An Approximation Algorithm for Optimal Consumption/Investment Problems*

Sanjiv Ranjan Das
Santa Clara University
Leavey School of Business
500 El Camino Real, Santa Clara, CA 95053
srdas@scu.edu

Rangarajan K. Sundaram
Stern School of Business
New York University
New York, NY 10012
rsundara@stern.nyu.edu

Versions: December 2001, May 2002.

*Special thanks to Luis Viceira and Raman Uppal for several sugestions. We are very grateful to the referees on the paper, as well as the editor for several useful comments which guided the paper into its improved form.

An Approximation Algorithm for Optimal Consumption/Investment Problems

Abstract

This article develops a simple approach to solving continuous-time portfolio choice problems. Portfolio problems for which no closed-form solutions are available may be handled by this technique, which substitutes the numerical solution of partial differential equations with a non-linear numerical algorithm approximating the solution. This paper complements the wide literature in economics on the solution of dynamic problems in discrete time using projection methods. Our approach extends the approximation function to power forms, which are shown to fit finance type problems well. The algorithm is parismonious, and is first illustrated by solving two basic examples, one, the standard Merton problem, and two, a jump-diffusion problem. Then, we demonstrate that the model is easy to implement on a larger scale, by optimizing a portfolio of 6 stock indexes, and stochastic volatility driven by two correlated state variables.

1 Introduction

The problem of optimal consumption and portfolio choice is one with a long history. Originally formulated in continuous time by Merton [30] [31], the problem has been extended substantially and several solution approaches have been developed. Barring the simplest problems, analytical solutions are difficult to come by. This paper provides a simple numerical approach to solving the optimal control problem using value function approximation.

For simpler problems, as in the original Merton formulation, closed-form solutions are achieved. The papers by Lehoczky, Sethi and Shreve [26], Karatzas, Lehoczky, Sethi and Shreve [23], Jacka [20], and Ocone and Karatzas [32] deal with explicit solutions, using the Bellman equation approach. The martingale approach of Cox and Huang [8] is also a well established one now, and has been extended to incomplete markets by He and Pearson [18] [19], Karatzas, Lehoczky, Shreve and Xu [24] and Cvitanic and Karatzas [9]. These problems are often further complicated by the choice of non-additive utility functions (see Duffie and Epstein [12], and more recently Dumas, Uppal and Wang [14]). Other complications arise when transactions costs are included in the analysis, as in Constantinides [7], Davis and Norman [10], and Dumas and Luciano [13].

In all these settings, simple versions admit either closed-form solutions or problems that are solved by applying simple numerical procedures. For example, in the case of the Bellman approach, if the number of state variables is low, the partial differential equation of optimality may be solved using finite-differencing methods. If the objective function is simple, then easily applied recursive methods may be used, as in Bertsimas, Kogan and Lo [1] where replication in incomplete markets is undertaken for a quadratic loss function, and the resulting system is quite tractable in low dimension. Bossaerts [2] examines a similar problem in an American option setting.

There are many approximation methods for the solution of these problems (see the review article by Taylor and Uhlig [35]). A wide variety of numerical approaches is applied such as iterating on the value function (Christiano [6]), quadrature methods (Tauchen [33] [34]), linear-quadratic approximations for the controls (Kydland and Prescott [25]), and parameterizations of the value function (Marcet [28], DeHaan and Marcet [11]). In this paper, we develop an analog of the parameterization approach in continuous-time, and demonstrate its application to a fairly general problem in continuous-time, that of a system driven by a mixed jump-diffusion stochastic process. Our model is a variant of known projection methods in the literature (see Judd [22]). More recent examples of approximation methods in the Finance literature include Campbell and Viceira [4], [5], Viceira [36], Brandt, Goyal and Santa-Clara [3], and Longstaff [27].

In continuous-time, solving the Bellman equation is more art than science. The usual approach involves making a clever guess as to the form of the value function, obtaining the optimal controls, and then verifying the solution after solving the Bellman PDE subject to the guess. Solutions have been obtained for some well-known and familiar utility functions, but whenever the number of state variables grows, or the stochastic processes chosen are not of the common geometric Brownian motion form, we are usually reliant on numerical schemes. This paper develops value function approximation as a method of extending the Bellman approach in a tractable way.

The basic idea is as follows. The optimal consumption-investment problem is set up as a Bellman control problem in the usual way. The first-order conditions provide the functional equations for the optimal controls, subject to solving for the value function. Rather than attempt to solve for the value function in closed form, we posit a very general polynomial form for the value function, extending the linear series forms that are prevalent in most models in computational economics. Thus the value function is described as a general function of a finite parameter set, denoted θ. Substituting this functional guess into the first-order conditions gives us the optimal controls as a function of θ. These optimal controls are then plugged back into the Bellman equation which should hold for all possible outcomes of the state variables. This will only be true when the guess for the value function coincides with the true value function, and for complex problems, this is unlikely. However, if the approximation to the true value function is a good one, then the distance between the approximate value function and the exact one should be small over all points in the state space. Thus, our solution comprises of minimizing a "distance function" between exact and approximate value functions, by means of finding the best-fit parameter set θ, subject to exactly satisfying the first-order conditions.

The algorithm works well, and we provide examples of its implementation in the paper. This approach has several benefits. First, it may be used to handle higher-dimensional problems, as well as problems with complex utility functions and stochastic processes. As an example, we solve a jump-diffusion model in the paper. We further extend the problem to one of higher-dimension, with solutions over a system of six assets and stochastic volatility driven by two state variables. Second, any standard minimizer routine may be applied for computational purposes making the problem computationally inexpensive. In fact, the illustrations in the paper use nothing more sophisticated than the optimizer in the Excel spreadsheet. All our implementations ran in under a minute, and were robust to different starting parameters. Third, general polynomial functions may be used to guess the value function. In other methods, such as weighted residuals methods, the approximation function is not nonlinear in the parameters, whereas our approach freely permits this. Fourth, as long as analytical derivatives of the approximated value function are obtainable, we are able to attain a rapid implementation. In other methods, such as finite-differencing to solve Bellman PDEs, numerical derivatives are taken, which impacts the speed and convergence of the algorithms used. Finally, examination of the numerical solution provides hints as to the form of the true value function, which may lead to an explicit analytical solution.

In the following sections, we provide the problem set up, and the formal presentation of the solution method. We present our approach within the didactic framework of projection methods. Numerical examples are also provided with appropriate discussion.

2 Stochastic Processes

Investors face a state space that is characterized by an infinite trading interval $T = [0, \infty)$. The uncertainty in the portfolio choice set emanates from a set of diffusion processes and Poisson jump processes, with probability spaces (Ω^Z, F^Z, Q^Z) and (Ω^N, F^N, Q^N) respectively. $\mathbf{Z_t} \in \mathbf{R^n}$ represents a vector of Wiener processes defined on (Ω^Z, F^Z, Q^Z) and $\mathbf{N_t} \in \mathbf{R^m}$ represents a vector of orthogonal jump processes defined on (Ω^N, F^N, Q^N), where $t \in T$ and $m, n > 0$. Each jump process is described by a sequence of random times $T^{(i)} \in T$ such that $N_{t,i} = \sum 1_{\{t \geq T^{(i)}\}}$. The Poisson jump arrival intensities are denoted $\lambda_i, i = 1...m$. We allow for the jump intensities to vary stochastically on T as well as functions of $\{\mathbf{Z_t}, \mathbf{N_t}\}$.

