Evaluating the optimal solution to a portfolio problem with transaction costs[1]

L.C.G. Rogers [2] and R.H. Stockbridge [3]

Abstract. Finding the optimal investment/consumption policy in a problem with transactions costs is not an easy matter. In their seminal paper, Davis & Norman (1990) take a simple portfolio optimisation problem which reduces to one dimension; finding that there is no closed-form solution, they propose an effective numerical scheme for computing the solution to the HJB equation. We consider a variant of this problem, in which the diffusion parameters and the riskless rate depend on a finite Markov chain. In this case, the numerical scheme of Davis & Norman appears to be incapable of generalisation, as the HJB equation becomes a coupled system. Nonetheless, we shall show how approximations to the solution can be calculated using occupation measures for controlled Markov processes.

Key words. linear programming, stochastic control, proportional transaction costs, HARA utility.

Abbreviated Title. Portfolio Optimization with Transaction Costs

AMS(MOS) subject classifications.

1 Introduction

In this paper we consider the problem of selecting a portfolio to maximize the expected discounted utility of consumption

$$E\left[\int_0^\infty e^{-\rho t}U(C_t)dt\right] \tag{1.1}$$

where the portfolio invests in a riskless bank account B and a risky asset S which satisfy

$$\begin{aligned} dB_t &= r_t B_t dt \\ dS_t &= S_t(\mu_t dt + \sigma_t dW_t). \end{aligned}$$

Here, $U(c) = c^{1-R}/(1-R)$ for some $R > 0$ different[4] from 1, W is a standard Brownian motion, and $\rho > 0$. We shall assume that r, μ and σ are all functions of

[1]This research is partially supported by NSF grant DMS 9803490, EPSRC grant GR/L58415, and a University of Bath Visiting Fellowship.

[2]Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, England

[3]Department of Statistics, University of Kentucky, Lexington, KY 40506–0027

[4]The case $R = 1$ corresponds to logarithmic utility, and could be tackled by similar methods.

some underlying Markov chain ξ independent of W, with finite state-space J and generator Q. Thus

$$r_t = r(\xi_t), \quad \mu_t = \mu(\xi_t), \quad \sigma_t = \sigma(\xi_t).$$

Changes in the portfolio incur proportional transaction costs; to buy shares costing 1 we have to pay $1+\delta > 1$ from the bank account, and selling shares worth 1 will give us $1-\varepsilon < 1$ in cash. Thus the dynamics of the amount x_t in the bank account at time t and the value y_t of the agent's holding of shares at time t will satisfy

$$dx_t = [r_t x_t - C_t]dt - (1+\delta)dK_t + (1-\epsilon)dL_t \tag{1.2}$$

$$dy_t = \mu_t y_t dt + \sigma_t y_t dW_t + dK_t - dL_t, \tag{1.3}$$

where K_t denotes the cumulative amount moved into shares by time t, L_t denotes the cumulative amount moved out of shares by time t, and C_t denotes the rate of consumption at time t. If the underlying Markov chain ξ took only one value, then the problem we consider would be exactly that of Davis & Norman (1990); they exploit scaling properties to reduce the HJB equation to one dimension and find that the proportion of wealth in the bank account

$$p_t \equiv \frac{x_t}{x_t + y_t}$$

has to be kept within some interval I^*, and the optimal buying and selling is the minimum required to keep p in I^*. The optimal processes K and L are continuous and increasing, and singular with respect to Lebesgue measure, in the manner of a local time.

In view of the Davis-Norman solution, we can expect that for each $j \in J$ there will be some interval I_j^* such that while $\xi = j$ the optimally-controlled proportion will have to be kept within I_j^* by similar local-time buying and selling. If this conjecture turns out to be correct, then there will inevitably be jumps in the optimal K and L, which happen when the state of ξ changes from j to q at an instant when the proportion p is in I_j^* but not in I_q^*; an instantaneous jump to the nearest point of I_q^* will be made.

