Comput Manag Sci (2009) 6:187–208
DOI 10.1007/s10287-008-0089-9

ORIGINAL PAPER

A stochastic programming approach for multi-period portfolio optimization

Alois Geyer · Michael Hanke · Alex Weissensteiner

Published online: 27 December 2008

Abstract This paper extends previous work on the use of stochastic linear programming to solve life-cycle investment problems. We combine the feature of asset return predictability with practically relevant constraints arising in a life-cycle investment context. The objective is to maximize the expected utility of consumption over the lifetime and of bequest at the time of death of the investor. Asset returns and state variables follow a first-order vector auto-regression and the associated uncertainty is described by discrete scenario trees. To deal with the long time intervals involved in life-cycle problems we consider a few short-term decisions (to exploit any short-term return predictability), and incorporate a closed-form solution for the long, subsequent steady-state period to account for end effects.

Keywords Life-cycle asset allocation · Stochastic linear programming · Scenario trees · VAR(1) process

1 Introduction

The classical treatments of strategic asset allocation can be traced back to Samuelson (1969) and Merton (1969, 1971). In the light of Markowitz' seminal papers

A. Geyer
WU Wirtschaftsuniversität Wien and Vienna Graduate School of Finance,
Heiligenstädter Straße 46, 1190 Vienna, Austria

M. Hanke · A. Weissensteiner (✉)
Department of Banking and Finance, University of Innsbruck,
Universitätsstraße 15, 6020 Innsbruck, Austria
e-mail: Alex.Weissensteiner@uibk.ac.at

on single-period portfolio selection, the early literature focused on conditions leading to the optimality of *myopic policies*, i.e., conditions under which portfolio decisions for multi-period problems coincide with those for single period problems. In addition, the lack of computing power lead to formulate models driven by the quest for closed-form solutions. To achieve these objectives, rather restrictive assumptions were made, and many of these models' results turned out to be inconsistent with conventional wisdom as expressed by the so-called Samuelson puzzle: whereas one of the main results from early multi-period portfolio models is that the fractions of risky assets are constant over time, this contradicts the advice obtained from many professionals in practice that investors should hold a share of risky assets which declines steadily as they approach retirement (often called the *age effect*). Since then, many researchers have tried to resolve this puzzle which is mainly rooted in some of the (simplifying) assumptions used in early models (fixed planning horizon, time-constant investment opportunities, no intermediate consumption, etc.).

Research in the area of life-cycle asset allocation models regained momentum in the early 1990s for two main reasons: first, a number of economic factors increased the number of people with sizeable wealth to invest (the "generation of heirs"), coupled with increased uncertainty about the security of public pension systems. Second, the enormous increase in computer power enabled the solution of models with more realistic assumptions. A number of additional features have been added to the classical models, in many cases with the goal of resolving the Samuelson puzzle: stochastic labor income, time-varying investment opportunities, parameter uncertainty (with and without learning), special treatment of certain asset classes (real estate), and habit formation, to name just the most important developments.

In contrast to other approaches in the literature using non-linear optimization (see, e.g., Blomvall and Lindberg 2002; Gondzio and Grothey 2007), we use multi-period stochastic linear programming (SLP) to solve the problem of optimal life-cycle asset allocation and consumption. This method has been explicitly chosen with the practical application of the approach in mind. Many features which are considered important for investment decisions in practice can be easily incorporated when using SLP. For example, personal characteristics of the investor can be taken into account (e.g., mortality risk, risk attitude, retirement, future cash flows for major purchases or associated with other life events). Combined with the availability of efficient solvers, this explains why the SLP approach has been successfully applied to a wide range of problems (see, e.g., Wallace and Ziemba 2005). To nest classical analytical results from this area within our model, we maximize expected utility of consumption over the investor's lifetime and expected utility of bequest rather than other objectives which can be implemented more easily (e.g., piecewise linear or quadratic penalty functions, or minimizing CVaR).

The paper is organized as follows: in Sect. 2 we provide a classification of the more recent life-cycle asset allocation models based on the type of available solutions. Section 3 describes the stochastic programming model, in particular the formulation of the objective, the optimization approach for its linearization, and the generation of scenarios. In Sect. 4 results from the SLP are compared to those in Campbell et al. (2003), and results for an extended setting are presented. Section 5 concludes.

2 Overview of solution methods

Many papers try to extend the classic Merton framework along different lines while maintaining analytical solutions (see, e.g., Bodie et al. 1992; Balvers and Mitchell 1997; Kim and Omberg 1996; Wachter 2002; Liu 2007). Analytical solutions are available for restrictive assumptions on the utility structure and the planning horizon. Another set of models obtain solutions which are exact only under (generally less stringent) assumptions, and approximately correct if these assumptions are not exactly met. In some cases, these approximate solutions are available in closed form, while others must be solved numerically. Approximate analytical solutions are provided by, e.g., Campbell and Viceira (1999, 2001, 2002), Campbell et al. (2004), and Chacko and Viceira (2005). Approximate numerical solutions can be found in, e.g., Schroder and Skiadas (1999) and Campbell et al. (2003).

To give some examples for the restrictive assumptions mentioned above, a number of the models from this category assume either a deterministic or an infinite planning horizon. Some of the finite-horizon models define utility over terminal wealth only. These assumptions are clearly problematic for individuals who face an uncertain lifetime and derive their utility mainly from what they consume *during* their lives, and not only from their bequest.

An important reference for the present paper is Campbell et al. (2003). They model asset returns and state variables as a first-order vector autoregression VAR(1) and consider Epstein–Zin utility with an infinite planning horizon. Additional assumptions include the absence of borrowing and short-sale constraints. Linearizing the portfolio return, the budget constraint, and the Euler equation, they arrive at a system of linear-quadratic equations for portfolio weights and consumption as functions of state variables. This system of equations can be solved analytically, yielding solutions which are exact only for a special case (very short time intervals and elasticity of intertemporal substitution equal to one), and accurate approximations in its neighborhood. In Sect. 4, we replicate their results as far as possible and subsequently exemplify the application of the SLP approach by investigating aspects beyond the scope of their setting, such as constraints on asset weights, transaction costs, and labor income.

Two main types of numerical solution methods can be found in the literature: One approach works via grid methods discretizing the state space, the other is based on Monte Carlo simulation. Grid discretizations are used in, among others, Brennan et al. (1997), Barberis (2000), Campbell et al. (2001), Cocco et al. (2005), and Gomes and Michaelides (2005). The main drawback of this approach is that the reduction in the state-space dimensionality, which is crucial for the solution in terms of computation time, requires to restrict the investment opportunity set (usually to one risky and one riskless asset). This may be inappropriate for many investors. Detemple et al. (2003) and Brandt et al. (2005) use simulation-based approaches. Detemple et al. approximate deviations from a closed-form solution, while Brandt et al. provide an approach that is inspired by the option pricing algorithm by Longstaff and Schwartz (2001).

The SLP used in the present paper has been applied successfully to a number of related problems. To cite only a few examples, there are applications in insurance (Cariño and Ziemba 1994, 1998; Cariño et al. 1998), and the pension fund industry (e.g., Gondzio and Kouwenberg 2001). Zenios (1999) surveys large-scale applications of SLP to fixed income portfolio management. General aspects of applying such models in a strategic asset allocation context are discussed in Ziemba and Mulvey (1998)), Pflug and Swietanowski (2000), Gondzio and Kouwenberg (2001), Wallace and Ziemba (2005), and Geyer and Ziemba (2007). Particular aspects that are relevant in a life-cycle portfolio context are discussed in Geyer et al. (2007).

3 Model description

We consider the consumption and investment decisions of an investor with uncertain lifetime. We start by introducing our notation and key variables. N is the number of assets the investor can choose from. t denotes stages (points in time) and runs from $t = 0$ (now) to $t = T$. T is the number of time intervals. τ_t is the number of years between stage t and stage $t + 1$ and the total number of years covered (the planning horizon) is given by $\tau = \tau_0 + \cdots + \tau_{T-1}$. Given the current age of the investor we define the planning horizon such that his maximum age is 101 years (the mortality tables we use assign a conditional probability of 1 for a person to die between age 100 and 101).

