

A MARKOVIAN APPROXIMATED SOLUTION TO A PORTFOLIO MANAGEMENT PROBLEM *

Jacek B. Krawczyk

*Victoria University of Wellington, PO Box 600, Wellington
New Zealand; fax +64-4-4955014
Email: Jacek.Krawczyk@vuw.ac.nz;
<http://www.vuw.ac.nz/~jacek>*

Version $\beta.1b$

Abstract: A portfolio management problem is approximated as a Markov decision chain. The weak Euler scheme is applied to discretise the time evolution of a portfolio and an inverse distance method is used to describe the transition probabilities. The approximating Markov decision problem is solved by value iteration. Numerical solutions of varying degrees of accuracy to a few specific financial portfolio problems are obtained. A sample of a fund manager's objective functions is analysed to tell which of them generates an acceptable Value-at-Risk.

Keywords: Computational economics; portfolio management; approximating Markov decision chains; weak Euler scheme.

JEL: C8, C63, D92, C87, G11

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

1. INTRODUCTION

The purpose of this paper is to describe a numerical method capable of solving a class of stochastic optimal control problems, which includes portfolio management. The paper draws the idea of solving a continuous finite-horizon stochastic optimal control problem as a Markov decision chain from [7] and [6]. In this paper, the weakly consistent Euler discretisation scheme, used for the approximation

of the stochastic process, and a scaling in the policy space have greatly improved the solutions, relative to those reported previously.

Optimal portfolio management can be modelled as a stochastic optimal control problem. One can usually solve a problem of this class by solving the Hamilton-Jacobi-Bellman (*HJB*) equation. This is a complex procedure in any case. Often this equation analytically insoluble and a numerical

* Research supported by VUW GSBGM and IGC. Helpful comments by my colleagues: Graeme Guthrie, Martin Lally and Leigh Roberts are gratefully acknowledged; all remaining errors are mine.

method has to be applied.¹ This involves a discretisation scheme.

The Kushner [9] approach is an efficient discretisation scheme, in which the state-space and time steps remain related. An implementation of this approach to infinite horizon decision problems has been successful even in the case of stochastic differential games [3].

In [8], [7] and [6] a simple approach was introduced that produced numerical solutions to a few *finite*-horizon stochastic optimal control problems. Instead of looking for a solution to the *HJB* equation, as in the Kushner approach, a Markov decision chain, discrete in time and space, was solved. This is a more elementary exercise: instead of looking for a numerical solution to a second-order partial differential equation (*HJB*), a first order difference equation (Bellman's) needs to be solved.

In this paper, the original continuous optimal control problem is discretised to produce a Markov decision chain. A method of approximating the continuous noise by a discretely valued noise is applied. Value iteration is used to solve the Bellman equation for the Markov decision chain thus obtained. This "simple" Markovian approximation method was directly applied in [8] and [7] to estimate the discounted profit of stochastic resource utilisation. Encouraging results were reported; in particular, a good level of agreement of numerical solutions with the existing solutions (see [11]) was achieved and a sensible degree of computational complexity of the method was observed. In [6], the same method was used for the solution to the classical portfolio management problem (see [2]). While the utility measures of the approximating and original problems were similar, there were some discrepancies in the policy shapes. These are overcome in this paper through a scaling in the policy space and a weakly consistent discretisation scheme of the Ito diffusion process.

The emphasis of the paper is on the solution method. However, a few financial engineering problems, difficult to solve analytically, will be solved numerically in this paper. In particular, rules will be computed for non stationary² and

¹ Computational methods have been used for financial optimisation for quite some time, see for example [11] and, for a review, [15].

² Analytical optimal portfolio rules are known for the HARA (Hyperbolic Absolute Risk-Aversion) utility func-

tioned³ portfolio problems. A bond pricing problem will be solved through a repetitive solution of the Markov decision chain.

The rest of this paper is organised as follows. In Section 2 the Markovian approximation method from [8] is modified through use of weakly consistent (Euler) 2-value noise discretisation, rather than an intuitively motivated 3-value noise discretisation scheme used in [8], [7] and [6]. The method is applied, in Sections 4-5, to a classical optimal portfolio selection problem from [2]. The portfolio problem is defined and solved analytically in Section 3. Numerical solutions of varying degree of computational effort are calculated in Section 4 and compared to the analytical solution. In Section 5, a few specific problems of financial engineering are solved. Concluding remarks close the paper.

2. A SIMPLE MARKOVIAN APPROXIMATION

2.1 *Optimal Stochastic Control*

Consider the stochastic system to be controlled

$$d\mathbf{X}(t) = f(\mathbf{X}(t), \mathbf{u}(t), t)dt + b(\mathbf{X}(t), \mathbf{u}(t), t)d\mathbf{W}(t) \quad (1)$$

where

$$\mathbf{X} = \{\mathbf{X}(t) \in X \subset \mathbb{R}^n, t \geq 0, \mathbf{X}(0) = \mathbf{x}_0 - \text{given}\}$$

is the state process, $\mathbf{u}(t) \in U \subset \mathbb{R}^m$ is the control, $\mathbf{W}(t)$ is a Wiener process, $f(\mathbf{X}(t), \mathbf{u}(t), t)$ is a drift, and $b(\mathbf{X}(t), \mathbf{u}(t), t)d\mathbf{W}(t)$ is diffusion. For the formal treatment of the optimally controlled diffusion process refer to [2]. The *optimal* control rule μ that determines the control \mathbf{u} is Markovian

$$\mathbf{u}(t) = \mu(t, \mathbf{X}(t)) \quad (2)$$

and chosen so as to maximise a functional J

$$\max_{\mathbf{u}} J(0, \mathbf{x}_0; \mathbf{u}) \quad (3)$$

where

tions include isoelastic, exponential and quadratic utility functions. family of utility functions, see [10]. However, the explicit solutions to some "practical" problems that would allow for time dependent model parameters are usually beyond the simple quadratures.

³ This is another class of analytically intractable yet sensible portfolio problems. In principle, constrained policies could be obtained through the Kuhn-Tucker conditions. In practice, their closed forms are unobtainable.

$$\mathbb{E} \left(\int_{\tau}^T g(\mathbf{X}(t), \mathbf{u}(t), t) dt + s(\mathbf{X}(T)) \mid \mathbf{X}(\tau) = x \right) \quad (4)$$

is the profit-to-go function. For Markovian feedback controls the maximum value of (4)

$$H(\tau, x) = \max_{\boldsymbol{\mu}(\cdot)} J(\tau, x; \boldsymbol{\mu}(\cdot))$$

satisfies the *HJB* equation

$$\max \{g(x, \mathbf{u}, t) + \mathcal{L}_u H(t, x)\} = 0 \quad (5)$$

with the boundary condition

$$H(T, \mathbf{X}(T)) = s(\mathbf{X}(T)) \quad (6)$$

where \mathcal{L}_u is the operator

$$\begin{aligned} \mathcal{L}_u &= \frac{\partial}{\partial \tau} + \sum_{i=1}^n f_i(\mathbf{x}, \mathbf{u}, \tau) \frac{\partial}{\partial x_i} \\ &+ \frac{1}{2} \sum_{i,j=1}^n B_{i,j}(\mathbf{x}, \mathbf{u}, \tau) \frac{\partial^2}{\partial x_i \partial x_j} \end{aligned} \quad (7)$$

and where f_i is the i -th component function of f and $B_{i,j}$ is the i, j -th entry of the covariance matrix $B = bb^T$.

