The Markov Chain Approximation Approach for Numerical Solution of Stochastic Control Problems: Experiences from Merton's Problem*

Claus Munk†

January 1998

Abstract. Many problems in modern financial economics involve the solution of continuous-time, continuous-state stochastic control problems. Since explicit solutions of such problems are extremely rare, efficient numerical methods are called for. The Markov chain approximation approach provides a class of methods that are simple to understand and implement. In this paper, we compare the performance of different variations of the approach on a problem with a well-known explicit solution, namely Merton's consumption/portfolio problem. We suggest a variant of the method, which outperforms the known variants, at least when applied to this specific problem. We document that the size of the contraction parameter of the control problem is of great importance for the accuracy of the numerical results. We also demonstrate that the Richardson extrapolation technique can improve accuracy significantly.

Keywords. Stochastic control, efficient numerical solution, Merton's consumption/portfolio problem.

JEL classification. C61, G11.

---

*Financial support from the Danish Research Councils for Natural and Social Sciences is gratefully acknowledged.

†Assistant Professor, Department of Management, Odense University, Campusvej 55, DK-5230 Odense M, Denmark. Phone: +45 6557 3257. Fax: +45 6593 0726. E-mail: `cmu@busieco.ou.dk`. Internet homepage: `http://www.busieco.ou.dk/dept/man/faculty/cmu.htm`

1 Introduction

Many problems in financial economics involve the solution of continuous-time, continuous-state stochastic control problems. The most prominent example is the optimal consumption/investment problem of a utility maximizing investor having access to continuous trade in assets, whose prices follow diffusion processes. Several simple versions of this problem was solved explicitly by Merton (1969, 1971), but explicit solutions of this type of stochastic control problems are rare. Therefore, efficient numerical solution techniques are needed.

While several types of numerical approaches have been suggested – see Tapiero and Sulem (1994) for an overview – the Markov chain approximation approach constitutes a particularly simple and comprehensible class of methods. The basic idea of the approach is to approximate the diffusion state variables of the continuous-time, continuous-state control problem with a discrete-time, discrete-state Markov chain. The original objective function for controlling the diffusion variables is then approximated by an objective function for controlling the Markov chain in such a way that the solution of the Markov chain control problem is easily computable and converges (in some sense) to the solution of the original controlled diffusion problem.

In this paper, we shall implement various versions of the Markov chain approximation approach on Merton's simple consumption/investment problem, which as already noted has a known, explicit solution. The purpose is to study the convergence properties and numerical efficiency of the approach, and to compare the performance of the different versions. We find that all the methods efficiently produce very accurate approximations to both the value function and the optimal controls of Merton's problem. We suggest a version of the Markov chain approximation approach, which outperforms the known versions, at least when applied to Merton's problem.

Since the state space of the original continuous-time, continuous-state control problem is unbounded from above, it is necessary to impose an upper bound on the state space for the approximating Markov chain. We find that the numerical results for all four versions of the method can be very inaccurate in a neighborhood of this upper bound, and that the size of this neighborhood is highly dependent on the contraction parameter in the dynamic programming equation, which, for Merton's problem, is identical to the time preference rate of the investor.

Finally, we demonstrate that Richardson extrapolation techniques can accelerate convergence significantly. Although well-known in the numerical analysis literature, not much attention has been paid to Richardson extrapolation in published work on numerical finance.[1]

The rest of the paper is organized in the following way. In Section 2, Merton's problem and

[1] Prominent exceptions are the American option pricing studies by Geske and Johnson (1984) and Bunch and Johnson (1992).

its solution are reviewed. Section 3 outlines the ideas of the Markov chain approximation approach and discusses the four versions we consider, and how they can be implemented on Merton's problem. In Section 4, we first study the experimental convergence properties of the method. Next, we compare the efficiency of the four versions. We show how the value of the contraction parameter influences the numerical performance of the methods, and discuss the use of Richardson extrapolation techniques. Some concluding remarks are given in Section 5.

2 Merton's Problem

In this section, we will briefly review Merton's problem and its solution. Consider an economic agent seeking to maximize her expected life-time utility of consumption given an initial wealth endowment $x > 0$. The agent has access to continuous trade in a riskless asset with a continuously compounded rate of return $r \geq 0$ and a risky asset with a price $P(t)$ evolving according to the geometric Brownian motion

$$dP(t) = P(t)\left[b\,dt + \sigma\,dw(t)\right],$$

where w is a one-dimensional standard Brownian motion, and the drift b and the volatility σ are assumed to be constant with $b > r$ and $\sigma > 0$. Let $\theta(t)$ denote the amount invested in the risky asset at time t, and let $c(t) \geq 0$ denote the rate of consumption of the investor at time t. The residual wealth (positive or negative) of the investor is invested in the riskless asset. Given an investment strategy $\theta = \{\theta(t)\}_{t\geq 0}$, a consumption strategy $c = \{c(t)\}_{t\geq 0}$, and the initial wealth x, the wealth $X_x^{\theta,c}$ of the investor satisfies the equation

$$dX_x^{\theta,c}(t) = \left[rX_x^{\theta,c}(t) + \theta(t)(b-r) - c(t)\right]\,dt + \theta(t)\sigma\,dw(t), \quad X_x^{\theta,c}(0) = x. \tag{2.1}$$

Assuming that the investor has an infinite time horizon and an additively time-separable preference structure with a constant rate of time preference, we can write her optimal consumption/investment problem as

$$V(x) = \sup_{(\theta,c)\in\mathcal{A}(x)} \mathrm{E}_x\left[\int_0^\infty e^{-\beta t}u(c(t))\,dt\right], \tag{2.2}$$

where u is a utility function for consumption, and $\mathcal{A}(x)$ is the set of admissible investment and consumption strategies (θ, c) given the initial wealth x. Except for technical integrability conditions, the only restriction on the set of strategies (θ, c) in Merton's set-up is that the wealth $X(t) = X_x^{\theta,c}(t)$ of the investor stays non-negative (almost surely) at any point in time. It is then easy to see that two-fund separation obtains, and, hence, the assumption of a single risky asset is without loss of generality.[2]

[2]See Merton (1971, Thm. II) or Karatzas, Lehoczky, Sethi, and Shreve (1986, Sect. 5) for proofs of this assertion.

Now consider the special case of a power utility function of consumption, $u(c) = c^\gamma$, $0 < \gamma < 1$. Let the constant A be defined as

$$A = \frac{\beta - r\gamma}{1-\gamma} - \frac{\gamma(b-r)^2}{2(1-\gamma)^2\sigma^2}.$$

It can then be shown that, if $A > 0$, the optimal control strategies are $c(t) = C(X(t))$ and $\theta(t) = \Theta(X(t))$, where

$$C(x) = Ax, \quad \Theta(x) = \frac{b-r}{\sigma^2(1-\gamma)}x, \tag{2.3}$$

and that the value function V is given by

$$V(x) = A^{\gamma-1}x^\gamma. \tag{2.4}$$

For a proof of this fact, see, e.g., Karatzas, Lehoczky, Sethi, and Shreve (1986), where a more or less explicit solution can be found for a wide range of utility functions. Merton (1971) found the solution (for a slightly different utility function) by solving the Hamilton-Jacobi-Bellman (HJB) equation associated with the problem, which is the fully non-linear second order ordinary differential equation

$$\beta V(x) = \sup_{\theta\in\mathbb{R}}\left\{\frac{1}{2}\theta^2 V''(x) + \theta(b-r)V'(x)\right\} + \sup_{c\geq 0}\left\{u(c) - cV'(x)\right\} + rxV'(x). \tag{2.5}$$

See also Duffie (1996, Chap. 9) or Munk (1997b) for more information on this problem. The general theory of stochastic control and the relation between HJB equations and value functions are thoroughly treated by Fleming and Soner (1993).

3 Implementing the Markov Chain Approximation Approach

In this section, we show how the Markov chain approximation approach can be implemented on Merton's problem reviewed in Section 2. The approach was developed by Kushner (1977, 1990) and is described in details in the book by Kushner and Dupuis (1992). Munk (1997a) extracts from these sources the points relevant for most stochastic control problems arising in the financial economics literature. In particular, note that the state variable of Merton's problem, i.e. the wealth of the investor, has a control-dependent volatility term. While Kushner and Dupuis (1992) only briefly address this situation – and the problems it can cause – Munk (1997a) allows for control-dependent volatilities throughout. Munk (1997a) also provides a comparison of the Markov chain approximation approach with standard finite difference type approximations of the Hamilton-Jacobi-Bellman equation associated with the control problem.

3.1 The General Principles

The state variable of the continuous-time, continuous-state problem, i.e. the wealth of the investor, follows the process (2.1) and can take values in $\mathbb{R}_+ \equiv [0, \infty)$. We approximate it by a discrete-time, discrete-state Markov process $\xi^h = (\xi_n^h)_{n\in\mathbb{Z}_+}$ on the state space $\mathcal{R}^h = \{0, h, 2h, \ldots, Ih\}$, where $\bar{x} \equiv Ih$ is an artificial upper bound.

The evolution of the Markov chain is given by transition probabilities $p^h(x, y \mid \theta, c)$ denoting the probability of the state of the Markov chain switching from $x \in \mathcal{R}^h$ to $y \in \mathcal{R}^h$ in one time-step, when the controls $\theta \in \mathbb{R}$ and $c \geq 0$ are currently applied. The possible controls of the Markov chain are pairs (θ^h, c^h) of functions of the state of the chain, i.e. $\theta^h : \mathcal{R}^h \to \mathbb{R}$ and $c^h : \mathcal{R}^h \to \mathbb{R}_+$. Such a pair (θ^h, c^h) is called admissible, if the process ξ^h indeed is a Markov chain on $\mathcal{R}^h$, when it is controlled by (θ^h, c^h). In particular, ξ^h stays non-negative almost surely. The set of admissible controls given the initial wealth $\xi_0^h = x$ are denoted by $\mathcal{A}^h(x)$.

Define an interpolation interval function $\Delta t^h : \mathcal{R} \times \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}$, such that $\Delta t^h(x, \theta, c)$ is the length of the time step of the Markov chain, when the state is x and the controls currently applied are θ and c. Let Δt_n^h denote $\Delta t^h(\xi_n^h, \theta^h(\xi_n^h), c^h(\xi_n^h))$ and let $t_n^h = \sum_{m=0}^{n-1} \Delta t_m^h$. Define the objective of the approximating Markov chain control problem as

$$V^h(x) = \sup_{(\theta^h, c^h)\in\mathcal{A}^h(x)} \mathrm{E}\left[\sum_{n=0}^{\infty} e^{-\beta t_n^h} c^h(\xi_n^h)^\gamma \Delta t_n^h \,\middle|\, \xi_0^h = x \right].$$

The Markov chain is to be chosen such that $V^h(x)$ converges to $V(x)$, when $h \to 0$,[3] and such that the computation of V^h is fairly simple. The function V^h can be found by exploiting the dynamic programming equation

$$V^h(x) = \sup_{\theta\in\mathbb{R}, c\geq 0} \left\{ c^\gamma \Delta t^h(x, \theta, c) + e^{-\beta\Delta t^h(x,\theta,c)} \sum_{y\in\mathcal{R}^h} p^h(x, y \mid \theta, c) V^h(y) \right\}, \quad x \in \mathcal{R}^h, \tag{3.1}$$

which can be solved by various methods, see, e.g., Rust (1996). We apply the *policy iteration method*, which works as follows.[4] Given some admissible control (θ_0^h, c_0^h), that is $(\theta_0^h(x), c_0^h(x))$ for all $x \in \mathcal{R}^h$. Compute a first approximation V_0^h of the value function V^h by solving the equation system

$$\begin{aligned} V_0^h(x) = {} & c_0^h(x)^\gamma \Delta t^h(x, \theta_0^h(x), c_0^h(x)) \\ & + e^{-\beta\Delta t^h(x,\theta_0^h(x),c_0^h(x))} \sum_{y\in\mathcal{R}^h} p^h(x, y \mid \theta_0^h(x), c_0^h(x)) V_0^h(y), \quad x \in \mathcal{R}^h. \end{aligned}$$

[3] Of course, this requires that $\bar{x} = Ih \to \infty$ as $h \to 0$.

[4] The value space method (also known as the method of successive approximations) was also implemented, but it converged very, very slowly. The value space method is an explicit search for a fixpoint of a contraction mapping. If the mapping is only "little contractive", the convergence will be slow. This is the case with our problem, since $\exp(-\beta\Delta t(x))$ is very close to 1 for small h.

