Life-cycle Asset Allocation and Optimal Consumption Using Stochastic Linear Programming

Preliminary draft; first version: December 21, 2006; this version: March 8, 2007

Abstract

We consider optimal consumption and (strategic) asset allocation of an investor with uncertain lifetime in the context of time-varying investment opportunities. To solve this problem we use a multi-stage stochastic linear programming (SLP) model. We consider aspects of the application of the SLP approach which arise in the context of life-cycle asset allocation, but are also relevant for other problems of similar structure. The objective is to maximize the expected utility of consumption over the lifetime and of bequest at the time of death of the investor. Since we maximize utility (rather than other objectives which can be implemented more easily) we provide a new approach to optimize the breakpoints required for the linearization of the utility function. Asset returns and state variables follow a vector autoregression 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 (which reflect return predictability), and incorporate a closed-form solution for the long, subsequent steady-state period. In our numerical examples we first show that available closed-form solutions can be accurately replicated with the SLP-based approach. Second, we add elements to the problem specification which are usually beyond the scope of closed-form solutions. We find that the asset allocation remains independent of age even if asset returns (and state variables) follow a vector autoregression, and short-sale constraints or transaction costs are included.

Key words: life-cycle asset allocation, stochastic linear programming, scenario trees

JEL classification: C61, G11

Acknowledgements

The authors gratefully acknowledge helpful comments, discussions and suggestions by Thomas Dangl, Pedro Santa-Clara, and Reimar Volkert. We also acknowledge financial support from the Austrian National Bank (Jubiläumsfondsprojekt Nr. 11962).

1 Introduction

One of the classical problems in finance is the optimal consumption and asset allocation over the life-cycle of a finitely-lived investor. This problem lies at the heart of the subfield of personal finance, and financial advisors as well as portfolio and pension fund managers throughout the world are faced with it every day. The interest and activities in this research area have grown in recent years, partly spurred by growing world-wide concern about the stability of public pension systems and the resulting trend towards private pensions. Classical treatments of this problem are Samuelson (1969) and Merton (1969, 1971) who formulate models in simplified settings, aiming at closed-form solutions. One of the main results from early multi-period portfolio models is that the fractions of risky assets are constant over the lifetime of an investor. This contradicts the advice obtained from many professionals in practice who recommend that the share of risky assets held by investors decline steadily as they approach retirement (often called the *age effect*). Since then, many researchers have tried to resolve this puzzle by incorporating more realistic assumptions. In many cases, this renders the models as analytically intractable, necessitating the use of numerical solution techniques.

In this paper, 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 our approach in mind. This distinguishes our work from literature which focuses more on gaining general insights into the dependencies between investment/consumption decisions and state variables but excludes a number of (possibly important) real-world aspects on purpose to preserve analytical tractability. For example, many models are confined to a small number of risky assets (often only one), do not allow for constraints on the asset allocation or ignore transactions costs. In contrast, many of these features which are considered important for investment decisions in practice can be easily incorporated when using SLP. Combined with the availability of efficient solvers for SLPs, this explains why the SLP approach has been successfully applied to a wide range of problems (see e.g. Ziemba and Mulvey 1998, Wallace and Ziemba 2005), in particular in the context of asset-liability management (see e.g. Cariño et al. 1998, Dempster et al. 2003, Zenios and Ziemba 2006, Geyer and Ziemba 2007).

Key elements of our model are: The expected utility of consumption over the investor's lifetime and expected utility of bequest is maximized, and the mortality risk of the investor is taken into account. Consumption and investment decisions are optimized jointly, time-varying investment opportunities are included, and personal characteristics of the investor can be taken into account (e.g. risk attitude, different utility functions for consumption and bequest, life expectancy, retirement, future cash flows for major purchases or associated with other life events).

Some aspects of the life-cycle problem we consider pose difficulties to the SLP approach (e.g. the long time span is in conflict with a manageable problem size). One of our objectives is to solve certain (technical) problems which are typically encountered in the formulation and implementation of SLPs for life-cycle asset allocation and consumption problems (which may be equally relevant for other problems of similar structure). A further goal is to obtain solutions under practically relevant settings, some of which are beyond the scope of other numerical techniques. Due to the complex nature of the problem it is difficult to isolate key factors and conditions which are potentially important (e.g. age and the associated mortality risk, or time-varying investment opportunities). The SLP approach provides a flexible tool to facilitate this assessment.

The present paper is – as far as we know – the first application of stochastic linear programming in a life-cycle asset allocation context. We maximize expected utility rather than other objectives which can be implemented more easily (e.g. piecewise linear or quadratic penalty functions, or minimizing CVaR). As far as we know no other papers in this area use utility functions in this way and for that purpose. The accuracy of the solution depends strongly on the way the utility function is linearized. Linearizing such functions can be rather difficult, in particular for high degrees of risk aversion. Therefore it is essential to optimize the breakpoints required for their linearization. We propose and implement a new approach for that purpose and show that closed-form solutions can be replicated with very high precision.