2.2 A Corresponding Markov Decision Chain

A Markov decision chain corresponding to the optimisation problem $\max(4)$ subject to (1) and (2) is obtained through the following three steps. First, the state equation (1) is discretised in time using the Euler-Maruyama approximation (see [5]). Then, the state space is restricted to a finite dimensional discrete state grid and, finally, the transition probabilities and rewards for these discrete states are specified.

Euler-Maruyama Approximation. An Euler-Maruyama approximation of process⁴ $\mathbf{X} \subset \mathbb{R}^1$ that satisfies equation (1) is a stochastic process

$$\mathbf{Y} = \{\mathbf{Y}_\ell \in X, \quad 0 \leq \ell \leq N\}$$

satisfying the equation (called the iterative scheme)

$$\begin{aligned} \mathbf{Y}_{\ell+1} &= \mathbf{Y}_\ell + f(\mathbf{Y}_\ell, \mathbf{u}_\ell, \tau_\ell)(\tau_{\ell+1} - \tau_\ell) + \\ &b(\mathbf{Y}_\ell, \mathbf{u}_\ell, \tau_\ell) (\mathbf{W}(\tau_{\ell+1}) - \mathbf{W}(\tau_\ell)) \end{aligned} \quad (8)$$

⁴ The approximation scheme is introduced for a one-dimensional process. The extension of the scheme to \mathbb{R}^n is obvious.

$\tau = \{\tau_\ell\}_{\ell=0}^N$, with $\tau_0 = 0$ and $\tau_N = T$ is a strictly increasing sequence of real numbers that partition the time interval $[0, T]$.

The indices run $\ell = 0, 1, 2, \dots, N-1$, the initial and subsequent values are, respectively

$$\mathbf{Y}_0 = \mathbf{X}(0) = x_0, \quad \mathbf{Y}_\ell = \mathbf{Y}(\tau_\ell). \quad (9)$$

For a time discretisation using a constant time step (where N is a positive integral number)

$$\tau_\ell = \ell \delta \quad \text{where} \quad \delta = \tau_{\ell+1} - \tau_\ell = \frac{T}{N}. \quad (10)$$

Notation. The discretisation scheme, while intuitively simple, overlays several layers of discretisation: of time, of state, and of noise. We adopt the following conventions.

- (1) Continuous-time variables: $x(t)$ (standard); variables in discrete time: x_ℓ .
- (2) Points of the discrete state space (“grid”) $\bar{x} \in \bar{X}_\ell$.
- (3) Stochastic processes: \mathbf{x} (bold).

Discrete State Space. Equidistant grids will be used for simplicity (see [8]). The discrete state space for stage ℓ is denoted by $\bar{X}_\ell \subset \mathbb{R}^1$. Let the upper and lower bounds of the state grid be

$$\bar{U}_\ell = \max \bar{X}_\ell \quad \text{and} \quad \bar{L}_\ell = \min \bar{X}_\ell.$$

respectively. A point $x \in X$ is defined to be within the grid \bar{X}_ℓ if $\bar{L}_\ell \leq x \leq \bar{U}_\ell$. The collection of the discrete state spaces for all the stages, $\{\bar{X}_\ell\}_{\ell=0}^N$, is denoted \bar{X} and called the discrete state space.

Adjacency. Heuristically, the scheme approximates a point of X at stage ℓ by the points of \bar{X}_ℓ which are “adjacent” to it.

- (1) Two states of \bar{X}_ℓ are **adjacent** if no other state of \bar{X}_ℓ lies between them⁵.
- (2) Given a point of the continuous state space, $x \in X$, a pair of states, $\bar{x}^\ominus \in \bar{X}_\ell$ and $\bar{x}^\oplus \in \bar{X}_\ell$, is **adjacent to** x if the states are adjacent and $\bar{x}^\ominus < x < \bar{x}^\oplus$.
- (3) Given $x \in X$ with $x \geq \bar{U}_\ell$ define \bar{U}_ℓ to be **adjacent to** x .
- (4) Given $x \in X$ with $x \leq \bar{L}_\ell$ define \bar{L}_ℓ to be **adjacent to** x .
- (5) Given $x \in X$ with $x \in \bar{X}_\ell$ define x to be **adjacent to** itself.

⁵ In \mathbb{R}^n two states are **adjacent** if their projections onto each of the n coordinate axes are adjacent in the sense just defined.

Transition Probabilities. Consider the stochastic process $\mathbf{Y} = \{\mathbf{Y}_\ell, \ell = 0, 1, 2, \dots, N\}$ where \mathbf{Y}_ℓ is defined through equation (8). For a given control sequence \mathbf{u}_ℓ and equidistant discretisation times, the iterative scheme (8) can be abbreviated to

$$\mathbf{Y}_{\ell+1} = \mathbf{Y}_\ell + \delta f_\ell + b_\ell \Delta \mathbf{W}_\ell \quad (11)$$

where f_ℓ and b_ℓ denote, respectively,

$$f_\ell = f(\mathbf{Y}_\ell, \mathbf{u}_\ell, \tau_\ell) \quad b_\ell = b(\mathbf{Y}_\ell, \mathbf{u}_\ell, \tau_\ell).$$

The increments

$\Delta \mathbf{W}_\ell = \mathbf{W}(\tau_{\ell+1}) - \mathbf{W}(\tau_\ell)$, for $\ell = 0, 1, 2, \dots, N$ refer to the Wiener process $\mathbf{W} = \{\mathbf{W}(t), t \geq 0\}$ and are known (cf [5]) to be independent Gaussian random variables with mean and variance:

$$\mathbb{E}(\Delta \mathbf{W}_\ell) = 0 \quad \mathbb{E}((\Delta \mathbf{W}_\ell)^2) = \delta.$$

The iterative scheme (11) thus defined is the simplest *strong* Taylor approximation of an Ito diffusion process (1), see [5]. Now, suppose that at some time τ_ℓ , $\mathbf{Y}_\ell = \bar{\mathbf{Y}}_\ell \in \bar{X}_\ell$.