Then compute a new control policy (θ_1^h, c_1^h) as

$$(\theta_1^h(x), c_1^h(x)) = \operatorname*{arg\,max}_{(\theta,c)\in\mathbb{R}\times\mathbb{R}_+} \left\{ c^\gamma \Delta t^h(x,\theta,c) + e^{-\beta\Delta t^h(x,\theta,c)} \sum_{y\in\mathcal{R}^h} p^h(x,y \mid \theta,c) V_0^h(y) \right\}, \quad x \in \mathcal{R}^h.$$

Continuing this way, we can compute a sequence of control policies (θ_k^h, c_k^h) and a sequence of value function approximations (V_k^h) by first performing a *policy evaluation*

$$\begin{aligned} V_k^h(x) &= c_k^h(x)^\gamma \Delta t^h(x, \theta_k^h(x), c_k^h(x)) \\ &\quad + e^{-\beta\Delta t^h(x,\theta_k^h(x),c_k^h(x))} \sum_{y\in\mathcal{R}^h} p^h(x,y \mid \theta_k^h(x), c_k^h(x)) V_k^h(y), \quad x \in \mathcal{R}^h, \end{aligned}$$

and then a *policy improvement*

$$\begin{aligned} (\theta_{k+1}^h(x), c_{k+1}^h(x)) = \operatorname*{arg\,max}_{(\theta,c)\in\mathbb{R}\times\mathbb{R}_+} \Big\{ & c^\gamma \Delta t^h(x,\theta,c) \\ & + e^{-\beta\Delta t^h(x,\theta,c)} \sum_{y\in\mathcal{R}^h} p^h(x,y \mid \theta,c) V_k^h(y) \Big\}, \quad x \in \mathcal{R}^h. \end{aligned} \tag{3.2}$$

3.2 Different Versions of the Markov Chain Approximation Approach

Given the discussion above, the problem that remains is to find transition probabilities and an interpolation interval function, such that the desired properties of the approximating problem are obtained. An intuitive and comprehensible derivation of one such set of transition probabilities together with some suggestions for minor modifications are given in Munk (1997a, Sect. 5 and 6). In the remainder of this paper, we shall suppress the h superscripts, where no confusion can arise.

Due to the form of the drift term of the state variable, $rX + \theta(b-r) - c$, cf. (2.1), we can take advantage of the "splitting the operator" technique. Instead of the terms $(rX + \theta(b-r) - c)^\pm$ appearing in the denominators of the transition probabilities, we get terms like $(rX)^\pm$, $[\theta(b-r)]^\pm$, and $c^\pm$. Since we know that r and X are non-negative, we have $(rX)^+ = rX$ and $(rX)^- = 0$. Since $b > r$, the investor will never take a short position in the risky asset, hence $[\theta(b-r)]^+ = \theta(b-r)$ and $[\theta(b-r)]^- = 0$. Finally, we know that $c \geq 0$. The "splitting the operator" method will simplify the necessary computations considerably.

With this modification, a very general transition probability scheme is given by

$$p(x, x+h \mid \theta(x), c(x)) = \frac{\frac{1}{2}\sigma^2\theta(x)^2 + h\left(rx + \theta(x)(b-r)\right)}{Q(x,\theta(x),c(x))}, \tag{3.3a}$$

$$p(x, x-h \mid \theta(x), c(x)) = \frac{\frac{1}{2}\sigma^2\theta(x)^2 + hc(x)}{Q(x,\theta(x),c(x))}, \tag{3.3b}$$

$$p(x, y \mid \theta(x), c(x)) = 0, \text{ for } y \notin \{x-h, x+h\}, \tag{3.3c}$$

for $x \in \{h, 2h, \ldots, (I-1)h\}$, where

$$Q(x,\theta,c) = \sigma^2\theta^2 + h\left(rx + \theta(b-r) + c\right).$$

Zero wealth is an absorbing state of the original problem. If the investor goes bankrupt, i.e. her wealth reaches zero, she is confined to zero consumption in all future, and since she gets zero utility from zero consumption, $V(0) = 0$.[5] Therefore, let $V^h(0) = 0$ and

$$p(0,0 \mid \theta, c) = 1, \tag{3.3d}$$

$$p(0,y \mid \theta, c) = 0, \quad y \neq 0, \tag{3.3e}$$

for all θ, c. At the upper boundary we take the transition probabilities to be[6]

$$p(\bar{x}, \bar{x} - h \mid \theta(\bar{x}), c(\bar{x})) = \frac{\frac{1}{2}\sigma^2\theta(\bar{x})^2 + hc(\bar{x})}{Q(\bar{x}, \theta(\bar{x}), c(\bar{x}))}, \tag{3.3f}$$

$$p(\bar{x}, \bar{x} \mid \theta(\bar{x}), c(\bar{x})) = 1 - \frac{\frac{1}{2}\sigma^2\theta(\bar{x})^2 + hc(\bar{x})}{Q(\bar{x}, \theta(\bar{x}), c(\bar{x}))}, \tag{3.3g}$$

$$p(\bar{x}, y \mid \theta(\bar{x}), c(\bar{x})) = 0, \quad \text{for } y \notin \{\bar{x} - h, \bar{x}\}. \tag{3.3h}$$

The interpolation interval is

$$\Delta t(x, \theta(x), c(x)) = \frac{h^2}{Q(x, \theta(x), c(x))}. \tag{3.4}$$

As can be seen from policy improvement step (3.2) of the solution of the dynamic programming equation (3.1), the computation of the approximating value function involves a maximization in the space of controls. With the transition probability scheme (3.3), this maximization is rather complex and has to be done numerically due to the presence of the controls in both the numerator and the denominator of the probabilities.[7] We can circumvent this by replacing $Q(x, \theta, c)$ in the denominators of the transition probabilities (3.3) and the interpolation interval (3.4) with

$$Q(x) = \sup_{\theta, c} Q(x, \theta, c)$$

and let $p(x, x \mid \theta(x), c(x)) = 1 - \sum_{y \neq x} p(x, y \mid \theta(x), c(x))$. But since θ and c are unbounded, $Q(x)$ will be infinite, making the resulting Markov chain degenerate. One way to get around this is to artificially bound the controls from above by requiring $\theta(x)$ and $c(x)$ to be less than or equal to Kx, where K is some positive constant.[8]

---

[5] The impact of other implications of bankruptcy is studied by, e.g., Karatzas, Lehoczky, Sethi, and Shreve (1986).

[6] The transition probabilities at the upper boundary correspond to assigning the Neumann boundary condition $V'(\bar{x}) = 0$. If $\bar{x}$ is chosen such that $\bar{x} \to \infty$ as $h \to 0$, the same limit for V^h is obtained no matter which boundary conditions are chosen. See Theorem IX.5.3 and the discussion on page 370 in Fleming and Soner (1993).

[7] An experiment shows that, with the parameter values used in Section 4, the function to be maximized in (3.1) will have several local extrema, so a numerical search for the global maximum will be extensive.

[8] Of course, we could allow for different bounds on the two controls, say, $\theta x \leq K_\theta x$ and $c(x) \leq K_c x$. Convergence will typically be faster, the faster the probability mass "spreads". Since a larger K will decrease the probability of leaving a state, K should not be too big a number.

Define

$$
\begin{aligned}
Q^*(x) &= \sup_{\theta(x)\le Kx, 0\le c(x)\le Kx} Q(x,\theta(x),c(x)) \\
&= Q(x,Kx,Kx) \\
&= \sigma^2K^2x^2 + h\left(rx + Kx(b-r) + Kx\right). \qquad (3.5)
\end{aligned}
$$

The resulting Markov chain is given by

$$\bar{p}(x, x+h \mid \theta(x), c(x)) = \frac{\frac{1}{2}\sigma^2\theta(x)^2 + h\left(rx+\theta(x)(b-r)\right)}{Q^*(x)}, \qquad (3.6a)$$

$$\bar{p}(x, x-h \mid \theta(x), c(x)) = \frac{\frac{1}{2}\sigma^2\theta(x)^2 + hc(x)}{Q^*(x)}, \qquad (3.6b)$$

$$
\begin{aligned}
\bar{p}(x, x \mid \theta(x), c(x)) &= 1 - \bar{p}(x, x+h \mid \theta(x), c(x)) - \bar{p}(x, x-h \mid \theta(x), c(x)) \\
&= \frac{Q^*(x) - \sigma^2\theta(x)^2 - h\left(rx + \theta(x)(b-r) + c(x)\right)}{Q^*(x)}, \qquad (3.6c)
\end{aligned}
$$

$$\bar{p}(x, y \mid \theta(x), c(x)) = 0, \text{ for } y \notin \{x-h, x, x+h\}, \qquad (3.6d)$$

for $x \in \{h, 2h, \ldots, (I-1)h\}$ and

$$\bar{p}(\bar{x}, \bar{x}-h \mid \theta(\bar{x}), c(\bar{x})) = \frac{\frac{1}{2}\sigma^2\theta(\bar{x})^2 + hc(\bar{x})}{Q^*(\bar{x})}, \qquad (3.6e)$$

$$
\begin{aligned}
\bar{p}(\bar{x}, \bar{x} \mid \theta(\bar{x}), c(\bar{x})) &= 1 - \bar{p}(\bar{x}, \bar{x}-h \mid \theta(\bar{x}), c(\bar{x})) \\
&= \frac{Q^*(\bar{x}) - \frac{1}{2}\sigma^2\theta(\bar{x})^2 - hc(\bar{x})}{Q^*(\bar{x})}, \qquad (3.6f)
\end{aligned}
$$

$$\bar{p}(\bar{x}, y \mid \theta(\bar{x}), c(\bar{x})) = 0, \text{ for } y \notin \{\bar{x}-h, \bar{x}\}, \qquad (3.6g)$$

with the interpolation interval

$$\overline{\Delta t}(x) = \frac{h^2}{Q^*(x)}. \qquad (3.7)$$

The corresponding dynamic programming equations are much easier to solve with this approximating Markov chain than the one given by the transition probability scheme (3.3) and the interpolation interval function (3.4).

Let us take a closer look at the computations involved in the policy iteration algorithm applied on the problem of controlling the Markov chain defined by the transition probabilities (3.6) and the interpolation interval (3.7). For a given control policy (θ_k, c_k), the policy evalution step amounts to solving the following system of equations

$$
\begin{aligned}
V_k^h(x) = \frac{1}{Q^*(x)}\Bigg\{ & e^{-\beta\overline{\Delta t}(x)}\Bigg[\left(\frac{1}{2}\sigma^2\theta_k(x)^2 + hc_k(x)\right)V_k^h(x-h) \\
& + \left(Q^*(x) - \sigma^2\theta_k(x)^2 - h\left(rx + \theta_k(x)(b-r) + c_k(x)\right)\right)V_k^h(x) \\
& + \left(\frac{1}{2}\sigma^2\theta_k(x)^2 + h\left(rx + \theta_k(x)(b-r)\right)\right)V_k^h(x+h)\Bigg] \\
& + h^2 c_k(x)^\gamma\Bigg\}, \quad x = h, 2h, \ldots, (I-1)h,
\end{aligned} \tag{3.8}
$$

$$
\begin{aligned}
V_k^h(\bar{x}) = \frac{1}{Q^*(\bar{x})}\Bigg\{ & e^{-\beta\overline{\Delta t}(\bar{x})}\Bigg[\left(\frac{1}{2}\sigma^2\theta_k(\bar{x})^2 + hc_k(\bar{x})\right)V_k^h(\bar{x}-h) \\
& + \left(Q^*(\bar{x}) - \frac{1}{2}\sigma^2\theta_k(\bar{x})^2 - hc_k(\bar{x})\right)V_k^h(\bar{x})\Bigg] \\
& + h^2 c_k(\bar{x})^\gamma\Bigg\}.
\end{aligned} \tag{3.9}
$$

With the notation $\mathbf{V}_k^h = \left(V_k^h(0), V_k^h(h), V_k^h(2h), \ldots, V_k^h(Ih)\right)^\top$, and similarly for $\boldsymbol{\theta}_k, \mathbf{c}_k$, we can write the equations as

$$
\mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)\mathbf{V}_k^h = \mathbf{d}(\mathbf{c}_k), \tag{3.10}
$$

where $\mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)$ is a tridiagonal $(I+1)\times(I+1)$ matrix with non-zero elements

$$
\begin{aligned}
\mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)_{i,i-1} &= \frac{1}{2}\sigma^2\theta_k(ih)^2 + hc_k(ih), \quad i = 1, \ldots, I, \\
\mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)_{i,i+1} &= \frac{1}{2}\sigma^2\theta_k(ih)^2 + h(rih + \theta_k(ih)(b-r)), \quad i = 1, 2, \ldots, I-1, \\
\mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)_{0,0} &= 1, \\
\mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)_{i,i} &= \left(1 - e^{\beta h^2/Q^*(ih)}\right)Q^*(ih) - \sigma^2\theta_k(ih)^2 - h(rih + \theta_k(ih)(b-r) + c_k(ih)) \\
&= \left(1 - e^{\beta h^2/Q^*(ih)}\right)Q^*(ih) - \mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)_{i,i-1} - \mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)_{i,i+1}, \quad i = 1, 2, \ldots, I-1, \\
\mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)_{I,I} &= \left(1 - e^{\beta h^2/Q^*(Ih)}\right)Q^*(Ih) - \frac{1}{2}\sigma^2\theta_k(Ih)^2 - hc_k(Ih) \\
&= \left(1 - e^{\beta h^2/Q^*(Ih)}\right)Q^*(Ih) - \mathbf{M}(\boldsymbol{\theta}_k, \mathbf{c}_k)_{I,I-1},
\end{aligned}
$$

and $\mathbf{d}(\mathbf{c}_k)$ is the $(I+1)$-dimensional vector with i'th element $-e^{\beta h^2/Q^*(ih)}h^2 c_k(ih)$ for $i = 1, \ldots, I$ and 0'th element equal to zero.