3.1 Variables

The following (decision) variables are used in the model formulation:

$C_0 \geq 0$, $\tilde{C}_t \geq 0$ $(t = 1,\ldots, T - 1)\ldots$ consumption in t; e.g., $\tilde{C}_2$ is the amount set aside in $t = 2$ for consumption between $t = 2$ and $t = 3$.
$\tilde{R}_t^i$ $(t = 1, \ldots, T; i = 1, \ldots, N)\ldots$ gross return of asset i for the period that ends in t.
$P_0^i \geq 0$, $\tilde{P}_t^i \geq 0$ $(t = 1, \ldots, T - 1; i = 1, \ldots, N)\ldots$ amount of asset i purchased in t.
$S_0^i \geq 0$, $\tilde{S}_t^i \geq 0$ $(t = 1, \ldots, T - 1; i = 1, \ldots, N)\ldots$ amount of asset i sold in t.
q_p^i and $q_s^i \ldots$ transaction costs for purchases and sales of asset i.
W_0^i, $\tilde{W}_t^i$ $(t = 1, \ldots, T - 1; i = 1, \ldots, N)\ldots$ total amount invested in asset i in t; e.g., $\tilde{W}_2^i$ is the amount invested in asset i in $t = 2$; in $t = 3$ the value of this investment will be $\tilde{W}_2^i \tilde{R}_3^i$.
$w_0^i \ldots$ initial value of asset i (before transactions).
$\tilde{B}_t \geq 0$ $(t = 1, \ldots, T)\ldots$ bequest in t given by $\tilde{B}_t = \sum \tilde{R}_t^i \tilde{W}_{t-1}^i$.
$\tau_t (t = 0, \ldots, T - 1)\ldots$ the number of years between stage t and stage $t + 1$.
$\varphi_y \ldots$ the (conditional) probability to survive the year following year y.
$\Phi(y_t, \tau_t)\ldots$ the probability to survive the period of length τ_t starting at stage t at an age of y_t years; $\Phi(y_t, \tau_t) = \prod_{k=y_t}^{y_t+\tau_t-1} \varphi_k$.
$\Lambda_t (t = 1, \ldots, T - 1)\ldots$ the probability to be alive at stage t (at an age of y_t); $\Lambda_t = \prod_{k=0}^{t-1} \Phi(y_k, \tau_k)$.
$\Theta_t (t = 1, \ldots, T)\ldots$ the probability to die between stage $t - 1$ and t; $\Theta_t = \Lambda_{t-1}[1 - \Phi(y_t, \tau_t)]$.

$L_t (t = 0, \dots, T-1) \dots$ labor income in t; e.g., L_2 is the present value of labor income received between $t = 2$ and $t = 3$.
$F_t (t = 0, \dots, T-1) \dots$ fixed cash flow paid or received in t; e.g., F_2 is the present value of cash flows paid or received between $t = 2$ and $t = 3$.
$r \dots$ the risk-free interest rate.
$\gamma \dots$ the coefficient of relative risk aversion.
$\delta \dots$ the investor's time preference rate, $d = \exp\{-\delta\}$ is the time discount factor, and D_t is the time discount factor applicable at stage t:

$$D_t = \exp\left\{-\delta \sum_{i=0}^{t-1} \tau_i\right\}.$$

The stochastic returns $\tilde{R}_t^i$ describe the uncertainty faced by the investor. The procedure to simulate their values and to construct the scenario tree is described in Sect. 3.5. C_0, $\tilde{C}_t$, W_0^i, $\tilde{P}_t^i$, $\tilde{S}_t^i$, $\tilde{W}_t^i$ and $\tilde{B}_t$ are the decision variables of the problem and their values are obtained from the optimal solution of the stochastic linear program.

Labor income is computed on the basis of initial labor income $\mathcal{L}_0$, the annual labor growth rate ℓ, the number of years until retirement y_r, and the fraction of income during retirement f_r. The annual stream of income before retirement is given by (the index y denotes years) $\mathcal{L}_y = \mathcal{L}_0 \exp\{y\ell\}$ $(y = 1, \dots, y_r)$ and by $\mathcal{L}_y = f_r \mathcal{L}_{y_r} \exp\{(y - y_r)\ell_r\}(y = y_r + 1, \dots, \tau)$ after retirement, where ℓ_r is the growth rate of labor income after retirement. The present value of labor income used in the budget constraints (see below) is defined as

$$L_t = \sum_{y=j_t}^{k_t} \mathcal{L}_y[(1 - \Phi(y_t, y-1)(1 - \varphi_y)) \exp\{-r(y - j_t + 1)\}], \tag{1}$$

where

$$j_t = 1 + \sum_{i=0}^{t-1} \tau_i \quad k_t = j_t + \tau_t - 1.$$

$\Phi(y_t, y-1)$ is the probability to survive until the beginning of year y given age y_t at stage t, and $(1 - \varphi_y)$ is the probability to die in the subsequent year. Labor income $\mathcal{L}_y$ is thus reduced by an amount that corresponds to the premium of a fairly priced life insurance (Richard 1975).

3.2 Constraints

The budget equations are given by

$$C_0 + \sum_{i=1}^{N} P_0^i (1 + q_p^i) = \sum_{i=1}^{N} S_0^i (1 - q_s^i) + L_0 + F_0$$

$$\tilde{C}_t + \sum_{i=1}^{N} \tilde{P}_t^i (1 + q_p^i) = \sum_{i=1}^{N} \tilde{S}_t^i (1 - q_s^i) + L_t + F_t \quad t = 1, \ldots, T - 1.$$

The value of investments accumulates according to the following equations:

$$\begin{aligned} W_0^i &= w_0^i + P_0^i - S_0^i \quad i = 1, \ldots, N \\ \tilde{W}_t^i &= \tilde{R}_t^i \tilde{W}_{t-1}^i + \tilde{P}_t^i - \tilde{S}_t^i \quad t = 1, \ldots, T - 1,\ i = 1, \ldots, N \\ \tilde{W}_T^i &= \tilde{R}_T^i \tilde{W}_{T-1}^i \quad i = 1, \ldots, N. \end{aligned}$$

To model restrictions on the portfolio composition we use the constraints

$$l_i \leq \frac{\tilde{W}_t^i}{\sum_{i=1}^{N} \tilde{W}_t^i} \leq u_i \quad t = 0, \ldots, T - 1, \tag{2}$$

where u_i is the maximum and l_i the minimum weight of asset i in the portfolio. Short sales can be excluded by $l_i = 0$ or limited by setting l_i equal to minus the maximum leverage of asset i. In general the decision variables $\tilde{W}_t^i$ can become negative. However, total wealth must be positive in all periods:

$$\sum_{i=1}^{N} \tilde{W}_t^i \geq 0 \quad t = 0, \ldots, T.$$

3.3 Objective

The objective is to maximize the expected utility of consumption over the lifetime and of bequest at the time of death of the investor:

$$U(C_0) + E\left[\sum_{t=1}^{T} \Theta_t U(\tilde{B}_t) D_t + \sum_{t=1}^{T-1} \Lambda_t U(\tilde{C}_t) D_t\right] \longrightarrow \max .$$

U is a power utility function with constant relative risk aversion γ. Θ_t is the probability to die between stage $t - 1$ and t, and Λ_t is the probability to be alive in t.

