

Optimal investment: bounds and heuristics

L. C. G. Rogers*
P. Zaczkowski†

Statistical Laboratory, University of Cambridge

June 6, 2014

Abstract

High-dimensional optimal investment/consumption problems are hard to deal with, not least because of the difficulty in characterizing the value function. This paper tries to offer ways to determine an approximately optimal *policy*, and to estimate its performance using duality methods. Though the value function is required as a concept in developing the theory, it plays no part in the computation, nor is it necessary to have global knowledge of the policy; it is enough to determine the policy along the realized sample path.

1 Introduction.

From the early work of Merton and his seminal papers [14] and [13], optimal investment tries to determine how to invest in financial markets with various constraints, objectives, and market imperfections; see [19] for a recent survey of the methods of the subject and a range of examples. The problems are all generic problems of controlling a diffusion¹, but with a special feature: the process to be controlled, the agent's wealth, is *scalar*. The first approach that springs to mind with such problems is to derive the Hamilton-Jacobi-Bellman (HJB) equation, and to try to solve it, and for simple enough problems this approach works well. But as the dimension of the state-space increases, this methodology becomes less and less effective. To see why, and to fix our ideas, we will consider throughout this paper the situation where there is a k -dimensional factor process X which is an autonomous diffusion, where k may be quite large. The n investable assets in the market have growth rates and covariances which are

*Corresponding author: Statistical Laboratory, Wilberforce Road, Cambridge CB3 0WB, United Kingdom.
lcgr1@cam.ac.uk

†Statistical Laboratory, Wilberforce Road, Cambridge CB3 0WB, United Kingdom.

¹We will work only in continuous time with controlled diffusions; more general processes could be considered, but with less explicit conclusions.

functions of X , so that the value function for the problem will depend on the current wealth, on the current value of X , and possibly also on time. This value function is not in general available in closed form, so if we want to obtain some approximation to it, we either have to characterize it by its values at some finite set of points and then interpolate; or we need to approximate it by some linear combination of basis functions. The first of these struggles to hold the values at enough points to give a reasonable understanding of what the value looks like, and the second similarly struggles to carry along enough members of the family of basis functions to make a reasonable approximation to the value function. Thus challenges of *storing* the approximation to the value function already look formidable, even before we think how it is to be calculated².

What else could we do? We could propose some recipe for what we think is a good control, and then evaluate its performance; and this is the general approach adopted in this paper. If we propose some recipe, we can always run many simulations of the evolution under this control, and this will give an estimate of the value of using this strategy, which in turn provides a lower bound for the true value of the optimal investment problem. But we would like to be able to provide an upper bound also, to assess how good the proposed control may be. To derive upper bounds, we use duality techniques, which also help us to propose sensible control policies.

There is already quite some literature on the use of duality methods in such problems. An early reference is Karatzas, Lehoczky, Shreve & Xu [12], who introduce the notion of a fictitious market completion to deal with optimization problems in incomplete markets. Following on from this, Cvitanic & Karatzas [7] use these techniques to handle constrained optimization problems. These studies are essentially theoretical in nature, but point the way to a possible technology. The paper of Haugh, Kogan & Wang [9] applies the ideas of [7] to provide bounds on the performance of candidate portfolio rules, and has similarities to what we do, but is different in at least one major respect, namely that it is assumed that the user is supplied with an explicit rule *and its associated value function* in closed form. As we explained above, such an assumption is a big requirement in a high-dimensional problem. We will also start from a candidate optimal policy, which does not need to be quite as explicit as in [9], but we will not require a global form for the value function. The reason is that the portfolio selection at a particular time and place needs only knowledge of the derivatives of the value function at that time and place, and this can be obtained from the *dual* value function and its derivatives. It turns out to be far more computationally viable to obtain these.

There is another strand of the duality literature where the notion of duality is understood in a different sense, and this starts with Rogers [18]. The essence of this approach is to permit non-adapted controls but to penalize the failure of adaptedness using a Lagrangian martingale. This approach also naturally leads to upper bounds on the value of the problem without needing a global specification of an approximate value function, but we see that this

²Calculating the value function will inevitably require calculation of expectations; and optimizing over controls. Neither task is particularly simple in high-dimensional problems.

approach is very different from those already described because it leads in principle to non-adapted choices, which none of the other methods do. There have been a few attempts to carry this approach forward, for example, Brown, Smith & Sun [3] and Brown & Smith [2]. The papers introduce a concept of Lagrangian relaxation in time, allowing the investor to have access to the future information at a certain price. The approach leads (like ours) to upper bounds on the value function of the problem, and offers great flexibility in the choice of Lagrangian penalty. Because of this flexibility, their approach does not lead directly to recommended trading strategies. Instead, the authors propose a couple of problem-specific heuristics and comment on the performance.

There are of course also numerous attempts to solve high-dimensional stochastic optimal control problems by approximating the value function in various ways. We do not attempt to provide a comprehensive literature review here, however we point an interested reader to Brandt, Goyal, Santa-Clara & Stroud [1] and Garlappi & Skoulakis [8]. In Brandt, Goyal, Santa-Clara & Stroud [1], the authors propose using Taylor expansion in the Bellman equation in order to make the optimisation with respect to the choice variable easier. A similar approach is used in Garlappi & Skoulakis [8], with additional features enhancing the applicability of Taylor series. The authors transform a value function in order to linearise it, and also decompose the choice variables from randomness in the following period. This allows them to compute conditional expectations in a much more efficient way. Both approaches are efficient when applied to the problems chosen by the authors. However, as remarked earlier, they are fundamentally value function methods, and will require clever parametrisations and grid choice in order to work in high dimensions.

This paper is structured as follows. In Section 2 we present the general problem and the methodology for solving it. Section 3 describes the algorithms used in the method. In Section 4 we give numerical evidence for the performance of the method, considering examples of the Merton problem; then non-constant relative risk aversion; a multi-dimensional incomplete market driven by a diffusion; and finally an example with correlated assets and a no-short-selling constraint. Section 5 concludes.

2 Continuous markets driven by a diffusion.

We shall present the methodology in the context of a finite-horizon optimal investment-consumption problem where the volatilities and drifts of the assets depend on some diffusion factor process. It will become evident that the general approach is not limited to such examples, but it is easier to explain in this more concrete setting. We shall also make various assumptions of boundedness on processes and global Lipschitz properties of coefficients which could be relaxed, but which simplify the exposition and proof: the aim is transparency, not maximality.

To begin with, suppose that X is an \mathbb{R}^k -valued diffusion process satisfying

$$dX_t = \sigma_X(X_t) dW_t + \mu_X(X_t) dt \equiv \sigma_X dW_t + \mu_X dt, \quad (2.1)$$

where W is a d -dimensional Brownian motion, and $\sigma_X : \mathbb{R}^k \rightarrow \mathbb{R}^k \otimes \mathbb{R}^d$ and $\mu_X : \mathbb{R}^k \rightarrow \mathbb{R}^k$ are globally Lipschitz coefficients.

We shall consider an investor who is allowed to invest in a market with a riskless asset yielding interest at rate $r_t \equiv r(X_t)$, and n stocks having volatility matrix $\sigma_t \equiv \sigma(X_t)$ and drift $\mu_t \equiv \mu(X_t)$. Here, $r : \mathbb{R}^k \rightarrow \mathbb{R}$, $\sigma : \mathbb{R}^k \rightarrow \mathbb{R}^n \otimes \mathbb{R}^d$, and $\mu : \mathbb{R}^k \rightarrow \mathbb{R}^n$ are bounded measurable functions. We assume non-degeneracy of the market, that is, $d \geq n$, and that the row rank of σ equals n . When $n = d$, the matrix σ is then invertible, and we have a special case of a complete market.

With these assumptions in place, we suppose that the investor's wealth w_t at time t evolves³ as

$$dw_t = w_t \{ r_t dt + \pi_t \cdot (\sigma_t dW_t + (\mu_t - r_t \mathbf{1}) dt) + g(\pi_t) dt \} - c_t dt, \quad (2.2)$$

where the n -vector process π_t represents the proportions of the overall wealth held in each of the stocks, c_t denotes the agent's consumption rate, and $g : \mathbb{R}^n \rightarrow [-\infty, \infty)$ is a concave penalty function. Most commonly, this function is taken to be identically zero, but the form (2.2) of the dynamics is the one used in Cuoco & Liu [6], and allows for interesting variants of the basic investment problem⁴, including constraints which require the portfolio proportions to remain in a convex set, and different rates for borrowing and lending. Thus in particular the theory we develop applies to the constrained portfolio problems considered by Haugh, Kogan & Wang [9]. As a piece of notation we define

$$\tilde{g}(z) \equiv \sup_{\pi} \{ g(\pi) - z \cdot \pi \}, \quad (2.3)$$

the convex dual function of g . In the special case where $g \equiv 0$, the convex dual \tilde{g} is equal to $+\infty$ everywhere except at 0, where it takes the value zero.

The agent's objective at time t is to achieve

$$\sup_{(c, \pi) \in \mathcal{A}} E \left[\int_t^T U(s, c_s) ds + \varphi(w_T) \middle| w_t = w, X_t = x \right] \equiv V(t, w, x), \quad (2.4)$$

where U and φ are strictly concave C^2 utility functions satisfying the Inada conditions⁵, and \mathcal{A} denotes the set of admissible consumption-portfolio pairs:

$$\mathcal{A} = \{(c, \pi) : c \text{ and } \pi \text{ are previsible, } c \geq 0, \text{ and for some } K < \infty, \|\pi_t\| \leq K\}. \quad (2.5)$$

REMARKS. (i) Notice that the function φ is defined on the whole of \mathbb{R} .

³We use the notations $a \cdot b$ for the scalar product of two vectors a and b , and $\mathbf{1}$ for the column vector of ones.

⁴See [17] for further discussion of this wealth dynamic and how to solve problems using it. In fact, [6] allow g to depend on time as well as π , though we omit this embellishment as it really changes nothing of substance.

⁵These are the conditions $\lim_{c \downarrow 0} U_c(t, c) = \infty = \lim_{w \downarrow -\infty} \varphi'(w)$, $\lim_{c \uparrow \infty} U_c(t, c) = 0 = \lim_{w \uparrow \infty} \varphi'(w)$.

(ii) The above definition of admissibility (2.5) is not the usual one⁶. Admissibility is imposed to eliminate doubling strategies, where wealth may go arbitrarily negative before time T , but ends up at a high value at time T . The assumptions made here rule this out; if we were to go to large negative wealth at some time in $(0, T)$, boundedness of σ , μ and π prevent us returning to positive wealth with certainty by time T , and the penalty imposed by the concave function φ then makes this a bad thing to do.

(iii) If the dimension k of the statespace of the factor diffusion X were not very small, it is not feasible to calculate and store the value function V . The approach we develop in this paper allows us to determine approximately optimal policies *without* the need to calculate V .

We shall require one technical condition on U , which is expressed as a condition on the inverse marginal utility I , defined by

$$U_c(s, I(s, z)) = z, \quad (z > 0). \quad (2.6)$$

We require

$$\text{ASSUMPTION: there exists } \alpha, A > 0 \text{ such that } I(t, z) \leq A(1 + z^{-\alpha}). \quad (2.7)$$

The inequality has to hold for all $z > 0$ and all $t \in [0, T]$.