The policy improvement step boils down to

$$\theta_{k+1}(0) = c_{k+1}(0) = 0,$$

(3.11)

$$\begin{aligned}(\theta_{k+1}(ih), c_{k+1}(ih)) = \underset{\theta \le Kih, 0 \le c \le Kih}{\arg\max} \Bigg\{ & e^{-\beta\overline{\Delta t}(ih)} \frac{\frac{1}{2}\sigma^2\theta^2 + hc}{Q^*(ih)} V_k^h((i-1)h) \\ & + e^{-\beta\overline{\Delta t}(ih)} \frac{Q^*(ih) - \sigma^2\theta^2 - h\left(rih + \theta(b-r) + c\right)}{Q^*(ih)} V_k^h(ih) \\ & + e^{-\beta\overline{\Delta t}(ih)} \frac{\frac{1}{2}\sigma^2\theta^2 + h\left(rih + \theta(b-r)\right)}{Q^*(ih)} V_k^h((i+1)h) \\ & + \overline{\Delta t}(ih)c^\gamma \Bigg\}, \quad i = 1, 2, \ldots, (I-1),\end{aligned}$$

and

(3.12)

$$\begin{aligned}(\theta_{k+1}(Ih), c_{k+1}(Ih)) = \underset{\theta \le KIh, 0 \le c \le KIh}{\arg\max} \Bigg\{ & e^{-\beta\overline{\Delta t}(Ih)} \frac{\frac{1}{2}\sigma^2\theta^2 + hc}{Q^*(Ih)} V_k^h((I-1)h) \\ & + e^{-\beta\overline{\Delta t}(Ih)} \frac{Q^*(Ih) - \frac{1}{2}\sigma^2\theta^2 - hc}{Q^*(Ih)} V_k^h(Ih) \\ & + \overline{\Delta t}(Ih)c^\gamma \Bigg\}.\end{aligned}$$

In this case, the Equations (3.11) and (3.12) can be solved explicitly, the solutions being

$$\theta_{k+1}(ih) = \min\left\{ -\frac{b-r}{\sigma^2} \frac{D^+V_k^h(ih)}{D^2V_k^h(ih)}, Kih \right\}, \quad i = 1, 2, \ldots, I-1, \tag{3.13}$$

$$c_{k+1}(ih) = \min\left\{ I\left(e^{-\beta h^2/Q^*(ih)} D^-V_k^h(ih)\right), Kih \right\}, \quad i = 1, 2, \ldots, I, \tag{3.14}$$

and

$$\theta_{k+1}(Ih) = 0, \tag{3.15}$$

where I is the inverse of u', i.e. $I(y) = (y/\gamma)^{1/(\gamma-1)}$, and we have used the finite difference operators

$$D^+V(x) = \frac{V(x+h) - V(x)}{h}, \quad D^-V(x) = \frac{V(x) - V(x-h)}{h},$$

and

$$D^2V(x) = \frac{V(x+h) - 2V(x) + V(x-h)}{h^2}.$$

To sum up, the algorithm is as follows. First guess initial policy vectors $(\boldsymbol{\theta}, \mathbf{c})$. Then (I) compute the vector $\mathbf{V}$ by solving the tridiagonal equation system (3.10) and (II) given the solution $\mathbf{V}$ compute new controls from the equations (3.13)-(3.15). The steps (I) and (II) are repeated until

$\mathbf{V}$ does not change very much from one iteration to the next. To be more precise, the algorithm is stopped when $\|\mathbf{V}_{(k+1)} - \mathbf{V}_{(k)}\|_\infty < \varepsilon$, where $\mathbf{V}_{(k)}$ is the solution from step (I) after k iterations, $\|\cdot\|_\infty$ is the maximum norm on $\mathbb{R}^I$, and ε is a predetermined tolerance.

A straightforward variation of the procedure outlined above is to use a constant denominator

$$\begin{aligned} Q' &= \sup_{\theta(x),c(x)\leq \tilde{K}Ih, 0\leq x\leq Ih} Q(x,\theta(x),c(x)) \\ &= Q(Ih,\tilde{K}Ih,\tilde{K}Ih) \\ &= \sigma^2\tilde{K}^2(Ih)^2 + h(rIh + \tilde{K}Ih(b-r) + \tilde{K}Ih), \end{aligned}$$

where $\tilde{K}$ is a constant, in the definition of the transition probabilities. Again, the probability of staying at a state is defined residually, such that the probabilities $p(x,\cdot \mid \theta, c)$ sum to one. If we take $K = \tilde{K}$, using Q^* instead of Q' gives a lower bound on the controls for all levels of wealth less than the maximum Ih. Also, the probability of leaving a given state is smaller with constant denominator Q' than with $Q^*(x)$, and, as previously mentioned, this will typically slow down convergence. The implementation is identical to the one outlined above, but, of course, the $Q^*(ih)$ terms in the tridiagonal matrix $\mathbf{M}$ and the vector $\mathbf{d}$ as well as in the equation (3.14) are replaced by Q'. In (3.13)-(3.15), Kih must be replaced by $\tilde{K}Ih$. In sum, the constant denominator simplifies the implementation, but only slightly.

3.3 Two Accelerated Markov-Like Approaches

In this subsection, we shall outline two variations of the basic Markov chain approach. For more details, we refer the reader to Munk (1997a). The constant denominator Q' scheme was applied by Fitzpatrick and Fleming (1991). They also suggested to use a denominator, which is the same for all wealth levels, but changes from one iteration of the policy space algorithm to the next. In the k'th iteration of (I)-(II) we use

$$Q'_k = \max_{1\leq i\leq I} \left\{\theta_k(ih)^2\sigma^2 + h\left(\theta_k(ih)(b-r) + rih + c_k(ih)\right)\right\},$$

where $(\boldsymbol{\theta}_{k-1}, \mathbf{c}_{k-1})$ is the control vector computed in the previous iteration.

Another variation, which was introduced by Munk (1997a), is to replace $Q^*(x)$ in the k'th iteration with

$$Q^*_k(ih) = \theta_{k-1}(ih)^2\sigma^2 + h\left(\theta_{k-1}(ih)(b-r) + rih + c_{k-1}(ih)\right).$$

This can be seen as an approximation of the original state- and control-dependent denominator $Q(x,\theta,c)$, but keeping the policy improvement step as simple as when using the $Q^*(x)$-based scheme. Compared to the state-independent Q'_k approach, the probability mass spreads even faster, and a time-consuming maximization in $\mathcal{R}^h$ is avoided.

Both the variations suggested above are simple modifications of the two Markov chain approaches discussed in the previous subsection. Note, however, that for both variations, the resulting transition "probabilities" are not guaranteed to be non-negative. In particular, $p(x, x \mid \theta, c)$ can be negative. Hence, the standard convergence argument does not apply.

4 Numerical Results

In the previous section, we outlined four different methods distinguished by the denominator applied in the transition scheme: The wealth-dependent denominator $Q^*(x)$, the constant denominator Q', the Fitzpatrick-Fleming acceleration given by Q'_k, and our acceleration given by $Q^*_k(x)$. To study and compare the numerical performance of each of these methods, we implement them on Merton's problem reviewed in Section 2. In the numerical experiment, we have taken the following parameter values:

$$\beta = 0.2, \quad b = 0.1, \quad r = 0.05, \quad \sigma = 0.3, \quad \gamma = 0.5,$$

in which case $A \approx 0.3222$ and the analytical solution in (2.3)–(2.4) becomes $V(x) \approx 1.7617\sqrt{x}$, $C(x) \approx 0.3222x$, and $\Theta(x) \approx 1.1111x$. For the policy space iterations, we take the tolerance level $\varepsilon = 0.001$. Smaller values of ε will only slightly improve precision and only for wealth levels near the artificial upper bound. In that end of the wealth interval, the numerical results, as will become clear below, are quite imprecise, and it is recommended to ignore those results. Therefore, smaller ε will simply result in more time-consuming and essentially futile policy space iterations.

4.1 Experimental Convergence Properties

First, we focus on the experimental convergence properties of the methods. In all experiments we have imposed an artificial upper bound on wealth at $\bar{x} = 100$ and put $h = \bar{x}/I$. Figure 1 and 2 depict the precision of the $Q^*(x)$ based approximation method for the optimal amount invested and the optimal consumption rate, respectively, where the measure of precision is the percentage deviation from the closed-form solution stated above. For relatively small or relatively large initial levels of wealth, the numerical results are quite imprecise, but for a wide intermediate range of wealth, the error is only a few percent using a refinement of $I = 400$, which can be handled even on standard PCs. As expected, the precision increases with the number of grid points. The improvement gained by using a finer grid is most pronounced for small and intermediate wealth levels, whereas a grid refinement does not significantly reduce the error for wealth levels near the artificial bound.

The same pattern can be seen from Figure 3 and 4, which show the precision of the numerically computed optimal controls using very fine grids. For a wide range of initial wealth levels, the error

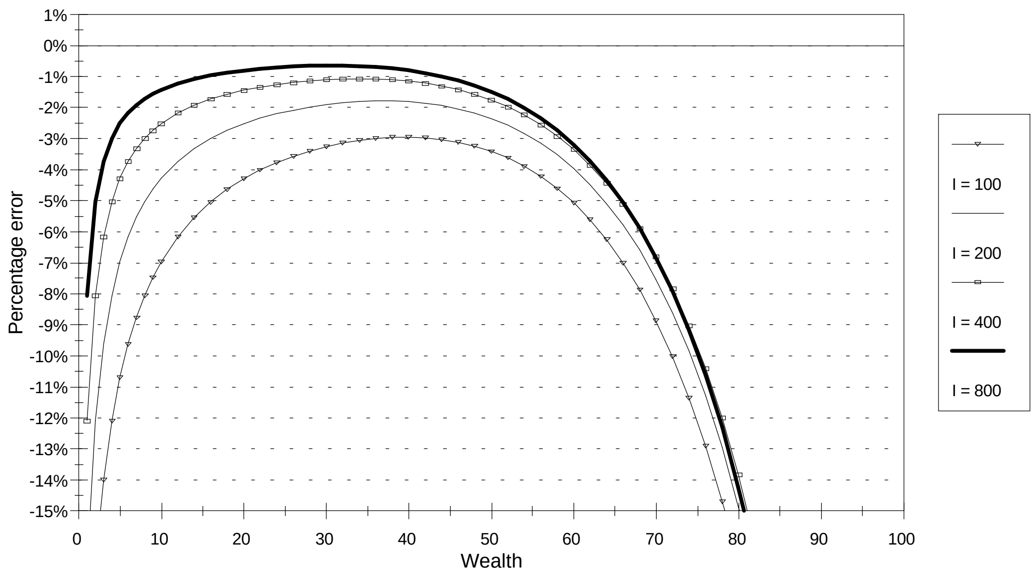

Figure 1: The percentage error of the numerically computed optimal risky investment for the $Q^*(x)$ method.