The long time intervals involved in life-cycle asset allocation problems pose a notoriously difficult problem. To keep the dimension of the problem manageable SLP-based approaches typically rebalance the portfolio very infrequently and define consumption plans for unrealistically long time intervals. We provide two contributions to this aspect. First, we incorporate a closed-form solution for optimal subperiod consumption during (long) time intervals in the objective. Second, the model formulation combines decisions for a few short time periods in the near future with a closed-form solution for the long, subsequent steady-state period. Both closed-form solutions take survival probabilities into account.

The SLP approach extends the range of available numerical methods for solving life-cycle problems.[1] One type of numerical methods works via grid methods discretizing the state space (see e.g. Brennan et al. 1997, Barberis 2000, Cocco et al. 2005, Gomes and Michaelides 2005). Another type derives a system of equations which is solved numerically (see Schroder and Skiadas 1999). The approach used by Brandt et al. (2005) combines Monte Carlo simulation and regression techniques, and is inspired by the option pricing algorithm of Longstaff and Schwartz (2001).

The linear nature of the SLP approach bears some similarities to Campbell et al. (2003) who linearize the portfolio return, the budget constraint, and the Euler equation, and arrive at a system of linear-quadratic equations for portfolio weights and consumption as functions of state variables. Asset returns and state variables are modelled as a first-order vector autoregression VAR(1). They consider Epstein-Zin utility and an infinite planning horizon. Additional assumptions include the absence of borrowing and short-sale constraints. The system of equations can

[1] A more comprehensive comparison with other numerical methods is very demanding and not the objective of this paper.

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. Since we also use a linear approach we take the stochastic setting of Campbell et al. (2003) as a starting point. Their data covers the period 1893 to 1997, and the three main asset classes T-bills, stocks and bonds. This provides an interesting and representative basis for our analysis. 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. However, we hesitate to derive general conclusions about the role of those aspects for life-cycle asset allocation from our results. We rather defer a more thorough investigation of the financial implications to a later paper.

The paper is organized as follows: Section 2 describes the stochastic programming model, in particular the formulation of the objective, the optimization approach for its linearization, and the generation of scenarios. In Section 3 results from the SLP are compared to cases where closed-form solutions are available, and results from life-cycle consumption and asset allocation decisions in a more realistic setting are presented. Section 4 concludes.

2 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 100% that a person dies between age 100 and 101). The choice of the length of the time intervals τ_t is described in Section 2.5.

2.1 Variables

The following (decision) variables are used in the model formulation; a tilde ˜ denotes stochastic (i.e. scenario dependent) decision variables:

$C_0\geq 0$, $\tilde{C}_t\geq 0$ (t=1,...,$T-1$) ...consumption in t; e.g. $\tilde{C}_2$ is the amount set aside in t=2 for consumption between t=2 and t=3.

$\tilde{R}^i_t$ (t=1,...,T; i=1,...,N) ...gross return of asset i for the period that ends in t.

$P^i_0\geq 0$, $\tilde{P}^i_t\geq 0$ (t=1,...,$T-1$; i=1,...,N) ...amount of asset i purchased in t.

$S^i_0\geq 0$, $\tilde{S}^i_t\geq 0$ (t=1,...,$T-1$; i=1,...,N) ...amount of asset i sold in t.

q^i_p and q^i_s ...transaction costs for purchases and sales of asset i.

W_0^i, $\tilde{W}_t^i$ (t=1,...,$T-1$; i=1,...,N) ... 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 ... initial value of asset i (before transactions).

$\tilde{B}_t \geq 0$ (t=1,...,T) ... bequest in t given by $\tilde{B}_t = \sum \tilde{R}_t^i \tilde{W}_{t-1}^i$.

τ_t (t=0,...,$T-1$) ... the number of years between stage t and stage t+1.

φ_y ... the (conditional) probability to survive the year following year y.

$\Phi(y_t, \tau_t)$ ... the probability to survive the period of length τ_t starting at stage t at an age of y_t years; $\Phi(y_t, \tau_t) = \prod_{k=y_t}^{y_t+\tau_t-1} \varphi_k$.

Λ_t (t=1,...,$T-1$) ... the probability to be alive at stage t (at an age of y_t); $\Lambda_t = \prod_{k=0}^{t-1} \Phi(y_k, \tau_k)$.

Θ_t (t=1,...,T) ... 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,...,$T-1$)... 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,...,$T-1$)... 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 ... the risk-free interest rate.

δ ... 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 Section 2.5. 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,...,y_r) and by $\mathcal{L}_y = f_r \mathcal{L}_0 \exp\{y\ell\}$ (y=y_r+1,...,τ) 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 (see Richard 1975).