We are now ready to state the main result of the paper, which allows us to derive effective Monte Carlo bounds on the value, and to find good sub-optimal strategies *pathwise*. The proof uses duality arguments similar to those presented in [5], [11], and later described in a more general setting in [10].

Theorem 2.1. *Suppose that κ is a bounded previsible process such that*

$$\tilde{g}(r_t \mathbf{1} - \mu_t + \sigma_t \kappa_t) = 0, \quad (2.8)$$

and that ζ solves the linear SDE

$$d\zeta_t = \zeta_t(-\kappa_t dW_t - r_t dt). \quad (2.9)$$

Define the function f by⁷

$$f(t, z, x) = E \left[\int_t^T \tilde{U}(s, \zeta_s) ds + \tilde{\varphi}(\zeta_T) \mid \zeta_t = z, X_t = x \right] \quad (2.10)$$

for $t \in [0, T]$, $z > 0$, $x \in \mathbb{R}^k$. Then for any $t \in [0, T]$, $z > 0$, $w \in \mathbb{R}$, $x \in \mathbb{R}^k$ we have the inequality

$$V(t, w, x) \leq f(t, z, x) + wz, \quad (2.11)$$

⁶One typically imposes a non-negativity constraint on the wealth process associated with the trading strategy π .

⁷The functions \tilde{U} , $\tilde{\varphi}$ are the convex dual functions, $\tilde{U}(t, z) \equiv \sup_x \{U(t, x) - zx\}$, $\tilde{\varphi}(z) \equiv \sup_x \{\varphi(x) - zx\}$.

If we suppose additionally that $\mathbf{g} \equiv \mathbf{0}$, then for any $t \in [0, T]$, $z > 0$, $w \in \mathbb{R}$, $x \in \mathbb{R}^k$, and bounded previsible π , we have the inequality

$$f(t, z, x) + wz - h(t, w, z, x, \pi) \leq V(t, w, x), \quad (2.12)$$

where

$$h(t, w, z, x, \pi) \equiv \mathbb{E} \left[\tilde{\varphi}(\zeta_T) - \varphi(w_T^\pi) + \zeta_T w_T^\pi \mid w_t = w, \zeta_t = z, X_t = x \right], \quad (2.13)$$

and the process w^π is the solution to the wealth evolution (2.2) with portfolio proportions π and consumption process

$$c_s = I(s, \zeta_s), \quad (s \geq t). \quad (2.14)$$

REMARKS. (i) When $g \equiv 0$, the condition (2.8) says that $\sigma_t \kappa_t = \mu_t - r_t \mathbf{1}$. In general the matrix σ is not even square, so not invertible, but we could try to find κ to satisfy (2.8) by taking the pseudo-inverse of σ :

$$\kappa_t = \sigma_t^T (\sigma_t \sigma_t^T)^{-1} (\mu_t - r_t \mathbf{1}). \quad (2.15)$$

This can be done if $(\sigma_t \sigma_t^T)^{-1}$ is bounded, in effect a uniform ellipticity condition of the kind commonly imposed in such problems. Notice that this particular choice of κ_t will be a function of X_t , and for the simulation methodology used later it turns out that this structure is essential. More generally, choice of the ‘market price of risk process’ κ corresponds to different choices of the martingale pricing measure.

(ii) From the definition of the convex dual function $\tilde{\varphi}$, it is clear that h is always non-negative. Since h dominates the gap between the lower and upper bounds, we should aim to make h as small as we can. Ideally, we would have that h was zero, which would require us to have

$$\varphi'(w_T) = \zeta_T. \quad (2.16)$$

If we demanded that this happens, then the problem becomes a BSDE with (2.16) as the terminal condition. However, it will only be solvable if w is an optimal wealth process and ζ is the state-price deflator which optimises the dual problem. Therefore, in some cases such an approach might be a possibility, but it will fail in the general case.

PROOF. (A) THE UPPER BOUND. The process ζ is determined by (2.8) and (2.9); in what follows, we shall suppose that c is determined from ζ by (2.14).

Consider the Itô expansion of $\zeta_T w_T$. We have:

$$\begin{aligned}
0 &= -\zeta_T w_T + \zeta_t w_t + \int_t^T (\zeta_s dw_s + w_s d\zeta_s + d[\zeta, w]_s) \\
&= -\zeta_T w_T + \zeta_t w_t + \int_t^T \zeta_s w_s (\pi_s \cdot \sigma_s - \kappa_s) dW_s \\
&\quad + \int_t^T \{ \zeta_s w_s (r_s + \pi_s \cdot (\mu_s - r_s \mathbf{1}) + g(\pi_s) - r_s - \pi_s \cdot \sigma_s \kappa_s) - \zeta_s c_s \} ds \\
&= -\zeta_T w_T + \zeta_t w_t + \int_t^T \zeta_s w_s (\pi_s \cdot \sigma_s - \kappa_s) dW_s - \int_t^T \zeta_s c_s ds \\
&\quad + \int_t^T \zeta_s w_s \{ g(\pi_s) - \pi_s (r_s \mathbf{1} - \mu_s + \sigma_s \kappa_s) \} ds
\end{aligned}$$

using (2.8) and (2.9).

We claim that the stochastic integral has zero mean, and in order to establish this, it is necessary to control the integrand. The processes κ , π , and σ are all bounded by hypothesis, so we need to have control on ζ and w . Since ζ satisfies the linear SDE (2.9) with bounded coefficients κ and r , it is not hard to establish a bound on $E[(\zeta_t^*)^p]$ for any $t > 0$, and for any $p \geq 2$, where $\zeta_t^* \equiv \sup_{0 \leq s \leq t} |\zeta_s|$; see, for example, Lemma V.11.5 of [20]. Similarly, we may bound $E[(\zeta_t^*)^{-p}]$ for any $t > 0$, and for any $p \geq 2$, by considering the linear SDE for ζ^{-1} . All that remains is to establish a similar bound for w_t^* , where w is given by (2.2). The only problematic part of this estimation is in controlling c , but this is where the Assumption (2.7) comes in, since ζ_t^{-1} is controlled as before, and c is bounded by some power of ζ .

We therefore conclude that

$$0 = E_t \left[-\zeta_T w_T + \zeta_t w_t - \int_t^T \zeta_s c_s ds + \int_t^T \zeta_s w_s \{ g(\pi_s) - \pi_s (r_s \mathbf{1} - \mu_s + \sigma_s \kappa_s) \} ds \right]. \quad (2.17)$$

We can add this equality to (2.4) to find⁸

$$\begin{aligned}
V(t, w, x) &= \sup_{(c, \pi) \in \mathcal{A}} E_t \left[\int_t^T \{ U(s, c_s) - \zeta_s c_s \} ds + \varphi(w_T) - \zeta_T w_T + \zeta_t w_t \right. \\
&\quad \left. + \int_t^T \zeta_s w_s \{ g(\pi_s) - \pi_s (r_s \mathbf{1} - \mu_s + \sigma_s \kappa_s) \} ds \middle| w_t = w, X_t = x, \zeta_t = \zeta \right] \quad (2.18) \\
&\leq E_t \left[\int_t^T \{ \tilde{U}(s, \zeta_s) + \zeta_s w_s \tilde{g}(r_s \mathbf{1} - \mu_s + \sigma_s \kappa_s) \} ds + \tilde{\varphi}(\zeta_T) + \zeta_t w_t \middle| w_t = w, X_t = x, \zeta_t = \zeta \right] \\
&= E_t \left[\int_t^T \tilde{U}(s, \zeta_s) ds + \tilde{\varphi}(\zeta_T) + \zeta_t w_t \middle| w_t = w, X_t = x, \zeta_t = \zeta \right] \\
&= E_t \left[\int_t^T \tilde{U}(s, \zeta_s) ds + \tilde{\varphi}(\zeta_T) \middle| w_t = w, X_t = x, \zeta_t = \zeta \right] + \zeta w \\
&= f(t, \zeta, x) + \zeta w. \quad (2.19)
\end{aligned}$$

⁸We use (2.14) at the first step.

This is the upper bound in (2.11).

(B) THE LOWER BOUND. Recall that we now assume that $g \equiv 0$, with the result that the condition (2.8) becomes simply that $r_t \mathbf{1} - \mu_t + \sigma_t \kappa_t = 0$. The argument reuses elements of the proof of the upper bound. The task this time is to propose some admissible (c, π) and deduce a lower bound from it.

Given the state-price density process ζ as in (2.9), our intention is to use the process c to be defined from it by (2.14). Doing this, we see that the first integral term appearing in the right-hand side of (2.18) is equal to

$$E \int_t^T \tilde{U}(s, \zeta_s) ds,$$

and the second integral on the right-hand side of (2.18) is absent thanks to the assumption that $g \equiv 0$. Moreover, (2.17) still holds by the same argument as before, but takes a simpler form because the second integral term is again absent. For any bounded previsible π , the pair (c, π) is admissible, so if we use that admissible pair we find as at (2.18) that

$$\begin{aligned} V(t, w, x) &\geq E \left[\int_t^T \tilde{U}(s, \zeta_s) ds + \varphi(w_T^\pi) - \zeta_T w_T^\pi + \zeta_t w_t^\pi \middle| w_t = w, X_t = x, \zeta_t = \zeta \right] \\ &= f(t, \zeta, x) + w\zeta - h(t, w, \zeta, x, \pi) \end{aligned} \quad (2.20)$$

when we recall the definitions (2.10) and (2.13) of f and h . □

REMARKS. (i) For any bounded previsible π and κ , Theorem 2.1 gives two-sided bounds on the value function. Importantly, the numerical values of f and h can be estimated *by forward simulation from current values*. It is no surprise that there has to be some kind of forward look, because the optimal controls must account somehow for what may happen in the future. It is also worth noting that the methodology does not require any ‘simulation within simulations’ which substantially increases the computation times; we will be evaluating the state-price density and the portfolio process along *just one trajectory*. All we need to do is to simulate sufficiently many sample paths to approximate the expectation operator in (2.10) and (2.13).

(ii) We need to have a measure for comparison between the bounds in (2.11), (2.12). Since affine transformations of utility functionals induce the same preferences, our measure needs to be invariant under those. Thus the difference between the upper and lower bounds is not informative.

We can however think of giving up a fraction of the initial wealth αw and look for the minimal α such that the upper bound corresponding to $(1 - \alpha)w$ initial wealth is at most as large as the lower bound for starting with wealth w . This α is of course:

$$\alpha(t, w, \zeta, X, \pi) \equiv \frac{h(t, w, \zeta, X, \pi)}{\zeta w}, \quad (2.21)$$

which will from now on be our efficiency measure. Notice that (2.21) is a dimensionless quantity. If, in a particular application, there is a possibility that the wealth might hit zero, we should instead only divide by ζ , expressing efficiency in the units of wealth.

(iii) The key issue for obtaining good bounds is of course the choice of the processes κ and π . The traditional way to approach solving the problem (2.4) would be to write down the HJB equation, derive the corresponding PDEs, and try to solve them. However, these PDEs are typically highly non-linear, and we only stand a chance of getting reasonably stable solutions in dimensions one or two.