We also allow for K state variables $\mathbf{x_t} \in \mathbf{R^K}$ which also evolve on the same jump-diffusion probability spaces defined above. These processes are defined as follows:

$$d\mathbf{x}(t) = \alpha_x(\mathbf{x}, t)dt + \sigma_x(\mathbf{x}, t)d\mathbf{Z}_x(t) + \mathbf{J}_x(\mathbf{x}, t)d\mathbf{N}(\lambda, t) \tag{1}$$

where $\alpha_x(\mathbf{x}, t) \in R^K$, $\sigma_x(\mathbf{x}, t) \in R^{K \times n}$, $\mathbf{J}_x(\mathbf{x}, t) \in R^{K \times m}$.

The economy offers a set of $L + 1$ traded assets, whose values evolve stochastically on T. These assets comprise L risky assets $(\mathbf{S})$ indexed $l = 1...L$, and a riskless asset (B) which earns a constant return r. Hence the stochastic process for this asset is

$$dB(t) = rB(t)dt. \tag{2}$$

The remaining assets earn a random rate of return and obey the following SDE:

$$\frac{d\mathbf{S(t)}}{\mathbf{S(t)}} = \alpha(\mathbf{S}, \mathbf{x}, \mathbf{t})\mathbf{dt} + \sigma(\mathbf{S}, \mathbf{x}, \mathbf{t})\mathbf{dZ} + \mathbf{J}(\mathbf{S}, \mathbf{x}, \mathbf{t})\mathbf{dN}(\lambda, \mathbf{t}) \tag{3}$$

where $\frac{d\mathbf{S(t)}}{\mathbf{S(t)}}, \alpha(\mathbf{S}, \mathbf{t}) \in \mathbf{R^L}$, $\sigma(\mathbf{S}, \mathbf{t}) \in \mathbf{R^{L \times n}}$ and $\mathbf{J}(\mathbf{S}, \mathbf{t}) \in \mathbf{R^{L \times m}}$.

3 Optimal Portfolio Choice

The investor seeks to implement a consumption (c_t) and portfolio plan $(\mathbf{w_t} \in \mathbf{R^L})$ so as to maximize his lifetime utility. The utility of consumption is given by the usual class of Von-

Neumann and Morgenstern utility functions, which we denote $U(c_t)$, satisfying the usual requirements of concavity, and other technical regularity conditions.

The portfolio plan of the investor is a choice of asset weights $w_l, l = 1...L$ such that the amount invested in the riskless asset is $w_0 = 1 - \mathbf{w}'\mathbf{1}$. At any time t, the investor chooses how much of his current wealth W_t to consume, and invests the balance in the riskless and risky assets. Thus, the stochastic process for wealth taking into account investment and consumption is as follows (see Merton [29] for details):

$$dW = \{W\left[\mathbf{w}'(\alpha - r\mathbf{1}) + r\right] - c\}\, dt + W\mathbf{w}'\sigma d\mathbf{Z} + W\mathbf{w}'\mathbf{J}d\mathbf{N} \tag{4}$$

where $\mathbf{1}$ is a unit vector. At the initial time $t = 0$, we are interested in solving for the investor's optimal consumption and investment program, so as to undertake the following maximization:

$$\max_{\{c,\mathbf{w}\}} E_0 \left\{ \int_0^\infty e^{-\rho s} U(c_s) ds \right\} \tag{5}$$

where ρ is the investor's time preference parameter, a scalar constant. The optimized function at any time t is denoted as the value function, defined recursively as a function of state variables $\{W, \mathbf{x}\}$

$$V(W, \mathbf{x}; \mathbf{t}) = \max_{\mathbf{c_t}, \mathbf{W_t}} \mathbf{E_t} \left\{ \mathbf{U(c_t)} + \mathbf{V(W, x; t + dt)} \right\}. \tag{6}$$

Using the method of stochastic dynamic programming (see Merton [29]), we arrive at the Bellman equation of optimality

$$0 = \max_{\{c,\mathbf{w}\}} H(W, \mathbf{x}) \tag{7}$$

which in full detailed form is:

$$\begin{aligned} 0 \;=\; & \max_{\{c,\mathbf{w}\}} \{U(c) - \rho V + V_W W\left[\mathbf{w}'\left(\alpha - r\mathbf{1}\right) + r\right] - cV_W \\ & + \frac{1}{2} V_{WW} W^2 \mathbf{w}'\sigma\sigma'\mathbf{w} + \sum_{j=1}^{m} E\left(\lambda_j\left[V\left(W + \mathbf{w}'\mathbf{J}_j W, \mathbf{x}\right) - V(W, \mathbf{x})\right]\right) \\ & + V_{\mathbf{x}}'\alpha_x + \frac{1}{2} E[(\sigma_{\mathbf{x}} \mathbf{dZ_x})' V_{\mathbf{xx}} (\sigma_{\mathbf{x}} \mathbf{dZ_x})] + W\ V_{W\mathbf{x}}' E[(\sigma_x \mathbf{dZ_x})(\sigma \mathbf{dZ})'\mathbf{w}] \\ & + \sum_{j=1}^{m} E\left(\lambda_j\left[V(W, \mathbf{x} + \mathbf{J}_x) - V(W, \mathbf{x})\right]\right) \Bigg\} \end{aligned} \tag{8}$$

where $\mathbf{J_j} \in \mathbf{R^L}$ is the jth column of matrix $\mathbf{J} \in \mathbf{R^{L \times m}}$. Subscripts denote partial derivatives, i.e. $V_w = \frac{\partial V}{\partial W}$, $V_{WW} = \frac{\partial^2 V}{\partial W^2}$. Likewise, $V_{\mathbf{x}} = \frac{\partial V}{\partial \mathbf{x}} \in R^K$, $V_{\mathbf{xx}} = \frac{\partial^2 V}{\partial \mathbf{x} \partial \mathbf{x}'} \in R^{K \times K}$ and $V_{W\mathbf{x}} = \frac{\partial V}{\partial W \partial \mathbf{x}'} \in R^K$. The solution method entails taking the first-order derivatives from the equation of optimality to arrive at the optimal controls $\{c, \mathbf{w}\}$, as functions of $V(W, \mathbf{x})$. Then, the substitution of these values into the Bellman equation provides a second order $(K+1)$-dimensional partial differential equation in $(W, \mathbf{x})$ which must be solved subject to suitably imposed boundary conditions.