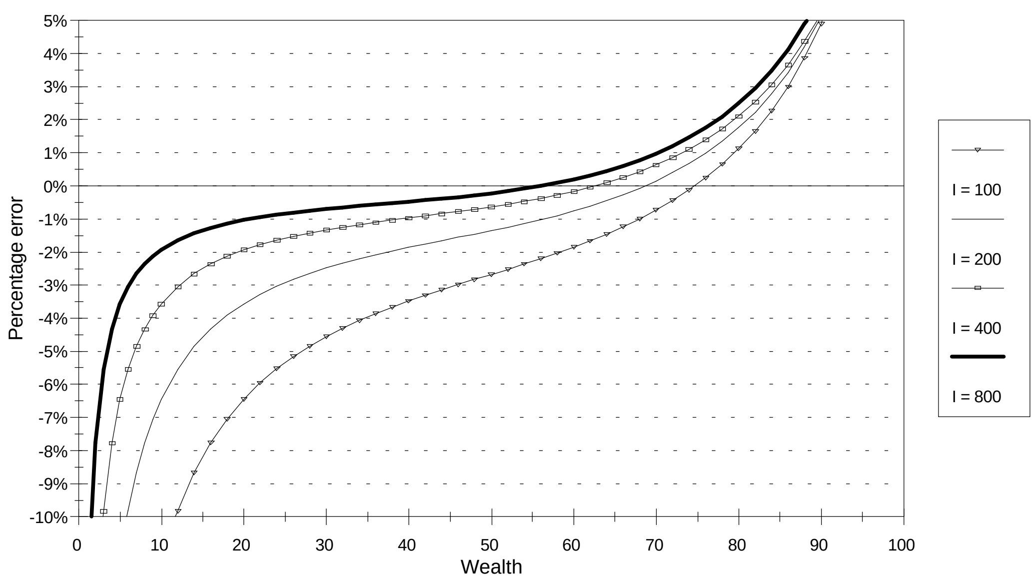

Figure 2: The percentage error of the numerically computed optimal consumption rate for the $Q^*(x)$ method.

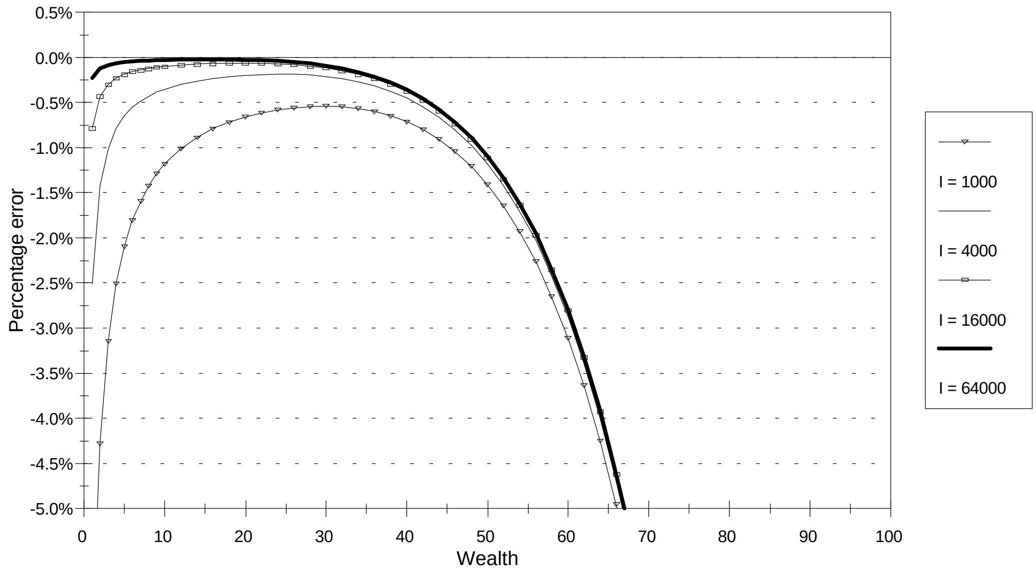

Figure 3: The percentage error of the numerically computed optimal risky investment for the $Q^*(x)$ method.

is less than one tenth of a percent with a refinement of $I = 64000$, but for high wealth levels the error is still substantial. This suggests that one should design the program such that the initial wealth levels of particular interest do not fall in the upper half of the interval $[0, \bar{x}]$. Of course, these results can differ from problem to problem and with the parameter values used. As we shall demonstrate in Section 4.3, the value of the time preference paremeter is important in this respect.

It is not surprising that the numerically computed optimal controls are that imprecise near the artificial bound on initial wealth. If wealth is near the upper bound, the investor cannot benefit much from a further increase in wealth. Of course, this moderates the incentive to invest in the risky asset, so it is quite natural that the numerically computed risky investment is much lower than the "true" optimal investment for wealth levels near the artificial upper bound.

It is well-known that the same preferences can be represented by different utility functions. This implies that no meaning can be given to a statement like "the precision of the value function is one percent". In order to have a *cardinal* measure of the value function and the precision of the value function, that is a measure which is independent of the particular choice of utility function, we can translate the value function into consumption. One way to do that is to compute what might be called the *constant consumption equivalent*, i.e., the consumption rate that, when held constant in the entire life-time, obtains the optimal expected utility. Let c denote the constant

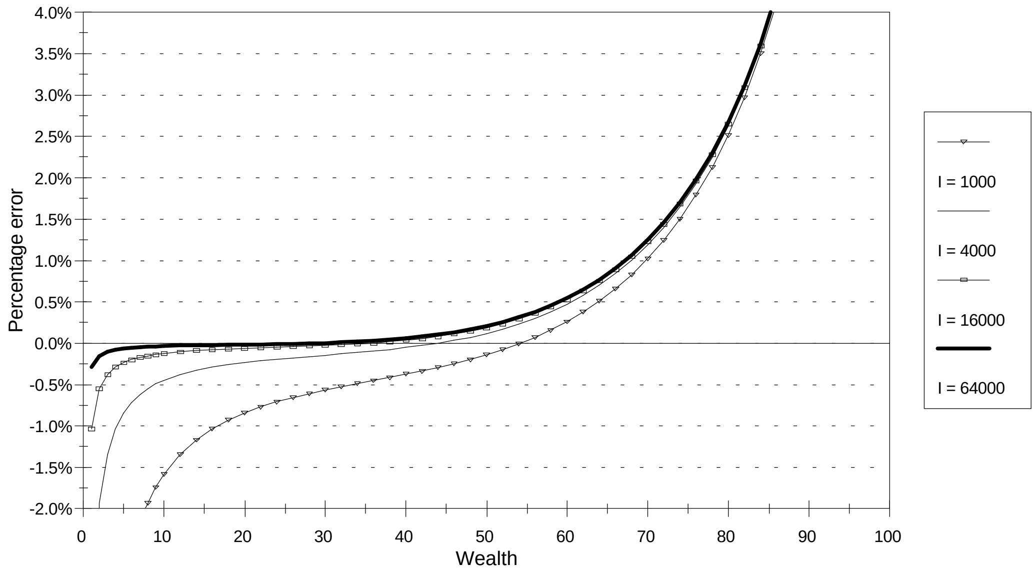

Figure 4: The percentage error of the numerically computed optimal consumption rate for the $Q^*(x)$ method.

consumption equivalent. Then

$$V(x) = A^{\gamma-1}x^{\gamma} = \int_0^{\infty} e^{-\beta t} c^{\gamma}\, dt,$$

i.e.,

$$c = (\beta V(x))^{1/\gamma} = \beta^{1/\gamma} A^{(\gamma-1)/\gamma} x.$$

The precision of the numerically computed constant consumption equivalent for the $Q^*(x)$ method is depicted in Figure 5 and 6. For example with $I = 400$, the error is less than two percent for wealth levels above 15. Even though the controls are quite imprecise for high levels of wealth, the value function is very precise in that end. Since the investor, of course, is interested in the expected utility, the imprecision of the optimal controls for high levels of wealth is not as destructive as one might have thought at first. For low wealth, the error can be quite substantial but is reduced considerably when I is increased.

Roughly the same pattern of convergence is found for the other methods, although there are differences in the precision of the various approaches, as we will demonstrate in the comparisons in the next subsection. However, the method we proposed, i.e. the $Q_k(x)$ method, has a somewhat different convergence pattern as can be seen from Figure 7–9. For the standard (i.e. non-accelerated) Markov approximation methods, the numerically computed value function converges to the true value function from below as in Figure 5. This is an implication of the policy iteration solution

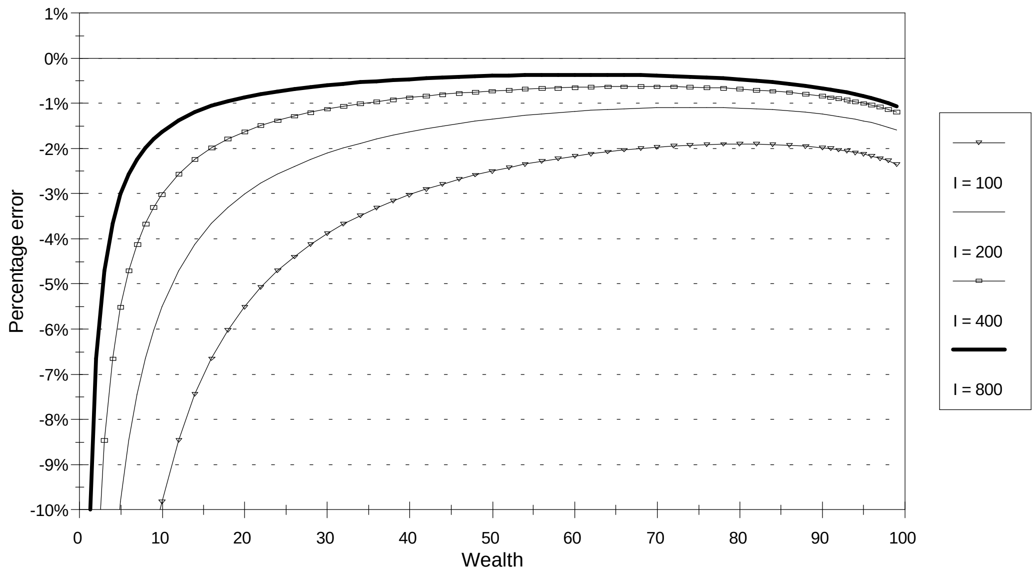

Figure 5: The percentage error of the numerically computed constant consumption equivalent for the $Q^*(x)$ method.

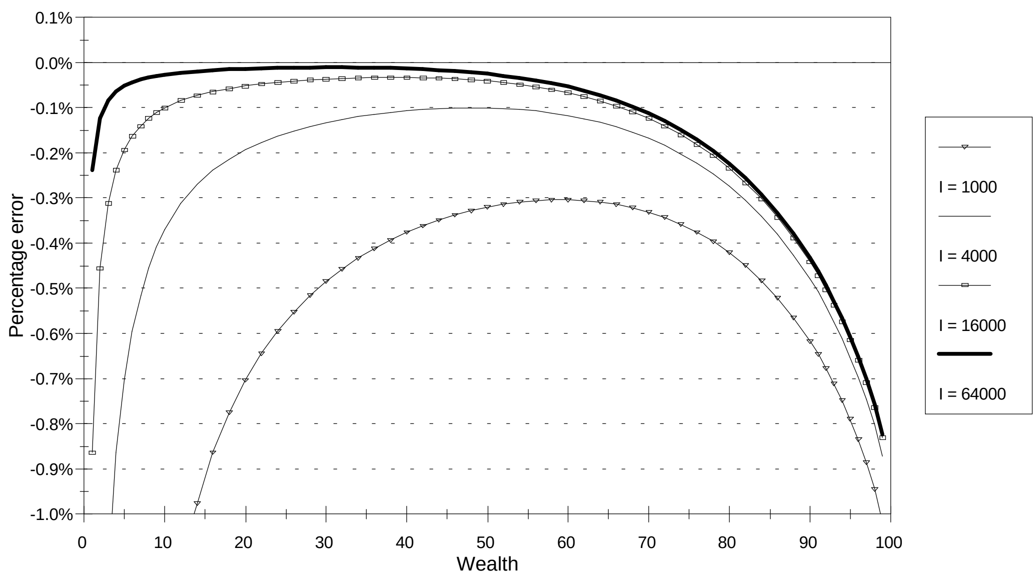

Figure 6: The percentage error of the numerically computed constant consumption equivalent for the $Q^*(x)$ method.