Nevertheless, we can deduce some worthwhile information from the HJB equation. For simplicity we limit the discussion to the main case $g \equiv 0$. Dropping the t subscript, and remembering the function V takes (t, w, X) as arguments, the HJB equation is

$$0 = \sup_{c, \pi} \left[U(t, c) + V_t + (rw + w\pi \cdot (\mu - r\mathbf{1}) - c) V_w + \mu_X \cdot V_X + \right. \\ \left. + \frac{1}{2} w^2 |\pi^T \sigma|^2 V_{ww} + w\pi \cdot \sigma \sigma_X^T V_{Xw} + \frac{1}{2} \text{tr}(\sigma_X \sigma_X^T V_{XX}) \right]. \quad (2.22)$$

Optimizing over c leads to the conclusion that $c_t = I(t, V_w)$, and optimizing over π tells us that we should have

$$w\pi = -(\sigma\sigma^T)^{-1} \{ (\mu - r\mathbf{1})V_w + \sigma\sigma_X^T V_{Xw} \} / V_{ww}. \quad (2.23)$$

Here $\sigma\sigma^T$ is invertible by our non-degeneracy assumptions on the market.

Assuming that V and f are dual (as we would expect from (2.11), (2.12)), that is,

$$V(t, w, x) = \inf_{\zeta} \{ f(t, \zeta, x) + w\zeta \}, \quad f(t, \zeta, x) = \sup_w \{ V(t, w, x) - w\zeta \}, \quad (2.24)$$

would lead us to the relations

$$w = -f_z(t, z, x), \quad \zeta = V_w(t, w, x). \quad (2.25)$$

Straightforward calculus then gives

$$V_{ww}(t, w, x) = -1/f_{\zeta\zeta}(t, \zeta, x). \quad (2.26)$$

These relations help us to make choices of κ and π .

If we ignore for the moment the cross-derivative term in (2.23), what we see is the equation

$$\pi = \frac{-V_w}{wV_{ww}} (\sigma\sigma^T)^{-1}(\mu - r\mathbf{1}) = \frac{-\zeta f_{\zeta\zeta}(t, \zeta, X)}{f_{\zeta}} (\sigma\sigma^T)^{-1}(\mu - r\mathbf{1}) \quad (2.27)$$

and we recognize $-wV_{ww}/V_w$ as the coefficient of relative risk aversion of the value function; if this were constant, then the portfolio recommendation (2.27) is just the Merton portfolio. If we were able to ignore the cross-derivative terms, then the message coming from (2.23), (2.27) is that *all we need to know in order to make the optimal portfolio choice is the coefficient of relative risk aversion of the value function at the present time and place.*

In general, we need to include the terms in the cross-derivatives V_{Xw} , but we notice from (2.25) that this just requires us to find the derivatives of ζ with respect to x . We intend to get hold of f by using forward simulation from the current state, and to estimate the derivative of ζ with respect to x we will have to repeat the calculation from x and from nearby starting points. This will of course be more cumbersome, and the required computational effort of our algorithm will increase linearly in the dimension of X .

For calculating the derivatives numerically, we have used a 5-point method, which has marginally outperformed the 3-point method.⁹

(iv) In practice, it will be clumsy to form an estimate of the term h in (2.12) if we are determining the portfolio process π according to the recipe just outlined, because if we are to simulate an evolution of (X, w) we will at each step need to identify derivatives of f , and this is a simulation within a simulation. We envisage the lower bound in (2.12) being used as a means to *assess* a particular portfolio rule which may be expressed explicitly as some function of (t, X, w) . In a high-dimensional problem, we do not expect the optimal portfolio rule to be something we can characterize, but we may well have some heuristic for some ‘good’ portfolio rule, and (2.12) gives us a way to tell how good that heuristic may be.

Summarising: Given an initial state (t, w, ζ) , we can follow the dynamics of w , ζ , and X , using (2.1), (2.2), (2.9), κ satisfying the relation (2.8), c given by (2.14), and π given by (2.27) (or perhaps (2.23)).

The key advantage of this formulation is that all we need to do now is to optimise the bound (2.11) for a one-dimensional starting value of the dual process ζ . This is a quick procedure numerically.

3 Algorithms.

We will now describe an algorithm for simulating the optimal path and controls for the problem (2.4), given *a particular realization of the Brownian motion*. That is, we do not attempt to recover the whole value function, as this is bound to fail in higher dimensions. Our method, which is effectively local, will follow a particular realization of the Brownian motion W and tell us how to invest and consume in that particular case. After all, we are predominantly interested in how to invest where we are now - we are less interested in what we would have done had we been somewhere different.

Algorithm 1 describes how to compute the best bounds numerically. The cost of running this algorithm will be $\mathcal{O}(N) \times \mathcal{O}(f)$, where N is the number of time grid points, and $\mathcal{O}(f)$ is the average cost of evaluation of the function f and h .

In Algorithm 1, we have not yet given the details of how to calculate the function f numerically (which will be the business of Algorithm 2). That is, we want to be able to numerically calculate the expectation in (2.10) and (2.13) for a given time $t = t_n$ on the discretization grid.

⁹We are grateful to Referee 1 for bringing this to our attention.

Algorithm 1: Computing the optimal path.

-
- Step 1: Initialization.* Pick starting values $w = w_0$, $X = X_0$ and a grid of time steps $0 = t_0 < t_1 < t_2 < \dots < t_N = T$ along which we want to know the solution. Simulate a realization of the Brownian motion W along which we want to calculate the optimal path.
- Step 2: Finding the optimal ζ_0 .* For any ζ , we can calculate $f(0, \zeta, X) + w\zeta_0$. This function is convex in ζ , so we can use the golden section search to find the minimum in (2.24). This gives us the value of $V(0, w_0, X_0)$ and the optimal starting value of the dual process ζ_0 .
- Step 3: Calculating the optimal path.* For each $n = 0, 1, \dots, N - 1$, we have $(t_n, \zeta_{t_n}, X_{t_n})$ available. We use (2.14) to work out c_{t_n} , (2.27) to work out π_{t_n} , and (2.25) to work out w_{t_n} . We then calculate κ_{t_n} with (2.8) and use the Euler scheme to move to time t_{n+1} using (2.1) and (2.9).
-

We approximate such expectation by sampling M paths of Brownian motion W for $t = t_n, t_{n+1}, \dots, t_N$ and hence simulating the system forwards along M different paths. We then average over the results.

In practice, we find that it might be necessary to use importance sampling in order to decrease the volatility of our estimates. In order to do that, define the change of measure martingale

$$dZ_s^{-1} = Z_s^{-1} \sigma_s^Z dW_s \text{ for } t \leq s \leq T, \quad Z_t = 1, \quad (3.1)$$

and set $\frac{d\mathbb{Q}}{d\mathbb{P}}|_{\mathcal{F}_t} = Z_t^{-1}$. Then we can rewrite (2.10) and (2.13) as

$$f(t, \zeta, X) = \mathbb{E}^{\mathbb{Q}} \left[\int_t^T Z_s \tilde{U}(s, \zeta_s) ds + Z_T \tilde{\varphi}(\zeta_T) \middle| \zeta_t = \zeta, X_t = X \right], \quad (3.2)$$

$$h(t, w, \zeta, X, \pi) = \mathbb{E}^{\mathbb{Q}} \left[Z_T \varphi(w_T^\pi) - Z_T w_T^\pi \zeta_T - Z_T \tilde{\varphi}(\zeta_T) \middle| w_t = w, \zeta_t = \zeta, X_t = X \right]. \quad (3.3)$$

with a new Brownian motion $W^{\mathbb{Q}}$ under \mathbb{Q} defined by

$$d\bar{W}_t = dW_t - \sigma_t^Z dt. \quad (3.4)$$

The idea now is to choose σ_Z in a way that the Ito expansion of the term $Z_T \tilde{\varphi}(\zeta_T)$ has no $d\bar{W}$ term. This has a variance reducing property. Writing \doteq whenever two sides of an equality differ only by integrals with respect to ds , we have

$$Z_T \tilde{\varphi}(\zeta_T) = Z_t \tilde{\varphi}(\zeta_t) + \int_t^T d(Z_s d\tilde{\varphi}(\zeta_s)) \doteq Z_t \tilde{\varphi}(\zeta_t) + \int_t^T (Z_s \tilde{\varphi}'(\zeta_s) d\zeta_s + dZ_s \tilde{\varphi}(\zeta_s)) \quad (3.5)$$

$$= Z_t \tilde{\varphi}(\zeta_t) + \int_t^T Z_s (-\kappa_s \tilde{\varphi}'(\zeta_s) \zeta_s - \sigma_s^Z \tilde{\varphi}(\zeta_s)) d\bar{W}_s. \quad (3.6)$$

Algorithm 2: Computing $f(t_n, \zeta, X)$ and $h(t_n, w, \zeta, X, \tilde{\pi})$

-
- Step 1: Initialization.* Recall $t = t_n$. Generate M paths of Brownian motion \bar{W}_t^i , $i = 1, 2, \dots, M$, with values evaluated at $t = t_n, t_{n+1}, \dots, t_N$. The corresponding paths for ζ , X , w and Z are denoted by ζ^i , X^i , w^i and Z^i with $\zeta_{t_n}^i = \zeta$, $X_{t_n}^i = X$, $Z_{t_n}^i = 1$ and $w_{t_n}^i = w$.
- Step 2: Simulation.* For $k = n, n+1, \dots, N-1$, update $\zeta_{t_{k+1}}^i$, $X_{t_{k+1}}^i$, $Z_{t_{k+1}}^i$ and $w_{t_{k+1}}^i$ as follows. Equations (3.7) and (3.4) give us the corresponding dW_{t_k} . We then use (2.9), (2.1), (3.1) and (2.2) to move to the next time point using the Euler scheme.
- Step 3: Averaging.* Having calculated paths ζ^i , X^i , Z^i and w^i corresponding to M paths of \bar{W}^i , we return the approximate values of f and h :

$$f(t_n, \zeta, X) \approx \frac{1}{M} \sum_{i=1}^M \left(\sum_{k=n}^{N-1} Z_{t_k}^i \tilde{U}(t_k, \zeta_{t_k}^i) + Z_{t_N}^i \tilde{\varphi}(\zeta_{t_N}^i) \right) \quad (3.8)$$

$$h(t_n, w, \zeta, X, \tilde{\pi}) \approx \frac{1}{M} \sum_{i=1}^M Z_{t_N}^i (\varphi(w_{t_N}^i) - w_{t_N}^i \zeta_{t_N}^i - \tilde{\varphi}(\zeta_{t_N}^i)). \quad (3.9)$$

Therefore, we set:

$$\sigma^Z \equiv -\kappa_s \frac{\zeta_s \tilde{\varphi}'(\zeta_s)}{\tilde{\varphi}(\zeta_s)}, \quad (3.7)$$

which cancels the $d\bar{W}$ term in (3.5), and in turn in (3.2).

With this in mind, we now present the numerical algorithm for calculating $f(t, \zeta, X)$.

The computational complexity of Algorithm 2 comes from (3.8), where we clearly see that we need $\mathcal{O}(N) \times \mathcal{O}(M)$ operations. Therefore, we deduce that $\mathcal{O}(f) = \mathcal{O}(MN)$.