The first-order condition for consumption,

$$U'(c) = V_W \tag{9}$$

implies the optimal consumption rule: $c^* = I\left(\frac{\partial V}{\partial W}\right), I = [U']^{-1}, U' = \frac{\partial U}{\partial c}$. Taking the first-order condition for optimal portfolio weights $\mathbf{w}^*$, we get the $(L \times 1)$ equation system

$$\begin{aligned} 0 \;=\; & V_W\left(\alpha - r\mathbf{1}\right)W + V_{WW}\sigma\sigma'\mathbf{w}W^2 \\ & + \sum_{j=1}^{m} E\left(\lambda_j\left[\frac{\partial}{\partial \mathbf{w}}V(W + \mathbf{w}'\mathbf{J}_j W) - V(W)\mathbf{1}\right]\right) \\ & + W\{V_{Wx}' E[(\sigma_{\mathbf{x}}\mathbf{dZ_x})(\sigma\mathbf{dZ})']\}' \end{aligned} \tag{10}$$

where $\frac{\partial}{\partial \mathbf{w}}V(W + \mathbf{w}'\mathbf{J}_j W) \in R^L$ and $\mathbf{1}$ is a unit vector.

4 The Approximation Algorithm

Exact solution of the problem in Section 3 is usually hard to achieve, except in the simplest of cases. The problem lies in the fact that the optimal controls are complicated functions of the state variables, and value function, which itself is the solution to a high-dimensional differential equation. The usual approach is to guess a functional form for the value function and then verify whether it satisfies the optimality conditions of the problem. Apart from a few well-known cases, this approach has proven rather fruitless. Alternatively, one could attempt to solve the differential equation using numerical methods such as finite-differencing, but achieving a stable numerical scheme with many state variables has proven to be a daunting task.

Here, we suggest an alternative approach which bypasses these problems. The idea is to posit a polynomial function of the state variables as an approximation to the value function. We choose a θ-parameterized function $V(\theta) \equiv V(W, \mathbf{x}; \theta)$, where $\theta \in R^{P+1}$ is a set of parameters $\{v_0, v_1 ... v_P\}$ which define the value function. If we are able to find the "best" possible value function $V^*(\theta)$, then we have automatically obtained the solution to the problem, since the controls $\{c, \mathbf{w}\}$ derive immediately. The exact solution will satisfy

$$0 = H(W, \mathbf{x}; \mathbf{V}^*) \tag{11}$$

subject to satisfying the constraints from the first-order conditions, in equations (9) and (10). If we are not able to find the optimal value function, we can find the best V in a set of value functions $\{V(\theta)\}$ which may be chosen arbitrarily. To do this we compute the following optimization program:

$$\min_{\{V(\theta)\}} \left\{ \min_{\theta} \sum_{u \in U} f\left(H[W(u), \mathbf{x}(u); V(W, \mathbf{x}; \theta)]\right) \right\} \tag{12}$$

subject to

$$\begin{aligned} w &= w^* \left[V(W(u), \mathbf{x}(u); \theta) \right], & \forall u \\ c &= c^* \left[V(W(u), \mathbf{x}(u); \theta) \right], & \forall u. \end{aligned}$$

Here, U represents a discrete set of choices of state variables, i.e. a chosen state-space for the problem. These values $u = \{W(u), \mathbf{x}(u)\}$ may be chosen to reflect the decision-makers envisaged outcomes of the state variables. The function $f(.)$ is the fitting function and may be chosen from a range of popular options. For example, a least-squares approach would set $f(H) = H^2$. Alternatively, a probability weighted function such as $f(H) = H^2 \times prob(u)$ may be used. A simple absolute valued function $f(H) = |H|$ is also possible. Optimization is undertaken by choosing a specific $V(\theta)$ and then optimizing. Searching over the set $\{V(\theta)\}$will produce the best value function.

This approach has certain advantages. First, it does not require the solution of a high-dimensional differential equation. Second, the complexity of form of the value functional $V(\theta)$ does not impact substantially the computational requirements of the algorithm. The number of points in the state space U does however increase the number of constraints to be satisfied in a linear way. But this was not found to be numerically difficult, and in fact, implementation with a spreadsheet optimizer works very well. In both the simpler examples in the following section, and the more complex one in the ensuing section thereafter, we were able to solve the models on a simple Excel spreadsheet. The simple models barely took a few seconds to solve. The more complex model, involving search over 10 parameters in the value function, took about a minute and roughly 50 iterations of the Newton method. Since the Excel optimizer is likely to be inferior to those found in formal optimization software, alternative implementations using superior software will speed up the time to generate the results even more.[1]

4.1 Context and didactic representation of the algorithm

The class of methods in which our algorithm falls is denoted as "projection methods" (See Fletcher [16], Gaspar and Judd [17], Judd [21], [22]). These methods offer a viable alternative to finite-differencing techniques. Projection methods are closely related to regression techniques, and our model, using a least-squares criterion is intuitively resident in this set of algorithms. While finite-differencing models seek to numerically determine the exact function by building up from the numerical derivatives, projection methods posit a parameterized approximation to the true function, substitute this into the problem, and then search over the parameter space for the best fit solution function. Finite-differencing solutions promise more

[1] We fully expected to require a formal optimization package for the more complex problem we attempted. However, we decided to attempt it first in Excel, and as it turns out, this resulted in surprisingly effective implementation. This suggests that the structure of the problem and our implementation set up does enhance the ease of implementing the algorithm. Hence, we did not even require the use of more sophisticated software for the solution.

accuracy, since they do not approximate the value function, but pay a price in convergence problems and computational complexity. In contrast, the cost of functional approximation buys the projection approach much more flexibility as well as better computational complexity.

We provide a brief, yet general and simple exposition of a projection method. We seek to solve for a function (denoted $f(x)$), embedded in a equation $g[x, f(x), \theta] = 0$, where x proxies for a set of state variables, and θ is a set of parameters in the function $g[.]$. The function $g[.]$ may arise from an optimization problem. For example, in the framework of this paper, $g[.]$ is the Bellman equation. In another setting it may be a differential equation on the function $f(x)$ emanating from a physical or economic problem. In this sense, power series solutions to differential equations may also be interpreted as a special form of projection methods. Intuitively, notice that

$$g[x, f(x), \theta] = 0, \quad \forall x \tag{13}$$

may be interpreted as a vector of moment conditions. The exact function $f(x)$ may be replaced by an approximating function $f'(x;\xi)$, where ξ is a set of parameters for the approximating function. Then the approach is to solve the following problem akin to the method of moments:

$$E\left(g[x, f'(x;\xi), \theta].h(x)\right) = 0, \quad \forall x \tag{14}$$

where $h(x)$ is a set of weighting functions. Hence, projection methods may be classified by (a) the choice of polynomial for the approximation function $f'(x;\xi)$, and (b) the choice of weights used.