Deterministic process. Assume, for the time being, that there is no noise in the process (11) so, for a given control value \mathbf{u}_ℓ , the process moves to $\mathbf{Y}_{\ell+1}$ which is defined by:

$$\mathbf{Y}_{\ell+1} = \bar{\mathbf{Y}}_\ell + \delta f_\ell. \quad (12)$$

If there is a pair of states of $\bar{X}_{\ell+1}$ adjacent to $\mathbf{Y}_{\ell+1}$ then the transition probabilities are assigned using an inverse distance method. Let $\bar{\mathbf{Y}}_{\ell+1}^\ominus < \bar{\mathbf{Y}}_{\ell+1}^\oplus$ be the pair of states adjacent to $\mathbf{Y}_{\ell+1}$. Define

$$h_\ell = \bar{\mathbf{Y}}_{\ell+1}^\oplus - \bar{\mathbf{Y}}_{\ell+1}^\ominus$$

and assign the following non-zero transition probabilities

$$p(\bar{\mathbf{Y}}_\ell, \bar{\mathbf{Y}}_{\ell+1}^\oplus | \mathbf{u}_\ell) = \frac{\mathbf{Y}_{\ell+1} - \bar{\mathbf{Y}}_{\ell+1}^\ominus}{h_\ell} \quad (13)$$

$$p(\bar{\mathbf{Y}}_\ell, \bar{\mathbf{Y}}_{\ell+1}^\ominus | \mathbf{u}_\ell) = \frac{\bar{\mathbf{Y}}_{\ell+1}^\oplus - \mathbf{Y}_{\ell+1}}{h_\ell}. \quad (14)$$

A Weak Taylor Approximation. If the Gaussian noise is present in (11) a value of $\mathbf{Y}_{\ell+1}$ is not deterministic. For this situation, the strong Euler scheme (11) will be replaced by a *weak* Euler scheme (see [5])

$$\mathbf{Y}_{\ell+1} = \mathbf{Y}_\ell + \delta f_\ell + b_\ell \Delta \tilde{\mathbf{W}}_\ell. \quad (15)$$

The difference is in $\Delta \tilde{\mathbf{W}}_\ell$, which is a ‘‘convenient’’ approximation of the random increments $\Delta \mathbf{W}_\ell$

of the Wiener process that has similar moment properties to those of $\Delta \mathbf{W}_\ell$. In the portfolio model, we will use an easily generated two point random variable taking values $\pm\sqrt{\delta}$ i.e.,

$$P(\Delta \tilde{\mathbf{W}}_\ell = \pm\sqrt{\delta}) = \frac{1}{2}. \quad (16)$$

This approximation of the continuously distributed perturbation $\Delta \mathbf{W}_\ell$ by a two-value noise is of course arbitrary. However, it is sufficient for the approximating solutions’ convergence. One can obviously use other more realistic discrete representations of $\Delta \tilde{\mathbf{W}}_\ell$ e.g., it can be modelled as a three-point distributed random variable T_ℓ with

$$P(T_\ell = \pm\sqrt{3\delta}) = \frac{1}{6} \quad P(T_\ell = 0) = \frac{2}{3}. \quad (17)$$

No matter how simple or complex these approximations are, they should preserve the original distribution’s first and second moments and depend on the partition interval’s length. The latter feature guarantees that, for all such approximations, the smaller δ the less diffuse the states become, to which the process transits.

For the noise representation (16), the definition of the transition probabilities in the stochastic case is only slightly different from (13), (14). Let $\mathbf{Y}_{\ell+1}$ be determined through (12). The noise discretisation method means that for $\delta > 0$ the process reaches, at $\ell + 1$:

$$\mathbf{Y}_{\ell+1}^- = \mathbf{Y}_{\ell+1} - b_\ell \sqrt{\delta} \quad \text{with prob. } \frac{1}{2} \quad (18)$$

$$\mathbf{Y}_{\ell+1}^+ = \mathbf{Y}_{\ell+1} + b_\ell \sqrt{\delta} \quad \text{with prob. } \frac{1}{2}. \quad (19)$$

If there are two adjacent states to $\mathbf{Y}_{\ell+1}^-$ and $\mathbf{Y}_{\ell+1}^+$ then apply the inverse distance method as in (13), (14) but weight the two probabilities by $\frac{1}{2}$. Thus, for example, if $\mathbf{Y}_{\ell+1}^- \notin \bar{X}_{\ell+1}$ but there exist $\bar{\mathbf{Y}}_{\ell+1}^\ominus < \bar{\mathbf{Y}}_{\ell+1}^\oplus$ in $\bar{X}_{\ell+1}$ adjacent to $\mathbf{Y}_{\ell+1}^-$ then the transition probabilities are

$$p(\bar{\mathbf{Y}}_\ell, \bar{\mathbf{Y}}_{\ell+1}^\oplus | \mathbf{u}_\ell) = \frac{1}{2} \frac{\mathbf{Y}_{\ell+1}^- - \bar{\mathbf{Y}}_{\ell+1}^\ominus}{h_\ell} \quad (20)$$

$$p(\bar{\mathbf{Y}}_\ell, \bar{\mathbf{Y}}_{\ell+1}^\ominus | \mathbf{u}_\ell) = \frac{1}{2} \frac{\bar{\mathbf{Y}}_{\ell+1}^\oplus - \mathbf{Y}_{\ell+1}^-}{h_\ell} \quad (21)$$

where $h_\ell = \bar{\mathbf{Y}}_{\ell+1}^\oplus - \bar{\mathbf{Y}}_{\ell+1}^\ominus$. If any of the states $\mathbf{Y}^\ominus, \mathbf{Y}^\oplus$, etc. overlaps another, the respective probabilities have to be summed up.

As evident, the above discretisation method is very simple and intuitive. However, as noted, it

preserves the first two moments of the original distribution so that the overall discretisation scheme is *weakly consistent* in the sense of [9]⁶.

Constraints. It has to be borne in mind that the discretisation of a constraint is always sensitive to the choice of the discretisation steps δ and h , and has to be dealt with “carefully”. *E.g.*, a local constraint $v(t) > a$ for $t \in [t_1, t_2]$ can be meaningfully translated as $v_\ell > a$ only if both δ is small in comparison to $[t_1, t_2]$ and the state grid step $h < |a|$. See inequality (34) for an example of how a state variable constraint can be represented in discrete time.

Transition Rewards. Let the control strategy be Markovian (2) and action at state ℓ computed as

$$\mathbf{u}_\ell = \mu(\ell, \bar{\mathbf{Y}}_\ell), \bar{\mathbf{Y}}_\ell \in \bar{\mathbf{X}}_\ell, \ell = 0, 1, \dots, N-1. \quad (22)$$

Recalling (4), note that for the approximating problem, the decision maker receives a reward that depends on the state at stage ℓ and on the action \mathbf{u}_ℓ

$$\gamma(\bar{\mathbf{Y}}_\ell, \mathbf{u}_\ell, \ell) = \delta g(\bar{\mathbf{Y}}_\ell, \mathbf{u}_\ell, \ell), \quad (23)$$

$\ell = 0, 1, 2, \dots, N-1$. The overall reward for the Markov decision chain \mathbf{Y} , starting from $\bar{\mathbf{Y}}_0 = x_0 \in \bar{\mathbf{X}}_0$ and controlled by $\mathbf{u} = \{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{N-1}\}$ can be determined as

$$J(0, \bar{x}_0; \mathbf{u}) = \mathbb{E} \left(\sum_{\ell=0}^{N-1} \gamma(\bar{\mathbf{Y}}_\ell, \mathbf{u}_\ell, \ell) + s(\bar{\mathbf{Y}}_N) \mid \bar{\mathbf{Y}}_0 = x_0 \right) \quad (24)$$

Finally, the problem:

$$\begin{cases} \max_{\mathbf{u}} J(0, \bar{x}_0; \mathbf{u}) \\ \text{subject to} \\ \mathbf{Y}_{\ell+1} = \mathbf{Y}_\ell + \delta f_\ell + b_\ell \mathbf{w}_\ell, \end{cases} \quad (25)$$

with the transition probabilities defined as above is the *Markov decision chain* approximating the

⁶ See [9], page 1002. Conditions “1” and “2” (about continuity of the Markov chain expected value and variance) are easy to prove. However, because in this discretisation method the state step is independent of the time step, condition “3” (about continuity of the Markov chain increments) can only be satisfied if the grid $\bar{\mathbf{X}}_\ell$ is allowed to become denser. Moreover, consistency fails along the boundary of the discrete state space so the scheme is *locally* weakly consistent. This is not surprising since it would be impossible for a system constrained to lie within a finite space to follow the behaviour of a system which is not similarly constrained at the points where the constraints become active. However, this feature is common to all approximation schemes of this kind.

original continuous-time optimisation problem of Section 2.1. For convenience, we use the notation $J(\delta, \bar{\mathbf{X}}) \equiv J(0, x_0; \hat{\mathbf{u}})$ where $\hat{\mathbf{u}}$ is a maximiser in the above optimisation.