The key to performance of the method is of course the accuracy of the Monte Carlo simulation. As we shall see in the following Section, the numerical results are promising. Even a fairly moderate number of Monte Carlo paths can provide a good approximation to the true value of f and h . With this in mind, we proceed to examine the numerical results for the performance of the method.

4 Numerical performance

In this Section, we shall compare the results of the Monte Carlo solutions with special cases of the problem (2.4) where we either know the solution in closed form, or we know highly accurate numerical schemes for approximating the solution.

We start off by analysing complete markets where some of the analysis in the previous Section simplifies. Recall that, in a complete market the asset volatility matrix σ is invertible. This means we have a unique¹⁰ state-price density for the problem, given by

$$\zeta_t = \zeta_0 \exp \left[- \int_0^t \kappa_s \cdot dW_s - \int_0^t \left(r_s + \frac{1}{2} |\kappa_s|^2 \right) ds \right], \quad (4.1)$$

where $\kappa_s \equiv \sigma_s^{-1}(\mu_s - r_s 1)$. Therefore, provided that (2.24) holds, our Monte Carlo method should be able to find the optimal path exactly, modulo numerical errors coming from Monte Carlo approximation of the expectation operator in (2.10), approximating the derivatives in (2.25) and (2.26), and finally the numerical optimisation over the (scalar!) value ζ_0 in (2.10). The positive side is that all these errors can be made small provided we use enough computational power.

With that in mind, we start off with two examples of problems dealing with complete markets where the benchmark answers are reliable; and finish by analysing runs in incomplete markets where we provide estimate error bounds, but where no other solutions methods are available.

4.1 The Merton problem

We start by comparing our results to the solutions of the Merton problem, which are available in closed form in multiple dimensions. Recall that the Merton problem assumes that functions r , μ and σ in (2.2) are constant, and the utility functions U and φ in (2.4) take a particular form:

$$U(t, c) = e^{-\rho t} u(c), \quad (4.2)$$

$$\varphi(w) = Au(w), \quad (4.3)$$

where A and ρ are positive constants, and u is a constant relative risk aversion utility:

$$u(c) = \frac{c^{1-R}}{1-R}, \quad (4.4)$$

for $R > 0$, $R \neq 1$. Constant A allows us to assign different weights to consumption and terminal utility, and ρ allows for the future utility decay rate. Then the optimal solution takes the form:

$$V(t, w, X) = F(t) u(w), \quad (4.5)$$

$$\theta_t = \pi_M w_t, \quad (4.6)$$

$$c_t = \gamma(t) w_t, \quad (4.7)$$

¹⁰Up to a multiplicative constant still to be found.

where

$$F(t) = \left\{ A^{1/R} e^{-b(T-t)} + \frac{e^{-\rho t/R}}{b + \rho/R} (1 - e^{-(b+\rho/R)(T-t)}) \right\}^R, \quad (4.8)$$

$$\pi_M = R^{-1}(\sigma\sigma^T)^{-1}(\mu - r\mathbf{1}), \quad (4.9)$$

$$\gamma(t) = e^{-\rho t/R} F(t)^{-1/R}, \quad (4.10)$$

where $b = (R - 1)(r + |\kappa|^2/2R)/R$; see [19], Section 2.1.

Figure 1 shows the results of the simulation runs for the 3-dimensional version of the problem using $M = 1000$ paths. The top left panel shows the running estimate of the value function $V_M(t, w_M(t))$ along a particular realization of Brownian motion W . The top right and bottom left panels show investment and consumption proportions, respectively. Finally, the bottom right panel depicts the estimated wealth process compared to the Merton wealth process.

As we see, all the graphs give a very satisfactory approximation to the Merton solution. This is especially remarkable taking into account that we are already in dimension 3, and we have used relatively few paths.

We now present the study of how the accuracy of the solutions to the Merton problem varies for different values of the number of simulations M and number of dimensions K . We found that the number of time steps N used to discretize the integral in (3.2) does not greatly influence the accuracy of the solutions.

We compare the estimates of the optimal starting ζ_0 found by the procedure (2.24) in Algorithm 1. For each test, we keep the initial data of *Step 1* fixed. We then run *Step 2* of Algorithm 1, each time approximating the function f with a different set of Monte Carlo paths. This way, we can investigate how sensitive our optimized values of ζ_0 are to the Monte Carlo procedure for approximating the expectation operator.

Table 3 and Table 4 present the results of the simulations for different number of Monte Carlo paths to calculate f , $M = 1000$ and $M = 10000$, respectively. We see that the numerical results work reasonably well for $K \leq 6$ when we choose to use 1000 Monte Carlo paths. The average ζ_0 is pretty close to the true value, and the volatility of the estimates stays modest. However, for larger values of K , we see that the estimates are either not as accurate, or become more volatile.

For $M = 10000$, the results look much better. For $K \leq 9$, we see a considerable drop in the volatility of the estimates, and all of them lie within two standard deviations of the true value, with most of them being less than one standard deviation away.

These results are very encouraging. They show that, even in dimensions up to 10, having a reasonably modest number of Monte Carlo paths of 10000 can provide satisfactory results when solving the Merton problem. This is particularly interesting since the traditional HJB approach would struggle in these dimensions unless the problem has a particular structure such that we can work out the value function explicitly.

One might think that the accuracy of the method relies on the special structure of the Merton problem. We now show that this is not the case. We consider departures from the basic problem where accurate numerical solutions are available.

	$K = 1$	$K = 2$	$K = 3$	$K = 4$	$K = 5$	$K = 6$	$K = 7$	$K = 8$	$K = 9$	$K = 10$
Merton	9.97	9.49	8.92	8.61	8.17	8.02	7.73	7.44	7.11	6.68
Average(ζ_0)	9.72	9.33	8.64	8.86	7.53	7.85	7.36	7.54	6.44	5.31
Stdev(ζ_0)	0.12	0.14	0.23	0.34	0.30	0.30	0.36	0.60	0.24	0.48
Time / run (min)	0.67	2.32	2.95	3.51	4.08	4.63	5.15	5.61	6.41	6.83

Table 3: Comparison of the ζ_0 for the Merton problem and the values found using the Monte Carlo method for different values of the dimension parameter K corresponding to the number of assets. The number of Monte Carlo paths each time was equal to $\mathbf{M} = \mathbf{1000}$. For each set of simulated Monte Carlo paths, we find the optimal implied value of ζ_0 . We then take the average as the estimate, and calculate its standard deviation. Here we take $r = 0.05$, $\rho = 0.03$, $R = 3$, $w_0 = 1$, $A = 1$, $N = 100$, $dt = 0.05$. The parameters μ and σ were generated randomly: μ had a $U[10\%, 50\%]$ distribution, once the entries of σ were drawn from $U[-1, 1]$ until the resulting matrix was positive definite.

	$K = 1$	$K = 2$	$K = 3$	$K = 4$	$K = 5$	$K = 6$	$K = 7$	$K = 8$	$K = 9$	$K = 10$
Merton	10.10	9.67	9.34	8.63	8.35	8.13	7.33	7.10	6.83	6.48
Average(ζ_0)	10.15	9.58	9.35	8.76	8.29	8.33	7.30	7.16	6.63	6.95
Stdev(ζ_0)	0.04	0.08	0.06	0.08	0.08	0.12	0.15	0.19	0.12	0.49
Time / run (min)	6.87	22.90	28.65	34.42	40.06	46.25	52.08	57.27	61.26	68.09

Table 4: Comparison of the ζ_0 for the Merton problem and the values found using the Monte Carlo method for different values of the dimension parameter K . The number of Monte Carlo paths each time was equal to $\mathbf{M} = \mathbf{10000}$. For each set of simulated Monte Carlo paths, we find the optimal implied value of ζ_0 . We then take the average as the estimate, and calculate its standard deviation. Here we take $r = 0.05$, $\rho = 0.03$, $R = 3$, $w_0 = 1$, $A = 1$, $N = 100$, $dt = 0.05$. The parameters μ and σ were generated randomly: μ had a $U[10\%, 50\%]$ distribution, once the entries of σ were drawn from $U[-1, 1]$ until the resulting matrix was positive definite.

4.2 Non-constant relative risk aversion

The example of the Merton problem has shown us that the Monte Carlo method can handle situations where we deal with a multi-dimensional Brownian motion. However, the multiplicative scaling property of the CRRA utility function u means that we are unable to assess the accuracy in predicting π . The remarkable accuracy in prediction in Figure 1 is caused by the fact that $f(t, \zeta, X) = \zeta^{1-1/R} \tilde{f}(t, X)$, for some function \tilde{f} , and the fact that the optimal π satisfies (2.27).

It will therefore be informative to consider an example where the proportion of money invested in the risky assets varies with wealth. This can be done, although the price to pay is dimensionality. In this Section, we assume that the financial market has constant coefficients and that there is only one asset in the market.

For $R_1 > 1 > R_2 > 0$, we define the agent's marginal utility as

$$I(t, y) = a_1^{1/R_1} e^{-\rho t/R_1} y^{-1/R_1} + a_2^{1/R_2} e^{-\rho t/R_2} y^{-1/R_2}, \quad (4.11)$$

$$I_\varphi(y) = b_1^{1/R_1} y^{-1/R_1} + b_2^{1/R_2} y^{-1/R_2}. \quad (4.12)$$

What this means is that, for small values of wealth w , the agent's relative risk aversion is close to R_1 and the agent behaves similarly to the Merton investor from Section (4.1) with $R = R_1$, $a = a_1$ and $b = b_1$, and value function $V_1(t, w)$. Conversely, the investor for large values of w is less risk averse, with risk aversion R_2 . He behaves like a Merton investor from Section (4.1) with $R = R_2$, $a = a_2$, and $b = b_2$, and value function $V_2(t, w)$.

In dimension one, there are two very effective methods for solving this problem: policy improvement and quantisation¹¹. We proceed by briefly describing each one of them, and then by comparing their performance with the Monte Carlo scheme we proposed earlier.

Policy improvement. We follow the approach described in Section 3.4 of [19]. The HJB equation for our problem is

$$0 = \sup_{c, \pi} [U(t, c) + V_t(t, w) + (rw + w\pi(\mu - r) - c)V_w(t, w) + \frac{1}{2}w^2\pi^2\sigma^2V_{ww}(t, w)], \quad (4.13)$$

and we are given the terminal value

$$V(T, w) = \varphi(w). \quad (4.14)$$

Given functions (4.11), functions U and φ , although not available in closed form, can be found efficiently using binary search.

We therefore give ourselves a grid of time points $0 < t_1 < t_2 < \dots < t_N = T$ and a grid of space points $w_1 < w_2 < \dots < w_M$ and we wish to find V evaluated at their mesh.

At the boundaries, we know that the solution resembles the Merton solutions:

$$V(t, w_1) = V_1(t, w_1), \quad V(t, w_N) = V_2(t, w_N). \quad (4.15)$$

¹¹Both of which are difficult to generalise to dimensions more than one, though.