We are forced to look for a different method from that used by Davis & Norman, because the shooting method they use to solve the one-dimensional HJB equation in their problem will not generalise to the coupled system which arises in the problem which we study here.

This problem is a generalization of the investment/consumption models of Merton (1971), Davis & Norman (1990) and Zariphopoulou (1992). Our model allows for the irregularity of paths as did Merton and Davis & Norman and also allows the dynamics of the stock price process to vary as did Zariphopoulou so that "bull" and "bear" markets can be accommodated, as can changes in the interest rate and volatility.

An extensive survey of models with transaction costs is provided by Cadenillas (2000). Our paper differs from the prior work in that we adopt a linear programming approach to the optimization problem in which the variables consist of measures on a state and control product space. As such the linear program is infinite-dimensional. The linear programming formulation for stochastic control problems has been established by Bhatt and Borkar (1996) and Kurtz and Stockbridge (1998) for a wide variety of optimality criteria. It has also been used to evaluate both controlled and uncontrolled processes in many areas of application (see e.g. Helmes, Röhl and Stockbridge (2001), Helmes and Stockbridge (2000), Mendiondo and Stockbridge (1998) Haurie (200?), and Dempster and Hutton (1999). Similar linear programs have been used for discrete time Markov decision processes (see e.g. Hernández-Lerma and Lasserre (1994). Numerical evaluation of the solution requires approximation by finite dimensional linear programs. We adopt a discretization approach and approximate the diffusion process by a continuous-time Markov chain. This approximation scheme has been well-studied (see e.g. Kushner and Dupuis (1992)). Convergence of such approximate linear programs is established by Mendiondo and Stockbridge (1998).

Here is a guide to the remainder of the paper. In Section 2, we show how scaling properties can be used to reduce the apparently two-dimensional problem in terms of (x, y) to a one-dimensional problem. We are in effect establishing a skew-product decomposition of the bivariate process (x, y), akin to the well-known skew-product decomposition of a multidimensional Brownian motion (see, for example, ?????); this part of the paper is a relatively straightforward application of the methods of stochastic calculus, with one little twist. Next in Section 3 we take the reduced problem and approximate the process by a finite-state Markov chain. We then explain in Section 4 how the problem is to be solved by the calculation of the occupation measure, and how this reduces to a (large) linear program. This method is illustrated by numerical examples in Section 5, where we see how well it performs. Section 6 concludes the paper.

2 Reducing the problem.

The dynamics of the process (x, y),

$$dx_t = [r_t x_t - C_t]dt - (1+\delta)dK_t + (1-\epsilon)dL_t \qquad (2.1)$$

$$dy_t = \mu_t y_t dt + \sigma_t y_t dW_t + dK_t - dL_t, \qquad (2.2)$$

are not totally unrestricted, otherwise we could take C to be arbitrarily large and the optimisation problem would be ill-posed. Following Davis & Norman (1990), we make the natural restriction that at all times the process (x_t, y_t) should lie in the *solvency region*

$$\mathcal{S} \equiv \{(x, y) : x + (1-\varepsilon)y \geq 0, x + (1+\delta)y \geq 0\}. \qquad (2.3)$$

The set $\mathcal{S}$ contains points for which $y < 0$, that is, where the agent is short the share; but at such points $x \geq -(1+\delta)y$, so the agent has at least enough in the bank account to buy the shares he is short. A similar interpretation holds for points in $\mathcal{S}$ where $x < 0$.