The variable $\tilde{C}_t$ is defined as the amount of money set aside for consumption between stage t and $t + 1$. If a small number of stages is used to keep the problem size manageable, the time intervals between stages must be very long to cover the lifetime of an investor. When stages cover more than 1 year, we assume that the amount $\tilde{C}_t$ is not consumed at once but in annual parts $\tilde{c}_{t,j},\ j = 0, 1, \ldots, \tau_t - 1$. To account for the fact that $U(\tilde{C}_t) \neq \sum_j U(\tilde{c}_{t,j})$, Geyer et al. (2007) derive a relation between $\tilde{C}_t$ and $\tilde{c}_{t,j}$ by optimizing sub-period consumption $\tilde{c}_{t,j}$ for any given $\tilde{C}_t$. This allows replacing the original decision variables C_0 and $\tilde{C}_t$ (which refer to consumption in an entire period) by the annual consumption variables $c_{0,0}$ and $\tilde{c}_{t,0}$. We define

$$\alpha_{t,j} = \exp\left\{\frac{r-\delta}{\gamma}\right\} \Phi(y_t, j)^{1/\gamma}$$

and

$$\alpha_t = \sum_{j=0}^{\tau_t - 1} \Phi(y_t, j)^{1/\gamma} \exp\left\{\frac{j(r(1-\gamma)-\delta)}{\gamma}\right\} = \sum_{j=0}^{\tau_t - 1} \alpha_{t,j} \exp\{-jr\}.$$

α_t can be interpreted as a kind of "annuity factor", taking into account the risk-free interest rate, mortality risk, and the optimal allocation of consumption to sub-periods. The budget constraints are now formulated as

$$\alpha_0 c_{0,0} + \sum_{i=1}^{N} P_0^i (1+q_p^i) = \sum_{i=1}^{N} S_0^i (1-q_s^i) + L_0 + F_0$$

$$\alpha_t \tilde{c}_{t,0} + \sum_{i=1}^{N} \tilde{P}_t^i (1+q_p^i) = \sum_{i=1}^{N} \tilde{S}_t^i (1-q_s^i) + L_t + F_t \quad t = 1, \ldots, T-1.$$

Utility of consumption in t is formulated in terms of $\tilde{c}_{t,0}$

$$U(\tilde{C}_t) = \sum_{j=0}^{\tau_t - 1} \Phi(y_t, j) \exp\{-j\delta\} U(\tilde{c}_{t,0}\alpha_{t,j}).$$

Since we consider non-linear objective functions (power utility of consumption and bequest), this may raise the question why we use SLP rather than non-linear optimization techniques. While good solvers for non-linear programs are quite expensive, several (stochastic) LP solvers are available for free (e.g., from the COIN project).

3.4 Choice of breakpoints

To be able to use linear programming solvers, non-linear objective functions need to be linearized. For that purpose a function $f(x)$ is approximated by m linear segments between the breakpoints b_j ($j = 0, \ldots, m$). The argument x is defined in terms of non-negative decision variables v_j associated with each segment:

$$x = \sum_{j=0}^{m+1} v_j$$

$$0 \le v_0 \le b_0 \quad v_{m+1} \ge 0$$

$$0 \le v_j \le b_j - b_{j-1} \quad j = 1, \ldots, m.$$

The slopes of the linear segments are given by

$$\Delta_j = \frac{f(b_j) - f(b_{j-1})}{b_j - b_{j-1}}$$

and $f(x)$ is approximated by

$$f(x) \approx \Delta_1 v_0 + \sum_{j=1}^{m} \Delta_j v_j + \Delta_m v_{m+1}.$$

The linearization of the objective function requires choosing the number and the position of breakpoints.

We define separate breakpoints for consumption and bequest to account for the different orders of magnitude of the two variables. In addition, these variables may show considerable variation across stages which requires using different breakpoints for each stage, too.

To define the minimum and maximum breakpoints for consumption we use closed-form solutions from (Ingersoll 1987, p. 238, 242) and (Duffie 2001, p. 210) as a guideline. To determine minimum and maximum breakpoints of bequest we consider a simplified version of the problem. For all nodes of a specific stage we assume that the fraction of consumed wealth and the asset allocation is the same. We use the same returns that are subsequently used to solve the SLP. Then we define a random grid of consumption-wealth ratios and asset allocations which obey leverage constraints and other bounds. We evaluate the objective function for each element of the grid, whereby we can use the exact form of the utility function. The optimal solution provides a rough guess for the order of magnitude and the dispersion of consumption and bequest in each stage. This guess is used to define the minimum and maximum breakpoints required for the linearization of the utility function.

For the remaining breakpoints we follow Geyer et al. (2007), who use the curvature of the utility function to position the breakpoints, allocating more breakpoints to areas with a greater curvature. In contrast to other alternatives considered, this approach is faster because it requires no optimization. The algorithm first divides the interval between b_t^0 and b_t^m into n equally wide segments separated by the points β_t^j $(j = 0, \ldots, n)$ where $\beta_t^0 = b_t^0$ and $\beta_t^n = b_t^m$. The curvature for each β_t^j is defined as (for details see Hanke and Huber 2008)

$$\kappa_t^j = \frac{U''(\beta_t^j)}{(1 + [U'(\beta_t^j)]^2)^{3/2}} \quad j = 0, \ldots, n.$$

In contrast to the second derivative, this measure is invariant to the orientation of a curve in the plane. Figure 1 illustrates the relation between two utility functions and their curvature.

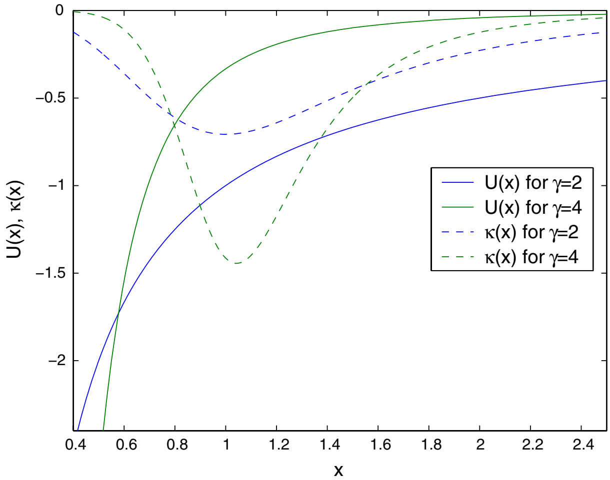

Fig. 1 Utility functions for risk aversions $\gamma = 2$ and $\gamma = 4$ together with their curvatures

The average curvature in each segment is the arithmetic mean of two consecutive curvatures

$$\overline{\kappa}_t^j = 0.5 \left(\kappa_t^{j-1} + \kappa_t^j \right) \quad j = 1, \ldots, n.$$

The relative average curvature is given by

$$\hat{\kappa}_t^j = \frac{\overline{\kappa}_t^j}{\sum_k \overline{\kappa}_t^k} \quad j = 1, \ldots, n$$

and is used to compute the number of breakpoints in each segment $n_t^j = [m \cdot \hat{\kappa}_t^j]$, where $[\cdot]$ denotes rounding to the nearest integer (surplus breakpoints can be ignored). The position of breakpoints b_t^i in the segment j is defined by

$$b_t^i = b_t^{i-1} + (\beta_t^{j+1} - \beta_t^j)/n_t^j \quad j = 1, \ldots, n; \; i = 1, \ldots, n_t^j.$$

3.5 Scenario generation and choice of intervals

The uncertainty associated with the consumption-investment problem and time-varying investment opportunities are modeled by a K-dimensional VAR(1) process as in Barberis (2000) or Campbell et al. (2003). The vector process consists of asset returns and other state variables (e.g., dividend yields or interest rate spreads). The multivariate return process evolves in discrete time, and the underlying probability distributions are approximated by discrete distributions in terms of a scenario tree (see Fig. 2). For that purpose, different approaches have been proposed in the literature. We consider two

Fig. 2 Scenario tree with node structure $2 \times 2 \times 3$

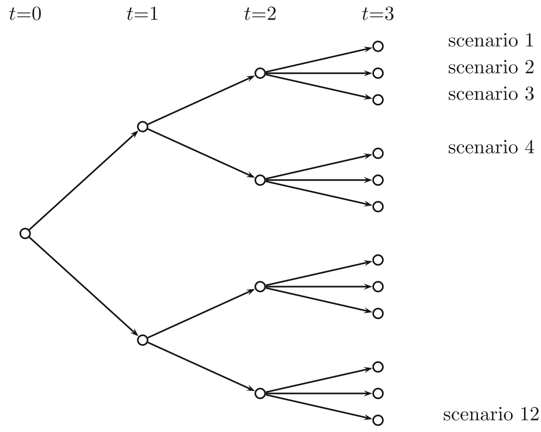

established methods, namely, scenario reduction (Pflug 2001; Heitsch and Römisch 2003; Dupačová et al. 2003) and moment matching (Høyland and Wallace 2001).