The usual form of the polynomial is

$$f'(x;\xi) = 1 + \sum_{i=1}^{n} \xi_i x^i. \tag{15}$$

This is the "linear" approximation form where the parameters apply linearly to the powers of the state variables. In this paper, we extend the specification to a more general form where the parameters enter as power coefficients of the state variables as well. Hence, for example, we specify:

$$f'(x;\xi,\delta) = \delta_0 + \sum_{i=1}^{n} \delta_i x^{\xi_i}. \tag{16}$$

The reason for this extension is simple. We noticed that the solutions of the Bellman equation in the few cases in which closed-form solutions were attained (for example, Merton [30], [31]) do possess such a function form for the value function. This form derives from the original power form of the direct utility function. Hence, we decided to adopt this as a modification to the classic polynomials used in projection methods. As will be shown, the approximation is well-founded. Other approximations may also be used, for example, the Galerkin method, which uses higher powers of x.

The other choice that is made in projection methods is the moment weighting function as applied to the "residuals". The classic choice is to set $h(.) = g(.)$, i.e. use a least-squares objective function. We adopt this criterion in our implementation of the algorithm. Many other weighting methods may be chosen as well. The objective function therefore, is usually a sum or integral of the residuals from the moment functions. Sometimes, parameters may be chosen to force the residuals to zero at a finite number of points on the state space grid, and this approach is called the collocation method.

Besides providing a new functional form for the value function, other points of difference are worth mentioning. In our setting, we are able to solve for the value function on a spreadsheet, even in complex settings as is shown in the final example in the paper. The problems are very nonlinear, yet yield a solution in a few seconds of computing time. And, just as in the finite-differencing model, we are able to provide the values of the controls, as well as the value function and all its derivatives at each point on the state-space grid. This allows generation of the plots for the solution in simple and easy way as will be seen in the figures provided for the last example in the paper. Even with a value function with as many as 10 parameters, the spreadsheet optimizer is able to converge easily in under a minute using the Newton's method with forward derivatives. Therefore, we did not face any severe root-finding problems in the implementation of the models. As will be seen from the plots, the value function is very smooth, and hence the properties of the problem itself would work in favor of this method. The following section provides illustrations of the model implementation.

5 Illustrative Examples

5.1 A simple implementation example

Consider the following simple problem. We begin with a single asset setting, where the asset follows a geometric Brownian motion.

$$dS = \alpha S dt + \sigma S dZ. \tag{17}$$

The notation follows from the previous section. Assume a power utility function over consumption where $U(c) = \frac{1}{\eta} c^{\eta}$. Analogous to equation (8) the Bellman equation of optimality will be:

$$0 = \max_{c,w} \left\{ U(c) - \rho V + V_W W[wR + r] - V_W c + \frac{1}{2} V_{WW} w^2 W^2 \sigma^2 \right\} \tag{18}$$

where $R = \alpha - r$ defines the equity premium/excess return. The first-order condition for consumption gives $U'(c) = V_W$ which implies that the optimal consumption is $c^* = (V_W)^{\frac{1}{\eta - 1}}$. Likewise, the first-order condition for the portfolio weights gives the optimal investment in the risky asset via the equation $V_W R + V_{WW} w W \sigma^2 = 0$, implying that $w^* = -\frac{R}{\sigma^2} \frac{V_W}{W V_{WW}}$. As is well known from Merton [29], the solution to this problem provides a value function of the form $V(W) = A \frac{W^{\eta}}{\eta}$, which implies that the optimal weights are $w^* = \frac{R}{\sigma^2} \frac{1}{(1-\eta)}$.

Now suppose we did not know the value function form in advance, and made a guess as to its form by choosing a somewhat more general function. Let $V(W) = V_0 + V_1 W^{V_2}$ where (V_0, V_1, V_2) are unknown scalar constants.[2] Our approach then entails substituting

[2]This approximation varies from other approaches using simple ploynomial forms, such as in the Galerkin method (see Judd [22], pg 373).

this posited value function into the Bellman equation and solving for the best fit values of (V_0, V_1, V_2) over the state space, which comprises a range of values of W.

Let the vector of N values of W be indexed by i, such that we have $W_i, i = 1...N$. This then defines a vector of values of the Bellman equation. Denote the optimized Bellman equation vector as $M_i, i = 1...N$. Thus, we have

$$M_i = U(c_i^*) - \rho V + V_W W_i[w_i^* R + r] - V_W c_i^* + \frac{1}{2} V_{WW} w_i^{*2} W_i^2 \sigma^2, \qquad \forall W_i. \tag{19}$$

Note here that the equation contains the optimized values (c_i^*, w_i^*) from the first-order conditions. Also, the derivatives V_W and V_{WW} are functions of W_i but the notation in the form $V_W(W_i)$ has been suppressed for expositional reasons.

Assume any unconditional distribution for W_i. Denote the probability of W_i as $f(W_i)$, and $\sum_{i=1}^{N} f(W_i) = 1$. For the optimal value function, it must be that $M_i = 0, \forall i$. However, since we are guessing the value function, and may not detect its exact form, the best that is possible is to choose (V_0, V_1, V_2) so as to make the values of $|M_i|, \forall i$ as small as possible. This suggests a range of objective functions of which an example is provided below.

$$\min_{\{V_0, V_1, V_2\}} \sum_{i=1}^{N} \left\{ M_i^2 f(W_i) \right\} \tag{20}$$

Under the assumption of equally weighted occurences of W_i, this is simply the least squares method.

We now explore the solution in more detail. First, we can compute the analytical derivatives of the value function. We get

$$\begin{aligned}
V_W &= V_1 V_2 W_i^{V_2 - 1} \\
V_{WW} &= V_1 V_2 (V_2 - 1) W_i^{V_2 - 2} \\
c_i^* &= [V_1 V_2 W_i^{V_2 - 1}]^{\frac{1}{\eta - 1}} \\
w_i^* &= \frac{R}{\sigma^2} \frac{1}{(1 - V_2)}.
\end{aligned}$$

The set of Bellman equations becomes (for every i)

$$M_i = \frac{1}{\eta} [V_1 V_2 W_i^{V_2 - 1}]^{\frac{\eta}{\eta - 1}} - \rho [V_0 + V_1 W^{V_2}]$$

$$
\begin{aligned}
&+V_1V_2W_i^{V_2-1}\left[\frac{R}{\sigma^2}\frac{1}{(1-V_2)}+r\right]\\
&-V_1V_2W_i^{V_2-1}[V_1V_2W_i^{V_2-1}]^{\frac{1}{\eta-1}}\\
&+\frac{1}{2}V_1V_2(V_2-1)W_i^{V_2-2}\left[\frac{R}{\sigma^2}\frac{1}{(1-V_2)}\right]^2 W_i^2\sigma^2, \qquad \forall i.
\end{aligned}
$$

The minimization problem may be simply stated as $\min_{\{V_0,V_1,V_2\}}\sum_{i=1}^N M_i^2$. Since we impose $V(W=0)=0$, we get that $V_0=0$. Further, it may easily be checked that the numerical minimization in fact leads to the solution $V_1>0, V_2=\eta$. This matches exactly the solutions in Merton [31]. Hence, the technique provides the known solution. We now proceed to a numerical illustration of a more complex model.