Define the processes $w_t \equiv x_t + y_t$ and $p_t \equiv x_t/(x_t+y_t)$, together with the normalised consumption process $c_t \equiv C_t/w_t$, the normalised purchase rate $d\kappa_t \equiv w_t^{-1}dK_t$ and the normalised sale rate $d\lambda_t \equiv w_t^{-1}dL_t$. Then a little Itô calculus gives us

$$\begin{aligned} dp_t &= -p_t(1-p_t)\sigma_t dW_t + (1-p_t)\Big[(r_t-\mu_t)p_t - c_t + \sigma_t^2 p_t(1-p_t)\Big]dt \\ &\qquad -(p_t+\delta(1-p_t))d\kappa_t + (p_t - \varepsilon(1-p_t))d\lambda_t \qquad (2.4) \\ dw_t &= w_t\Big[(1-p_t)\sigma_t dW_t + \{(1-p_t)\mu_t + r_t p_t - c_t\}dt - \delta d\kappa_t - \varepsilon d\lambda_t\Big], \qquad (2.5) \end{aligned}$$

where $dk_t = w_t^{-1}dK_t$, $dl_t = w_t^{-1}dL_t$. The process w can thus be expressed explicitly in terms of p as

$$w_t/w_0 = \exp\Big[M_t + \int_0^t g_s ds - \delta\kappa_t - \varepsilon\lambda_t\Big], \qquad (2.6)$$

where $M_t \equiv \int_0^t (1-p_s)\sigma(\xi_s)dW_s$, and $g_t \equiv (1-p_t)\mu_t + r_t p_t - c_t - \frac{1}{2}\sigma_t^2(1-p_t)^2$. We may therefore rewrite the payoff of the agent as

$$\begin{aligned} E\Big[\int_0^\infty e^{-\rho t}U(C_t)dt\Big] &= E\Big[\int_0^\infty e^{-\rho t}\, w_t^{1-R}\, c_t^{1-R}\,\frac{dt}{1-R}\Big] \\ &= w_0^{1-R}\, E\Big[\int_0^\infty \exp\{-\rho t + (1-R)[M_t + \int_0^t g_s ds - \delta\kappa_t - \varepsilon\lambda_t]\}\; c_t^{1-R}\,\frac{dt}{1-R} \end{aligned} \qquad (2.7)$$

Now we can reinterpret the term $\exp\{(1-R)M_t\}$ in terms of a change of measure. Indeed, if we let

$$Z_t = \exp\{(1-R)M_t - \frac{1}{2}\int_0^t (1-R)^2(1-p_s)^2\sigma_s^2 ds\},$$

then using Z as the change-of-measure martingale (which takes us from probability P to probability $\hat{P}$) converts the Brownian motion into a drifting Brownian motion

$$dW_t = d\hat{W}_t + (1-R)(1-p_t)\sigma(\xi_t)dt$$

where $\hat{W}$ is a $\hat{P}$-Brownian motion. Thus we can rewrite the objective (1.1) as

$$w_0^{1-R}\hat{E}\Big[\int_0^\infty \exp\{\int_0^t b(p_s,\xi_s)ds - (1-R)\int_0^t c_s ds - (1-R)(\delta\kappa_t + \varepsilon\lambda_t)\}\; c_t^{1-R}\,\frac{dt}{1-R}\Big], \qquad (2.8)$$

where

$$b(p_t,\xi_t) = -\rho + (1-R)\Big[(1-p_t)\mu(\xi_t) + r(\xi_t)p_t - \frac{1}{2}R\sigma(\xi_t)^2(1-p_t)^2\Big].$$

Under the transformed measure, the equation satisfied by p becomes

$$\begin{aligned} dp_t &= -p_t(1-p_t)\sigma_t d\hat{W}_t + (1-p_t)\Big[(r_t-\mu_t)p_t + R\sigma_t^2 p_t(1-p_t)\Big]dt \\ &\quad -(1-p_t)c_t dt - (p_t+\delta(1-p_t))d\kappa_t + (p_t-\varepsilon(1-p_t))d\lambda_t \\ &\equiv \Sigma(p_t,\xi_t)d\hat{W}_t + \beta(p_t,\xi_t)dt - (1-p_t)c_t dt \\ &\quad -(p_t+\delta(1-p_t)))d\kappa_t + (p_t-\varepsilon(1-p_t))d\lambda_t, \end{aligned} \tag{2.9}$$

say. To summarise the situation: the dynamics of the problem are completely represented by (2.9), which is the dynamics of a controlled one-dimensional diffusion, controlled by the processes c, κ and λ, with drift and diffusion coefficients depending on the state of the independent Markov chain ξ; and the payoff (2.8) is a functional of the controlled diffusion, the controls, and the Markov chain.