2.3 Computational Complexity

There are two crucial parameters for the solution method outlined above: the number of states and the number of time steps. One expects that increasing these numbers would improve the solution’s accuracy. However, the computation time also increases. Recent papers [13] and [14] report mitigating the curse of dimensionality for a certain subclass of Markov decision chains through use of randomisation. The Markovian approximation defined in this paper leads to a similar conclusion. Following [7] a result of this nature will be proved.

Claim 1. The computation time needed for the solution of the Markov decision chain (25) increases approximately linearly in both the number of states and the number of time steps.

Proof. Suppose a solution is computed by backward induction for a state in stage ℓ and that the solution from stage $\ell+1$ onwards has already been determined. The time required to compute the optimal decision for the current state is *largely* independent of both the number of time steps and the number of states. Its independence of the number of states is a consequence of the approximation scheme scanning only the adjacent states in the next stage. Doubling the number of states means that twice the time is taken for each stage and the computation time doubles. Doubling the number of time steps leaves the computation time for each stage fixed but doubles the number of stages and hence the computation time doubles. This exact linear relationship reported in [7], see Figure 1, is spoiled by the vagaries of the computation time of the numerical maximisation, as well as the load dependence of the computer performance (see Figure 1 right panel). \diamond

An array of test problems was solved in [7] (and [8]) using a similar Markovian approximation method⁷. The approximating solutions closely

⁷ The noise approximation method used in [7], [8] was different from the weak Taylor approximation used here. So, only deterministic problems’ solutions there reported are directly relevant for this paper.

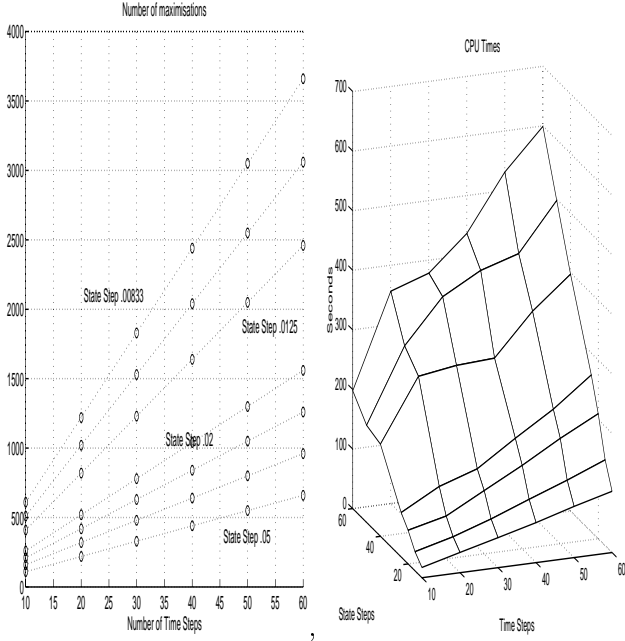


Fig. 1. Computational complexity.

followed the optimal ones, see [8]. However, the test problems were “easy” in that they contained a quadratic cost component (they were not linear-quadratic though). Here, the method will be used to solve a classical portfolio selection problem (see [2]) which is a generically stochastic and non linear-quadratic problem.

3. A PORTFOLIO SELECTION MODEL

3.1 The Model

A simplified version of Merton’s [10] model of optimal portfolio selection is analytically solved in [2], pp. 160-161.

The stock portfolio consists of two assets, one “risky” and the other “risk free”. If the price per share of the risky asset $p(t)$ changes according to

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

while the price q per share for the risk free asset changes according to

$$dq = qr dt$$

then the wealth $x(t)$ at time $t \in [0, T]$ changes according to the following stochastic differential equation

$$dx = (1 - u_1)rx dt + u_1x(\alpha dt + \sigma dw) - U_2 dt. \quad (26)$$

Here, w is a one-dimensional standard Brownian motion, and r, α, σ are constants with $r < \alpha$ and $\sigma > 0$. Symbol $u_1(t)$ denotes the fraction of the wealth invested in the risky asset at t and $U_2(t)$ is the consumption rate. The agent’s objective is to find an optimal two-dimensional strategy $u = [u_1(x), U_2(x)]$, such that

$$0 \leq u_1(t) \leq 1, \quad \text{and} \quad U_2(t) \geq 0, \quad (27)$$

and such that maximises expected discounted total utility

$$J(0, x(0); u) = \mathbb{E} \left(\int_0^T e^{-\rho t} [U_2(t)]^\gamma dt \mid x(0) = x_0 \right) \quad (28)$$

given discount rate $\rho > 0$ and assuming that $[U_2(t)]^\gamma$ is the agent’s utility function. Here no value is assigned to wealth at T while x_0 is the wealth at the initial time 0. The problem to maximise (28) subject to (26) is clearly one of the class described in Section 2.1.

3.2 The Optimal Solution

The Hamilton-Jacobi-Bellman equation can be solved for the optimal value function in the following form

$$H(\tau, x) = g(\tau)x^\gamma. \quad (29)$$

Function $g(\tau)$ can be integrated and equals

$$g(\tau) = e^{-\rho\tau} \left[\frac{1 - \gamma}{\rho - \nu\gamma} \left(1 - e^{-\frac{(\rho - \nu\gamma)}{1 - \gamma}(T - \tau)} \right) \right]^{1 - \gamma} \quad (30)$$

where

$$\nu = \frac{(\alpha - r)^2}{2\sigma^2(1 - \gamma)} + r.$$

The optimal investment and consumption strategies \hat{u}_1 and \hat{U}_2 can be computed as

$$\hat{u}_1 = \frac{\alpha - r}{\sigma^2(1 - \gamma)} \quad (31)$$

$$\hat{U}_2(\tau, x) = [e^{\rho\tau} g(\tau)]^{\frac{1}{\gamma - 1}} x. \quad (32)$$

In this example, only \hat{U}_2 is a (linear) function of wealth while \hat{u}_1 is constant. Notice also that the above solution is “internal” in that both constraints (27) will be satisfied for some parameter set. In particular $\hat{u}_1 \leq 1$ if $\alpha - r \leq \sigma^2(1 - \gamma)$.

4. A CALIBRATED MODEL

4.1 The reference solution

Suppose an agent with an original wealth of $x_0 = \$100,000$ wants to maximise their satisfaction during the coming $T = 10$ years. The instantaneous satisfaction is measured by $\sqrt{U_2(t)}$. The risk free asset price drift is $r = .05$ and that of the risky asset is $\alpha = .11$ with the volatility $\sigma = .4$. The agent's discount rate is $\rho = .11$.