In this method, we need to specify the state-space grid over which the value function is computed. The first choice to be made is the range of the state variable. This is easy to determine in the case of consumption/investment problems, as we start our investor off with initial wealth $W_0=1$. Given this, the terminal wealth may be chosen to reside over a range from 0 to 2, where the upper support refers to a return of 100% per period, which is substantial. Once the overall support was chosen, we did not find the model solution to be too sensitive to the number of grid points used. The reason for this is that unlike other approaches to solve the PDEs emanating from Bellman equations, we do not need to compute derivatives numerically in our model, as we can compute them analytically from the approximated value function. For instance, in the finite-difference method to compute a grid over the value function, derivatives are taken numerically and hence are very sensitive to mesh size on the grid. In our case, we do not encounter this problem. Hence, mesh size is not an issue. This is an often overlooked benefit of the approach we develop here.

5.2 An example with jump processes

We extend the process driving the risky asset to including jumps with stochastic intensity. Thus,

$$dS = \alpha S dt + \sigma S dZ + JSdN(\lambda)$$

$$d\lambda \quad = \quad k(\theta - \lambda)dt + \delta\sqrt{\lambda}dY.$$

Thus, the jump intensity λ follows a mean-reverting square-root diffusion, and we assume that dZ, dY are orthogonal diffusions. The value function is now extended to cover the new state variable λ in addition to wealth, so that we write it as $V(W, \lambda)$. From (8) we get the Bellman equation:

$$\begin{aligned} 0 \quad = \quad & \max_{c,w} \Big\{ U(c) - \rho V(W,\lambda) + V_W W[wR + r] - V_W c + \frac{1}{2} V_{WW} w^2 W^2 \sigma^2 \\ & + V_\lambda k(\theta - \lambda) + \frac{1}{2} V_{\lambda\lambda} \delta^2 \lambda + \lambda E\left[V(W + wWJ, \lambda) - V(W,\lambda)\right] \Big\} \end{aligned}$$

We guess the following functional form for the value function

$$V(W, \lambda) = V_0 + V_1 W^{V_2} + V_3 \lambda^{V_4} + V_5 W \lambda \tag{21}$$

which yields the following terms

$$\begin{aligned} V_W &= V_1 V_2 W^{V_2 - 1} + V_5 \lambda \\ V_{WW} &= V_1 V_2 (V_2 - 1) W^{V_2 - 2} \\ V_\lambda &= V_3 V_4 \lambda^{V_4 - 1} + V_5 W \\ V_{\lambda\lambda} &= V_3 V_4 (V_4 - 1) \lambda^{V_4 - 2} \end{aligned}$$

From the first-order condition for consumption, we obtain the optimal value of c

$$c^* = [V_1 V_2 W^{V_2 - 1} + V_5 \lambda]^{\frac{1}{\eta - 1}}. \tag{22}$$

The first-order condition for the risky asset weights is

$$V_W W R + V_{WW} w \sigma^2 W^2 + \lambda E\left[\frac{\partial}{\partial w} V[W + wWJ]\right] = 0. \tag{23}$$

Thus, w^* is implicitly defined as the solution to the above equation. The last term requires

$$V[W + wWJ] = V_0 + V_1 W^{V_2}(1 + wJ)^{V_2} + V_3 \lambda^{V_4} + V_5 W(1 + wJ)\lambda, \tag{24}$$

which provides

$$\frac{\partial}{\partial w} V[W + wWJ] = V_1 V_2 W^{V_2}(1 + wJ)^{V_2 - 1} J + V_5 W \lambda J. \tag{25}$$

For the purposes of this example we assume a binary form for the jump, i.e

$$J = \begin{cases} +j, & \text{w/prob } \frac{1}{2} \\ -j, & \text{w/prob } \frac{1}{2} \end{cases} \tag{26}$$

which leads to

$$\begin{aligned}
E[V(W,\lambda)] &= V_0 + V_1 W^{V_2} + V_3\lambda^{V_4} + V_5 W\lambda \\
E\{V[W + wWJ, \lambda]\} &= V_0 + V_1 W^{V_2} \times \frac{1}{2}[(1+wj)^{V_2} + (1-wj)^{V_2}] \\
&\quad + V_3\lambda^{V_4} + V_5 W\lambda \\
E\left\{\frac{\partial}{\partial w} V[W + wWJ, \lambda]\right\} &= \frac{1}{2}[V_1 V_2 W^{V_2}(1+wj)^{V_2-1} j \\
&\quad - V_1 V_2 W^{V_2}(1-wj)^{V_2-1} j].
\end{aligned}$$

The first-order condition for portfolio weights may now be written as

$$\begin{aligned}
0 &= \left[V_1 V_2 W^{V_2-1} + V_5\lambda\right] WR + \left[V_1 V_2 (V_2 - 1) W_i^{V_2-2}\right] w\sigma^2 W^2 \\
&\quad + \lambda \frac{1}{2}[V_1 V_2 W^{V_2}(1+wj)^{V_2-1} j - V_1 V_2 W^{V_2}(1-wj)^{V_2-1} j].
\end{aligned} \tag{27}$$

The vector of Bellman equations is now written as (i now indexes the joint space over state variables W, λ):

$$\begin{aligned}
M_i &= U(c_i^*) - \rho V(W_i, \lambda_i) \\
&\quad + \left(V_1 V_2 W_i^{V_2-1} + V_5\lambda\right) W_i[w_i^* R + r] - \left(V_1 V_2 W_i^{V_2-1} + V_5\lambda\right) c_i^* \\
&\quad + \frac{1}{2}\left(V_1 V_2 (V_2 - 1) W_i^{V_2-2}\right) w_i^{*2} W_i^2 \sigma^2 \\
&\quad + \left(V_3 V_4 \lambda_i^{V_4-1} + V_5 W_i\right) k(\theta - \lambda_i) \\
&\quad + \frac{1}{2}\left(V_3 V_4 (V_4 - 1)\lambda_i^{V_4-2}\right)\delta^2\lambda_i \\
&\quad + \lambda_i\left\{V_0 + V_1 W_i^{V_2} \times \frac{1}{2}[(1+w_i^* j)^{V_2} + (1-w_i^* j)^{V_2}] + V_3\lambda_i^{V_4} + V_5 W_i\lambda_i\right\} \\
&\quad - \lambda_i\left\{V_0 + V_1 W_i^{V_2} + V_3\lambda_i^{V_4} + V_5 W_i\lambda_i\right\}.
\end{aligned}$$

Since the first-order condition for portfolio weights w is an implicit equation, we solve the following optimization problem to obtain the values $(V_0, V_1, V_2, V_3, V_4, V_5)$:

$$\min_{\{V_0,V_1,V_2,V_3,V_4,V_5\}} \sum_{i=1}^{N} M_i^2$$

subject to equations (22),(27).

This problem is solved numerically, and the results are provided below. Additionally, since the jump-diffusion based problem nests the pure-diffusion model, we can examine the features of the first solution from Section 5.1 as well.

5.3 Numerical results

The following parameters were chosen as a base case for the jump-diffusion model.