These methods differ with respect to calculating the distance between the true continuous probability distribution and its discrete approximation. Whereas scenario reduction methods explicitly compute probability metrics such as the Kantorovich (or Wasserstein) distance, moment matching methods, as their name implies, minimize distance functions based on the first few (co-)moments of these two distributions. Hochreiter and Pflug (2007) propose a scenario reduction method aimed at approximating the entire probability distribution and show that pure moment matching may lead to strange results. However, Rasmussen and Clausen (2007) point out that arbitrage opportunities may arise when using scenario reduction in a financial portfolio optimization problem.

Another approach for generating scenario trees (in the spirit of Høyland and Wallace 2001) matches the first few moments and the correlations of the simulated processes. More nodes facilitate the matching of moments but increase the number of scenarios. In the examples presented in Sect. 4 we set the branching factor $n_t = 2K + 2, \forall t$ to match the first four (co)moments within reasonable time. The numerical results in Geyer et al. (2007) show that available closed-form solutions can be replicated very well with this method and this choice of n_t. To exclude arbitrage opportunities in the simulated returns, we apply the procedure proposed by Klaassen (2002).

The number of scenarios in the tree grows at a rate $O(n^t)$, where t is the stage index and n is the constant branching factor. Given the long period of time covered by a life-cycle model, it is computationally infeasible to work with annual decision (rebalancing) intervals over the entire lifetime of an investor. To keep the total number of scenarios practically manageable (e.g., several thousand scenarios) only a rather small number of stages (e.g., three to six) and a small number of nodes is usually considered. For example, in Dempster et al. (2003) the first revision of the portfolio is made after 1 year since the initial decisions are considered to be most important. The remaining time intervals are much longer and serve to *approximate* the fact that further portfolio revisions are possible until the planning horizon is reached. This approach implies that the investor is 'locked in' in the chosen asset allocation for a considerable amount of time—possibly much longer than the planned or anticipated rebalancing interval. This problem can be partly alleviated by using more stages and shorter time intervals, but more scenarios and longer solution times would be required.

Therefore we also consider a different approach which consists of a sequence of short (e.g., 1-year) periods followed by a long, steady-state period which lasts until the maximum lifetime of the investor. This design accounts for the short-term dynamics of the VAR model in the first few years, and the possibility of frequent rebalancing. For modeling the subsequent steady state period we have considered three candidate models: Richard (1975), Campbell et al. (2003), and Jurek and Viceira (2005). Whereas the latter two account for time-variation in investment opportunities, Richard (1975) assumes a multivariate geometric Brownian motion for the risky assets. However, this model accounts for mortality risk and intertemporal consumption, whereas the other two ignore these issues by assuming either an infinite horizon Campbell et al. 2003 or a fixed, finite horizon Jurek and Viceira 2005 with utility defined over terminal wealth. Since none of these alternatives fits our purpose exactly we consider the relative importance of their assumptions. From the perspective of period $t = 0$ the predictability of asset returns implied by the VAR model several periods ahead may have little effect, given the rather weak short-term temporal dependence empirically found for asset returns (we will look at this feature more closely in Sect. 4). Since we consider mortality risk to be conceptually important in a life-cycle context, we use the model of Richard (1975) to derive the utility from optimal consumption and investment decisions in the steady-state period. This amounts to reformulate the objective function as follows:

$$U(C_0) + E\left[\sum_{t=1}^{T-1} \Theta_t U(\tilde{B}_t) D_t + \sum_{t=1}^{T-2} \Lambda_t U(\tilde{C}_t) D_t \right.$$
$$\left. + \Lambda_{T-1} J(\tilde{W}^+_{T-1}, y_{T-1}) D_{T-1}\right] \longrightarrow \max. \tag{3}$$

$J(\tilde{W}^+_{T-1}, y_{T-1})$ is the value function (4) defined in Appendix 5. It depends on available wealth $\tilde{W}^+_{T-1}$ (which includes the present value of future labor income or other cash flows) and the age of the investor y_{T-1} at the beginning of the steady-state period. As described in Appendix 5 the value function is derived in a continuous-time setting. It accounts for optimal consumption and trading, the investor's survival probability, and it is based on geometric Brownian motions for the risky assets and power utility. To implement the steady-state solution according to Richard (1975) we need to define the tangency portfolio. To be consistent with his continuous-time setting the drifts of the assets are defined as $\boldsymbol{\mu} + 0.5\text{diag}(\mathbf{C})$ (where $\boldsymbol{\mu}$ is the vector of mean log returns and $\mathbf{C}$ is their covariance matrix). An asset which earns the risk-free rate is added to the set of traded assets. Its (constant) return r is also included in the check for arbitrage opportunities in the scenario generation.

Using analytical results from a continuous-time framework in the discrete-time optimization model has obvious advantages.We avoid the unrealistic implications associated with long rebalancing intervals, and we can reduce the number of stages and the size of the scenario tree. It has to be admitted, however, that the value function does not account for restrictions on asset weights or transaction costs. There is also an inconsistency associated with combining 1-year decision intervals in discrete time with continuous consumption and trading. In our opinion, however, the advantages outweigh these drawbacks by far.

4 Numerical results

The framework described has been tested extensively against known closed-form solutions in Geyer et al. (2007). Even a small number of scenarios with only three stages and a low branching factor yield solutions which are very close to their analytic counterparts. Any remaining deviations can in part be attributed to a small bias from model error (most closed-form solutions have been derived in a continuous-time framework, whereas the SLP model is formulated in discrete time), and sampling error (given the fact that the tree representation is not unique).

Given that closed-form solutions can be replicated very well, we now analyze optimal consumption and asset allocation over the life-cycle given time-varying investment opportunities modeled by the K-dimensional VAR(1) process from Campbell et al. (2003). They consider three asset return series (ex-post real T-bill rate, excess stock returns, and excess bond returns) and three state variables (dividend-price ratio, nominal T-bill yield, and yield spread). Using their annual data set covering the period 1893–1997 (rather than 1890–1998 as stated in their paper) we can replicate their parameter estimates. We admit that there may be some finite-sample bias but share the viewpoint of Campbell et al. (2003) who take estimated VAR coefficients as given and known by the investor. Likewise we explore the implications of time-varying investment opportunities for optimal portfolios. We are well aware of the potential effects associated with parameter uncertainty but do not explicitly address them here (see, e.g., Barberis 2000).

The impulse-response function derived from the VAR parameters shows that after about three or four years the impact of shocks on the asset returns has practically vanished. The main response takes place after 1 year. This can be taken as evidence to justify our approach of using a few 1-year periods followed by a long steady-state period. Shocks to the state variables and stocks, however, remain statistically significant for up to 10 years. Although the associated effects may be economically small we investigate the sensitivity of the SLP solution to the steady-state assumption.