For these parameter values, the agent's expected discounted total utility is $\hat{J} = g(0)\sqrt{100000} = 723.09$. Figure 2 presents the optimal strategies; the expected wealth and consumption rate time profiles are shown in Figure 3.




Fig. 2. Optimal (reference) strategies.

Figure 4 shows ten wealth and strategy time profiles, which correspond to ten noise realisations $dw(t)$, $t \in [0, T]$, obtained from a random number generator. The average total discounted utility of these 10 portfolios is 751, which is more than the theoretical $\hat{J} = 729$. However, it is evident that the optimal portfolios' performance has a large variance⁸.

⁸ An estimate of the optimal utility standard deviation was computed for the same integration parameter set as in Section 4.2. (I.e., there were 600 noise realisations and the integration step was .025.) The mean utility was 719.7, which was 99.5 % of the theoretical expected optimal performance; the corresponding standard deviation was 165.

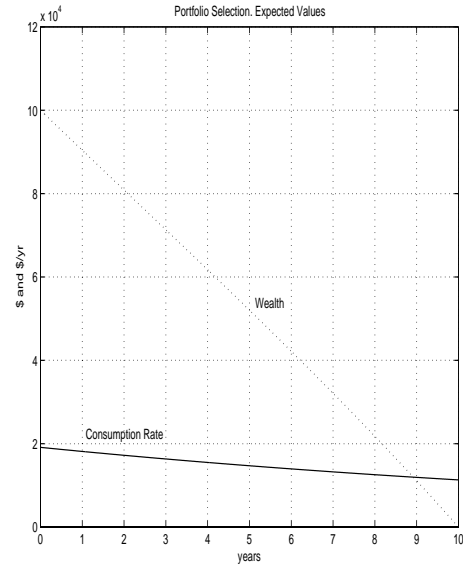


Fig. 3. Optimal expected wealth and consumption rate time profiles.




Fig. 4. Optimal wealth, investment and consumption rate realisations.

4.2 Numerical Solutions

The SOCSol [18] suite of Matlab functions was used to optimise a portfolio from Section 3. Most of the problem transformation (e.g., from a continuous time and space formulation to a discrete model) is taken care of by the software.

Constraints. The constraints have to be dealt with “manually”. The local instantaneous constraints on controls $u_1(t)$ and $U_2(t)$ (27) can be immediately expressed in discrete time as

$$0 \leq u_{1,\ell} \leq 1, \quad \text{and} \quad U_{2,\ell} \geq 0. \quad (33)$$

The portfolio admissibility condition, which in the continuous version amounts to $x(t) \geq 0$, $\forall t \in [0, T]$, [4], cannot however be directly replaced by $x_\ell \geq 0$. This is (mainly) because x_ℓ is expressible in terms of $u_{1,\ell-1}$ and $U_{2,\ell-1}$ yet we need a condition valid for time ℓ . It appears (from [10]) that the best discrete-time counterpart of $x(t) \geq 0$ is

$$x_\ell(1 + \delta(r + u_{1,\ell}(\alpha - r))) - \delta U_{2,\ell} \geq 0 \quad (34)$$

where δ is the time discretisation step (see Section 2.2). It has to be borne in mind that (34) is an *approximation* to the portfolio admissibility condition and that it depends on the time discretisation step.

Technical hints. A cautionary remark about numerical optimisation is in place. Most optimisation methods work (much) more efficiently if the solution vector components are of comparable magnitudes. This is not the case of the control variables u_1 and U_2 . Indeed, u_1 is bounded between 0 and 1 but, U_2 is practically unbounded from above, see Figure 2. This caused some (not insuperable) difficulties in [6] in obtaining accurate approximating solutions. In this paper, such difficulties were avoided through *re-scaling* of the model. It follows from (32) that $U_2(t)$ is linear in the state $x(t)$. All $U_2(t)$ were then replaced by $u_2(t)x(t)$ in the optimisation problem \max (28) subject to (26). Consequently, the numerical routines were looking for $u_2(t)$ that was not greater than 12 for most of the cases solved. Moreover, because of the above transformation the strategy graphs will no longer be linear as in Figure 2 low panel but horizontal (as in Figure 7).

Important software control parameters are the time discretisation step δ and the state space grid width h . To get an idea of their range values, necessary for an accurate approximation, a deterministic portfolio control problem ($T = 2$) was solved: first analytically, then, the discretised model solutions were computed. Figure 5 shows the results.

The plot coordinates are the time discretisation step δ and a utility measure. The horizon is $T = 2$; the remaining model parameters are as in Section 4.1.

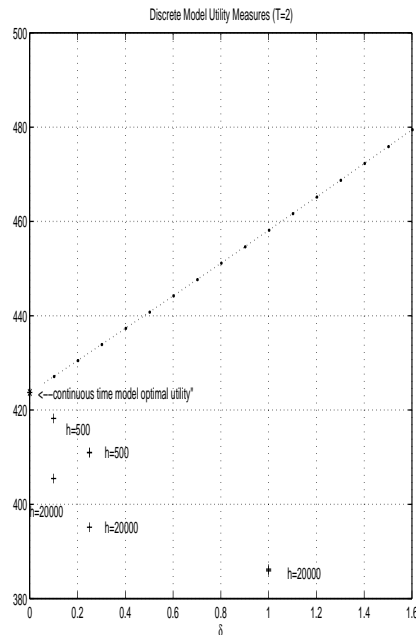


Fig. 5. Discretised model utility realisations.

The point denoted “*” (0,422) is the continuous model optimal utility. The *discrete* time model utility values converge toward this point as $\delta \rightarrow 0$. Notice that they are greater than the continuous model utility. This is because (in the rectangular method) the integration error grows in δ . The points denoted “+” correspond to utility realisations of a model discretised both in time and space (Markov chain). It is clear from the figure that reasonable utility approximations can only be obtained for $\delta \leq .1$ and $h \leq 500$.

The impact of the length of the time step δ on the solution accuracy in a stochastic model is shown in Figure 6.

Consider time ℓ ($\ell = 0, 1 \dots N - 1$) and $u_{1,\ell}$ to be applied at this time. Assuming that the choice of $U_{2,\ell}$ is made optimal

$$\begin{aligned} u_{1,\ell} &= \arg \max \left(\delta \sqrt{U_{2,\ell}} + e^{-\rho \delta} g(\ell + \delta) \mathbb{E} \sqrt{x_{\ell+\delta}} \right) \\ &= \arg \max \left(\mathbb{E} \sqrt{x_{\ell+\delta}} \right), \end{aligned} \quad (35)$$

see (29). The expected value in (35) was computed using a Taylor series (second order) expansion and presented as a function of strategy u_1 in Figure 6. The strategy domain was “extended” beyond the feasible range $[0 \ 1]$ to show the utility measure shapes. Notice that, for the feasible $u_1 \in [0, \ 1]$, the utility measures would all look flat.




Fig. 6. Strategy convergence.

The vertical line $u_1 = .75$ shows where the utility maximum “should” be, for it is known from Figure 2 that the optimal strategy is $u_1 = .75$, independent of time. It is clear from the figure that the discretised model strategy $u_{1,\ell}(\delta)$ converges to the optimal strategy as $\delta \rightarrow 0$. Again, any reasonable approximation requires $\delta \leq .1$.