2.2 Constraints

The budget equations are given by

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

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

The value of investments accumulates according to the following equations:

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

To model restrictions on the portfolio composition we use the constraints

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

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:

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

2.3 Objective

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

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

U_b is the utility function of bequest and U_c is the utility function of consumption. Θ_t is the probability to die between stage $t{-}1$ and t, and Λ_t is the probability to be alive in t. In principle we can use any time-additive utility function, but in the numerical examples of Section 3 we only consider power utility functions (which imply constant relative risk aversion γ). The risk aversion associated with U_c and U_b need not be the same, unless we use the analytical solution based on Richard (1975) for the final period (see Section 2.5).

The utility function is linearized by defining m linear segments between the breakpoints b_t^j (j=0,…,m). The choice of b_t^0 and b_t^m is described below. Note that different breakpoints are used in each stage. The same linearization procedure is used for $\tilde{B}_t$, C_0 and $\tilde{C}_t$ which is now described using the generic variable X_t. X_t is defined in terms of non-negative decision variables V_t^j associated with each segment:

$$X_t = \sum_{j=0}^{m+1} V_t^j$$

$$0 \leq V_t^0 \leq b_t^0 \qquad V_t^{m+1} \geq 0 \qquad t = 0, \ldots, T$$

$$0 \leq V_t^j \leq b_t^j - b_t^{j-1} \qquad j = 1, \ldots, m \qquad t = 0, \ldots, T.$$

The slopes of the linear segments are given by

$$\Delta_t^j = \frac{U(b_t^j) - U(b_t^{j-1})}{b_t^j - b_t^{j-1}}$$

and the utility of X_t is approximated by

$$U(X_t) \approx \Delta_t^1 V_t^0 + \sum_{j=1}^{m} \Delta_t^j V_t^j + \Delta_t^m V_t^{m+1}.$$

We linearize the utility of bequest and consumption by defining the constraints

$$\tilde{B}_t = \sum_{j=0}^{m+1} \tilde{V}_{bt}^j \qquad C_0 = \sum_{j=0}^{m+1} V_{c0}^j \qquad \tilde{C}_t = \sum_{j=0}^{m+1} \tilde{V}_{ct}^j.$$

C_0 and $\tilde{C}_t$ refer to consumption during the subsequent time *interval* of length τ_t which may be longer than one year. We assume that the amount $\tilde{C}_t$ is not consumed at once but in annual parts. Accordingly, the utility of consuming those parts differs from the utility of consuming $\tilde{C}_t$ at once. As a simple approximation the utility of consumption over a period of length τ_t can be defined as $\tau_t U_c(\tilde{C}_t/\tau_t)$, which can be improved by replacing $1/\tau_t$ with an annuity factor. However, we prefer to use a more refined approach which is based on considering the following optimization problem. We derive our solution by considering stage 0 and three subperiods (years) (i.e. τ_0=3) and generalize later. Assume, for the time being, that C_0 is known. We define the optimal *annual* consumption levels in the three subperiods c_0, c_1 and c_2 such that the utility of consumption in the period is maximized

$$U(c_0) + \Phi_1 \exp\{-\delta\} U(c_1) + \Phi_2 \exp\{-2\delta\} U(c_2) \longrightarrow \max$$

subject to the constraint c_0+$\exp\{-r\}c_1$+$\exp\{-2r\}c_2$=C_0. To simplify notation in this derivation Φ_j=$\Phi(y_0, j)$ is the conditional probability to survive the next j years (i.e. to be alive at the beginning of subperiod j+1), given a current age of y_0. The constraint assumes that c_1 and c_2 are invested at the risk-free rate. For power utility $U(c)$=$c^{1-\gamma}/(1-\gamma)$ we have

$$\begin{aligned} L &= \frac{c_0^{1-\gamma}}{1-\gamma} + \Phi_1 \exp\{-\delta\} \frac{c_1^{1-\gamma}}{1-\gamma} + \Phi_2 \exp\{-2\delta\} \frac{c_2^{1-\gamma}}{1-\gamma} + \\ & +\lambda \left(C_0 - c_0 - \exp\{-r\}c_1 - \exp\{-2r\}c_2 \right) \longrightarrow \max \end{aligned}$$

which leads to

$$c_1 = c_0 \exp\left\{ \frac{r-\delta}{\gamma} \right\} \Phi_1^{1/\gamma} \qquad c_2 = c_0 \exp\left\{ 2\frac{r-\delta}{\gamma} \right\} \Phi_2^{1/\gamma}.$$

Generalizing for τ_t subperiods we obtain (if $\gamma \neq 0$ and $c_j > 0 \ \forall j$)

$$C_0 = c_0 \sum_{j=0}^{\tau_t - 1} {\Phi_j}^{1/\gamma} \exp\left\{ \frac{j(r(1-\gamma)-\delta)}{\gamma} \right\} \qquad (\Phi_0 = 1).$$

Based on this result we can reformulate the model such that optimal consumption in the subperiods is taken into account. This is accomplished without increasing the number of decision variables. We replace 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 and $\tilde{c}_t$. We define

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

and

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

The budget constraints are formulated as

$$\alpha_0 c_0 + \sum_{i=1}^{N} P_0^i (1 + q_p^i) = \sum_{i=1}^{N} S_0^i (1 - q_s^i) + L_0 + F_0$$

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

Utility of consumption in t is formulated in terms of $\tilde{c}_t$

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

which is linearized in the same way as described above.