3 Approximating the controlled diffusion.

The controlled diffusion process p takes values in [0,1], and we need to approximate it by a Markov chain Z which will take values in the set $\mathcal{X} \equiv \{x_0, x_1, \dots, x_N\} \subseteq [0,1]$, arranged in increasing order. To compute the jump intensities starting from some state x_i, $(0 < i < N)$, we consider the diffusion p (respectively, the chain Z) until the first time τ (respectively, τ') when it reaches $\{x_{i-1}, x_{i+1}\}$. We assume that the control being used and the state of the underlying Markov chain ξ remain unaltered until this stopping time, and we further assume that the diffusion is well approximated by a drifting Brownian motion

$$Y_t = x_i + vW_t + at,$$

where the drift and the variance are given by the values at x_i and the relevant control and ξ values. Applying the Optional Stopping Theorem to the two martingales

$$\begin{aligned} M_t &= Y_t - x_i - at \\ N_t &= M_t^2 - v^2 t \end{aligned}$$

at the stopping time τ, we have that

$$\begin{aligned} 0 &= E(Y_\tau - x_i - a\tau), \\ 0 &= E[(Y_\tau - x_i - a\tau)^2 - v^2\tau], \end{aligned}$$

so we pick the intensities with which Z jumps from x_i, at rate λ_i up to x_{i+1} and rate μ_i down to x_{i-1} so that the analogues of these two equalities hold:

$$0 = E(Z_{\tau'} - x_i - a\tau'), \tag{3.1}$$
$$0 = E[(Z_{\tau'} - x_i - a\tau')^2 - v^2\tau']. \tag{3.2}$$

It is easy to work out that these two equalities imply the following relations for λ_i and μ_i:

$$\frac{\lambda_i}{\lambda_i+\mu_i}(x_{i+1}-x_i)+\frac{\mu_i}{\lambda_i+\mu_i}(x_{i-1}-x_i) = \frac{a}{\lambda_i+\mu_i} \tag{3.3}$$

$$\frac{\lambda_i}{\lambda_i+\mu_i}(x_{i+1}-x_i)^2+\frac{\mu_i}{\lambda_i+\mu_i}(x_{i-1}-x_i)^2 = \frac{v^2}{\lambda_i+\mu_i} \tag{3.4}$$

These simultaneous linear equations are easily solved for λ_i and μ_i; writing $\alpha = x_{i+1}-x_i$ and $\beta = x_i - x_{i-1}$, we obtain

$$\lambda_i = \frac{v^2+\beta a}{\alpha(\alpha+\beta)}, \tag{3.5}$$

$$\mu_i = \frac{v^2-\alpha a}{\beta(\alpha+\beta)}. \tag{3.6}$$

These expressions may fail to be non-negative, but as α and β get smaller, the contribution in the numerators from the v^2 term eventually prevails. If one of (3.5) or (3.6) is negative, it is because of a large value of a. We deal with this by replacing the negative value by 0, and determining the other value from the equation (3.1). The effect of this is to alter the solutions (3.5) and (3.6) to

$$\lambda_i = \max\Big\{\frac{v^2+\beta a}{\alpha(\alpha+\beta)}, \frac{a}{\alpha}, 0\Big\} \tag{3.7}$$

$$\mu_i = \max\Big\{\frac{v^2-\alpha a}{\beta(\alpha+\beta)}, \frac{a}{-\beta}, 0\Big\}. \tag{3.8}$$

Finally, we shall make the end points x_0 and x_N absorbing. This is in fact a harmless assumption, since the optimal control will keep the process away from the endpoints.