Another point to remember is that the time step δ should be *large* enough for the process to move beyond the adjacent state *i.e.*, $\delta|f_i| > h$. On the other hand, as shown above, δ should be small for high accuracy of the approximating solutions.

The convergence. There are various ways in which the goodness of a numerical solution can be evaluated. The most “objective” one would perhaps be to look at the average discounted total utility \hat{J} generated by the application of an approximated optimal numerical solution to the continuous model. However, as evident from Figure 4, the portfolio performance is very “volatile” and the standard deviation of utility distribution is large. Therefore, using \hat{J} thus computed is difficult to judge which solution is best.

We will first evaluate the convergence by comparing the approximating policy profiles (Figures 7 - 12), to the optimal ones (Figure 2).⁹ Then, we

⁹ Remember that because of the model re-scaling, if U_2 was linear (Figure 2) u_2 will be horizontal.

will generate a few realisation profiles to compare them to those of Figure 4 and, eventually, we will compute the corresponding utility distribution.




Fig. 7. Approximating strategies for $t = 0$ ($\delta = .2$).

Examine the policy rules shown in Figures 7 - 12. The bold dotted lines correspond to optimal strategies (compare Figure 2). One can see that the strategy convergence is more difficult to achieve for later times ($t = 9$) than at the beginning of the horizon ($t = 0$).




Fig. 8. Approximating strategies for $t = 0$ ($\delta = .1$).




Fig. 9. Approximating strategies for $t = 9$ ($\delta = .2$).




Fig. 11. Approximating strategies for $t = 9$ ($\delta = .05$).




Fig. 10. Approximating strategies for $t = 9$ ($\delta = .1$).




Fig. 12. Approximating strategies for $t = 9$ ($\delta = .02$ “small”).

Indeed, the gap between the optimal strategy and the approximating strategies for $t = 0$ is narrow for $\delta = .2$ and closes for $\delta = .1$ (for reasonably small h) whereas, for $t = 9$, it narrows down only for smaller δ s, see Figures 11 and 12. This is to be expected because the optimal $\hat{U}_2(T) = \infty$, and $\hat{u}_2(T) = \infty$ (see (32) and (30)), which is impossible to reproduce numerically.

Figure 13 shows the wealth and strategy realisations for $\delta = .05$ and $h = 100$. They look

very similar to the optimal ones in Figure 4. The simulation of 2000 noise realisations and the application of the approximating policy rules computed for the same parameters (*i.e.*, $\delta = .05$ and $h = 100$) resulted in the utility distribution (integrated with the time simulation step equal to .025) shown in Figure 14.

The mean discounted utility is $\hat{J} = 715.4$ (98.9 % optimal) and the corresponding standard deviation is 161. However, the portfolio performance




Fig. 13. Approximated wealth, investment and consumption rate realisations.

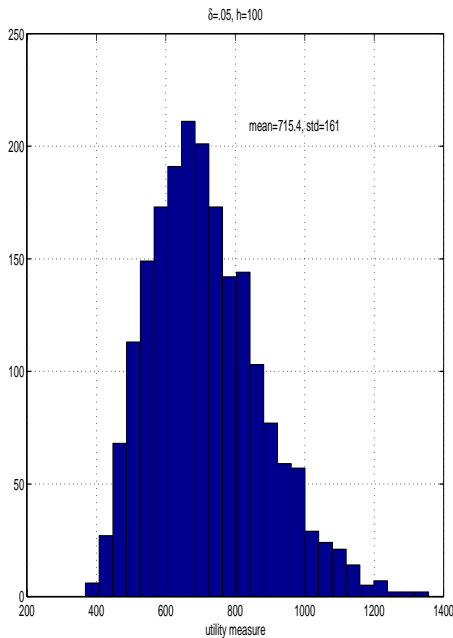


Fig. 14. Utility realisation distribution.

as judged by index \hat{J} for other approximating rules (e.g., $\delta = .05, .02$ and $h = 500, 100$) was comparable ($\hat{J} \in [701, 720]$ with standard devia-

tions $\in [159, 162]$).¹⁰ Averaging the utility over more realisations and diminishing the simulation step could help to improve the utility variance estimate. However, the improvement would not be substantial, as the portfolio performance is highly “volatile” whether optimal (Figure 4) or not (Figure 13).

5. PORTFOLIO MODEL MODIFICATIONS

5.1 Time Varying Parameters

Suppose now that the agent expects the volatility coefficient σ to vary as follows

$$\sigma(t) = \underline{\sigma} \left(1 - 0.09 \cos \left(\frac{2t}{\pi} \right) \right) \quad (36)$$

where $\underline{\sigma} = .4$. This means that the volatility σ used in Sections 3 and 4 was an average value. Now, the volatility will rise from 36.4 % to 43.6 % in the middle of T , and then drop. The agent would like to know whether this information should change their investment strategy or not.

Figure 15 reveals the modified strategy obtained as a solution to the discretised portfolio problem with $\delta = .05$ and $h = 500$. As expected¹¹, the agent will invest more (i.e., above .75) when the volatility is low. The solution is “exact” in that we can read how much one has to invest at each time. Interestingly, the optimal consumption strategy remains unchanged.

The wealth and strategy sample paths for the cosine volatility problem are shown in Figure 16. The discounted total utility is 664 (std=158).

In summary, the agent should modify their investment strategy once the information about a volatility scenario becomes available. In a similar way, a portfolio problem with a time dependent interest rate (or other parameters) can be solved.

¹⁰ However, using $\delta = .5, h = 10000$ that is clearly a “bad” policy, resulted in $\hat{J} = 124.13$ with the standard deviation equal 100.

¹¹ See footnote 19.




Fig. 15. Cosine modified strategies.




Fig. 16. Cosine modified realisation paths.

5.2 Constrained Policies

A portfolio manager might have an *a priori* belief that their investment should not exceed a certain wealth percentage. Such a constraint can easily be allowed for in the above numerical optimisation procedure.

Suppose that the permissible investment level is $\bar{u}_1 = .5$. Figure 17 reveals the modified strategy. Not surprisingly, the agent is expected to invest

at the constraint. However, this does not affect the consumption strategy. The new discounted total utility is 717 and the standard deviation equals 108, which is substantially less than under the unconstrained regime. The wealth and strategy sample paths are shown in Figure 18. The conservative investment strategy results in much less “volatility” in consumption and wealth as observed by eye and measured by the utility standard deviation.




Fig. 17. Constrained strategies.




Fig. 18. Constrained strategy realisation paths.

An impact of the investment constraint on the expected performance can be observed in more detail in Figure 19. Two histograms of the discounted total utility realisations are presented. The dark shadowed one corresponds to the unconstrained policies (compare Figure 14). The light histogram represents the constrained policy performance. Evidently, the constrained policy guarantees more “secure” performance (standard deviation = 108 *vis-à-vis* 166 of the unconstrained policy). However, the unconstrained policy brings a (marginally) higher utility value. More computations of that kind would generate the “efficient boundary”.