2.4 Choice of breakpoints

The linearization of the objective function requires choosing the number and the position of breakpoints. This choice determines the standard errors of the SLP optimization results as shown in Section 3. Since our model can accommodate different utility functions and degrees of risk aversion for consumption and bequest we define separate breakpoints for the two components of utility. However, even if the same utility and risk aversion were used, separate breakpoints for consumption and bequest would be necessary 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 in 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 (1996, p. 198) as a guideline. To find 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 functions. 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.

To obtain optimal positions of the remaining breakpoints we have considered and tested the following alternatives:

1. minimize the vertical distance between the piecewise linear function and the nonlinear utility function at the midpoints of two adjacent breakpoints,
2. minimize the area between the utility function and the piecewise linear function, or
3. use the curvature of the utility function.

We have found that all three methods yield similar results in terms of the standard errors of the solution. In the applications presented in Section 3 we use the curvature-based approach since it requires no optimization, and it is also much faster than the other methods. 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, …, n) where $\beta_t^0 = b_t^0$ and β_t^m=b_t^m. The curvature for each β_t^j is defined as (for details see Hanke and Huber 2006)

$$\breve{U}(\beta_t^j) = \frac{U''(\beta_t^j)}{(1 + [U'(\beta_t^j)]^2)^{3/2}}.$$

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

$$\overline{\breve{U}}(\beta_t^j) = 0.5\left(\breve{U}(\beta_t^{j-1}) + \breve{U}(\beta_t^j)\right) \qquad j = 1, \ldots, m.$$

The relative average curvature is given by

$$\mathcal{U}_t^j = \frac{\overline{\breve{U}}(\beta_t^j)}{\sum_j \overline{\breve{U}}(\beta_t^j)} \qquad j = 1, \ldots, m$$

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

$$b_t^j = b_t^{j-1} + (\beta_t^{j+1} - \beta_t^j)/n_t^j \qquad j = 1, \ldots, m-1.$$

2.5 Scenario generation and choice of intervals

The uncertainty associated with the consumption-investment problem and time-varying investment opportunities are modelled 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). To avoid arbitrage opportunities in the simulated returns (which would be exploited by the optimization algorithm) we apply the procedure proposed by Klaassen (2002). We follow the approach by Høyland and Wallace (2001), Høyland et al. (2003) and Kaut (2003) to match the first four moments (including the correlations) of the simulated processes. Details on the VAR process and its simulation are described in Appendix B.

The VAR process evolves in discrete time, and the underlying probability distributions are approximated by discrete distributions in terms of a scenario tree. Decisions are made at each node of the tree and depend on the current state which reflects previous decisions and uncertain future paths. Non-anticipatory constraints are imposed to guarantee that a decision made at a specific node is identical for all scenarios leaving that node. The tree is defined by the number of stages and the number of arcs leaving nodes at a particular stage (the branching factor n_t). For instance, the branching structure n_1=6, n_2=4 and n_3=2 (denoted by 6×4×2) corresponds to a total number of 6·4·2=48 scenarios. The tree always starts with a single node which corresponds to the present stage (t=0). A single scenario s_t is a trajectory that corresponds to a unique path leading from the single node at t=0 to a single node at t. Two scenarios s'_t and s''_t are identical until t and differ in subsequent stages t+1,…,T. The scenario assigns specific values to all uncertain parameters (mainly returns) along the trajectory. For a K-dimensional VAR process a branching factor of $2K$ is necessary to match the first four (co)moments within reasonable time. More nodes facilitate the matching of moments but increase the number of scenarios. In the examples presented in Section 3 we set n_t=$2K$+2, $\forall t$.

The number of scenarios in the tree grows exponentially with the number of stages (at which decisions are made) and the number of scenarios following each node. 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 one 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 consideration is particularly relevant in case of time-varying investment opportunities. 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 one-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. During the steady-state period we assume that asset returns are serially uncorrelated. This assumption requires that the effects of shocks in the VAR process disappear rather quickly, which seems justified given the rather weak short-term temporal dependence empirically found for asset returns (we will look at this feature more closely in Section 3). The assumption implies that after a few periods simulated asset returns will be mainly driven by their unconditional moments (consistent with the VAR model parameters). Therefore we can use the analytical solution obtained by 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:

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

$J(\tilde{W}^+_{T-1},y_{T-1})$ is the value function (4) defined in Appendix A. 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 A 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 unconditional means of the assets are defined as $\boldsymbol{\mu}+0.5\mathrm{diag}(\mathbf{C})$ (using the notation from Appendix B). 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.[2] 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 one-year decision intervals in discrete time with continuous consumption and trading. In our opinion, however, the advantages outweigh these drawbacks by far. To provide evidence on this view we analyze the validity and implications of the steady-state assumption in Section 3.