The tangency portfolio required for the steady-state solution according to Richard (1975) is based on the unconditional drifts $\boldsymbol{\mu} + 0.5\text{diag}(\mathbf{C})$ (using the notation from Appendix 5). The risk-free rate is set equal to the unconditional mean of real T-bill returns without adding the variance term. The analysis in Campbell et al. (2003) is based on the properties of stock and bond returns in excess of the real T-bill rate. In our setting comparable results are obtained by estimating the VAR coefficients using raw (as opposed to excess) stock and bond returns. The weight of the risk-free asset is added to the weight of T-bills, and the resulting asset is labeled ‘cash’. Comparability also requires to start simulating returns for period 1 using the unconditional means $\boldsymbol{\mu}$.

We use scenario trees with four stages, starting with three 1-year periods, and followed by a long period which lasts until the maximum age of 101 years. The node structure is $14 \times 14 \times 14$ which amounts to 2,744 scenarios. The linearization of the objective function is based on the curvature approach using 40 breakpoints.

Table 1 shows SLP-based results for optimal consumption and asset allocation in $t = 0$ for various assumptions about the investor's current age and degree of risk aversion. We find that consumption is decreasing in risk aversion, which is to be expected. The more risky an investor's asset allocation, the higher her expected returns, allowing for more consumption today. Consumption is uniformly higher for older investors who can afford to set aside less for their (shorter) future life span. We find that consumption is very precisely measured, and the standard errors of asset weights are comparatively large. As shown in Geyer et al. (2007), however, these could be reduced by using more scenarios and/or more breakpoints.

Campbell et al. (2003) use Epstein–Zin utility and an infinite horizon. They obtain numerical solutions based on linearizing the Euler equation and the budget constraint. Although our setting is quite similar, it differs in the following aspects: We use time-additive power utility and a finite horizon accounting for survival probabilities. The only case which should yield comparable results is a very young investor with log utility (Epstein–Zin and power utility coincide for $\gamma = 1$ if the elasticity of intertemporal substitution—the second parameter of Epstein–Zin utility—is equal to one). Results for cases which are comparable to Campbell et al. (2003) using the same time discount factor $d = 0.92$ show hardly any differences (see Table 1). This not only supports the use of our approach but also provides a sound basis to investigate cases beyond the scope of their setting, such as constraints on asset weights, transaction costs, and labor income.

Similar to Campbell et al. (2003) we find large stock and bond holdings financed by short positions in T-bills (see Table 1). Their asset allocation for $\gamma = 1$ shows 220% in stocks, 242% in bonds and −361% in cash, which is very close to our results (we obtain 215, 246 and −361%, respectively). The short positions decrease with risk aversion, which is again to be expected and in line with their results.

Table 1 further shows that the asset allocation does not significantly change with age. This seems to contradict results from the recent literature on asset allocation with time-varying investment opportunities (see Barberis 2000; Wachter 2002; Brandt et al. 2005). However, these models have two features in common: first, shocks to the dividend yield are strongly negatively correlated to shocks associated with equity returns. Second, the dividend yield itself (which is used as a state variable) exhibits high autocorrelation. Barberis (2000) notes that the combination of predictability in asset returns with extreme persistence in the dividend yield causes a substantial hedging demand. In Barberis (2000) and Brandt et al. (2005) these parameters are estimated from monthly or quarterly data, while a residual correlation of −1 is assumed in the analytical treatment of Wachter (2002). In contrast, we use the parameters estimated from annual data given in Campbell et al. (2003), where persistence of the dividend yield is weaker, and the correlation between shocks to the dividend yield and stock returns is only −0.725. In addition, Wachter (2002) illustrates the impact of the utility structure used in the optimization process: For utility defined over consumption (as used here), time-horizon effects are much smaller than for utility defined over terminal

Table 1 Optimal consumption and asset allocation in $t = 0$ for a man at age 20, 40, or 60, with risk aversion $\gamma = 1, 2, 5$ and 10, and time discount factor $d = 0.92$

Age of investor	20	40	60
$\gamma = 1$			
Consumption	8.2(0.0)	8.6(0.0)	10.7(0.0)
Cash	−360.7(10.4)	−358.5(10.4)	−362.2(10.4)
Stocks	214.6(2.4)	213.8(2.4)	215.6(2.4)
Bonds	246.0(8.7)	244.7(8.7)	246.6(8.7)
$\gamma = 2$			
Consumption	7.1(0.0)	7.4(0.0)	9.0(0.0)
Cash	−150.8(5.7)	−147.9(5.6)	−150.4(5.7)
Stocks	113.5(1.3)	112.1(1.2)	113.2(1.2)
Bonds	137.3(4.6)	135.8(4.6)	137.3(4.6)
$\gamma = 5$			
Consumption	4.0(0.0)	4.4(0.0)	6.0(0.0)
Cash	1.8(2.1)	2.4(2.1)	2.0(2.1)
Stocks	42.8(0.5)	42.8(0.5)	42.7(0.5)
Bonds	55.4(1.8)	54.9(1.8)	55.3(1.8)
$\gamma = 10$			
Consumption	2.6(0.0)	3.2(0.0)	4.6(0.0)
Cash	51.7(1.2)	51.4(1.2)	51.5(1.2)
Stocks	20.7(0.3)	20.8(0.3)	20.7(0.3)
Bonds	27.6(1.1)	27.8(1.1)	27.8(1.1)

The parameters for the VAR process driving asset returns and state variables are from Campbell et al. (2003, p. 58). We use scenario trees with four stages where the first three periods are each 1 year long. The node structure is 14 × 14 × 14 which amounts to 2,744 scenarios. The results are presented in terms of means and standard errors (in parentheses) from 100 solutions of the problem. The same 100 scenario trees are used for each pair of age and γ

wealth (as in Barberis 2000; Brandt et al. 2005). Finally, Barberis (2000) argues that the time-horizon effects found in the asset allocation may no longer exist if the investor has a broader range of asset classes to choose from. This is the case for our model with three risky assets compared to models with only one risky asset as in Barberis (2000) or Brandt et al. (2005).

Although the aspects discussed in the previous paragraph provide plausible explanations for the age-independence found in Table 1, one might argue that this result is mainly due to the steady-state assumption for the final period. During this period the risky assets are assumed to follow a geometric Brownian motion, and we use the analytical solution of Richard (1975) which implies an age-independent asset allocation. The utility from this very long period dominates the utility from the few initial 1-year periods. This could possibly bias the asset allocation in $t = 0$ towards the one which holds in the steady-state. A first piece of evidence against this objection is the similarity of our results to those of Campbell et al. (2003): despite using the analytical approximation by Richard (1975) (which assumes a geometric Brownian motion) for the final stage, this does not bias our results away from those in Campbell et al. (2003) (based on an infinite-horizon VAR model) for long planning horizons.

Second, we consider the case that the stochastic variables evolve according to the unconditional moments implied by the VAR model from the very beginning. The

Table 2 Optimal consumption and asset allocation based on the unconditional moments implied by the VAR model for a male investor of age 20 and $d = 0.92$

	Consumption	Cash	Stocks	Bonds
$\gamma = 1$	7.9(0.0)	−206.3(7.7)	187.9(2.1)	118.4(6.1)
$\gamma = 2$	6.7(0.0)	−71.2(2.8)	104.2(0.7)	67.1(2.3)
$\gamma = 5$	3.7(0.0)	31.4(0.9)	41.7(0.2)	26.9(0.7)
$\gamma = 10$	2.5(0.0)	66.1(0.4)	20.60.1)	13.3(0.3)

Results are presented in terms of means and standard errors (in parentheses) from 100 solutions of the problem

results from the unconditional case in Table 2 can be compared to the results from the VAR model for stage $t = 0$ from Table 1. There is a substantial difference in the asset allocation which reflects the impact of time-varying investment opportunities. We note that investigating the unconditional case may also be justified from a different perspective. The issue of return predictability and its impact on asset allocation decisions has found considerable attention in the literature (see, e.g., Barberis (2000) and the references cited there). Even though the empirical evidence cannot simply be ignored, it has to be admitted that predicting asset returns out-of-sample is by no means an easy task. Thus, it is not implausible that an investor is sceptical about the precision of short-term forecasts obtained from a VAR model. There may be a certain risk of having found spurious short-term dependence. In that case she may consider investing according to the unconditional results presented in Table 2.