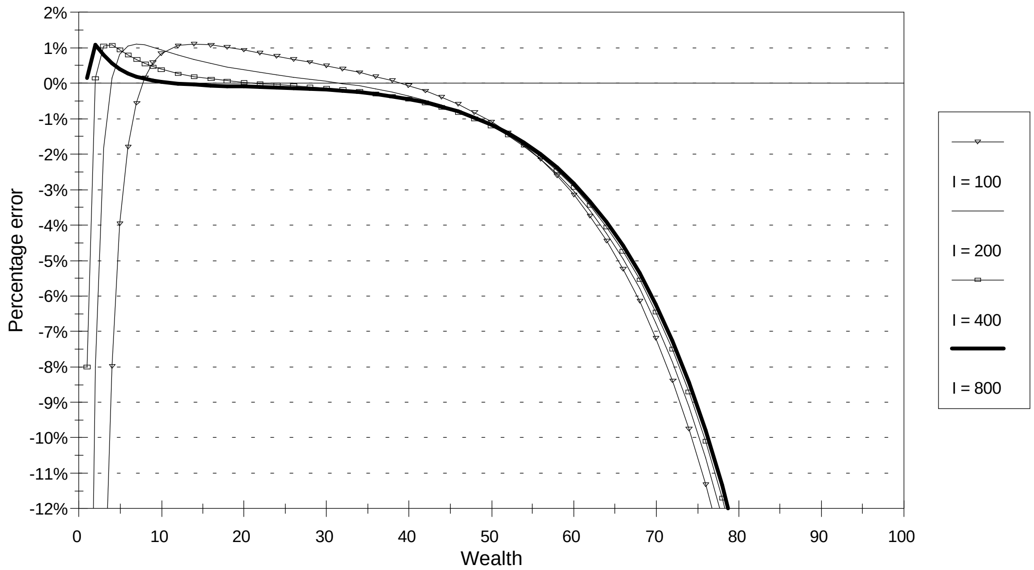

Figure 7: The percentage error of the numerically computed optimal risky investment for our $Q_k(x)$ method.

approach to the discretized problem, since that approach is based on iterative improvements of the controls. With our approximation scheme, this feature is not present, and indeed Figure 9 shows that the numerically computed value function does not uniformly converge in a monotonic way.

Next, we estimate the order of convergence of the methods. This concept is defined in most introductory textbooks on numerical analysis. We remind the reader that the order of convergence, m, can be estimated from three computed values of the number sought. For example, if we have computed the optimal risky investment $\theta^h(x)$ at some wealth level for $I = 50$, $I = 100$, and $I = 200$, that is $h = 2$, $h = 1$, and $h = 0.5$, respectively, we get the following estimate of the order of convergence of $\theta^h(x)$ to $\theta(x)$:

$$\tilde{m}(\theta(x)) = \frac{\ln \frac{\theta^{h=2}(x) - \theta^{h=1}(x)}{\theta^{h=1}(x) - \theta^{h=0.5}(x)}}{\ln 2}.$$

In Figure 10, we have graphed estimates of the convergence order for both the controls and the value function for the $Q^*(x)$ method. Note that the horizontal axis has a log scale. As indicated above, the experimental convergence order depends on the wealth level at which we compute the controls and the value function. We get a similar picture for a wide range of low and intermediate wealth levels. For large wealth levels, the experimental convergence orders are very unstable as the grid refinement is varied, especially for the controls and for rather fine grids. This indicates that, although the numerical results do converge to the true results, the convergence is not smooth. From Figure 10, it seems reasonable to conclude that the method is of order one, i.e., it converges

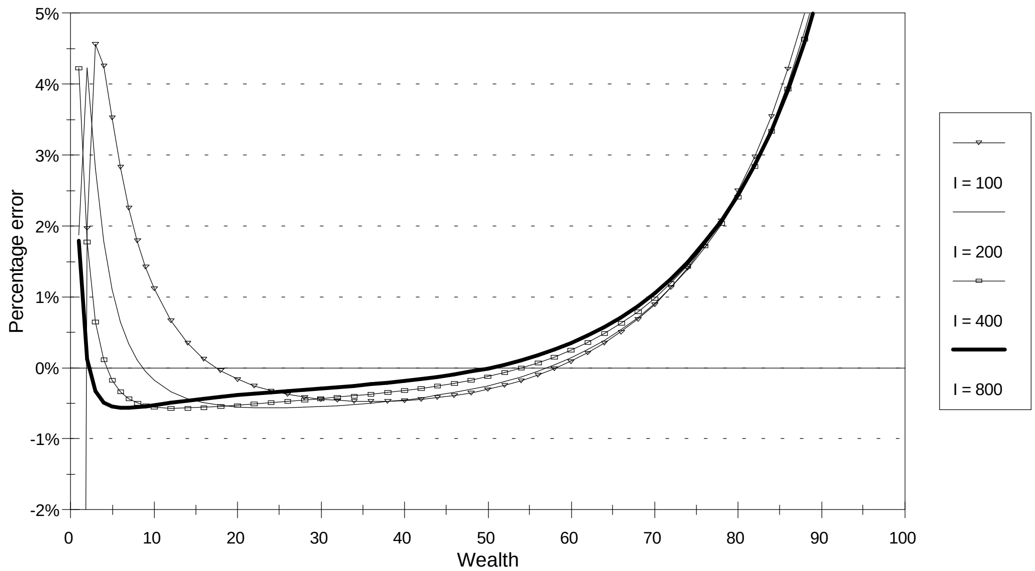

Figure 8: The percentage error of the numerically computed optimal consumption rate for our $Q_k(x)$ method.

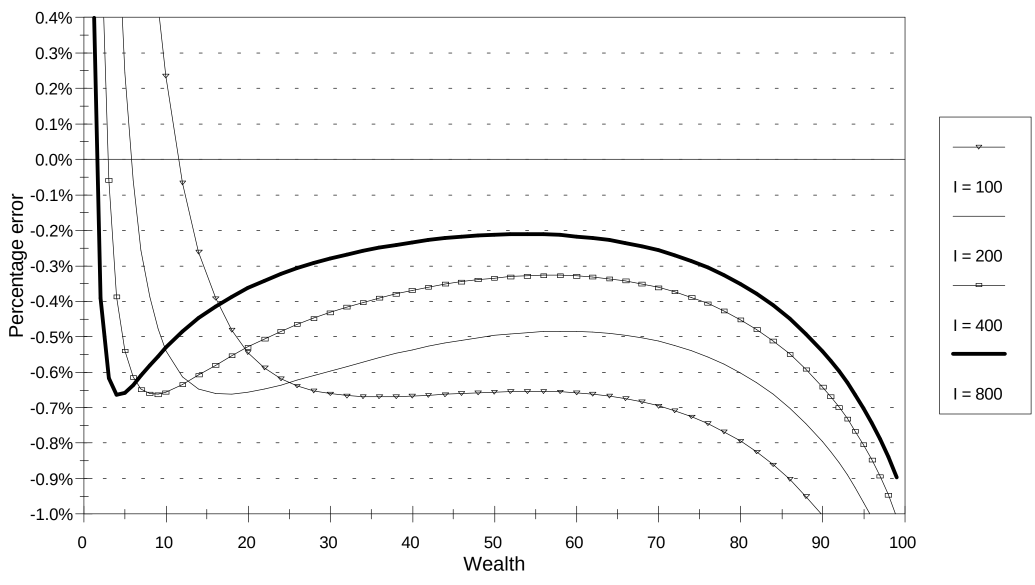

Figure 9: The percentage error of the numerically computed constant consumption equivalent for our $Q_k(x)$ method.

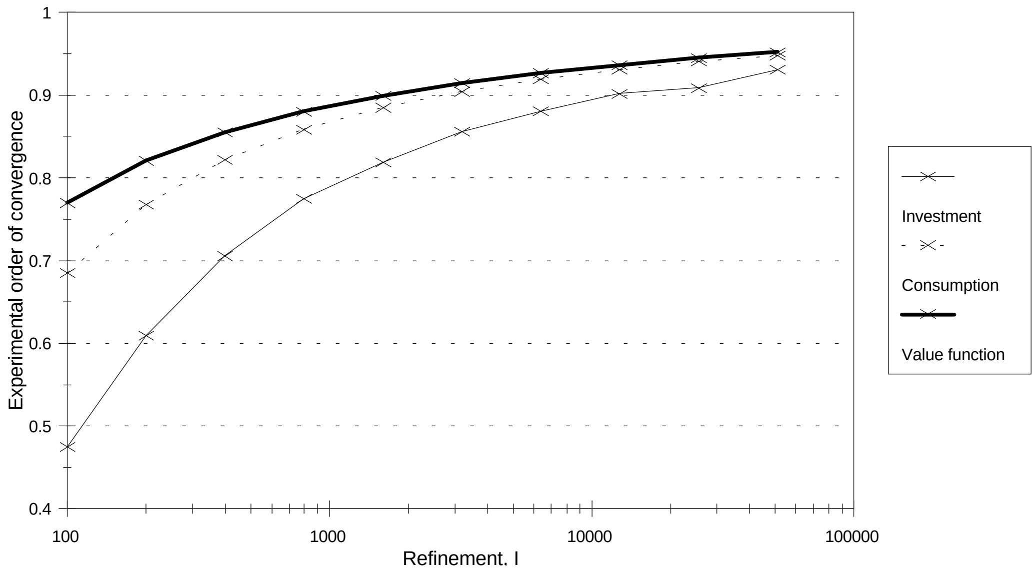

Figure 10: The experimental convergence order for the $Q^*(x)$ method computed for an initial wealth of $x = 25$.

linearly, in the "relevant range" of wealth.[9]

For the constant denominator Q' scheme, we get a similar picture for coarse grids, but the convergence using relatively fine grids seems somewhat dubious, implying useless experimental convergence orders both for the controls and the value function. For the Fitzpatrick-Fleming Q'_k scheme, the convergence pattern, and hence the computed experimental convergence orders, are almost identical to that of the wealth-dependent scheme.

As indicated above, the $Q_k(x)$ method suggested in Munk (1997a) has a different convergence pattern, and this is also reflected in the experimental convergence orders. For very coarse grids, the experimental convergence orders are meaningless for both the controls and the value function, but for fine grids the orders are close to one. As discussed in Section 4.3, the value of the time preference rate β, which is the contraction parameter of the problem, has a great impact on the convergence pattern. For higher values of β, the experimental order of convergence for our method is higher and more stable.

[9] Munk (1997a) demonstrates that the Markov chain approximation approach is equivalent to a non-standard finite difference solution of the HJB equation associated with the control problem. Since the finite differences involved are one-sided, an experimental convergence order of one is not surprising.

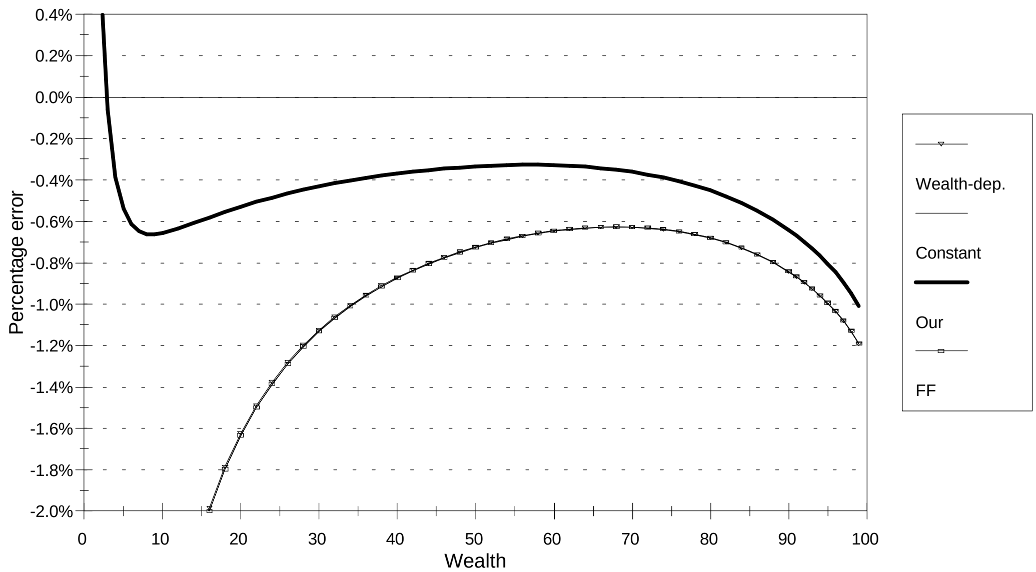

Figure 11: The accuracy of the numerically computed constant consumption equivalent for each of the four methods ($I = 400$). The graphs for the wealth-dependent scheme, the constant scheme, and the Fitzpatrick-Fleming acceleration are nearly indistinguishable.

4.2 A Comparison of the Different Methods

Below we compare the four different methods. Of course, we seek a fast and accurate method, but as usual there is a trade-off between the two objectives. First, we compare the results of the methods, when the same number (I) of wealth intervals has been used for all methods. In Figure 11 and 12, the precision of the numerically computed value function is depicted for each of the four methods – in Figure 11 for a refinement of $I = 400$ and in Figure 12 for $I = 10000$. Clearly, our $Q_k(x)$ method is by far the most accurate, and it is particularly advantageous for grids of a size manageable on a PC. The accuracy of the other three methods is almost indistinguishable with a coarse grid ($I = 400$), but with a fine grid ($I = 10000$) the precision of the constant denominator scheme is distinctively different from the wealth-dependent scheme and the Fitzpatrick-Fleming scheme.