3 Numerical results

We optimize the stochastic linear program on the basis of routines from the open source project COIN-OR (see http://www.coin-or.org). The problem is formulated using the Stochastic Mathematical Programming System (SMPS) input format for multi-stage stochastic programs (see Gassmann and Schweitzer 2001, King et al. 2005) in terms of three input files: the core-, stoch- and time-files. The core-file contains information about the decisions variables, constraints, right-hand-sides and bounds. It contains all fixed coefficients and dummy entries for random elements. The stoch-file reflects the node structure of the scenario tree and contains all random elements, i.e. asset returns and probabilities. The time-file assigns decision variables and constraints to stages. The solution of a problem with four stages, three assets, 2744 scenarios and 40 breakpoints requires less than three minutes using a Pentium 4/640 processor with 3.2 GHz and 1 GB RAM. A corresponding problem with only three stages and 196 scenarios solves in less than 5 seconds.

To test the scenario generation and the SLP formulation we first present results from cases where closed-form solutions are available using artificial data. Thereafter we consider life-cycle consumption and asset allocation decisions in a more realistic setting using data from Campbell

[2]We do not consider scenario reduction techniques (see Heitsch and Römisch 2003, Rasmussen and Clausen 2006) for two reasons. First, we can replicate closed-form solutions very precisely with only a few scenarios (see Section 3). Second, it is not clear how well scenario reduction performs in the context of VAR models.

et al. (2003).

Closed-form solutions are obtained from Ingersoll (1987), Duffie (1996) and Richard (1975) (see Appendix A) and compared to the SLP-based results in terms of consumption and asset allocation decisions. In case of uncertain lifetime we use survival probabilities for Austrian men estimated in 2005. For cases with certain lifetime the survival probabilities φ are set to 1 (except for the final one). We consider two risky assets with a drift rate of 0.06 and a risk-free asset with a return of 0.04. The correlation among risky assets is 0.5 and their volatility is 0.2.[3] Since these assumptions imply that asset returns are in steady-state we use the corresponding objective function (3). The node structures considered are 6×6 and 36×12, and the linearization of the objective is based on 40 and 80 breakpoints, respectively. We present results for stage t=0 in terms of means and standard errors from sampling and solving the model 100 times. The same 100 scenario trees are used in each setting (e.g. for varying risk aversion or time preference). We consider small-scale problems with only a few scenarios (36 and 432, respectively) which can be solved in only a few seconds. By solving such problems repeatedly we can quickly obtain standard errors of the optimal solutions. This is preferable to solving large-scale problems once, which provide better but unknown precision.

Table 1 shows that closed-form results can be replicated with relatively high precision even if only a few scenarios are used. The standard errors for optimal consumption are negligibly small. The slight discrepancies found for optimal consumption may be explained by our mixture of working in discrete time and using the steady-state solution (which implies continuous consumption and rebalancing). The precision of portfolio weights is lower than for consumption. However, the accuracy of the results can be improved by increasing number of scenarios, and/or using more breakpoints in the linearization of the utility function.

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 modelled 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 to 1997 (rather than 1890 to 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, being well aware of the potential effects associated with parameter uncertainty.

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 one year. This justifies our approach of using a few one-year periods

[3]This choice yields an equally weighted portfolio for an investor with log utility. With power utility and γ>1 the weights of the risky assets are reduced by the factor $1/\gamma$.