Let $\mathcal{L}(c, \pi)$ be a functional acting on smooth test functions $\psi(t, w)$ as

$$\mathcal{L}(c, \pi)\psi(t, w) = (rw + w\pi(\mu - r) - c)\psi'(t, w) + \frac{1}{2}w^2\pi^2\sigma^2\psi''(t, w). \quad (4.16)$$

\mathcal{L} here is the infinitesimal generator of the wealth process.

Noticing that

$$\psi'(t, w) \approx \frac{\psi(t, w_{i+1}) - \psi(t, w_{i-1})}{\Delta_+ + \Delta_-} \quad (4.17)$$

$$\psi''(t, w) \approx \frac{\Delta_-(\psi(t, w_{i+1}) - \psi(t, w_i)) - \Delta_+(\psi(t, w_i) - \psi(t, w_{i-1}))}{\Delta_+\Delta_-(\Delta_+ + \Delta_-)}, \quad (4.18)$$

where $\Delta_+ = w_{i+1} - w_i$ and $\Delta_- = w_i - w_{i-1}$, it is possible to approximate \mathcal{L} acting on $\psi(t, \cdot)$ by a sparse triangular matrix $L(c, \pi)$ acting on a column vector $\psi(t, w_i), i = 2, \dots, M-1$, using approximations (4.17) plugged into (4.16)¹².

We now discretize the differential operator appearing in the HJB equation (4.13) on the chosen time and space grid. By letting $V_i^n = V(t_n, w_i)$ and $V^n = (V_i^n)_{i=1,2,\dots,M}$, we obtain:

$$0 = \sup_{c, \pi} \left[\frac{V^{n+1} - V^n}{t_{n+1} - t_n} + \alpha(L(c_n, \pi_n)V^n + U(t_n, \cdot, c_n)) + (1 - \alpha)(L(c_{n+1}, \pi_{n+1})V^{n+1} + U(t_{n+1}, \cdot, c_{n+1})) \right]. \quad (4.19)$$

We took $\alpha = 0.5$, giving the Crank-Nicholson method. We define L to act on the boundary points w_1 and w_M in such a way that (4.19) yields boundary solutions given by (4.15).

Given (c, π) and the forward solution at grid points with index $n+1$, (4.19) is then a sparse set of linear equations which we solve for $(V^n)_{i=1,2,\dots,M}$. We then improve on (c, π) by maximisation in (4.19), given the found V . We iterate the process until convergence.

Figure 2 shows the results of the policy improvement algorithm for $t = 0$. As we see, we were able to recover the whole value function using the method described above. It is worth pointing out, though, that the method is tricky to implement even in one dimension, and higher dimensions are almost certainly out of question. However, once V has been found in one dimension, working out the optimal consumption and investment around a sample path of Brownian motion are immediate.

Quantization. We proceed to a method which builds on the observations from Section 2, but avoids using the Monte Carlo method for approximating the expectation operator in (2.10). Instead, quantisation proposes approximating the expectation of the Brownian functional by a deterministic sum. Here we follow the details from the website [16] and related papers [15] and [4]. The idea is to use the Karhunen-Loeve expansion of Brownian motion $(W_t)_{0 \leq t \leq T}$:

$$W_t = \sum_{k=1}^{\infty} \xi_k e_k(t), \quad (4.20)$$

¹²Where we consider $w = (w_1, w_2, \dots, w_M)^T$ as a column vector, with the corresponding controls $(c_1, c_2, \dots, c_M)^T$ and $(\pi_1, \pi_2, \dots, \pi_M)^T$.

where $(\xi_n)_{n \geq 1} \sim N(0, \lambda_n)$ is a sequence of independent normal random variables with variance λ_n . Here the decomposition functions are

$$e_n(t) = \sqrt{\frac{2}{T}} \sin\left(\frac{\pi t}{T} \left(n - \frac{1}{2}\right)\right), \quad (4.21)$$

$$\lambda_n = \left(\frac{T}{\pi(n - \frac{1}{2})}\right)^2. \quad (4.22)$$

Brownian motion W is then firstly approximated by choosing the first d terms in the sum (4.20). We can then think of $\xi = (\xi_n)_{n=1,2,\dots,d}$ and $e(t) = (e_n(t))_{n=1,2,\dots,d}$ as d -dimensional vectors, and the Brownian motion as being approximated by the dot product

$$W_t \approx \xi \cdot e(t) \quad (4.23)$$

We then quantise the random d -dimensional vector ξ by a random variable X taking n distinct values $x_1, x_2, \dots, x_n \in \mathbb{R}^d$ with respective probabilities p_1, p_2, \dots, p_n , and giving us the final approximation

$$W_t \approx X \cdot e(t). \quad (4.24)$$

Now, if we need to calculate an expected value of a functional

$$\mathbb{E} \left[\int_0^T q(t, W_t) dt + F(W_T) \right], \quad (4.25)$$

we can now approximate it by a deterministic sum

$$\sum_{k=1}^n p_k \left[\int_0^T q(t, x_k \cdot e(t)) dt + F(x_k \cdot e(t)) \right]. \quad (4.26)$$

The effectiveness of this application depends on the number of terms n taken in the expansion (4.20), as well as the placing of the points and weights x_i and p_i . Files of the points and weights for many different values of n and for dimension up to 10 may be freely downloaded from the website [16]. These points and weights are optimal quantizations of the standard Gaussian distribution, in a sense explained in detail there. For a chosen number of n , we can therefore load up the optimal x_i and p_i from these files. For our runs, we use $n = 10160$.

The important thing is that, for the current problem, the calculations we need to perform are of the particular form (4.25). Indeed, in a complete market with one asset, we have

$$f(t, z) = \mathbb{E} \left[\int_t^T \tilde{U}(t, \zeta_s) ds + \tilde{\varphi}(\zeta_T) \middle| \zeta_t = z \right], \quad (4.27)$$

compare it with (2.10). Here ζ has a closed-form expression

$$\zeta_s = \zeta_t \exp \left[-\kappa(W_s - W_t) - \left(r + \frac{1}{2}\kappa^2\right)(s - t) \right] \text{ for } s \geq t, \quad (4.28)$$

which is of the required form (4.25).

Having laid out the problem setup and the accurate numerical methods for solving the problem, we now show the numerical results of our calculations.

Comparison of the methods. Figure 3 shows the results of the simulation runs. It is clear that all the methods proposed give virtually the same answers; with Monte Carlo being only marginally different from the two benchmark methods of policy improvement and quantisation. The most reassuring message here is that the Monte Carlo methodology also does a very good job on approximating the investment proportions for the problem as in (2.27) and (2.14). This is the part the the Merton problem example was unable to reveal due to the special structure.

The time taken to get the answer for the policy improvement was approximately 10 minutes, most of which was taken on the calculation of the value function (evaluating the solution along a chosen path is extremely fast). In comparison, quantisation has taken roughly 4 minutes, and Monte Carlo took 8 minutes.

Of course, each of the methods has their costs and benefits. The value function takes a time-investment at the start, but is very fast regardless of how many sample paths we would like to evaluate. This is not the case for quantisation and the Monte Carlo method. Quantisation is the overall speed-winner here, however we must remember that this is mainly due to the preloaded files which we used to quantise the Brownian motion.

Overall, we conclude that the Monte Carlo method performs very well on the complete market problems, as it should. After all, as mentioned before, the only errors we are incurring are numerical: approximating the derivatives and the expectation operator. With sufficient computational power, it should be possible to make these small.

4.3 Incomplete markets driven by a diffusion

Further, we consider an example where no benchmark methods are available, and the bounds derived in Theorem 2.1 are the only sensible indicator for how well our method is doing. We consider an example that is as challenging as possible: an incomplete market driven by a diffusion.

As a specific example, we consider a market composed of 4 stocks driven by a 5-dimensional Brownian motion. The same Brownian motion drives the 5-dimensional factor process X , which we take as a 5-dimensional OU process with independent components. We then generate a volatility matrix σ_0 such that σ_0 has rank 4, and $\sigma_0^T \sigma_0$ is invertible. We then set $\sigma_t = \sigma_0(1 + \exp(-1 \cdot X_t))$ which we then take to be the volatility process of the assets. Other parameters of the problem are kept constant. This guarantees that the market in question is incomplete.

Finally, we take a CRRA utility function, with a number of Monte Carlo paths being equal to $M = 1000$.

The results of the optimisation run are depicted on Figure 4. The run time here took 23h.¹³ The details regarding the parameters are displayed below the panel.

¹³This is due to the fact that we insist on evaluating the lower bound function h in (2.13) at each point

As we can see from the first two panels on the top, the upper and lower bounds stay reasonably close during the sample runs, with the error measure defined in (2.21) between 10% and 22%, and generally decreasing as we near the end of investment.

This is a positive result, especially in light of the dimensionality of the problem. Notice that the market is incomplete, and that the value function for this problem would need to be 7-dimensional (1 dimension for wealth, 1 for time, and 5 for the factor process X). Hence, any other method for approaching this problem would really struggle.

We could of course try to improve on the performance of this algorithm. We lose efficiency when we use the approximation of κ in (2.15), and also when we truncate the expression (2.23) for π . However, our main goal of the paper has already been achieved here: we have illustrated how to use our method on a very difficult problem, and derived satisfactory bounds on the efficiency.

4.4 Correlated assets with no short-selling constraints

Our last example illustrates how the method deals with portfolio constraints through allowing a non-zero function g in (2.2). We shall assume that the agent is not allowed to short the stocks, i.e. $\pi_t \geq 0$ for all t . This translates to taking

$$g(\pi) = \begin{cases} 0 & \text{if } \pi \geq 0 \\ -\infty & \text{otherwise,} \end{cases} \quad (4.29)$$

which conveniently gives

$$\tilde{g}(z) = \begin{cases} 0 & \text{if } z \geq 0 \\ \infty & \text{otherwise,} \end{cases} \quad (4.30)$$

In order to use Theorem 2.1, we need to propose a process κ such that (2.8) holds or, equivalently, $r\mathbf{1} + \mu + \sigma\kappa \geq 0$. At the same time, we will show how to construct a corresponding feasible process π_t so that the lower bound in Theorem 2.1 can also be used.

In order to do that, set

$$\bar{\pi} \equiv \max \{ (\sigma\sigma^T)^{-1}(\mu - r\mathbf{1}), 0 \}. \quad (4.31)$$

We then have that the maximiser of the HJB equation over $\pi \geq 0$ in (2.22) is $\pi = \frac{-V_w}{wV_{ww}} \bar{\pi}$. Further, setting $\kappa = \sigma^T \bar{\pi}$ we get

$$r\mathbf{1} - \mu + \sigma\kappa = \sigma\sigma^T \left((\sigma\sigma^T)^{-1}(r\mathbf{1} - \mu) + \max\{(\sigma\sigma^T)^{-1}(\mu - r\mathbf{1}), 0\} \right) \geq 0,$$

since $\sigma\sigma^T$ is positive-definite, and $(\sigma\sigma^T)^{-1}(r\mathbf{1} - \mu) + \max\{(\sigma\sigma^T)^{-1}(\mu - r\mathbf{1}), 0\} \geq 0$ component-wise.

on our trajectory, which in turn requires an evaluation of a forward expectation of the proposed investment strategy. If we were satisfied with an initial estimate of the error bound at time $t = 0$ and only wanted the proposed investment strategy, the run would have taken a couple of minutes.