Despite the differences implied by using conditional and unconditional moments the asset weights based on the VAR model may still be biased towards the unconditional results. To justify the steady-state assumption it is essential that the conditional moments have approached the unconditional moments before the steady-state period starts. Whether this is the case can be judged upon the impulse-response function of the VAR model. Rather than making this judgement on the basis of statistical significance (as done above) we prefer to inspect the economic consequences in terms of the SLP solution. We consider an investor at age 98 in a problem with three stages, and solve the problem with and without the steady-state assumption. These two settings only differ with respect to the properties of returns in the last period. Whereas the former case is based on the conditional moments in the third period, the latter uses unconditional moments. Table 3 shows that the results from the two problems are rather similar, in particular regarding the weights of stocks and bonds. Therefore we can assume that the difference in the moments is economically insignificant, and that the returns from the VAR model have *practically* converged to their unconditional moments. Overall we find no evidence to conclude that the results are biased by the steady-state assumption.

In Sect. 3.5 we have argued that choosing rather long time intervals between stages may distort results compared to the case of short-period rebalancing. In Table 4 we show results for different rebalancing intervals which can be compared to the results for 1-year intervals in Table 1. We analyze the effects for one example only (age 40 and $\gamma = 5$). We find only a slight drop in annual consumption. However, there are

Table 3 Optimal consumption and asset allocation in $t = 0$ for a man at age 98 for $\gamma = 5$ and $d = 0.92$

Steady-state	Consumption	Cash	Stocks	Bonds
yes	30.8(0.0)	5.6(1.5)	39.8(0.4)	54.6(1.3)
no	30.3(0.0)	3.1(1.9)	42.0(0.4)	54.9(1.7)

The table compares results with ('yes') and without ('no') the steady-state assumption. Results are presented in terms of the means and standard errors (in parentheses) from 100 solutions of the problem

Table 4 Optimal consumption and asset allocation in $t = 0$ for various choices of rebalancing intervals. Consumption is expressed in annual terms

Rebalancing intervals	Consumption	Cash	Stocks	Bonds
Annual	4.4(0.0)	2.4(2.1)	42.8(0.5)	54.9(1.8)
Annual and bi-annual (twice)	4.5(0.0)	−5.4(1.7)	52.3(0.5)	53.1(1.5)
Bi-annual	4.5(0.0)	33.6(1.7)	50.7(0.5)	15.7(1.5)
Three years	4.5(0.0)	33.5(1.7)	54.6(0.8)	11.9(1.3)

Results are presented in terms of means and standard errors (in parentheses) from 100 solutions of the problem. We consider a man at age 40 with uncertain lifetime with $\gamma = 5$ and $d = 0.92$

significant changes in the asset allocation which reflect a complicated interplay of various effects (e.g., being locked-in in the asset allocation for varying periods of time and the annualization of consumption). These results indicate that the economic implications associated with time-varying investment opportunities will not be correctly reflected if the rebalancing intervals used to construct the scenario tree are longer than the interval represented by the underlying VAR process. Note, however, that results from increasing the rebalancing intervals even further need not converge to those from using unconditional moments in Table 2 which are based on annual rebalancing.

Closed-form solutions are usually derived by allowing for short sales and excluding transaction costs. Very little is known about the effects of those aspects in the context of time-varying investment opportunities (e.g., Barberis (2000) precludes short sales). For the average investor extreme short positions as obtained in the Campbell et al. (2003) setting have limited practical relevance. Since debt-financed stock investments are usually strongly restricted or impossible for private investors we also consider the case of excluding short sales altogether. In addition, we include 0.5% transaction costs for buying and selling stocks or bonds. Table 5 shows that optimal consumption levels are not affected by either of these aspects. Long positions in cash are obtained if short sales are excluded, but the weight of cash is strongly increased at the expense of bonds by adding transaction costs. In all cases, however, the asset allocation remains rather unaffected by changing the age of the investor. Bodie et al. (1992) and Chen et al. (2006) find a significant impact of human capital on asset allocation decisions over the life cycle. We therefore also investigate the importance of labor income on the age dependence of asset allocation decisions. As opposed to their models we have to treat labor income as deterministic (unrelated to assets and state variables) to make

Table 5 Optimal consumption and asset allocation implied by the VAR model under various assumptions about short sales and transactions costs (0.5% for buying or selling stocks and bonds)

Age	Consumption	Cash	Stocks	Bonds
Short sales excluded; no transaction costs				
20	3.9(0.0)	55.7(0.2)	44.3(0.1)	0.0(0.0)
40	4.2(0.0)	55.7(0.2)	44.3(0.1)	0.0(0.0)
60	5.7(0.0)	55.8(0.2)	44.2(0.2)	0.0(0.0)
Short sales excluded; transaction costs				
20	3.8(0.0)	59.6(0.1)	40.4(0.1)	0.0(0.0)
40	4.3(0.0)	59.5(0.1)	40.5(0.1)	0.0(0.0)
60	5.6(0.0)	59.7(0.1)	40.3(0.1)	0.0(0.0)
Short sales allowed; transaction costs				
20	4.0(0.0)	52.7(0.9)	39.2(0.3)	8.1(0.8)
40	4.4(0.0)	52.8(0.9)	39.2(0.3)	8.0(0.8)
60	5.9(0.0)	52.9(0.9)	38.9(0.3)	8.1(0.8)

Results are presented for a male investor with risk aversion $\gamma = 5$ and $d = 0.92$ in terms of means and standard errors (in parentheses) from 100 solutions of the problem

use of Richard (1975) closed-form solution. However, the uncertainty associated with survival probabilities is accounted for as described in Eq. (1).

Table 6 shows optimal consumption and asset weights for various assumptions about labor income. Compared to Table 1 there is a distinct age effect: the short positions in cash and the long positions in stocks decrease with age. Overall the short positions in cash (and long positions in stocks) are far more extreme than in Table 1 or the bottom panel of Table 5 where labor income is ignored. Excluding short sales leads to 100% investments in stocks (and zero in the other assets). Higher labor income leads to more consumption and makes the distribution of portfolio weights more uneven. These results can be explained by the hedging effect associated with the certain stream of income. The decreasing share of stocks with increasing age is consistent with the results in Bodie et al. (1992) (we have replicated their results to the extent possible given the differences in the two settings). They consider cases where initial labor income is about 10–30% of initial wealth, and their results are also characterized by extreme short positions in the risk-free asset. Despite the fact that age plays a role as soon as labor income is included, we also observe a rather stable ratio of stocks to bonds. This ratio is rather independent of age and slightly increases with labor income. However, we hesitate to derive far reaching conclusions from this particular case, since it depends on many aspects whose role has yet to be investigated more thoroughly.

In summary, in the context of time-varying investment opportunities, constant relative risk aversion, and uncertain lifetime, the fractions invested in risky assets are independent of age. This is the case even if short sales are excluded, and transaction costs are included. Age-dependence is only found if labor income is taken into account. We defer a closer examination of this finding to future research where we also intend to include stochastic labor income.