	consumption	asset A	asset B	risk-free
log utility, d=1, certain lifetime				
closed-form solution	1.61	33.33	33.33	33.33
40 breakpoints, 6×6 nodes	1.66	33.05	32.98	33.97
	(0.00)	(0.13)	(0.11)	(0.15)
80 breakpoints, 6×6 nodes	1.66	33.05	32.95	34.00
	(0.00)	(0.07)	(0.07)	(0.09)
40 breakpoints, 36×12 nodes	1.66	33.00	32.94	34.06
	(0.00)	(0.06)	(0.06)	(0.03)
80 breakpoints, 36×12 nodes	1.66	33.05	32.92	34.03
	(0.00)	(0.06)	(0.06)	(0.01)
log utility, d=0.92, certain lifetime				
closed-form solution	8.05	33.33	33.33	33.33
40 breakpoints, 6×6 nodes	8.02	33.05	33.18	33.77
	(0.00)	(0.11)	(0.11)	(0.14)
80 breakpoints, 6×6 nodes	8.02	32.99	33.01	34.00
	(0.00)	(0.07)	(0.06)	(0.08)
40 breakpoints, 36×12 nodes	8.02	33.09	32.93	33.98
	(0.00)	(0.06)	(0.06)	(0.03)
80 breakpoints, 36×12 nodes	8.02	33.05	32.93	34.01
	(0.00)	(0.06)	(0.05)	(0.02)
log utility, d=0.92, uncertain lifetime				
closed-form solution	8.82	33.33	33.33	33.33
40 breakpoints, 6×6 nodes	8.71	32.95	33.23	33.81
	(0.00)	(0.11)	(0.11)	(0.14)
80 breakpoints, 6×6 nodes	8.72	32.95	33.01	34.03
	(0.00)	(0.06)	(0.05)	(0.08)
40 breakpoints, 36×12 nodes	8.72	33.00	32.98	34.02
	(0.00)	(0.06)	(0.06)	(0.03)
80 breakpoints, 36×12 nodes	8.71	33.03	32.96	34.01
	(0.00)	(0.06)	(0.06)	(0.02)
power utility, γ=4, d=0.92, certain lifetime				
closed-form solution	5.18	8.33	8.33	83.33
40 breakpoints, 6×6 nodes	5.37	8.22	8.29	83.49
	(0.00)	(0.09)	(0.10)	(0.14)
80 breakpoints, 6×6 nodes	5.37	8.30	8.30	83.40
	(0.00)	(0.06)	(0.07)	(0.10)
40 breakpoints, 36×12 nodes	5.36	8.24	8.24	83.52
	(0.00)	(0.03)	(0.03)	(0.04)
80 breakpoints, 36×12 nodes	5.36	8.26	8.22	83.52
	(0.00)	(0.02)	(0.02)	(0.03)

Table 1: Optimal consumption and asset allocation from closed-from solutions and the stochastic linear program in t=0. SLP 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 entry of the table. The investor is assumed to be 40 years old. For cases with uncertain lifetime we use Austrian mortality tables for men.

followed by a long steady-state period. Shocks to the state variables and stocks, however, remain statistically significant for up to ten 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 means $\boldsymbol{\mu}$+0.5diag($\mathbf{C}$) (using the notation from Appendix B). The risk-free rate is set equal to the unconditional mean of real T-bill returns without adding the variance term. The analysis of 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 labelled 'cash'. Comparability also requires to start simulating returns for period 1 using the unconditional means $\boldsymbol{\mu}$.[4] As a final measure to maintain comparability we restrict short selling to a maximum leverage of 500% (i.e. l_i=−5 in equation 2). This is based on the results of Campbell et al. (2003), where the most extreme short position is about 360%.

We use scenario trees with four stages, starting with three one-year periods, and followed by a long period which lasts until the maximum age of 101 years. The node structure is 14×14×14 which amounts to 2744 scenarios. The linearization of the objective function is based on the curvature approach using 40 breakpoints.

Table 2 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. Similar to Table 1 we find that consumption is very precisely measured, and the standard errors of asset weights are comparatively large. As shown in Table 1, 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 γ=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 2). 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 2). Their asset allocation for γ=1 shows 220% in stocks, 242% in

[4]We could also use observed returns and state variables to show how the conditioning information affects the decisions over time, but this is not our focus here.

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 2 further shows that the asset allocation does not significantly change with age. This is contrary to popular belief which expects (or proposes) that older investors should uniformly invest less in risky assets than younger investors. A constant fraction of risky assets is an expected result if assets are assumed to follow a geometric Brownian motion. It comes as a surprise, however, that accounting for time-varying investment opportunities by way of the VAR model also leads to age-independent asset allocations.

One might argue that this age-independence 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 one-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). Second, we consider the case that the stochastic variables evolve according to the unconditional moments implied by the VAR model from the very beginning.[5] The results from the unconditional case in Table 3 can be compared to the results from the VAR model for stage t=0 from Table 2. There is a substantial difference in the asset allocation which reflects the impact of time-varying investment opportunities.

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 4 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

[5]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 therein). 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 3.

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

Table 2: Optimal consumption and asset allocation in t=0 for a man at age 20, 40, or 60, with risk aversion γ=1, 2, 5 and 10, and time discount factor d=0.92. 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 one year long. The node structure is 14×14×14 which amounts to 2744 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.