Parameter Description	Notation	Value
Relative risk aversion	η	0.5
Mean return on risky asset	α	0.07
Riskless rate	r	0.03
Subjective discount rate	ρ	0.03
Volatility coefficient for risky asset	σ	0.3
Mean reversion for jump intensity (λ)	k	0.5
Mean level of λ	θ	7.5
Volatility coefficient of λ	δ	5
Jump amplitude	j	0.1

In order to solve the problem we need to choose a grid of points $(W_i, \lambda_i), \forall i$. We used a range of values of $W \in [0, 10]$ and for jump intensity we assumed a two-state model where $\lambda \in \{5, 10\}$, i.e. low and high jump states. The algorithm was implemented on an Excel spreadsheet, and converges in a few seconds. The optimal values $(V_0, V_1, V_2, V_3, V_4, V_5)$ are:

Optimal Value Function Parameters (Jump-diffusion model)					
V_0	V_1	V_2	V_3	V_4	V_5
0	14.1290	0.4995	-1.0231	-0.0198	-0.0021

The signs of the parameters are exactly as expected. Note that V_1, V_2 are greater than zero, since indirect utility is increasing in the level of wealth. Likewise V_3, V_4, V_5 are less than zero, since utility declines when jump risk increases. For the purposes of comparison, we switched off the jump process to reduce the problem to the pure-diffusion model. To do so, we set the parameters as follows: $(k, \theta, \delta, j, \lambda)$ to zero. In this setting, the value function is:

Optimal Value Function Parameters (Pure-diffusion model)					
V_0	V_1	V_2	V_3	V_4	V_5
0	18.0889	0.5000	0	0	0

This solution corresponds exactly to that of the known solution in Merton [29]. Notice that the value of $V_2 = \eta$ as required in theory. In the following table we present some of the qualitative results from the two models, and undertake a comparison of outcomes.

Optimal Consumption and Investment Values									
	Pure-diffusion model			Jump-diffusion model					
	$\lambda = 0$			$\lambda = 5$			$\lambda = 10$		
W	c^*	w^*	$V(W, \lambda)$	c^*	w^*	$V(W, \lambda)$	c^*	w^*	$V(W, \lambda)$
0.1	0.0012	0.8888	5.72	0.0020	0.5702	3.24	0.0020	0.4200	3.25
0.9	0.0110	0.8888	17.16	0.0181	0.5697	12.16	0.0182	0.4192	12.17
2.5	0.0306	0.8888	28.60	0.0505	0.5691	21.07	0.0507	0.4184	21.06
4.5	0.0550	0.8888	38.37	0.0910	0.5687	28.67	0.0916	0.4177	28.64
10.1	0.1235	0.8888	57.48	0.2052	0.5678	43.52	0.2071	0.4164	43.42

The table presents results from the pure-diffusion model, and the jump-diffusion model. For varying levels of the state variables W and λ, we examine three values of interest: optimal consumption, investment in the risky asset, and the value function.

First, we note that as the level of wealth increases, so does the value function. Second, as jump intensity increases, investor utility decreases since additional risk is borne. The only exception occurs when wealth is at a very low level and the jump intensity increases from 5 to 10. This may be on account of the fact that at low levels of wealth, additional jumps cannot harm the investor given a floor level of zero on wealth. Plus, at high levels of jump intensity, mean reversion will lower jump risk. Third, as wealth increases, the investor consumes more. Fourth, as jump risk increases the investor also consumes more, since investing becomes less attractive, and consumption from the future is shifted to the present. Fifth, as jump risk increases, the investor correspondingly invests less in the risky asset. Sixth, when there is no jump risk, the amount invested in the risky asset is independent of wealth, as is known from the Merton model. However, when jump risk exists, the hedging term in equation (27) comes into play, and the choice of risky assets is no longer independent of wealth level. Finally, jump risk has a greater effect on the investor decision when jump risk is small and increasing than when it is large and increasing.

Computationally, the following results apply. First, the model converges to the same solution for many different starting values that we tried. We did not specially select any initial values. Second, the time taken for this optimization was always less than 10 seconds. Hence, the spreadsheet optimizer does not face any root-finding problems. Third, we examined the residuals from the Bellman equation. We did this both, in-sample and out-of-sample with reference to the state-space grid. While we ran wealth from 0-10 on the in-sample grid, we then tried it out-of-sample on the 10-12 region. The residuals statistics are as follows:

Residual Statistics

Sample	Mean	Std. deviation
In-sample	-1.58×10^{-7}	1.7×10^{-3}
Out-of-sample	2.4×10^{-3}	1.9×10^{-3}

As can be seen, the residuals are extremely small, and attest to an accurate solution. As is to be expected the in-sample residuals are smaller than those out-of-sample, but in all cases the residuals are not large. We also re-optimized the model with the additional out-of-sample points in the state space to check if the parameters of the value function changes drastically and found the changes to be small. The statistics for the residuals after this optimization are still very small in mean and standard deviation.

6 Higher Dimensional Problems

In the previous sections, we considered a problem where there was only one risky asset and a single state variable driven by a jump process. The approach we presented in the previous section is easily extendable to more securities and more state variables.

In order to demonstrate this, we extended the choice set of risky assets to six stock national indexes. The countries we considered are: United States, United Kingdom, Japan, Germany, Switzerland and France. We extracted index date for the period January 1982 to February 1997. From this data we computed the mean vector of returns (α), and the covariance matrix of returns ($\mathbf{\Sigma}$). We denoted this as the "base" covariance matrix. Uncertainty over the indexes is driven by a vector of diffusions ($\mathbf{dz}_s$).

To make the problem complex as well as realistic, we allowed the covariance matrix to vary from period to period, driven by two positive state variables ($\mathbf{x} = \{x_1, x_2\}$), which multiplicatively impact the covariance matrix in a linear way. Thus the actual covariance matrix was set to be equal to $\mathbf{\Sigma}(\mathbf{x}) = x_1 x_2 \mathbf{\Sigma}$. The two state variables are assumed to obey the following stochastic processes:

$$\mathbf{dx} = \alpha_x dt + \sigma_x \mathbf{dz}_x \tag{28}$$

where $\alpha_x \in R^2$ and $\mathbf{dz}_x \in R^2$. Also note that $\mathbf{dz}_x \mathbf{dz}'_x = \mathbf{\Omega} dt$, and $\mathbf{dz}_s, \mathbf{dz}_x$ are assumed to be orthogonal. The two factors driving volatility are an open choice for the purposes of calibration. Any two state variables could be utilized.

Let the value function be an extension of the prior simpler form, i.e. $V(W, x_1, x_2)$. If we denote $\mathbf{dy} = \sigma_{\mathbf{x}} \mathbf{dz}_{\mathbf{x}} \in R^2$, then we may write the Bellman equation over 6 assets and 2 stochastic volatility state variables as follows (same as before with additional terms):

$$\begin{aligned} 0 &= \max_{\mathbf{w},c} : U(c) - \rho V + V_W W[\mathbf{w}'(\alpha - r\mathbf{1}) + r] - cV_W \\ &+ \frac{1}{2} V_{WW} W^2 \mathbf{w}' \mathbf{\Sigma}(\mathbf{x}) \mathbf{w} \\ &+ \mathbf{V_x}' \alpha_{\mathbf{x}} + \frac{1}{2} E[\mathbf{dy}' \mathbf{V_{xx}} \mathbf{dy}] \end{aligned} \tag{29}$$

where $\mathbf{V_x}$ is the gradient matrix of the value function with respect to the state variables, and $\mathbf{V_{xx}}$ is the hessian matrix.