We have therefore suggested a κ and π which are a natural extension of the unconstrained market equivalents in a way which lets us use Theorem 2.1.

As a specific numerical example, we take the Merton problem with two assets, and allow the correlation between the assets to vary in a stochastic way. That is, we keep the same notation as in section 4.1, but let the correlation of the assets to be

$$\rho_t = \tanh(X_t), \quad (4.32)$$

where X is a one-dimensional OU process following dynamics

$$X_0 = x, \quad (4.33)$$

$$dX_t = dW_t - \lambda(X_t - x)dt, \quad (4.34)$$

for constants $x \geq 0$, $\lambda > 0$. The idea is that the behaviour of the investor will be changing depending on how much the assets are correlated.

The results of the simulations are displayed on Figure 5. We can see the behaviour of the investor under three different conditions: base case correlation, high positive correlation, and high negative correlation. We can see that the investor wants to hold around 20% of his wealth in asset 1 (which has a higher average return) when the two assets are around 30% correlated. Asset 2 appears to be useful as a hedging instrument. When the correlation gets higher, e.g. at time $t = 0.2$ or $t = 0.7$, the investor wants to short asset 1, which we can see looking at the corresponding investment proportions. Due to the constraints we imposed on the problem, this is not allowed and the investor only holds one asset at times of high positive correlation. Finally, when the assets are negatively correlated, such as just before $t = 0.1$ or $t = 0.4$, it is beneficial to hold both of them at the same time to reduce volatility. In fact, we see that, at such times, the investor is willing to buy more of each asset than (s)he would at the base case scenario.

The behaviours described are certainly sensible, and our method also delivers bounds on the value function along the sample path. As we can see, the error measure defined in (2.21) is between 2.0% – 5.5% along the run, which is satisfactory.

For comparison, we also attach the performance of the myopic Merton portfolio in this problem. That is, at time t , we choose the investment proportions and consumption proportions to be the Merton ones, with input coefficients being equal to the current values of μ_t, r_t, σ_t . We can see that, in general, this strategy performs worse than ours, with corresponding error rates ranging from 6% – 11%. The wealth processes appear to be quite volatile and difficult to compare. However, looking at the investment proportions and consumption proportions we see that the Merton strategy tends to be riskier in this case, investing more in the asset 2, especially later on in the simulation. It also tends to consume at a slightly higher rate. The Merton strategy tends to expose the investor to more market risk than the solution using our simulation.

5 Conclusions

This paper presents a way of finding upper bounds on the value of using some reasonably explicit portfolio rule in an optimal investment/consumption problem that can be of very high dimension, such as may occur where the characteristics of the available investable assets depend on some high-dimensional diffusion driver. The approach works for convex-constrained problems as well as for unconstrained problems. We have illustrated with numerical examples how effective the technology can be. The key point is that we do not attempt to characterize the value function of the problem, since this is not needed to allow us to decide what the investment/consumption decisions should be at any time - all that is needed is derivatives of the value function, and these can be numerically approximated by derivatives of the dual value which are easier to get hold of by simulation. For an explicit investment/consumption rule such as those considered in this paper, we provide a dual methodology for deriving lower bounds on the value; of course, given an explicit policy, we could always find a lower bound on the value by directly simulating the paths of the controlled process.

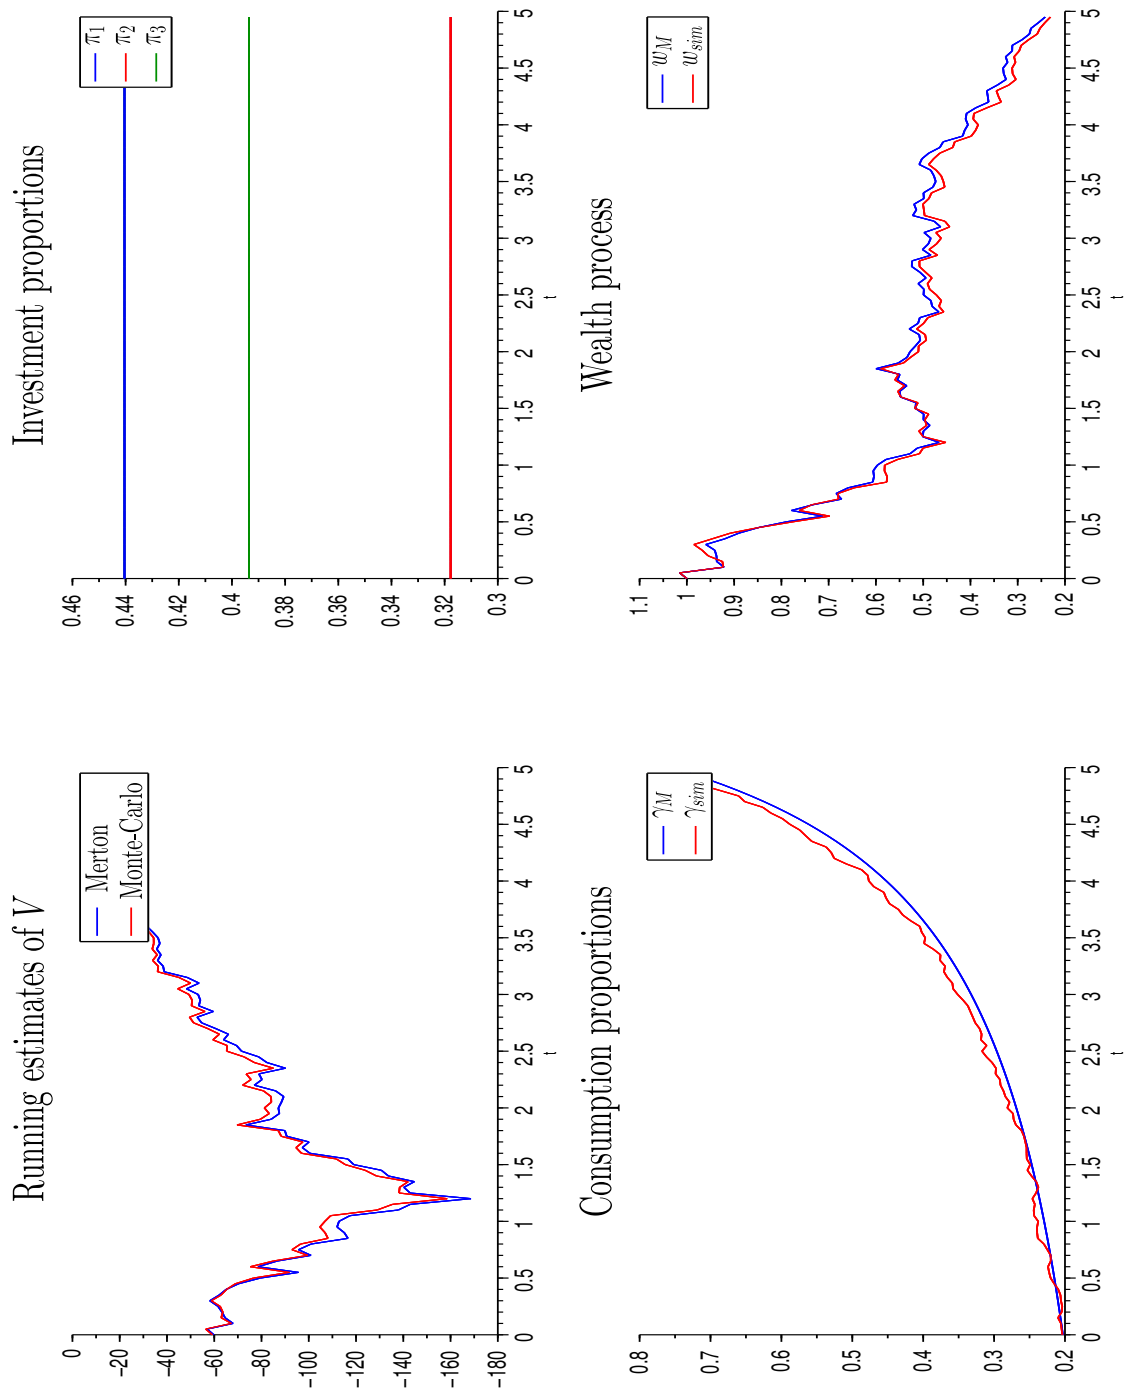


Figure 1: **Monte Carlo solution to the Merton problem.** Here we take the number of assets $k = 3$, $r = 0.05$, $\rho = 0.03$, $R = 3$, $A = 2$, $N = 100$, $dt = 0.05$, $M = 1000$, $w_0 = 1$, $\mu = [0.09; 0.26; 0.16]$, $\sigma = [0.12, 0.01, 0.03; 0.01, 0.45, 0.01; 0.03, 0.01, 0.27]$. Note that, since $\tilde{U}(t, z)$ for the Merton problem is a power function in z , the suggested proportions (2.27) will automatically match the Merton investment proportions.