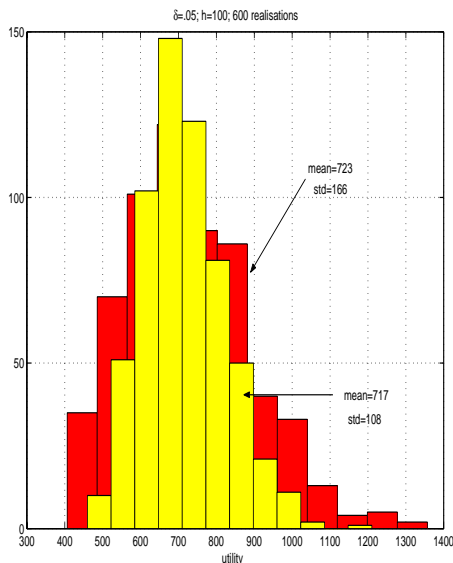


Fig. 19. Utility realisation distributions.

In a similar way, another portfolio problem, in which some minimal (or maximal) consumption rate is given could be solved.

5.3 Pension funds

A practical problem of financial engineering is one in which an agent pays an amount x_0 to a pension fund, to be repaid by a quantity \bar{x}_T at time T . The latter is a result of an investment policy $u_1(x)$ adopted by the fund’s manager.

The manager’s policy depends on his or her objective function, which could be the maximisation of an expected value, the minimisation of risk to obtain a target amount, *etc.* Once the objective function is revealed, the manager’s policy can

be computed as a solution to a stochastic optimal control problem associated with the objective function. The problem solution will routinely comprise an optimal decision rule $u_1(x), U_2(x)$ and a Monte-Carlo simulated distribution of x_T . Knowing the former is crucial for the manager to control the portfolio. The latter is “practical” in that it tells the pension buyer what they can, or should, expect as \bar{x}_T .

Knowing the distribution of x_T also helps the manager. It gives them an idea of what probabilities, or risks, are associated with obtaining particular realisations of the objective function. For example, the distribution may suggest that, for every x_0 there is a probable terminal value \bar{x}_T , which the manager may choose to advertise as the pension target.

We will first solve a pension fund problem for the expected value criterion, as follows. In (4), set $g(\mathbf{X}(t), \mathbf{u}(t), t) = 0$, $s(T(t)) = x(T)$ ¹² and suppose that the management fee is $2\%x(t)$. This means that we need to solve an optimisation problem in $u_1(x)$ with $U_2(t) = .02x(t)$.

Using the Markovian approximation approach as in Section 4.2, with the same model parameters *i.e.*, $T = 10, r = .05$, *etc.*, generates a rather trivial optimal strategy: $u_{1,t} = 1, U_{2,t} = .02x_t$ for positive states and times. Applying the strategy to different initial outlays x_0 generates the following final fund yield spread and location measures¹³ at $T = 10$, see Figure 20.

The figure tell us, among other things, that an initial deposit of \$40,000 corresponds to the expected pension value of about \$100,000. However, the median fund’s yield for this objective function is significantly below the mean. This indicates that the fund distribution is skewed, which is evident from Figure 21 (upper panel). The histogram shows us too that the probability of earning less than the “secure” revenue

$40,000 \exp\{(r - \text{“management fee”})10\} = 53,994$ is more than .5.¹⁴ It is even fairly probable (with probability $>.4$) to earn less than the initial outlay $x_0 = 40,000$. Evidently, using a policy that maximises the expected yield is a very risky strategy of managing a portfolio.

¹² Here, the objective function is not HARA.

¹³ Averaged over 1000 realisations.

¹⁴ To prove this and the subsequent claims integrate the area under the histogram from zero to 53,994 and 40,000, respectively.




Fig. 20. Yield location and spread measures.

To quantify a risk level associated with this policy, two popular risk measures, Value-at-Risk (VaR) and Conditional Value-at-Risk¹⁵ (CVaR) can be approximately calculated from the histogram. The β -VaR of a portfolio is the lowest amount α such that, with probability β the loss will not exceed α ; here

$$.9\text{-VaR} \approx 30,000.$$

The β -CVaR is the conditional expectation of losses above the amount α ; here

$$.9\text{-CVaR} \approx 34,000.$$

These measures disqualify the policy of maximizing the expected yield as a fund manager's objective function. No manager would accept such a high risk in controlling a pension fund.

A different strategy has to be considered. Suppose that the manager will use a "constrained" strategy: $u_{1,t} = .5$, $U_{2,t} = .02x_t$ (i.e., non optimal with respect to the expected value criterion). We can see from Figure 21 (lower panel) that, for the same initial outlay $x_0 = \$40,000$, the yield distribution (represented by the light histogram) is more concentrated and less skewed than the unconstrained one (upper panel). The mean for this portfolio is $\$84,100$ (median= $58,710$) and the standard deviation diminishes to $\$45,563$ from $\$168,000$ for the unconstrained policy. The dark histogram represents the final fund's yield distribution for the constrained strategy applied to $x_0 = 60,000$. In this case, the mean value is

¹⁵ See [1] for a static portfolio analysis based on VaR and CVaR.




Fig. 21. Fund yield spread.

$\$109,013$ (median= $93,000$) and the standard deviation $\$63,322$. Overall, the risk of performing worse than investing in the secure asset alone is much less under a constrained strategy. The risk of scoring less than x_0 is, for $x_0 = 40,000$

$$.9\text{-VaR} \approx 13,000 \text{ and } .9\text{-CVaR} \approx 17,000;$$

whereas for $x_0 = 60,000$,

$$.9\text{-VaR} \approx 21,000 \text{ and } .9\text{-CVaR} \approx 30,000.$$

Suppose now that the fund manager would like to advertise their pension fund as paying an amount \bar{x}_T for an initial outlay x_0 .¹⁶ It is clear from the histograms in Figure 21 that an expected value maximisation policy (constrained or not) cannot be used for this purpose. Instead, we will examine the policy determined as a solution to the stochastic optimal control problem with the objective function given as

$$J(0, x(0); u) = \mathbb{E} \left(h(x_T) \mid x(0) = x_0 \right) \quad (37)$$

where

$$h(x_T) = \begin{cases} (x_T - \bar{x}_T)^{\frac{1}{2}} & \text{if } x_T \geq \bar{x}_T, \\ -(x_T - \bar{x}_T)^2 & \text{otherwise.} \end{cases} \quad (38)$$

This criterion reflects the manager's wish to dispose of sufficient funds to meet the target \bar{x}_T . At the same time, it does not prompt the manager to accumulate (much) more than needed.

¹⁶ In other words, the manager will sell a ten year "bond" \bar{x}_{10} for x_0 .

The optimal investment policy¹⁷ u_1 , called “cautious”, is shown in Figure 22. The target is $\bar{x}_{10} = 100,000$.




Fig. 22. A cautious policy.

The solution is intuitively explicable: if there is a shortage of funds, or time, the manager’s actions become risky. In other words, the policy lines get higher and steeper as x_0 falls. Also, for the part of the strategy graph where it is impossible to meet the target x_T by investing in the riskless asset only (*i.e.*, before each line hits 0.), each policy line representing a later time dominates the lines that correspond to earlier times.

The usefulness of the cautious strategy for pension fund management can be assessed from Figure 23.