When it comes to the precision of the numerically computed optimal controls, Figure 13 and 14 show that our method again outperforms the three other methods. As the other methods, however, our method provides imprecise control values in a wide neighborhood of the upper wealth bound.

The computer time for two different methods can differ even if the same grid is used, first of all due to the fact that the number of iterations required in the policy space algorithm differ, but also because the complexity of each iteration varies from method to method. In Table 1, we have

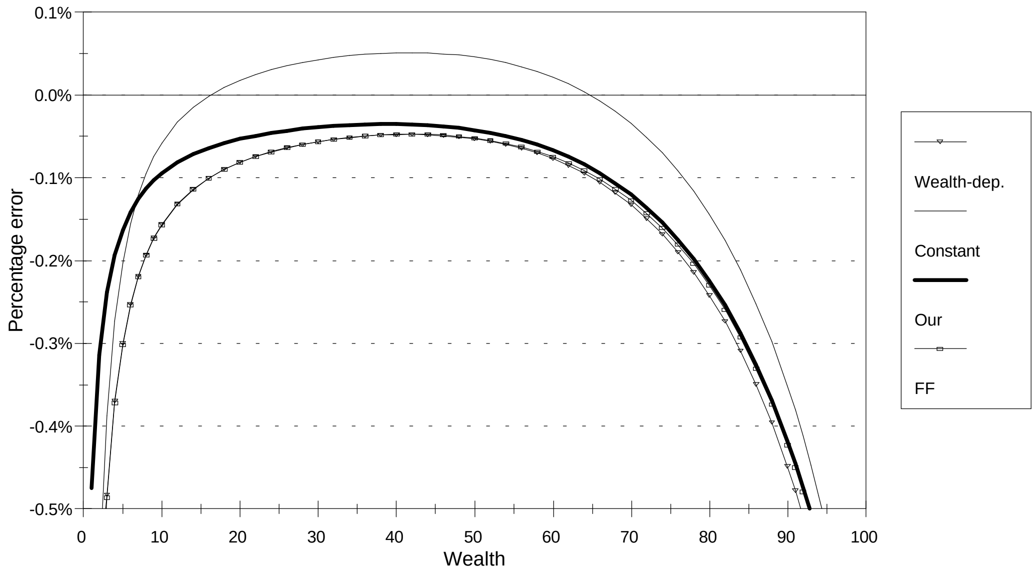

Figure 12: The accuracy of the numerically computed constant consumption equivalent for each of the four methods ($I = 10000$). Over a wide range of wealth, the graphs for the wealth-dependent scheme and the Fitzpatrick-Fleming acceleration are nearly indistinguishable.

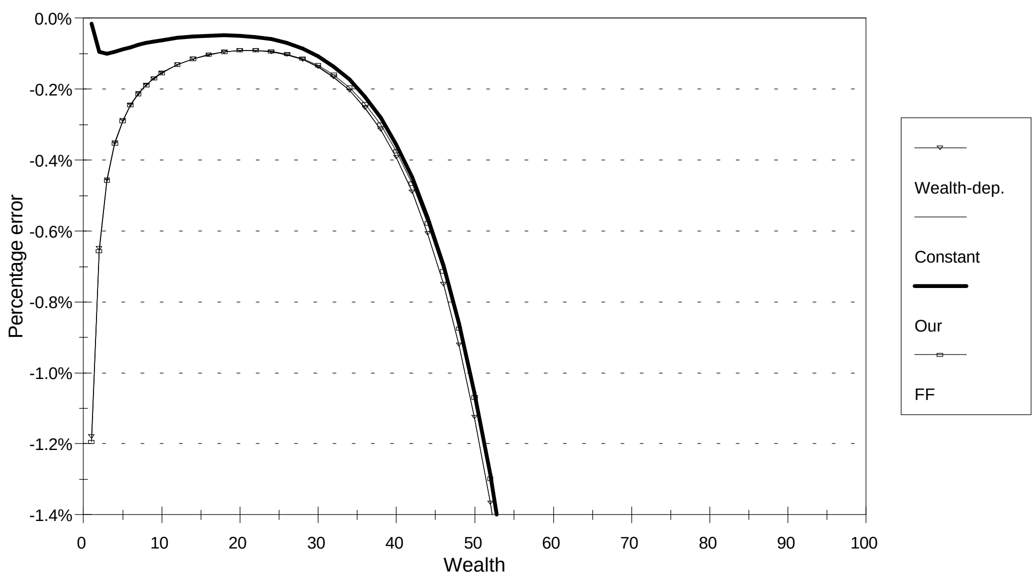

Figure 13: The accuracy of the numerically computed risky investment for each of the four methods ($I = 10000$). Over a wide range of wealth, the graphs for the wealth-dependent scheme, the constant scheme, and the Fitzpatrick-Fleming acceleration are nearly indistinguishable.

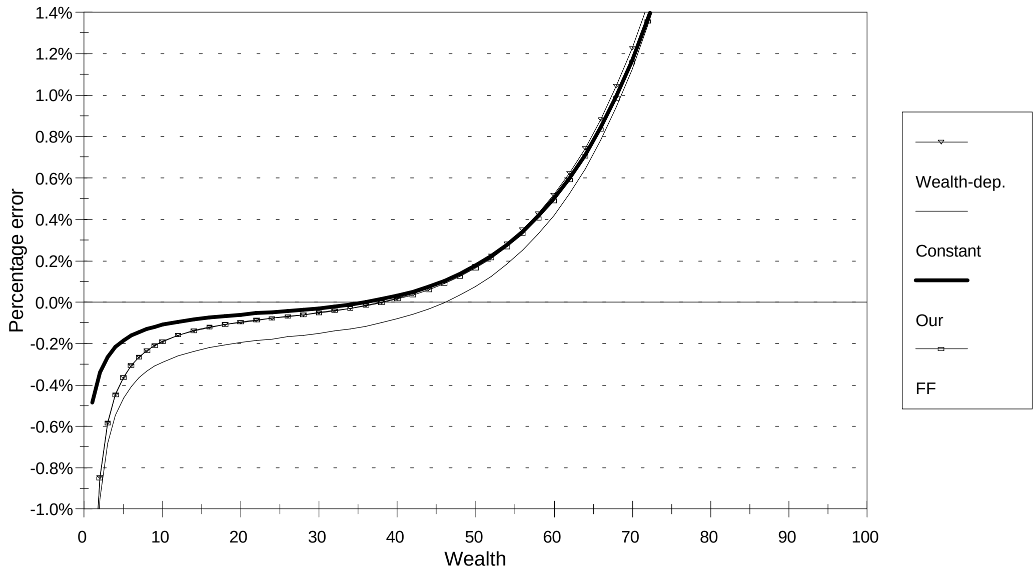

Figure 14: The accuracy of the numerically computed optimal consumption rate for each of the four methods ($I = 10000$). Over a wide range of wealth, the graphs for the wealth-dependent scheme and the Fitzpatrick-Fleming acceleration are nearly indistinguishable.

tabulated the run time of the methods on an HP9000/E35 128MB computer, as well as the number of iterations necessary for the policy space algorithm to meet the stopping criterion. Clearly, the two standard Markov chain approximation approaches (the wealth-dependent $Q^*(x)$ scheme and the constant Q' scheme) need fewer iterations and run faster than their two modifications (our $Q_k(x)$ scheme and the Fitzpatrick-Fleming Q'_k scheme).

The most fair comparison of the efficiency of the methods is to compute the precision for a fixed run time. Of course, the run time is not an input to the program, so one has to vary I until the given run time is reached. In Table 2, we give the grid refinement I and the number of policy space iterations for the constant denominator (Q') scheme, the Fitzpatrick-Fleming (Q'_k) scheme, and our ($Q_k(x)$) scheme, such that they demand (approximately) the same computer time as the wealth-dependent denominator ($Q^*(x)$) scheme with $I = 10000$.

In Figure 15, the precision of the methods is compared for the I values in Table 2, i.e., for approximately the same computer time. Obviously, our method remains the most precise, even though we use a significantly coarser grid compared to the wealth-dependent denominator Markov chain approximation approach. Furthermore, the latter now clearly outperforms the two remaining methods. Similar conclusions hold for the precision of the optimal controls.

When comparing various methods, it must be kept in mind that their performance depends on the problem being solved. The ranking of the methods indicated through our comments above

method	number of iterations	run time in seconds
wealth-dep.	7	2.68
constant	8	2.62
our	10	3.65
FF	10	3.28

Table 1: The number of policy space iterations and the run time in seconds for the four methods for $I = 10000$.

method	grid refinement I	number of iterations	run time in seconds
wealth-dep.	10000	7	2.68
constant	10200	8	2.70
our	7300	10	2.67
FF	8100	10	2.67

Table 2: The grid refinement, the number of policy space iterations, and the run time in seconds for the computations, whose results are depicted in Figure 15.

could be different, if the methods were implemented on some other control problem. On the other hand, our problem does not seem to be particularly favorable for any of the methods.

4.3 Dependence on the Contraction Parameter

In the previous subsections, we have seen that all of the four numerical methods provide rather inaccurate results for the value function and especially for the optimal controls in a wide neighborhood of the artificial upper bound on the state variable. In this section, we demonstrate that these findings are very sensitive to the value of the contraction parameter β, which in Merton's problem has the interpretation of a time preference rate for consumption.

Figure 16 and 17 show the precision of the numerically computed optimal risky investment and optimal consumption rate, respectively, for four different values of β. These results are produced with the wealth-dependent Markov chain approximation scheme, but similar conclusions can be drawn for the other three methods. The graphs show that the range of the state variable over which the numerically computed optimal controls are reliable increases dramatically as the contraction parameter β is increased. Therefore, the fixing of the upper bound relative to the interesting range

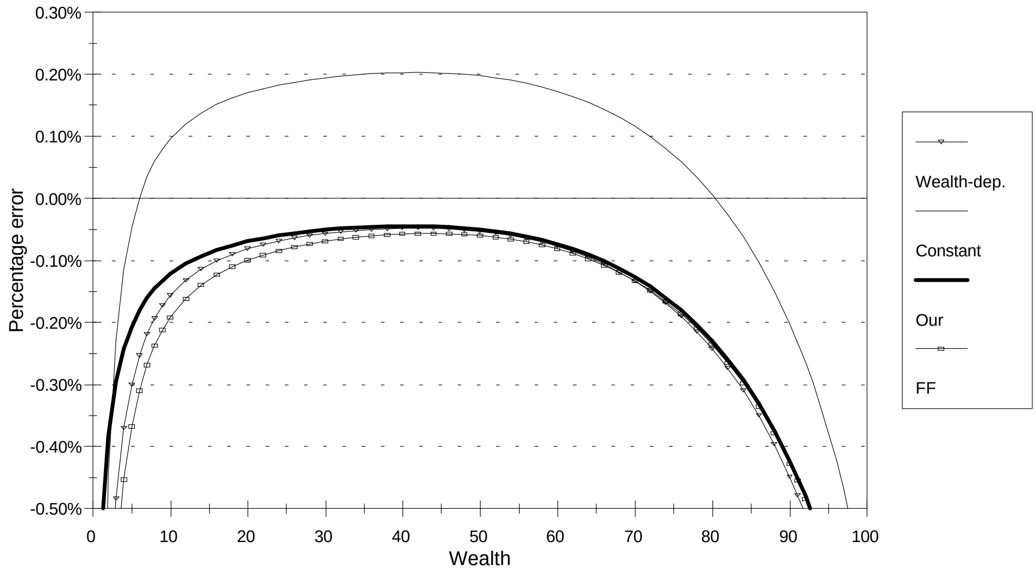

Figure 15: The accuracy of the numerically computed constant consumption equivalent with approximately identical run time for all methods.

of values of the state variable is highly dependent on the value of β. Even with a very high value of β, the computed optimal controls – in particular the optimal risky investment – are quite imprecise very close to the upper bound for the reasons given in Section 4.1.

In Figure 18, the impact of the contraction parameter value on the precision of the numerically computed value function is investigated. Recall that while the numerically computed optimal controls in the earlier sections were found to be very inaccurate near the upper bound, the computed value function was rather precise. Still, we see from Figure 18 that the accuracy of the numerically computed value function near the upper bound improves significantly, when the contraction parameter β is increased. For lower values of the state variable the precision is slightly reduced.