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 Section 2.5 we have argued that choosing rather long time intervals between stages may distort results compared to the case of short-period rebalancing. In Table 5 we show results for different rebalancing intervals which can be compared to the results for one-year intervals in Table 2. We analyze the effects for one example only (age 40 and γ=5). We find only a slight drop

	consumption	cash	stocks	bonds
γ=1	7.9	−206.3	187.9	118.4
	(0.0)	(7.7)	(2.1)	(6.1)
γ=2	6.7	−71.2	104.2	67.1
	(0.0)	(2.8)	(0.7)	(2.3)
γ=5	3.7	31.4	41.7	26.9
	(0.0)	(0.9)	(0.2)	(0.7)
γ=10	2.5	66.1	20.6	13.3
	(0.0)	(0.4)	(0.1)	(0.3)

Table 3: 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. Results are presented in terms of means and standard errors (in parentheses) from 100 solutions of the problem.

steady-state	consumption	cash	stocks	bonds
yes	30.8	5.6	39.8	54.6
	(0.0)	(1.5)	(0.4)	(1.3)
no	30.3	3.1	42.0	54.9
	(0.0)	(1.9)	(0.4)	(1.7)

Table 4: Optimal consumption and asset allocation in t=0 for a man at age 98 for γ=5 and d=0.92. 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.

in annual consumption. However, there are significant changes in the asset allocation[6] 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.

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 6 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

[6] Note, however, that results from increasing the rebalancing intervals even further need not converge to those from using unconditional moments in Table 3 which are based on annual rebalancing. In addition to this incompatibility, there are effects associated with differences in survival probabilities.

rebalancing intervals	consumption	cash	stocks	bonds
annual	4.4	2.4	42.8	54.9
	(0.0)	(2.1)	(0.5)	(1.8)
annual and bi-annual (twice)	4.5	−5.4	52.3	53.1
	(0.0)	(1.7)	(0.5)	(1.5)
bi-annual	4.5	33.6	50.7	15.7
	(0.0)	(1.7)	(0.5)	(1.5)
three years	4.5	33.5	54.6	11.9
	(0.0)	(1.7)	(0.8)	(1.3)

Table 5: Optimal consumption and asset allocation in t=0 for various choices of rebalancing intervals. Consumption is expressed in annual terms. 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 γ=5 and d=0.92.

of bonds by adding transaction costs. In all cases, however, the asset allocation remains rather unaffected by changing the age of the investor.

age	consumption	cash	stocks	bonds
	short sales excluded; no transaction costs			
20	3.9	55.7	44.3	0.0
	(0.0)	(0.2)	(0.1)	(0.0)
40	4.2	55.7	44.3	0.0
	(0.0)	(0.2)	(0.1)	(0.0)
60	5.7	55.8	44.2	0.0
	(0.0)	(0.2)	(0.2)	(0.0)
	short sales excluded; transaction costs			
20	3.8	59.6	40.4	0.0
	(0.0)	(0.1)	(0.1)	(0.0)
40	4.3	59.5	40.5	0.0
	(0.0)	(0.1)	(0.1)	(0.0)
60	5.6	59.7	40.3	0.0
	(0.0)	(0.1)	(0.1)	(0.0)
	short sales allowed; transaction costs			
20	4.0	52.7	39.2	8.1
	(0.0)	(0.9)	(0.3)	(0.8)
40	4.4	52.8	39.2	8.0
	(0.0)	(0.9)	(0.3)	(0.8)
60	5.9	52.9	38.9	8.1
	(0.0)	(0.9)	(0.3)	(0.8)

Table 6: 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). Results are presented for a male investor with risk aversion γ=5 and d=0.92 in terms of means and standard errors (in parentheses) from 100 solutions of the problem.

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

Table 7 shows optimal consumption and asset weights for various assumptions about labor income. Compared to Table 2 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 2 or the bottom panel of Table 6 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)[7]. They consider cases where initial labor income is about 30% to 40% 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.

4 Conclusion

We have presented a stochastic linear programming 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. We have first shown that available closed-form solutions can be accurately replicated with the SLP-based approach. Key requirements are exactly matching the moments of the (conditional) distributions of asset returns and state variables, and the linearization of the utility functions of consumption and bequest. 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 one-year periods followed by a long, steady-state period. Thereby the short-term dynamics of the VAR model and frequent rebalancing in the

[7]We have replicated their results to the extent possible given the differences in the two settings.

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

Table 7: Optimal consumption and asset allocation implied by the VAR model for various assumptions about current, annual labor income L_0. 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 γ=5 and d=0.92 in terms of means and standard errors (in parentheses) from 100 solutions of the problem.

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.

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

$$H_t = \int_{y_t}^{\bar{\tau}} \mathcal{L}(s) \frac{S(s)}{S(y_t)} \exp\{-(s - y_t)r\} ds, \tag{5}$$

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 Section 3 are based on the discrete-time version of labor income defined in Section 2.1, equation (1).

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

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

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

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

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

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 Section 2.1 are related to the i-th 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

Barberis, Nicholas C. 2000. Investing for the long run when returns are predictable. *Journal of Finance* **55** 225–264.

Bodie, Zvi, Robert C. Merton, William F. Samuelson. 1992. Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control* **16** 427–449.