Table 6 Optimal consumption and asset allocation implied by the VAR model for various assumptions about current, annual labor income L_0

Income	Age	Consumption	Cash	Stocks	Bonds
	20	11.0(0.0)	−35.5(2.3)	114.9(0.5)	20.7(2.1)
$L_0 = 5$	40	11.0(0.0)	−18.4(2.0)	99.7(0.4)	18.7(1.9)
	60	11.7(0.0)	6.9(1.6)	78.2(0.4)	15.0(1.5)
	20	18.3(0.0)	−110.0(3.6)	178.5(0.8)	31.5(3.2)
$L_0 = 10$	40	17.1(0.0)	−82.5(3.2)	155.0(0.6)	27.5(2.8)
	60	17.4(0.0)	−38.9(2.4)	117.6(0.5)	21.3(2.1)
	20	31.8(0.0)	−187.6(6.1)	245.8(1.8)	41.8(5.2)
$L_0 = 20$	40	29.3(0.0)	−162.3(5.1)	227.0(1.4)	35.4(4.5)
	60	28.3(0.0)	−112.8(3.6)	182.8(0.8)	30.0(3.2)

Initial wealth is $w_0 = 100$. Transactions costs are 0.5% for buying or selling stocks and bonds. The investor is assumed to retire at age 65. After retirement he receives 65% of his pre-retirement income. Results are presented for a male investor with risk aversion $\gamma = 5$ and $d = 0.92$ in terms of means and standard errors (in parentheses) from 100 solutions of the problem

5 Conclusion

We have presented a SLP approach to obtain optimal consumption and life-cycle asset allocation of an investor with uncertain lifetime in the context of a VAR model of asset returns and state variables. The SLP approach is based on a discrete scenario tree with only a few stages. To cover the very long time span required in a life-cycle context we work with a few 1-year periods followed by a long, steady-state period. Thereby the short-term dynamics of the VAR model and frequent rebalancing in the first few years can be accounted for. The results of this approach compare well to existing results from the literature. The SLP approach is a flexible tool that may also be used to assess the importance of aspects such as time-varying investment opportunities, short-sale constraints, transaction costs, and labor income. An interesting finding is that the asset allocation seems to be independent of age even if asset returns and state variables follow a vector autoregression model. To confirm this numerically derived result analytically calls for further research, as well as the age-dependence we find if labor income is taken into account.

Acknowledgments The authors gratefully acknowledge financial support from the Austrian National Bank (Jubiläumsfondsprojekt Nr. 11962). The authors are grateful for helpful comments, discussions and suggestions by participants of the May 2007 meeting of the EURO working group on financial modeling in Rotterdam, the 2007 Computational Methods in Finance conference (Waterloo), the 2007 EFMA annual meeting (Vienna), the 2007 International Conference on Stochastic Programming (Vienna), and the 2007 Workshop on Credit and Operational Risks in the Financial Sector (Bolzano).

Appendix A: Closed-form solutions in case of uncertain lifetime

Richard (1975) obtains a closed-form solution for the consumption and investment decisions of an uncertain lived investor in a continuous time model. He assumes geometric Brownian motions for the risky assets, one riskless asset, and power utility for consumption and bequest of an investor whose current age is y_t. Provided that relative risk aversion γ is the same in both utility functions, the closed-form solution for the value function J is given by

$$J(W_t, y_t) = \frac{a_{y_t}}{1-\gamma}(W_t + H_t)^{1-\gamma}. \tag{4}$$

The value function is based on the following definitions:

$$a_{y_t} = \left(\int_{y_t}^{\bar{\tau}} k(\theta) \frac{S(\theta)}{S(y_t)} \exp\left\{ \frac{1-\gamma}{\gamma}(v+r)(\theta - y_t) \right\} d\theta \right)^{\gamma}$$

with

$$k(\theta) = [h(\theta)m(\theta)]^{1/\gamma} + m(\theta)^{1/\gamma} \quad m(\theta) = \exp\{-\delta(\theta - y_t)\} \quad v = \frac{(\nu_p - r)^2}{2\gamma\sigma_p^2}.$$

ν_p and σ_p are drift and standard deviation of the tangency portfolio (which only consists of risky assets). $S(y_t)$ is the survival function defined as

$$S(y_t) = P(\theta \geq y_t) = \int_{y_t}^{\bar{\tau}} \vartheta(\theta) d\theta \qquad \int_{0}^{\bar{\tau}} \vartheta(\theta) d\theta = 1.$$

$h(\theta)$ is the conditional probability density for death conditional upon the investor being alive at age θ, so that $h(\theta) = \vartheta(\theta)/S(\theta)$.

H_t is the present value of labor income received until the final age of the underlying mortality table $\bar{\tau} = 101$. H_t assumes an actuarially fair life insurance of labor income and is given by

$$H_t = \int_{y_t}^{\bar{\tau}} \mathcal{L}(s) \frac{S(s)}{S(y_t)} \exp\{-(s - y_t)r\} ds,$$

where $\mathcal{L}(s)$ is continuous labor income and $S(s)/S(y_t)$ is the conditional probability density to be alive at time s conditional upon the investor being alive at age y_t. This definition of H_t agrees with the continuous-time formulation of Richard (1975). The results presented in Sect. 4 are based on the discrete-time version of labor income defined in Sect. 3.1, Eq. (1).

Since we work with discrete mortality tables where age is integer-valued we can simplify the definition of a_{y_t} as follows:

$$a_{y_t} = \left(\sum_{\theta=y_t}^{\bar{\tau}-1} k(\theta) \frac{S(\theta)}{S(y_t)} \int_{\theta}^{\theta+1} \exp\left\{c(u - y_t)\right\} du \right)^{\gamma} \qquad c = \frac{1-\gamma}{\gamma}(v + r)$$

$$a_{y_t} = \left(\sum_{\theta=y_t}^{\bar{\tau}-1} k(\theta) \frac{S(\theta)}{S(y_t)} \left[\exp\{c(\theta - y_t)\}(\exp\{c\} - 1)/c\right] \right)^{\gamma}.$$

Appendix B: VAR model for asset returns

To describe time-varying investment opportunities as in Barberis (2000) or Campbell et al. (2003) we use a VAR(1) model

$$\mathbf{Y}_s = \mathbf{c} + \mathbf{A}\mathbf{Y}_{s-1} + \mathbf{e}_s \quad \mathbf{e}_s \sim N(0, \mathbf{C}_e),$$

where $\mathbf{Y}_s$ is a $K \times 1$ vector of asset returns and state variables, $\mathbf{c}$ is a vector of constants, $\mathbf{A}$ is a $K \times K$ matrix of autoregressive coefficients and $\mathbf{e}_s$ is a vector of uncorrelated normal disturbances. Assuming a normal distribution seems justified given the long time intervals we consider.

$\mathbf{C}_e$ is related to the correlation matrix of disturbances $\mathbf{R}_e$ and the standard errors $\mathbf{s}_e$ by $\mathbf{C}_e = \mathbf{R}_e \cdot (\mathbf{s}_e \mathbf{s}_e')$. Mean and covariance of $\mathbf{Y}_s$ are given by $\boldsymbol{\mu} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{c}$ and

$$\mathbf{C} = \sum_{i=0}^{\infty} \mathbf{A}^i \mathbf{C}_e \mathbf{A}'^i$$

(see Lütkepohl 1993, p. 11). Asset returns are observed at a relatively high frequency and parameter estimates refer to that data frequency. However, in our model asset allocation decisions are made at only a few points in time which may be one or several years apart. Therefore we have to consider the properties of multi-period returns, i.e. the sum of $\mathbf{Y}_s$ over h periods $\mathbf{Y}_s^h = \mathbf{Y}_{s+1} + \mathbf{Y}_{s+2} + \cdots + \mathbf{Y}_{s+h}$. $\mathbf{Y}_s^h$ can be shown (see Barberis 2000, p. 241) to be normally distributed with mean (conditional on $\mathbf{Y}_s$)

$$\boldsymbol{\mu}^h = \sum_{i=0}^{h-1} (h - i)\mathbf{A}^i \mathbf{c} + \left(\sum_{i=1}^{h} \mathbf{A}^i \right) \mathbf{Y}_s$$

and covariance

$$\begin{aligned} \mathbf{C}^h = \ & \mathbf{C}_e + (\mathbf{I} + \mathbf{A})\mathbf{C}_e(\mathbf{I} + \mathbf{A})' \\ & + (\mathbf{I} + \mathbf{A} + \mathbf{A}^2)\mathbf{C}_e(\mathbf{I} + \mathbf{A} + \mathbf{A}^2)' \\ & + \cdots + (\mathbf{I} + \mathbf{A} + \cdots + \mathbf{A}^{h-1})\mathbf{C}_e(\mathbf{I} + \mathbf{A} + \cdots + \mathbf{A}^{h-1})'. \end{aligned}$$

For each time interval of length τ_t we simulate a sample of log returns $\tilde{\mathbf{Y}}^{\tau_t}$ such that the τ_t-period moments $\boldsymbol{\mu}^{\tau_t}$ and $\mathbf{C}^{\tau_t}$ are matched, and their skewness is zero and kurtosis is three. The gross returns $\tilde{R}_t^i$ of asset i defined in Sect. 3.1 are related to the ith element of the τ_t-period simulated returns by $\tilde{R}_t^i = \exp\{\tilde{Y}_i^{\tau_t}\}$.