The cautious strategy was used to manage two initial payments x_0 of \$60,000 and \$73,500. The corresponding yield spreads are represented by the two central histograms in Figure 23. It is easy to see that the risk indices VaR and CVaR will be very small for the first outlay and virtually zero for the second. For comparisons, the background “back-to-back” histogram represents the result of managing the second payment using the expected value maximisation strategy. Once again we can see that results worse than a “secure” investment outcome (\$99,215 in this case) are very probable.

¹⁷The rest of the problem parameters are as before *i.e.*, $T = 10$, $u_2 = .02$, $r = .05$, *etc.*

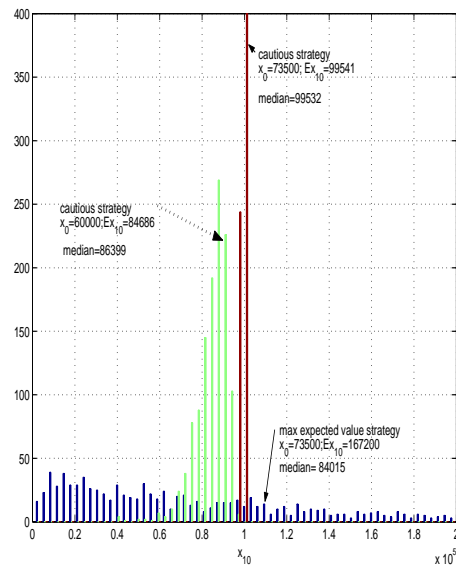


Fig. 23. A policy comparison.

Overall the cautious strategy appears an attractive portfolio management policy. However, to use it for pension advertising, or for pricing the x_{10} “bond”, the objective function should reward the manager for exceeding a target (here, \$100,000) more than the square root term¹⁸.

An optimal control problem with $g(\mathbf{X}(t), \mathbf{u}(t), t) \neq 0$ and $s(\mathbf{X}(T)) = e^{-\rho T} x(T)$ *i.e.*, one in which a combination of the final wealth and the utility from consumption is maximised, can also be solved using the above approximation method.

6. CONCLUSION

A discretisation method useful for Markovian approximations of finite-horizon continuous-time stochastic optimal-control problems has been described. An optimising algorithm has been developed (see [18]) and applied to solve a portfolio selection problem. For the calibrated models, an overall agreement between the analytical (where obtainable) and approximating solutions was noticed. However, a large variance of utility realisations was observed. It was also noticed that the variance diminished for constrained optimal solutions.

In the example with variable volatility, no difference was reported in consumption patterns between periods of low and high volatilities.

¹⁸For example, $(x_T - \bar{x}_T) \frac{\rho}{10}$ could be tried.

There were, however, differences in the investment schedules: as expected¹⁹, the volatility troughs triggered higher investment levels. A pension fund management example illustrated the use of the method for a non HARA objective function; it also highlighted the insufficiency of a mean value as an optimisation criterion. A cautious strategy was computed with very low VaR and CVaR.

All solutions were practical in that they could be applied to real life situations describable by the portfolio model. The method of obtaining them is ready to be applied to other scenarios of the model parameters (including a variable discount rate, *etc.*).

Some optimisation runs (on a Pentium II PC) lasted up to 30 hours (for $\delta = .02, h = 100$). However, an “intelligent” state space search (*à la* [12] or [16]) could be implemented to accelerate the algorithm convergence.

7. REFERENCES

- [1] ANDERSSON, F., S. URYASEV, 1999, “Credit Risk Optimization with Conditional Value-at-Risk Criterion”, *working paper available from authors*.
- [2] FLEMING, W. H., R. W. RISHEL, 1975, *Deterministic & Stochastic Optimal Control*, Springer-Verlag, New York *etc.*
- [3] HAURIE, A., J. B. KRAWCZYK, M. ROCHE, 1994, “Monitoring Cooperative Equilibria in a Stochastic Differential Game”, *JOTA*, Vol. 81, No 1 (April), pp. 73-95.
- [4] KARATZAS, I., S. E. SHREVE, 1992, *Methods of Mathematical Finance*, Springer-Verlag, New York *etc.*
- [5] KLOEDEN, P. E., E. PLATEN, 1992, *Numerical Solution of Stochastic Differential Equations*, Springer-Verlag, Berlin *etc.*
- [6] KRAWCZYK, J. B., 1999, “Approximated Numerical Solutions to a Portfolio Management Problem”. Paper prepared for, and presented at, *Stanford Institute for Theoretical Economics 1999 Summer Workshop on Computational Economics & Economic Theory*.
- [7] KRAWCZYK, J. B., 1998, “Finite-Horizon Continuous-Time Stochastic Optimisation via a Markovian Chain Approximation”, Preprints of the Society for Computational Economics 1998 Conference on *Computation in Economics, Finance and Engineering: Economic Systems*, (CEFES'98), Cambridge, England, paper No 33, 7 pages.
- [8] KRAWCZYK, J. B. & A. WINDSOR, 1997, *An Approximated Solution to Continuous-Time Stochastic Optimal Control Problem Through Markov Decision Chains*, Economic Working Papers Archive, [comp/971001](#)
- [9] KUSHNER, H.J., 1990, “Numerical Methods for Stochastic Control Problems in Continuous Time”, *SIAM J. Contr. & Optim.*, Vol. 28, No. 5, pp. 999-1048.
- [10] MERTON, R. C., 1971, “Optimum Consumption and Portfolio Rules in a Continuous-Time Model”, *J. of Economic Theory*, 3, pp. 373-413.
- [11] MORCK, R., E. SCHWARZ & D. STANGELAND, 1989, “The Valuation of Forestry Resources under Stochastic Prices and Inventories”, *J. of Financial and Quantitative Analysis*, Vol. 24, No. 4, pp. 473-487.
- [12] PEREIRA, M.V.F. & L.M.V.G PINTO, 1991, “Multi-stage Stochastic Optimization Applied to Energy Planning”, *Mathematical Programming*, Series B, 52, pp. 359-375.
- [13] RUST, J, 1997, “Using Randomization to Break the Curse of Dimensionality”, *Econometrica*, Vol. 65, No. 3 (May), pp. 487-516.
- [14] RUST, J, 1997, A Comparison of Policy Iteration Methods for Solving Continuous-State, Infinite-Horizon Markovian Decision Problems Using Random, Quasi-random, and Deterministic Discretizations, Economic Working Papers Archive, [comp/9704001](#).
- [15] TAPIERO, C., 1998, *Applied Stochastic Models and Control for Finance and Insurance*, Kluwer, Boston, *etc.*
- [16] TOLWINSKI, B. & R. UNDERWOOD, 1991, *An Algorithm to Estimate the Optimal Evolution of an Open Pit Mine*, Work. Paper MCS-91-05, Colorado School of Mines, Golden, Colorado.
- [17] SARKAR S., 2000, “On the investment-uncertainty relationship in a real options model”, *J. of Economic Dynamics and Control*, 24, pp. 219-225.
- [18] WINDSOR, A. & J. B. KRAWCZYK, 1997, *SOCsol-1: a Matlab Package for Approximating the Solution to a Continuous-Time Stochastic Optimal Control Problem*, Economic Working Papers Archive, [comp/9701002](#).

¹⁹There are situations where the conclusion could have been opposite, see [17].