Obviously, the value of the contraction parameter β has a great influence on the accuracy of the Markov chain approximation methods. For stochastic control problems with the objective of maximizing an additively time-separable utility function (as in Merton's problem and other consumption/investment problems), the contraction parameter is the time preference rate of the investor. Since the time preference rate typically is assumed to be quite small, this implies that the Markov chain approximation scheme must be designed such that the numerically computed values in a rather wide neighborhood of the imposed upper bound on the state variables are ignored or at least treated with great care.

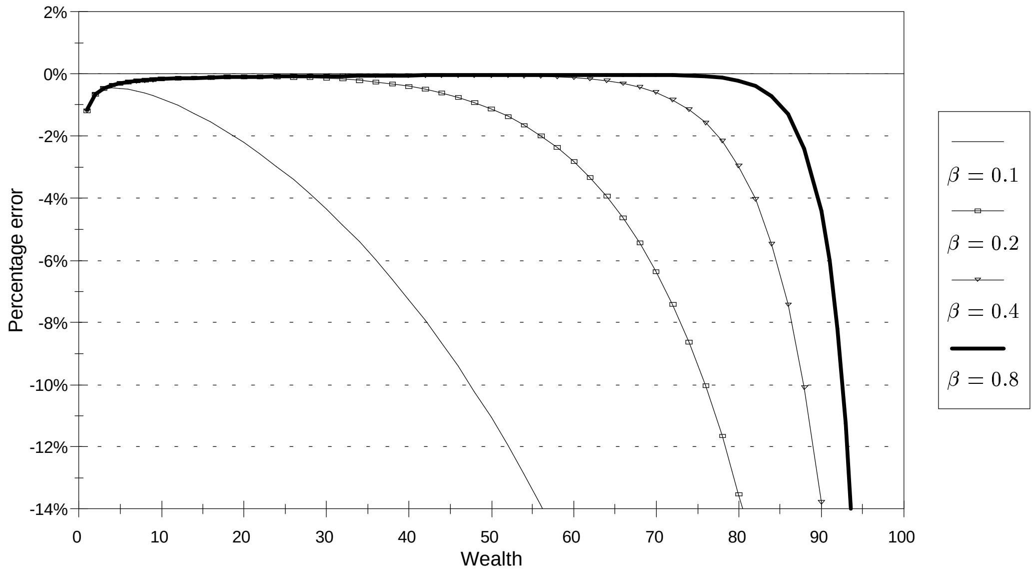

Figure 16: The accuracy of the numerically computed risky investment as a function of the time preference rate β. Results for $I = 10000$.

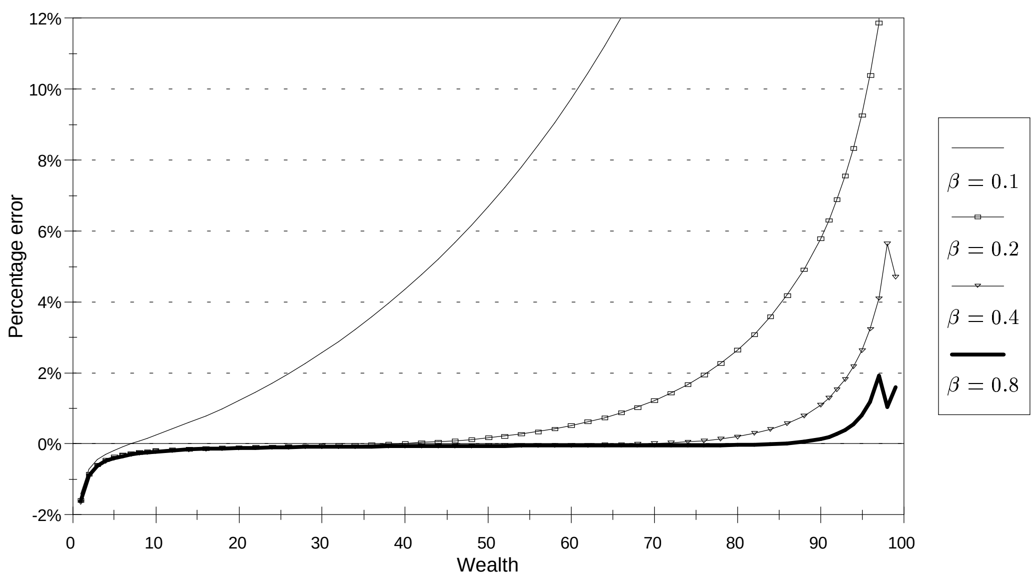

Figure 17: The accuracy of the numerically computed consumption rate as a function of the time preference rate β. Results for $I = 10000$.

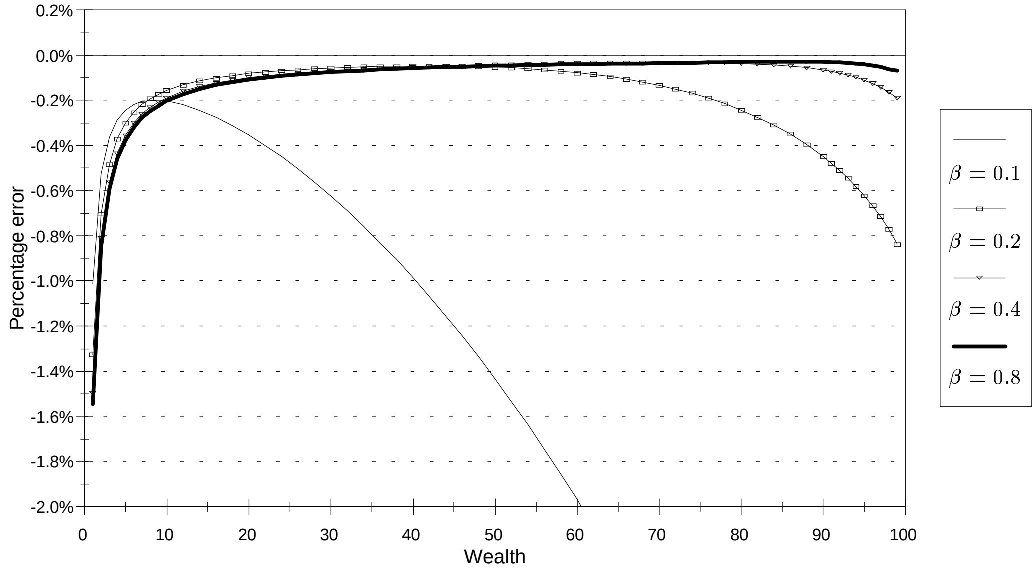

Figure 18: The accuracy of the numerically computed constant consumption equivalent as a function of the time preference rate β. Results for $I = 10000$.

4.4 Extrapolation

The numerical methods discussed in this paper involve the computation of an approximation to some wanted number. This approximation depends on a single parameter – the grid size parameter h. Hence, we can exploit the Richardson extrapolation technique, which is described in most intermediate-level textbooks on numerical analysis, e.g. Buchanan and Turner (1992). The Richardson extrapolation technique is, loosely speaking, a way to accelerate convergence or, in other words, to improve the precision of the computed approximations without involving additional computations. Note that in our solution of Merton's problem, we approximate a lot of numbers simultaneously, namely $V(x)$, $\theta(x)$, and $c(x)$ by $V^h(x)$, $\theta^h(x)$, and $c^h(x)$ for all $x \in \mathcal{R}_h$. The Richardson extrapolation method is applied to each of these approximations.

To perform the extrapolation, we must know the order of convergence of the numerical approximation method. In Section 4.1, we estimated the convergence order of the different methods to be one for both the optimal controls and the value function, cf. Figure 10. If we have computed, say, an estimate $V^h(x)$ of $V(x)$ for two different h-values h_1 and $h_2 = h_1/2$, we can compute the order one Richardson extrapolated value as $2V^{h_2}(x) - V^{h_1}(x)$.

Consider first the wealth-dependent method, i.e., the Markov chain approximation approach with $Q^*(x)$ as the denominator in the transition probabilities. In Figure 19, the curve tagged 'ri(100,200)' shows the results of a Richardson extrapolation of the values computed for $I = 100$, i.e. $h = 1$, and $I = 200$, i.e. $h = 0.5$. Similarly for the other curves tagged 'ri(·)'. From Figure 19, we see that the 'ri(100,200)' value function is much more precise than the $I = 800$, and approximately

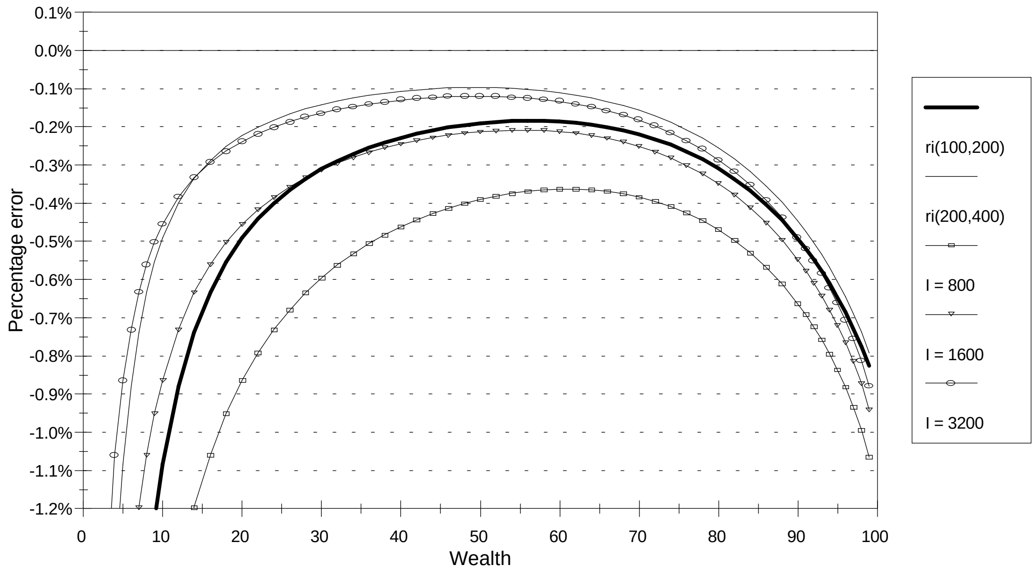

Figure 19: The accuracy of the numerically computed constant consumption equivalent for the $Q^*(x)$ method using one Richardson extrapolation.

as precise as the $I = 1600$ value functions, whereas the accuracy of the 'ri(200,400)' values is comparable to that of the $I = 3200$ results. The run time is approximately proportional to the refinement I, so the Richardson extrapolated method achieves the same precision as the original method much faster (approximately a factor 1:5). Another important point is that, due to memory limitations, it may be impossible to implement the original algorithm with a refinement greater than some level of I, depending on the computer and the programming software. In that case, the Richardson extrapolation technique offers additional precision under the given restrictions.

The effects of the Richardson extrapolation on the numerically computed optimal controls are depicted in Figure 20 and 21. Roughly the same conclusions can be drawn from these pictures as from the value function graph discussed above. In addition, we note that extrapolation does not substantially improve the precision of the controls for high wealth levels.

Encouraged by the success of the Richardson extrapolation procedure discussed above, we try another extrapolation, that is, we extrapolate the extrapolated results.[10] To do that consistently, we first estimate the convergence order of the extrapolated results. This is depicted in Figure 22. Although the convergence order is certainly higher than without extrapolation, cf. Figure 10, the order is still close to one. Therefore, we perform another Richardson extrapolation to the order one.

[10]For details on repeated Richardson extrapolations, we refer the interested reader to Christiansen and Petersen (1989).

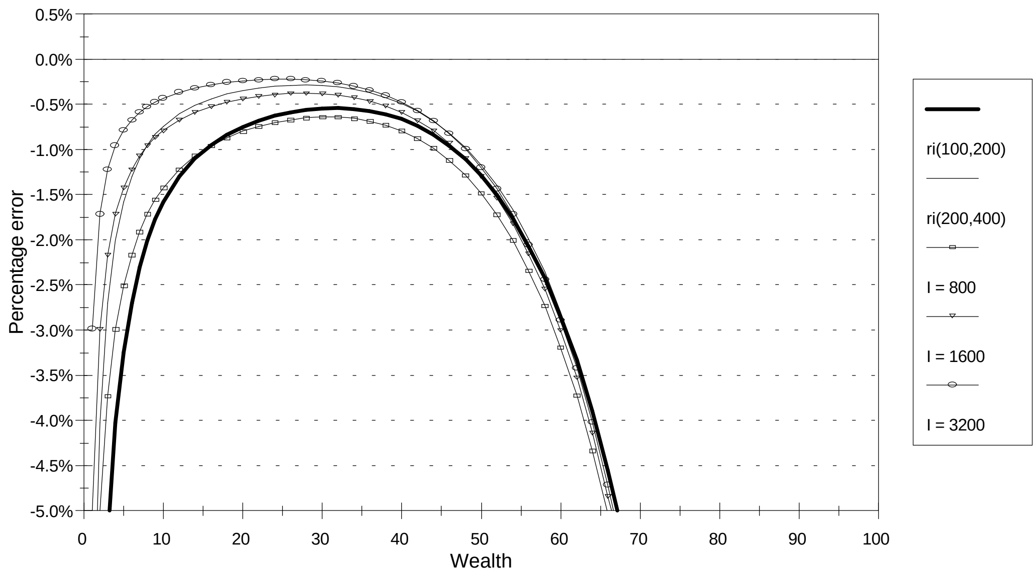

Figure 20: The accuracy of the numerically computed risky investment for the $Q^*(x)$ method using one Richardson extrapolation.