The simulated returns for period 1 are based on the unconditional means $\boldsymbol{\mu}$. These simulated returns provide the conditioning information for subsequent periods. The number of samples drawn depends on the node structure of the scenario tree, and determines the actual dimension of the (stacked) vector of simulated returns.

References

Balvers RJ, Mitchell DW (1997) Autocorrelated returns and optimal intertemporal portfolio choice. Manage Sci 43(11):1537–1551

Barberis NC (2000) Investing for the long run when returns are predictable. J Finance 55:225–264

Blomvall J, Lindberg PO (2002) A Ricatti-based primal interior point solver for multistage stochastic programming. Eur J Oper Res 143:452–461

Bodie Z, Merton RC, Samuelson WF (1992) Labor supply flexibility and portfolio choice in a life cycle model. J Econ Dyn Control 16:427–449

Brandt MW, Goyal A, Santa-Clara P, Stroud JR (2005) A simulation approach to dynamic portfolio choice with an application to learning about return predictability. Rev Financ Stud 18(3):831–873

Brennan MJ, Schwartz ES, Lagnado R (1997) Strategic asset allocation. J Econ Dyn Control 21:1377–1403

Campbell JY, Viceira LM (1999) Consumption and portfolio decisions when expected returns are time varying. Q J Econ 114:433–495

Campbell JY, Viceira LM (2001) Who should buy long-term bonds? Am Econ Rev 91(1):99–127

Campbell JY, Viceira LM (2002) Strategic asset allocation. Oxford University Press, USA

Campbell JY, Cocco J, Gomes F, Maenhout PJ, Viceira LM (2001) Stock market mean reversion and the optimal asset allocation of a long-lived investor. Eur Finance Rev 5:269–292

Campbell JY, Chan YL, Viceira LM (2003) A multivariate model of strategic asset allocation. J Financ Econ 67:41–80

Campbell JY, Chacko G, Rodriguez J, Viceira LM (2004) Strategic asset allocation in a continuous-time VAR model. J Econ Dyn Control 28:2195–2214

Cariño DR, Ziemba WT (1994) The Russell–Yasuda–Kasai model: an asset/liability model for a Japanese insurance company using multistage stochastic programming. Interfaces 24(1):29–49

Cariño DR, Ziemba WT (1998) Formulation of the Russell–Yasuda–Kasai financial planning model. Oper Res 46:433–449

Cariño DR, Myers DH, Ziemba WT (1998) Concepts, technical issues, and uses of the Russell-Yasuda-Kasai financial planning model. Oper Res 46:450–462

Chacko G, Viceira LM (2005) Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets. Rev Financ Stud 18((2):1369–1402

Chen P, Ibbotson RG, Milevsky MA, Zhu KX (2006) Human capital, asset allocation, and life insurance. Financ Anal J 62:97–109

Cocco JF, Gomes FJ, Maenhout PJ (2005) Consumption and portfolio choice over the life cycle. Rev Financ Stud 18(2):491–533

Dempster MAH, Germano M, Medova EA, Villaverde M (2003) Global asset liability management. Br Actual J 9:137–216

Detemple JB, Garcia R, Rindisbacher M (2003) A Monte Carlo method for optimal portfolios. J Finance 58:401–446

Duffie D (2001) Dynamic asset pricing theory. Princeton University Press, USA

Dupačová J, Gröwe-Kuska N, Römisch W (2003) Scenario reduction in stochastic programming: an approach using probability metrics. Math Program Ser A 95:493–511

Geyer A, Ziemba WT (2007) The Innovest Austrian pension fund financial planning model InnoALM. Oper Res 56(4):797–810

Geyer A, Hanke M, Weissensteiner A (2007) Life-cycle asset allocation and optimal consumption using stochastic linear programming. Working paper, Vienna Graduate School of Finance and University of Innsbruck. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=952956
Gomes F, Michaelides A (2005) Optimal life-cycle asset allocation: understanding the empirical evidence. J Finance 60:869–904
Gondzio J, Grothey A (2007) Solving nonlinear portfolio optimization problems with the primal-dual interior point method. Eur J Oper Res 181:1019–1029
Gondzio J, Kouwenberg R (2001) High performance computing for asset-liability management. Oper Res 49:879–891
Hanke M, Huber S (2008) Curvature, not second derivative. Mathe Spectr (in press)
Heitsch H, Römisch W (2003) Scenario reduction algorithms in stochastic programming. Comput Optim Appl 24:187–206
Hochreiter R, Pflug GC (2007) Financial scenario generation for stochastic multi-stage decision processes as facility location problems. Ann Oper Res 152:257–272
Høyland K, Wallace SW (2001) Generating scenario trees for multistage decision problems. Manage Sci 47(2):295–307
Ingersoll JE (1987) Theory of financial decision making. Rowman & Littlefield, USA
Jurek JW, Viceira LM (2005) Optimal value and growth tilts in long-horizon portfolios. Working paper
Kim TS, Omberg E (1996) Dynamic nonmyopic portfolio behavior. Rev Financ Stud 9:141–161
Klaassen P (2002) Comment on generating scenario trees for multistage decision problems. Manage Sci 48:1512–1516
Liu J (2007) Portfolio selection in stochastic environments. Rev Financ Stud 20:1–39
Longstaff F, Schwartz E (2001) Valuing American options by simulation: a simple least-square approach. Rev Financ Stud 14(1):113–147
Lütkepohl H (1993) Introduction to multiple time series analysis. Springer, Heidelberg
Merton RC (1969) Lifetime portfolio selection under uncertainty: the continuous-time case. Rev Econ Stat 51:247–257
Merton RC (1971) Optimum consumption and portfolio rules in a continuous-time model. J Econ Theory 3:373–413
Pflug GC (2001) Optimal scenario tree generation for multiperiod financial planning. Math Program 89(2):25–271
Pflug GC, Swietanowski A (2000) Asset-liability optimization for pension fund management. In: Operations Research Proceedings, Springer, Heidelberg
Rasmussen KM, Clausen J (2007) Mortgage loan portfolio optimization using multi-stage stochastic programming. J Econ Dyn Control 31(3):742–766
Richard SF (1975) Optimal consumption, portfolio and life insurance rules for an uncertain lived individual in a continous time model. J Financ Econ 2:187–203
Samuelson PA (1969) Lifetime portfolio selection by dynamic stochastic programming. Rev Econ Stat 51:239–246
Schroder M, Skiadas C (1999) Optimal consumption and portfolio selection with stochastic differential utility. J Econ Theory 89(1):68–126
Wachter J (2002) Portfolio and consumption decisions under mean-reverting returns: an exact solution. J Financ Quant Anal 37(1):63–91
Wallace SW, Ziemba WT (eds) (2005) Applications of stochastic programming. MPS-SIAM book series on optimization
Zenios SA (1999) High-performance computing for financial planning: the last ten years and the next. Parallel Comput 25:2149–2175
Ziemba WT, Mulvey J (1998) Worldwide asset and liability modeling. Cambridge University Press, London