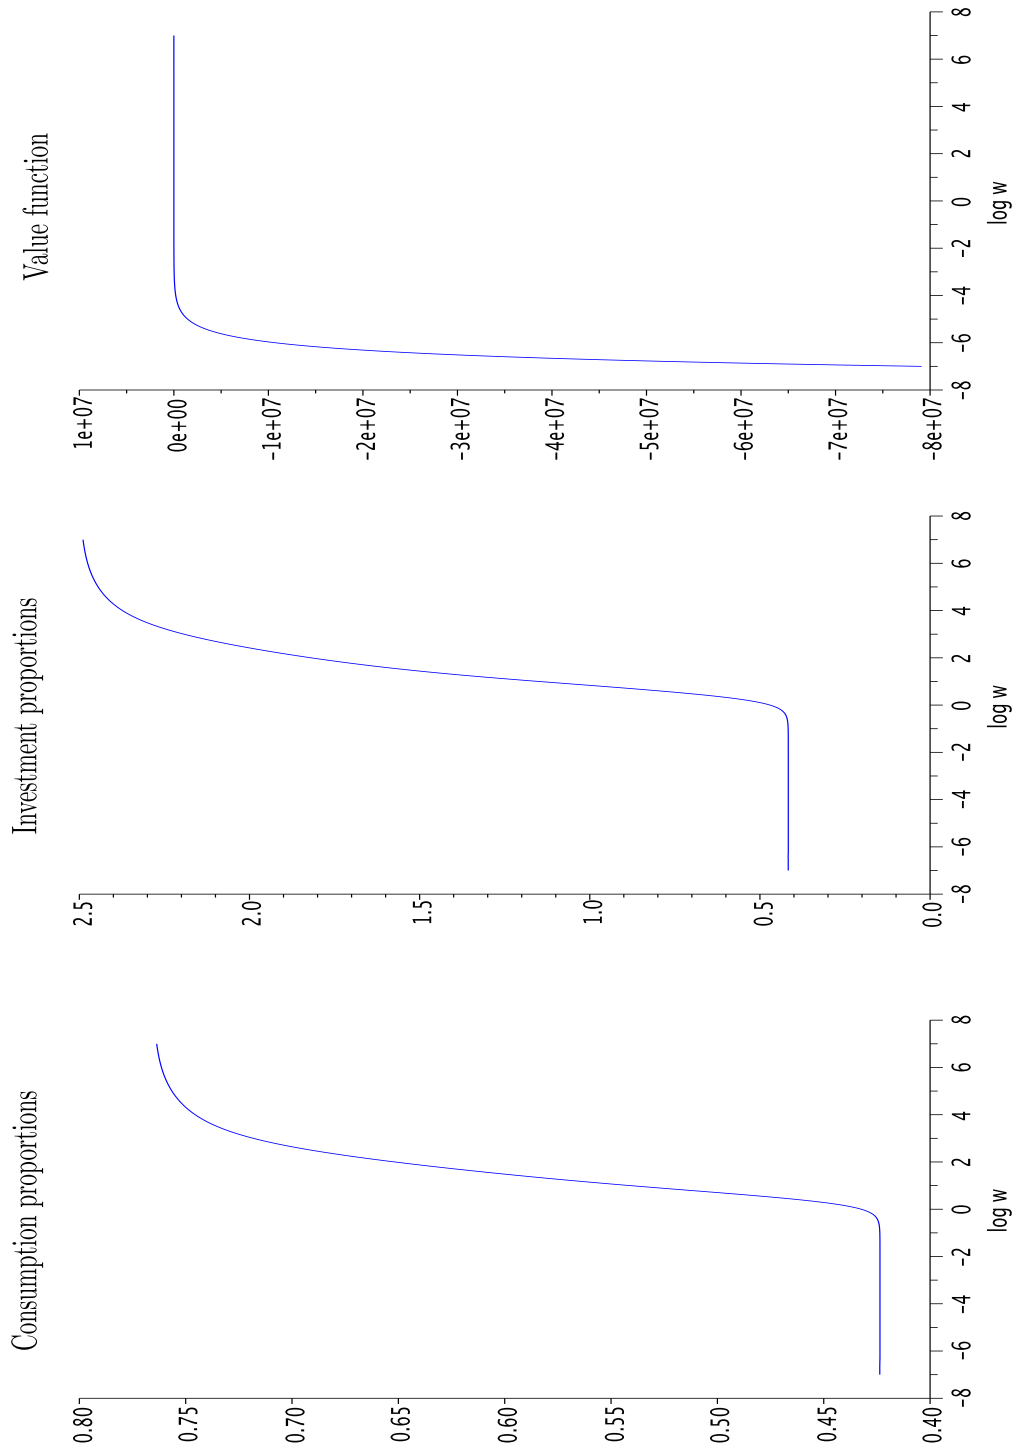


Figure 2: **Value function found using policy improvement algorithm.** Here we take $w_0 = 2$, $\mu = 0.10$, $\sigma = 0.20$, $r = 0.05$, $\rho = 0.03$, $a_1 = 10$, $a_2 = 20$, $b_1 = 30$, $b_2 = 10$, $R_1 = 3$, $R_2 = 0.5$, $T = 1$, $N = 100$.

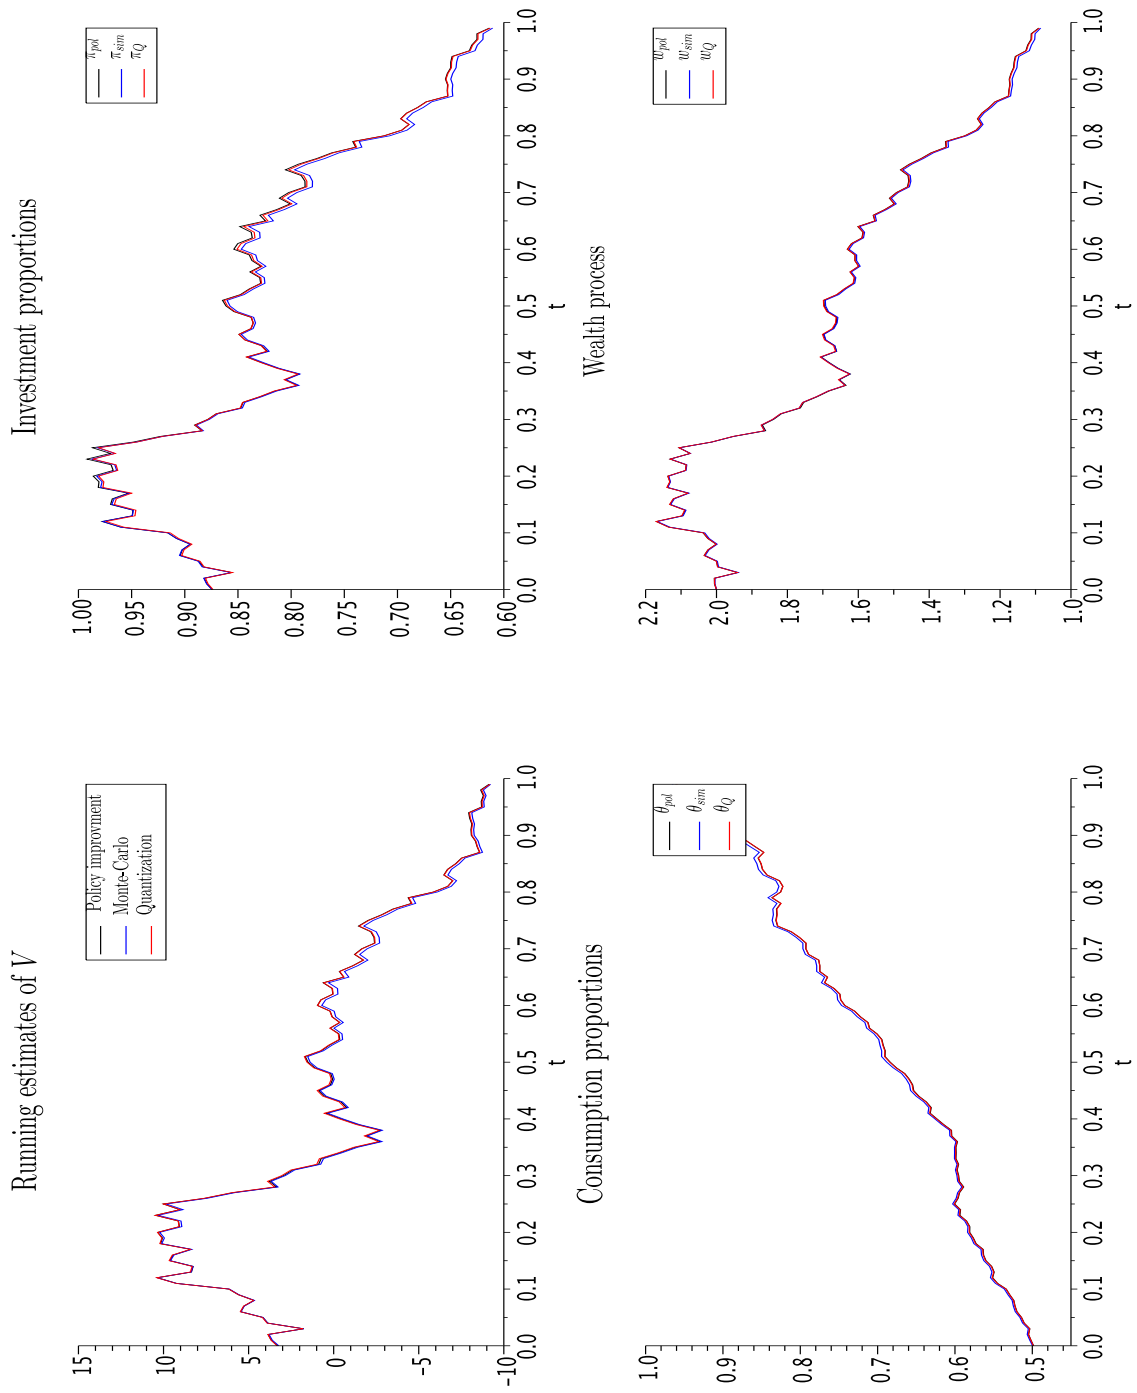


Figure 3: **Comparison of different methods for the non-constant relative risk aversion example.** Here we take $w_0 = 2$, $\mu = 0.10$, $\sigma = 0.20$, $r = 0.05$, $\rho = 0.03$, $a_1 = 10$, $a_2 = 20$, $b_1 = 30$, $b_2 = 10$, $R_1 = 3$, $R_2 = 0.5$, $T = 1$, $N = 100$. The number of Monte Carlo paths we took is $M = 10000$.