This is a much more complex problem than those presented previously, since the implementation requires manipulation of fairly sizeable matrices. However, it is still tractable on a spreadsheet. And, of course, the value function approximation is also extended to a 10 parameter form as follows:

$$V(W, x_1, x_2) = V_0 + V_1 W^{V_2} + V_3 x_1^{V_4} + V_5 x_2^{V_6} + V_7 W x_1 + V_8 W x_2 + V_9 x_1 x_2. \tag{30}$$

The approach taken is the same as before. We take the necessary derivatives of the value function, and substitute them into the Bellman function, and its first order conditions, computed over a grid containing all three state variables. Unlike in the past, when computation on the spreadsheet only required a few seconds, the size of this problem required about 50 iterations in Excel, and took up to a minute in computational time. However, this is still a computationally economical procedure.[3]

[3] We used the optimizer in Microsoft excel over a range of values of W from 0 to 2.0, values of x_1, x_2 were taken to be 0.5, 1.0 and 1.5. The settings of the optimizer were as follows: convergence criterion was 0.0001, the precision setting was at 0.000001, the maximum number of iterations set to 500 (never reached), tolerance at 5%. The estimates were based on the tangent method, with forward derivatives. Search was implemented using Newton's method. Finally, the results and time did not vary much at all despite a wide range of choices of starting parameters, even naive choices such as setting the initial values to all be 0.1 or 0.5, for example.

6.1 Numerical results for higher-dimensional problem

In this section we present the results for the extended problem outlined above. The means for the six country stock index (monthly) returns were as follows:

US	UK	JP	GE	SW	FR
0.0102	0.0084	0.0080	0.0120	0.0102	0.0113

The "base" covariance matrix of monthly returns is:

US	UK	JP	GE	SW	FR
0.0018	0.0014	0.0010	0.0010	0.0012	0.0013
0.0014	0.0032	0.0018	0.0016	0.0016	0.0019
0.0010	0.0018	0.0048	0.0015	0.0016	0.0020
0.0010	0.0016	0.0015	0.0033	0.0020	0.0023
0.0012	0.0016	0.0016	0.0020	0.0026	0.0019
0.0013	0.0019	0.0020	0.0023	0.0019	0.0037

The state variables for stochastic volatility are assumed to evolve as follows:

$$\begin{pmatrix} dx_1 \\ dx_2 \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} + \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{pmatrix} \begin{pmatrix} dz_1 \\ dz_2 \end{pmatrix} \tag{31}$$

and $dz_1 dz_2 = Corr(dz_1, dz_2)dt$. The remaining input parameters were as follows:

Parameters	η	ρ	r	α_1	α_2	σ_{11}	σ_{12}	σ_{21}	σ_{22}	$Corr(dz_1, dz_2)$
	0.5	0.005	0.005	0.2	0.2	0.1	0.1	0.1	0.1	0.5

We ran the algorithm on this problem and obtained the ten parameters of the value function. These are as follows:

V_0	V_1	V_2	V_3	V_4	V_5	V_6	V_7	V_8	V_9
10.1137	3.4771	0.4281	23.1806	0.0538	23.1817	0.0538	0.6441	0.6441	-17.1148

The results of this model are presented using plots, as this will make for clearer exposition. To start with, we examine the resulting value function from the model. The value function $V(W, x_1, x_2)$ is hard to plot in four dimensions. Hence, we decided to present it for varying values of W, and for values of the product $(x_1 \times x_2)$, which is the term that enters the stochastic volatility process. Our graph is two-dimensional, with terminal wealth W on the x-axis, and the value function on the y-axis. Figure 1 presents the plots of the results. We plot different lines for varying values of the volatility factors. First, notice that the value function (i.e. indirect utility) is increasing in wealth as expected. Second, the numerically computed indirect utility function derives its shape from the original direct one, and is therefore concave and increasing monotonically. Third, as the volatility of the assets increases, indirect utility drops as one would expect. However, in the presence of risk-aversion, the growing volatility factor results in greater percentage reductions in utility.

We also present (in Figure 2) the plots depicting optimal consumption as terminal wealth changes and for different levels of the volatility state variables. The relationship is as expected. As wealth increases, consumption also increases but at a decreasing rate. As volatility increases, amount of consumption declines as the amount invested shifts to provide a greater investment buffer in ensuing risky periods.

7 Concluding Comments

This paper develops a simple numerical approach to solving optimal consumption and investment problems when analytic solutions are not achievable. The Bellman problem is translated into an econometric one, where we minimize a parameterized objective function of the Bellman equation over the state space, based on a polynomial guess for the value function. The method is tractable, and converges quickly. The approach offers a means to numerically guessing the form of the value function which, if fortuitous, may lead to an analytic solution.

The algorithm may be extended to solving problems in other domains in finance, such as asset-liability management, optimal replication of derivative securities, equilibrium problems with incomplete markets and market micostructure games.

References

[1] BERTSIMAS, D., L. KOGAN, and A. LO (1997) "Pricing and Hedging Derivative Securities in Incomplete Markets: An $\in$ - Arbitrage Approach," working paper No. 6250, NBER.

[2] BOSSAERTS, P. (1989) "Simulation Estimators of Optimal Early Exercise," working paper, Carnegie Mellon University.

[3] BRANDT, M., A. GOYAL, and P. SANTA-CLARA (2000). "Dynamic Portfolio Choice with Simulation Methods," working paper, UCLA.

[4] CAMPBELL, J., and L. VICEIRA (1996) "Consumption and Portfolio Decisions when Expected returns are Time Varying," NBER working paper no. 5857.

[5] CAMPBELL, J., and L. VICEIRA (1996) "Who should buy long-term bonds?" unpublished paper, Harvard University.

[6] CHRISTIANO, L.J. (1990) "Solving the Stochastic Growth Model by Linear-Quadratic Approximation and by Value-Function Iteration," *Journal of Business and Economic Statistics*, v8(1), 23-26.

[7] CONSTANTINIDES, G. (1986) "Capital Market Equilibrium with Transactions Costs," *Journal of Political Economy*, v94, 842-862.

[8] COX, D., and C-F. HUANG (1989) "Optimal Consumption and Portfolio Policies when Asset Prices follow a Diffusion Process," *Journal of Economic Theory*, v49, 33-83.

[9] CVITANIC, J., and I. KARATZAS (1991) "Convex Duality in Constrained Portfolio Optimization," working paper, Columbia University.

[10] DAVIS, M., and A. NORMAN (1990) "Portfolio Selection with Transaction Costs," *Mathematics of Operations Research*, v15, 676-713.