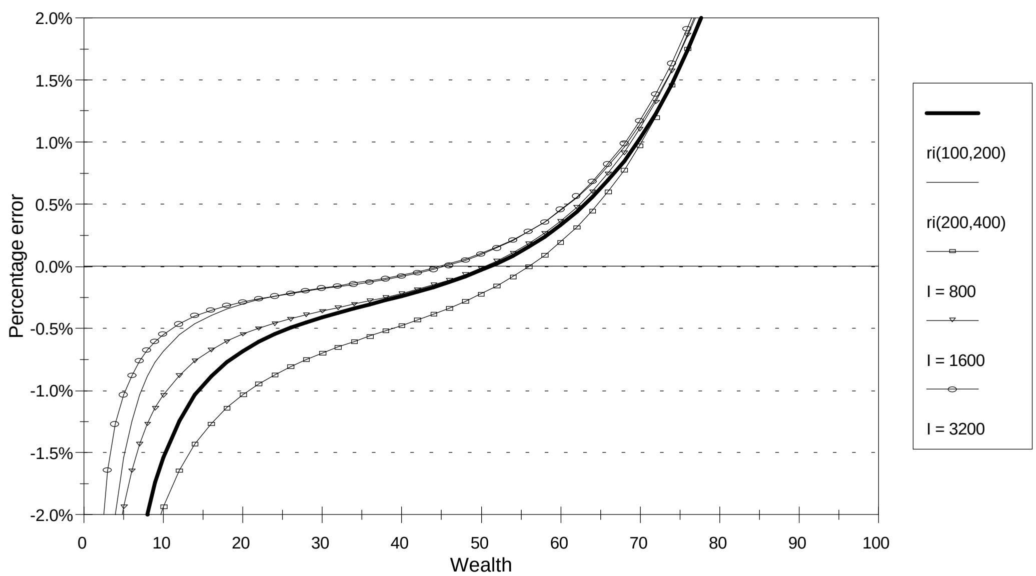

Figure 21: The accuracy of the numerically computed consumption rate for the $Q^*(x)$ method using one Richardson extrapolation.

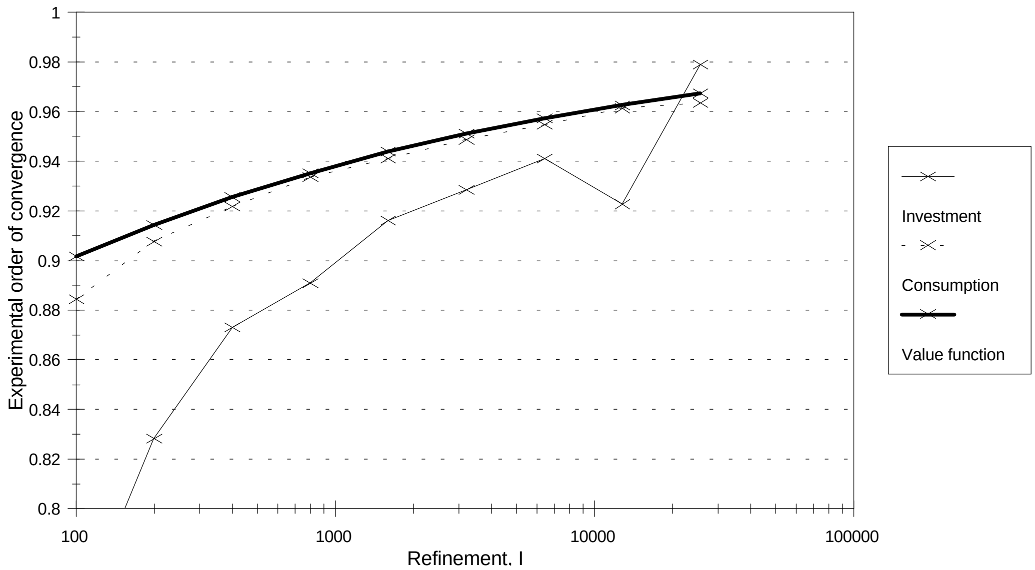

Figure 22: The experimental convergence order for the $Q^*(x)$ method after a first order Richardson extrapolation. Computed for an initial wealth of $x = 50$.

Figure 23 graphs the precision of the constant consumption equivalent computed by the $Q^*(x)$ method, this time with a curve representing doubly extrapolated results. The curve tagged 'ri(100,200,400)' is the results from an order one extrapolation of (i) an order one extrapolation of the $I = 100$ and $I = 200$ results, and (ii) an order one extrapolation of the $I = 200$ and $I = 400$ results. The accuracy of the value function computed with the 'ri(100,200,400)' method is comparable to that of the $I = 12800$ and $I = 25600$ results, although there are fewer computations, and hence less computer time is required, as can be seen by comparing $100 + 200 + 400 = 700$ with 12800 or 25600.

The success of the Richardson extrapolation technique depends highly on the stability of the estimated experimental convergence orders and hence on the smoothness of the convergence of the numerical method applied. Above, we considered extrapolation of the $Q^*(x)$ method, which displays a very smooth convergence pattern, as demonstrated in Section 4.1. In that section, we also discussed the convergence pattern and experimental convergence orders of the other three methods. Since the Fitzpatrick-Fleming method shows a convergence pattern almost identical to the $Q^*(x)$ method, superimposing an order one Richardson extrapolation will provide almost the same gain in precision as for the $Q^*(x)$ method. For the constant denominator scheme, Richardson extrapolation would be most succesful for coarse grids, where the convergence is smoother than for fine grids. Conversely for our method, which shows smooth convergence for fine grids. Therefore, a Richardson extrapolation of the results obtained with our method using coarse grids cannot reduce

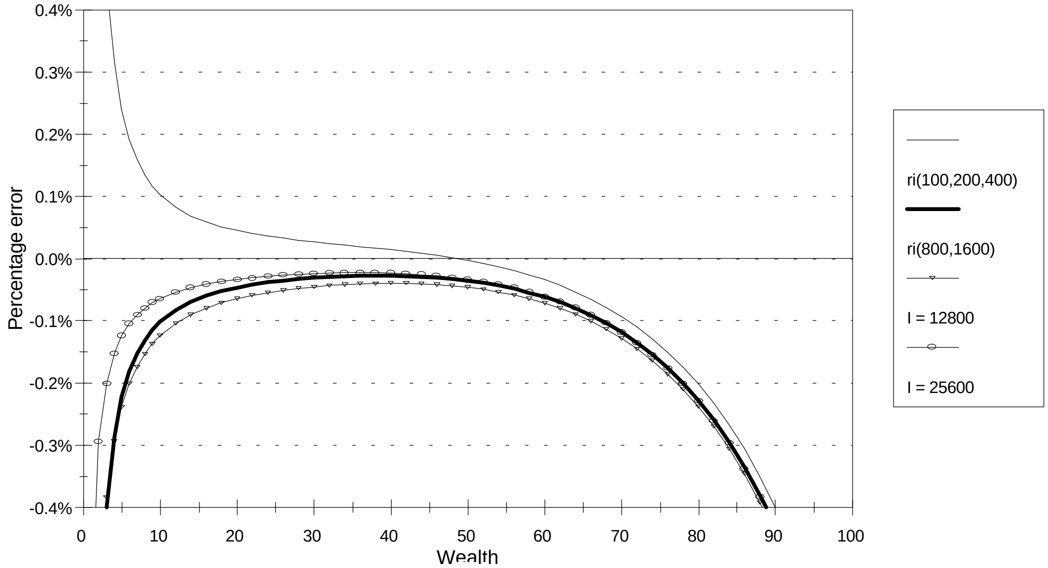

Figure 23: The accuracy of the numerically computed constant consumption equivalent for the $Q^*(x)$ method using two Richardson extrapolations.

the error substantially. Recall however that our method was significantly more accurate than the other methods.

5 Concluding Remarks

We have studied the experimental performance of two Markov chain approximation approaches and two closely related alternatives on a problem with a known solution, namely Merton's consumption/investment problem. We have found that the methods yield efficient approximations of the solution to the underlying continuous-time stochastic control problem. For all the methods considered, the numerically computed optimal controls are generally very precise, except near both the lower bound and especially near the artificial upper bound, whereas the numerically computed value function is imprecise only at the lower bound. While the precision of the methods near the lower bound can be improved substantially be refining the grid, this does not help near the upper bound. This is very important to have in mind, when the methods are applied to problems without a known solution. Of the four methods studied, the method suggested by us is the most efficient. We have also demonstrated how the application of Richardson extrapolation techniques is very successful in reducing the error of the numerical solution without increasing the computation time.

References

Buchanan, J. L. and P. R. Turner (1992). *Numerical Methods and Analysis*. McGraw-Hill.

Bunch, D. and H. Johnson (1992). A Simple and Numerically Efficient Valuation Method for American Puts Using a Modified Geske-Johnson Approach. *The Journal of Finance 47*, 809–816.

Christiansen, E. and H. G. Petersen (1989). Estimation of Convergence Orders in Repeated Richardson Extrapolation. *BIT 29*, 48–59.

Duffie, D. (1996). *Dynamic Asset Pricing Theory* (Second ed.). Princeton, New Jersey, USA: Princeton University Press.

Fitzpatrick, B. G. and W. H. Fleming (1991). Numerical Methods for an Optimal Investment-Consumption Model. *Mathematics of Operations Research 16*(4), 823–841.

Fleming, W. H. and H. M. Soner (1993). *Controlled Markov Processes and Viscosity Solutions*, Volume 25 of *Applications of Mathematics*. New York: Springer-Verlag.

Geske, R. and H. E. Johnson (1984). The American Put Valued Analytically. *The Journal of Finance 39*(5), 1511–1524.

Karatzas, I., J. P. Lehoczky, S. P. Sethi, and S. E. Shreve (1986). Explicit Solution of a General Consumption/Investment Problem. *Mathematics of Operations Research 11*(2), 261–294.

Kushner, H. J. (1977). *Probability Methods for Approximations in Stochastic Control and for Elliptic Equations*, Volume 129 of *Mathematics in Science and Engineering*. New York: Academic Press.

Kushner, H. J. (1990). Numerical Methods for Stochastic Control Problems in Continuous Time. *SIAM Journal on Control and Optimization 28*(5), 999–1048.

Kushner, H. J. and P. G. Dupuis (1992). *Numerical Methods for Stochastic Control Problems in Continuous Time*, Volume 24 of *Applications of Mathematics*. New York: Springer-Verlag.

Merton, R. C. (1969). Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case. *Review of Economics and Statistics 51*, 247–257. Reprinted as Chapter 4 in Merton (1992).

Merton, R. C. (1971). Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory 3*, 373–413. Erratum: Merton (1973). Reprinted as Chapter 5 in Merton (1992).

Merton, R. C. (1973). Erratum. *Journal of Economic Theory 6*, 213–214.

Merton, R. C. (1992). *Continuous-Time Finance.* Padstow, UK: Basil Blackwell Inc.

Munk, C. (1997a). Numerical Methods for Continuous-Time, Continuous-State Stochastic Control Problems. Publications from Department of Management 11, Odense University.

Munk, C. (1997b). *Optimal Consumption/Portfolio Policies and Contingent Claims Pricing and Hedging in Incomplete Markets.* Ph. D. thesis, School of Business and Economics, Odense University, DK-5230 Odense M, Denmark.

Rust, J. (1996). Numerical Dynamic Programming in Economics. In H. M. Amman, D. A. Kendrick, and J. Rust (Eds.), *Handbook of Computational Economics*, Volume I. North-Holland.

Tapiero, C. S. and A. Sulem (1994). Computational Aspects in Applied Stochastic Control. *Computational Economics 7*, 109–146.