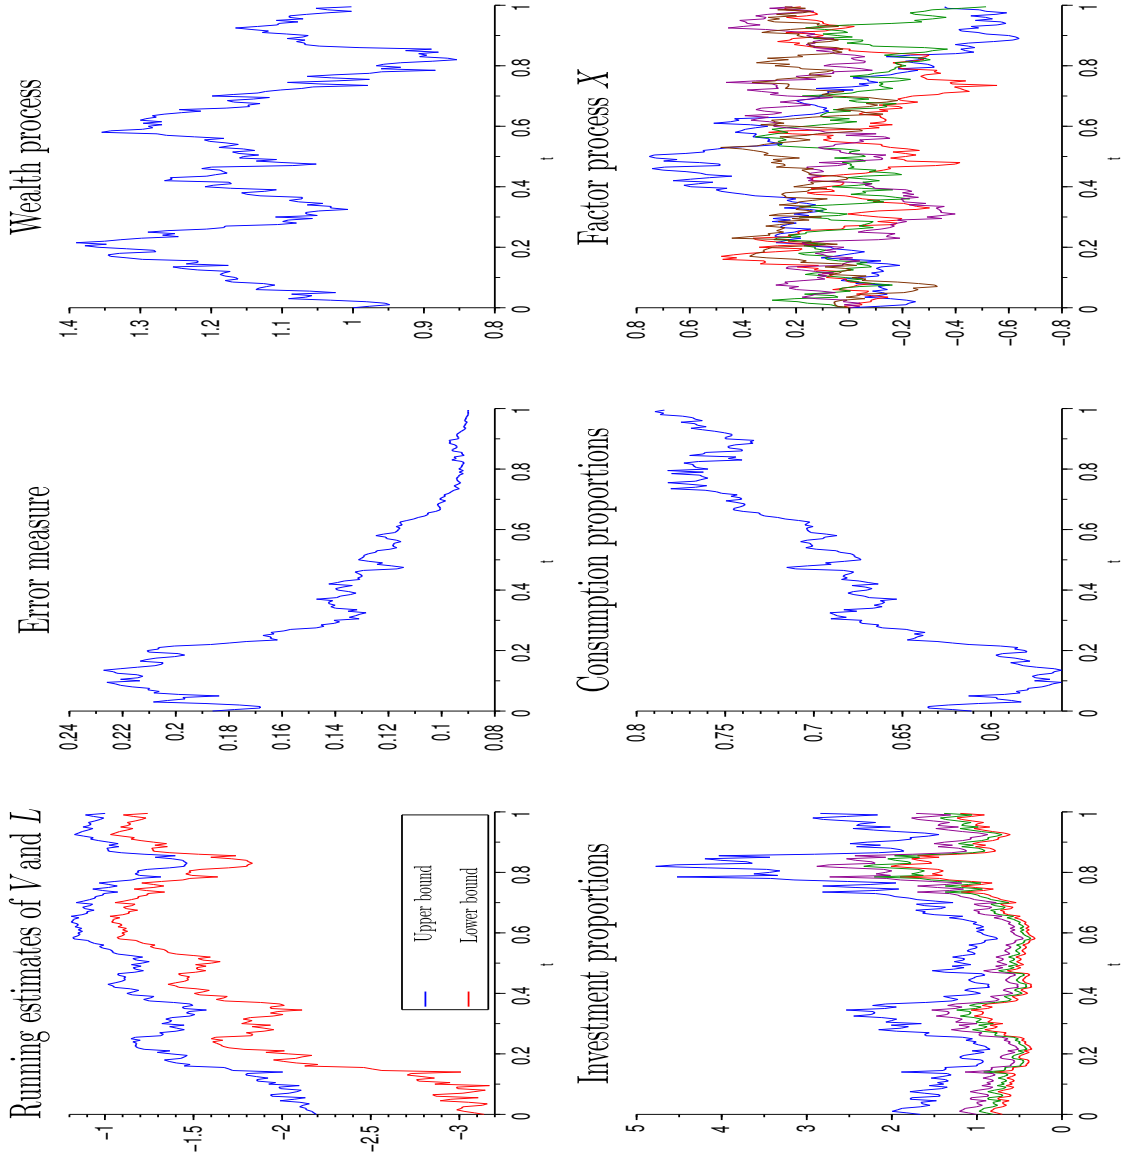


Figure 4: **An incomplete market driven by a stochastic factor.** Here we take an incomplete market 5-dimensional Brownian motion with 4 independent assets and $M = 1000$ Monte Carlo paths in the approximation of f . We start with $w_0 = 1$, $\rho = 0.03$, $T = 1$, $N = 100$. We take the CRRA utility function with parameters $A = 2$, $R = 3$. X is taken to be an OU process with randomly-generated volatility matrix and mean-reverting drift. The market interest rate r and μ are constant and randomly generated. Market volatility σ is a random 4×5 matrix multiplied by a stochastic scaling factor $1 + \exp(-\mathbf{1} \cdot X_t)$. The randomisation is done by drawing relevant parameters from $U[-1, 1]$ distribution until the regularity conditions imposed by the paper are met (i.e. $\mu \geq r \geq 0$, σ has rank 4 and $\sigma^T \sigma$ is invertible). The top panels represent the running upper and lower bounds on the objective as defined in (2.11) and (2.12), error rate as in (2.21), and the corresponding wealth process from (2.25). The bottom panel shows the investment and consumption proportions, and the factor process X .

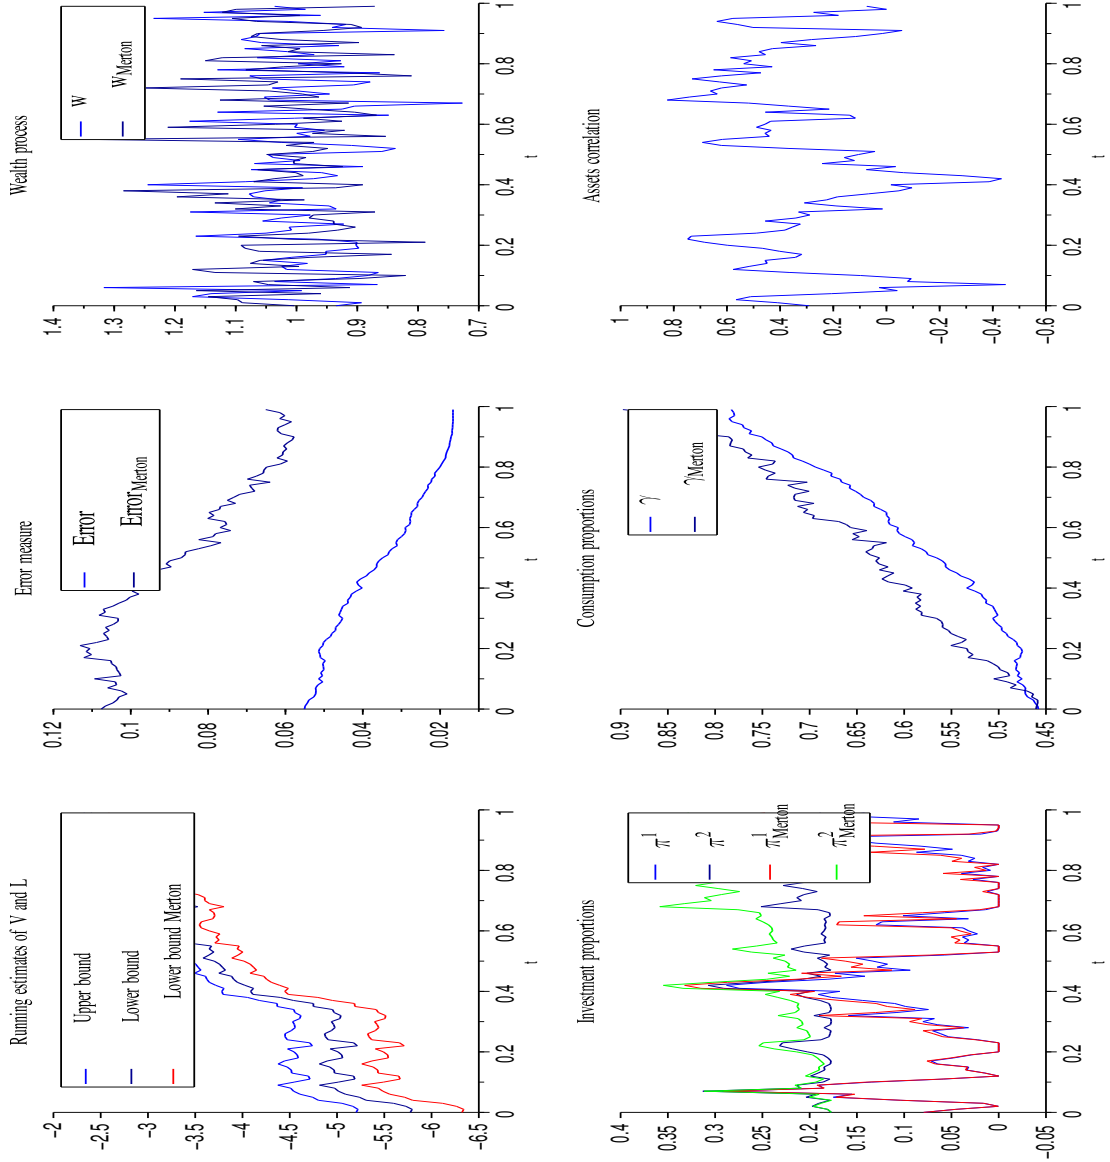


Figure 5: **Stochastic asset correlation and no short-selling constraints.** Here we take a market with two assets of constant volatilities, but whose correlation $\rho_t = \tanh(X_t + x)$ is stochastic, where X_t is a one-dimensional OU process with $dX_t = dW_t - \lambda X_t dt$. We also require that no short sales are allowed so that $\pi_t \geq 0$. The parameter values we took are $w_0 = 1$, $\rho = 0.03$, $T = 1$, $r = 0.06$, $A = 2$, $R = 3$, $x = 0.3$, $\mu_1 = 0.12$, $\mu_2 = 0.18$, $\sigma_1 = 0.35$, $\sigma_2 = 0.45$, $\lambda = 15$. The top panels represent the running upper and lower bounds on the objective as defined in (2.11) and (2.12) error rate as in (2.21), and the corresponding wealth process from (2.25). The bottom panel represents the investment and consumption proportions, and the resulting correlation process ρ_t . For comparison, we attach the corresponding heuristic strategy which attempts to follow the Merton strategy at time t , treating the coefficients as if they were constant.

References

- [1] BRANDT, M. W., GOYAL, A., SANTA-CLARA, P., AND STROUD, J. R. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Review of Financial Studies* 18, 3 (2005), 831–873.
- [2] BROWN, D. B., AND SMITH, J. E. Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Science* 57, 10 (2011), 1752–1770.
- [3] BROWN, D. B., SMITH, J. E., AND SUN, P. Information relaxations and duality in stochastic dynamic programs. *Operations Research* 58 (2010), 785–801.
- [4] CORLAY, S. *Some aspects of optimal quantization and applications to finance*. PhD thesis, Université Pierre et Marie Curie, 2011.
- [5] COX, J. C., AND HUANG, C.-F. Optimal consumption and portfolio policies when asset prices follow a diffusion process. *Journal of economic theory* 49, 1 (1989), 33–83.
- [6] CUOCO, D., AND LIU, H. A martingale characterization of consumption choices and hedging costs with margin requirements. *Mathematical Finance* 10 (2000), 355–385.
- [7] CVITANIĆ, J., AND KARATZAS, I. Convex duality in constrained portfolio optimization. *The Annals of Applied Probability* (1992), 767–818.
- [8] GARLAPPI, L., AND SKOULAKIS, G. Solving consumption and portfolio choice problems: The state variable decomposition method. *Review of Financial Studies* 23, 9 (2010), 3346–3400.
- [9] HAUGH, M. B., KOGAN, L., AND WANG, J. Evaluating portfolio policies: A duality approach. *Operations Research* 54 (2006), 405–418.
- [10] KARATZAS, I. Optimisation problems in the theory of continuous trading. *Control Optim* 27 (1989), 1221–1259.
- [11] KARATZAS, I., LEHOCZKY, J. P., AND SHREVE, S. E. Optimal portfolio and consumption decisions for a “small investor” on a finite horizon. *SIAM journal on control and optimization* 25, 6 (1987), 1557–1586.
- [12] KARATZAS, I., LEHOCZKY, J. P., SHREVE, S. E., AND XU, G.-L. Martingale and duality methods for utility maximization in an incomplete market. *SIAM Journal on Control and Optimization* 29 (1991), 702–730.
- [13] MERTON, R. Lifetime portfolio selection under uncertainty: the continuous-time model. *Rev. Econ. Statist.*, 51 (1969), 247–257.
- [14] MERTON, R. Optimum consumption and portfolio rules in a continuous-time model. *J. Econ. Theory*, 3 (1971), 373–413.

- [15] PAGÈS, G., AND PRINTEMS, J. Optimal quadratic quantization for numerics: the Gaussian case. *Monte Carlo Methods and Applications 9* (2003), 135–166.
- [16] PAGÈS, G., PRINTEMS, J., AND CORLAY, S. The optimal quantization web site. http://www.quantize.maths-fi.com/gaussian_process_database; accessed 20-March-2013.
- [17] ROGERS, L. C. G. Duality in constrained optimal investment and consumption problems: a synthesis. In *Paris-Princeton Lectures on Mathematical Finance 2002*. Springer, 2003, pp. 95–131.
- [18] ROGERS, L. C. G. Pathwise stochastic optimal control. *SIAM Journal on Control and Optimization* 46 (2007), 1116–1132.
- [19] ROGERS, L. C. G. *Optimal Investment*. Springer Briefs in Quantitative Finance. Springer-Verlag, 2013.
- [20] ROGERS, L. C. G., AND WILLIAMS, D. *Diffusions, Markov Processes and Martingales*, vol. 2. Cambridge University Press, 2000.