[11] DEN HAAN, W.J., and A. MARCET (1990) "Solving the Stochastic Growth Model by Parameterizing Expectations," *Journal of Business and Economic Statistics*, v8(1), 31-34.

[12] DUFFIE, D., and L. EPSTEIN (1992) "Stochastic Differential Utility," *Econometrica*, v60, 353-394.

[13] DUMAS, B., and E. LUCIANO (1989) "An Exact Solution to a Dynamic Portfolio Choice Problem with Transations Costs," working paper, Wharton School, University of Pennsylvania.

[14] DUMAS, B., R. UPPAL and T. WANG (1997) "Efficient Intertemporal Allocations with Recursive Utility," working paper, HEC School of Management, and University of British Columbia.

[15] FLEMING, W., and R. RISHEL (1975) *Deterministic and Stochastic Optimal Control*, Berlin: Springer-Verlag.

[16] FLETCHER, C.A.J (1984). "Computational Galerkin Techniques," Springer-Verlag, New York.

[17] GASPAR, J., and K. JUDD (1997). "Solving Large-Scale Rational Expectations Models," *Macroeconomic Dynamics*, v1, 45-75.

[18] HE, H., and N. PEARSON. (1991) "Consumption and Portfolio Policies with Incomplete Markets: The Infinite-Dimensional Case," *Journal of Economic Theory*, v54, 259-305.

[19] HE, H., and N. PEARSON. (1991) "Consumption and Portfolio Policies with Incomplete Markets: The Finite-Dimensional Case," *Mathematical Finance*, v1, 1-10.

[20] JACKA, S. (1984) "Optimal Consumption of an Investment," *Stochastics*, v13, 45-60.

[21] JUDD, K. (1992). "Projection Methods for Solving Aggregate Growth Models," *Journal of Economic Theory*, v58, 410-452.

[22] JUDD, K. (1998) "Numerical Methods in Economics," MIT Press, Cambridge, MA.

[23] KARATZAS, I., J. LEHOCZKY, S. SETHI and S. SHREVE (1986) "Explicit Solution of a General Consumption/Investment Problem," *Mathematics of Operations Research*, v11, 261-294.

[24] KARATZAS, I., J. LEHOCZKY, S. SHREVE and G-L. XU (1991) "Martingale and Duality Methods for Utility Maximization in Incomplete Markets," *SIAM Journal of Control and Optimization*, v29, 702-730.

[25] KYDLAND, F., and E. PRESCOTT (1982) "Time to Build and Aggregate Flucutations," *Econometrica*, v50, 1345-1370.

[26] LEHOCZKY, J., S. SETHI, and S. SHREVE (1983) "Optimal Consumption and Investment Policies allowing Consumption Constraints and Bankruptcy," *Mathematics of Operations Research*, v8, 613-636.

[27] LONGSTAFF, F (2000) "Optimal Portfolio Choice and the Valuation of Illiquid Securities," forthcoming *Review of Financial Studies.*

[28] MARCET, A. (1988) "Solving Non-Linear Models by Parameterizing Expectations," working paper, Graduate School of Industrial Administration, Carnegie-Mellon University.

[29] MERTON, R.C. (1992) "Continuous-Time Finance," Basil Blackwell, New York.

[30] MERTON, R.C. (1969) "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," *Review of Economics and Statistics*, v51, 247-257.

[31] MERTON, R.C. (1971) "Optimum Consumption and Portfolio Rules in a Continuous Time Model," *Journal of Economic Theory*, v3, 373-413.

[32] OCONE, D., and I. KARATZAS (1991) "A Generalized Clark Representation Formula, with Application to Optimal Portfolios," *Stochastics and Stochastics Reports*, v34, 187-220.

[33] TAUCHEN, G. (1987) "Quadrature-Based Methods for Obtaining Approximate Solutions to Non-linear Asset Pricing Models," working paper, Duke University, Dept of Economics.

[34] TAUCHEN, G. (1990) "Solving the Stochastic Growth Model by using Quadrature Methods and Value-Function Iterations," Journal of Business and Economic Statistics, v8(1), 49-51.

[35] TAYLOR, J.B., and H. UHLIG (1990) "Solving Non-Linear Stochastic Growth Models: A Comparison of Alternative Solution Methods," *Journal of Business and Economic Statistics*, v8(1), 1-17.

[36] VICEIRA, L. (1997) "Optimal Portfolio Choice for Long-horizon Investors with Non-tradeable Income," unpublished paper, Harvard University.

Figure 1: **The Value Function**

The value function $V(W, x_1, x_2)$ is presented for values of W, and for values of the product $(x_1 \times x_2)$, which is the term that enters the stochastic volatility process. Our graph is two-dimensional, with terminal wealth W on the x-axis, and the value function on the y-axis.

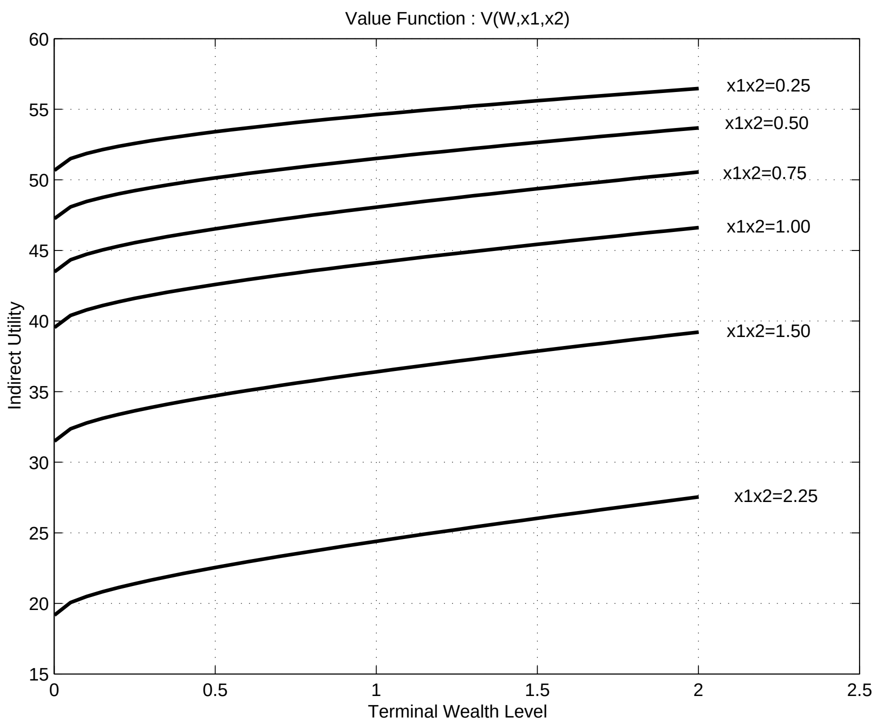

Figure 2: **The Consumption Function**

The consumption function $c(W, x_1, x_2)$ is presented for values of W, and for values of the state variables, (x_1, x_2), which are the terms that enter the stochastic volatility process. Our graph is two-dimensional, with terminal wealth W on the x-axis, and the value function on the y-axis.