The controls c, κ, and λ all take values in the interval $[0,\Lambda]$, and this interval must be discretised. We do this by allowing a finite mesh of c values, and by supposing that κ and λ take values in the set $\{0,\Lambda\}$.

To summarise the approximation, we build a Markov chain on the finite set $\mathcal{X}\times J$, with jumps from state (x_i,j) to (x_{i+1},j) and (x_{i-1},j) at rates (dependent on (c,κ,λ)) calculated as just described, and with jumps from (x_i,j) to (x_i,k) at rate q_{jk}.

4 The occupation measure method.

To explain the main idea of our approach, we suppose that we have a controlled finite-state Markov chain X, whose jump intensities depend on the current state, and on the current value of the control u, which also takes values in a finite set. The chain begins with law μ_0, and the aim of the controller is to

$$\max\ E^{\mu_0}\int_0^\infty \exp(-A_t)g(u_t,X_t)dt, \tag{4.1}$$

where

$$A_t \equiv \int_0^t f(u_s, X_s) ds.$$

ISSUES OF FINITENESS OF THE OBJECTIVE NEED TO BE DEALT WITH. If we let u^* denote the optimal control, and X^* the optimally-controlled process, we can define the occupation measure m^* of the optimally-controlled process via

$$\int\int \varphi(u,x) m^*(du,dx) \equiv E^{\mu_0} \int_0^\infty \exp(-A_t^*) \varphi(u_t^*, X_t^*) dt,$$

where $A_t^* \equiv \int_0^t f(u_s^*, X_s^*) ds$, and φ is any suitable test function. Knowledge of the measure m^* gives us full information on the solution, by taking the regular conditional m^*-distribution of u given x; what typically happens is that the optimal control is simply a function of the underlying state, so from m^* we can deduce what that function is. More generally, the optimal control may be a 'mixed' control in some (or all) states, and m^* will again tell us what the optimal mixture will be in such states. Thus we convert the (somewhat complicated) objective (4.1) into the (much simpler *linear*) objective

$$\int\int g(u,x) m(du,dx), \tag{4.2}$$

where now the measure m is to be the occupation measure for some controlled version of X. The question is therefore, 'How do we characterise such measures?' The answer turns out to be remarkably simple. Suppose we pick some feedback control γ, and let X^γ denote the controlled process. For a test function ψ, we write $\mathcal{G}^\gamma \psi(x) \equiv (\mathcal{G}\psi)(\gamma(x), x)$ as alternative notations for the action of the controlled generator on ψ.

Next we consider the martingale

$$\psi(X_t^\gamma) \exp(-A_t) - \int_0^t e^{-A_s} (\mathcal{G}\psi - f\psi)(\gamma(X_s^\gamma), X_s^\gamma) ds.$$

Assuming the first term goes to zero in L^1, the optional sampling theorem gives

$$\begin{aligned} -\int \psi(x) \mu_0(dx) &= E^{\mu_0} \int_0^\infty e^{-A_s} (\mathcal{G}\psi - f\psi)(\gamma(X_s^\gamma), X_s^\gamma) ds \\ &\equiv \int\int (\mathcal{G}\psi - f\psi)(u,x) m^\gamma(du,dx). \end{aligned} \tag{4.3}$$

The point of this is that the occupation measures m in terms of which the objective (4.2) is expressed can be characterised by the statement that *for all (suitable) test functions* ψ,

$$-\int \psi(x) \mu_0(dx) = \int\int (\mathcal{G}\psi - f\psi)(u,x) m(du,dx). \tag{4.4}$$

Thus we have a *linear programming problem:* the problem has the linear objective (4.2) (linear in the unknown m) constrained by the linear constraints (4.4). While it may appear that the size of the problem is going to be unmanageable in any real application (after all, the control will have to be discretised into finitely many possible values, and then the number of variables is the number of states of X times the number of discretised values of u, which is likely to get very big), it nevertheless turns out that one can often do quite well.

5 Numerical examples.

6 Conclusions.