Brandt, Michael W., Amit Goyal, Pedro Santa-Clara, Jonathan R. Stroud. 2005. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *The Review of Financial Studies* **18**(3) 831–873.

Brennan, Michael J., Eduardo S. Schwartz, Ronald Lagnado. 1997. Strategic asset allocation. *Journal of Economic Dynamics and Control* **21** 1377–1403.

Campbell, John Y., Yeung L. Chan, Luis M. Viceira. 2003. A multivariate model of strategic asset allocation. *Journal of Financial Economics* **67** 41–80.

Cariño, David R., David H. Myers, William T. Ziemba. 1998. Concepts, technical issues, and uses of the Russell-Yasuda-Kasai financial planning model. *Operations Research* **46** 450–462.

Chen, Peng, Roger G. Ibbotson, Moshe A. Milevsky, Kevin X. Zhu. 2006. Human capital, asset allocation, and life insurance. *Financial Analysts Journal* **62** 97–109.

Cocco, João F., Francisco J. Gomes, Pascal J. Maenhout. 2005. Consumption and portfolio choice over the life cycle. *Review of Financial Studies* **18 (2)** 491–533.

Dempster, Michael A. H., Matteo Germano, Elena A. Medova, Michael Villaverde. 2003. Global asset liability management. *British Actuarial Journal* **9** 137–216.

Duffie, Darrel. 1996. *Dynamic Asset Pricing Theory*. Princeton University Press.

Gassmann, Horand I., Eithan Schweitzer. 2001. A comprehensive input format for stochastic linear programs. *Annals of Operations Research* **104**(1-4) 89–125.

Geyer, Alois, William Ziemba. 2007. The Innovest Austrian pension fund financial planning model InnoALM. Working paper (submitted to *Operations Research*).

Gomes, Francisco, Alex Michaelides. 2005. Optimal life-cycle asset allocation: Understanding the empirical evidence. *Journal of Finance* **60** 869–904.

Hanke, Michael, Sonja Huber. 2006. Curvature, not second derivative. *Mathematical Spectrum* Forthcoming.

Heitsch, Holger, Werner Römisch. 2003. Scenario reduction algorithms in stochastic programming. *Computational Optimization and Applications* **24** 187–206.

Høyland, Kjetil, Michal Kaut, Stein W. Wallace. 2003. A heuristic for moment-matching scenario generation. *Computational Optimization and Applications* **24** 169–185.

Høyland, Kjetil, Stein W. Wallace. 2001. Generating scenario trees for multistage decision problems. *Management Science* **47** 295–307.

Ingersoll, Jonathan E. 1987. *Theory of financial decision making*. Rowman & Littlefield.

Kaut, Michal. 2003. Updates to the published version of 'A Heuristic for Moment-matching Scenario Generation by K. Høyland, M. Kaut, and S.W. Wallace'. URL http://work.michalkaut.net/CV_and_study/SG_Heuristic_updates.htm. Working paper.

King, Alan J., Stephen E. Wright, Gyana R. Parija, Robert Entriken. 2005. *Applications of Stochastic Programming*, chap. The IBM Stochastic Programming System. MPS SIAM - Series in Optimization, 21–37.

Klaassen, Pieter. 2002. Comment on 'Generating Scenario Trees for Multistage Decision Problems'. *Management Science* **48** 1512–1516.

Longstaff, Francis, Eduardo Schwartz. 2001. Valuing American options by simulation: A simple least-square approach. *Review of Financial Studies* **14 (1)** 113–147.

Lütkepohl, Helmut. 1993. *Introduction to Multiple Time Series Analysis*. Springer.

Merton, Robert C. 1969. Lifetime portfolio selection under uncertainty: the continuous-time case. *Review of Economics and Statistics* **51** 247–257.

Merton, Robert C. 1971. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* **3** 373–413.

Rasmussen, Kourosh Marjani, Jens Clausen. 2006. Mortgage loan portfolio optimization using multi-stage stochastic programming. *Journal of Economic Dynamics and Control* **31** 742–766.

Richard, Scott F. 1975. Optimal consumption, portfolio and life insurance rules for an uncertain lived individual in a continous time model. *Journal of Financial Economics* **2** 187–203.

Samuelson, Paul A. 1969. Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics* **51** 239–246.

Schroder, Mark, Costis Skiadas. 1999. Optimal consumption and portfolio selection with stochastic differential utility. *Journal of Economic Theory* **89 (1)** 68–126.

Wallace, Stein W., William T. Ziemba. 2005. *Applications of Stochastic Programming*. MPS-SIAM Book Series on Optimization.

Zenios, Stavros A., William T. Ziemba. 2006. *Handbook of Asset and Liability Management: Theory and Methodology*. North-Holland.

Ziemba, William T., J.M. Mulvey. 1998. *Worldwide Asset and Liability Modeling*. Cambridge University Press.