Network for Studies on Pensions, Aging and Retirement

Theo Nijman

Bas Werker

Ralph Koijen

Appendix to: When Can Life-cycle Investors Benefit from Time-varying Bond Risk Premia?

Discussion Paper 2006 - 017

May 2007 (replacing version of January, 2006)

Appendix Describing the Numerical Method Used in

When Can Life-cycle Investors Benefit from Time-varying Bond Risk Premia?*

Ralph S.J. Koijen[†] Theo E. Nijman[‡] Bas J.M. Werker[§]

This version: November 2006

Abstract

We discuss the numerical approach for the above-mentioned paper in detail. The methodology is based on Brandt, Goyal, Santa-Clara, and Stroud (2005) (*Review of Financial Studies*) and Carroll (2006) (*Economics Letters*). Next to combining these numerical techniques, we suggest two extensions. First, the approach of Brandt, Goyal, Santa-Clara, and Stroud (2005) approximates the conditional expectations encountered in optimizing the utility function via polynomial expansions in the state variables. The coefficients in these expansions are estimated using the cross-sectional regressions across a set of simulated trajectories of returns and state variables. In order to facilitate fast optimization over the portfolio weights, we develop an accurate approximation of these regression coefficients. This allows us to deal with a large number of decision variables without relying on iterative procedures. Second, to approximate the conditional expectations that lead to the optimal consumption strategy, we ensure that the approximation remains strictly positive, while keeping the approximation computationally tractable.

*We are grateful to Michael Brandt, Hans Schumacher, and Jules van Binsbergen for useful discussions on the numerical approach.

[†]Finance Group, CentER, Tilburg University and Netspar, Tilburg, the Netherlands, 5000 LE. Phone: 31-13-4663238. Email: R.S.J.Koijen@TilburgUniversity.nl.

[‡]Finance and Econometrics Group, CentER, Tilburg University and Netspar, Tilburg, the Netherlands, 5000 LE. Phone: 31-13-4662342. Email: Nyman@TilburgUniversity.nl.

[§]Finance and Econometrics Group, CentER, Tilburg University and Netspar, Tilburg, the Netherlands, 5000 LE. Phone: 31-13-4662532. Email: Werker@TilburgUniversity.nl.

1 Summary of the life-cycle problem

We briefly summarize the life-cycle investment problem before turning to the details of the numerical method used. The investor derives utility from real intermediate consumption (C_t/Π_t) and real terminal wealth (W_T/Π_T). More formally, the investor solves[1]

$$\max_{(C_t,x_t)\in\mathcal{K}_t} \mathbb{E}_0 \left(\sum_{t=1}^{T-1} \frac{\beta^t}{1-\gamma} \left(\frac{C_t}{\Pi_t}\right)^{1-\gamma} + \frac{\phi\beta^T}{1-\gamma} \left(\frac{W_T}{\Pi_T}\right)^{1-\gamma} \right), \tag{1}$$

where ϕ governs the relative importance of terminal wealth versus intermediate consumption. The investment opportunity set at time t is summarized by a vector of state variables, X_t. The simple, nominal asset returns are denoted by R_t and the nominal return on the single period cash account are indicated by R_t^f. The investor is entitled to a stream of labor income, denoted by Y_t at time t in nominal terms. The dynamics of wealth (cash-on-hand), W_t, is then given by

$$W_{t+1} = (W_t - C_t)\left(x_t'\left(R_{t+1} - \iota R_t^f\right) + R_t^f\right) + Y_{t+1}. \tag{2}$$

Section 2 of Koijen, Nijman, and Werker (2006) specifies the financial market and income dynamics. For notational convenience, we formulate the problem in real terms, with small letters indicating real counterparts, i.e.,

$$c_t = \frac{C_t}{\Pi_t},\ w_t = \frac{W_t}{\Pi_t},\ r_t = \frac{R_t\Pi_{t-1}}{\Pi_t},\ r_t^f = \frac{R_{t-1}^f\Pi_{t-1}}{\Pi_t},\ y_t = \frac{Y_t}{\Pi_t}. \tag{3}$$

This implies that the budget constraint is given in real terms by

$$w_{t+1} = (w_t - c_t)\left(x_t'\left(r_{t+1} - \iota r_{t+1}^f\right) + r_{t+1}^f\right) + y_{t+1}. \tag{4}$$

The relevant state variables at time t are given by (X_t, y_t, w_t) and the control variables by (c_t, x_t), i.e., the consumption and investment choice. The set $\mathcal{K}_t = \mathcal{K}(w_t)$ summarizes the constraints on the consumption and investment policy. We assume that the investor is liquidity constrained, i.e.,

$$c_t \leq w_t, \tag{5}$$

[1] We normalize the price index at time $t = 0$ to one.

which implies in turn that the investor cannot borrow against future income to increase today's consumption. Second, we impose standard borrowing and short-sales constraints

$$x_{st} \geq 0 \ (s = 1, \ldots, n) \text{ and } \iota' x_t \leq 1. \tag{6}$$

Formally, we have

$$\mathcal{K}(w_t) = \{(c, x) : c \leq w_t, \, x \geq 0, \text{ and } \iota' x \leq 1\}. \tag{7}$$

Note that the investor cannot default within the model as a result of these constraints.

We now discuss the dynamic problem the individual solves in detail. The optimal policies are determined by means of dynamic programming. We start in the final period, in which it is optimal to consume all wealth available, w_T, as there are no successor states in which the individual derives utility from wealth. The value function at time t is indicated by $J_t(w_t, X_t, y_t)$. Since the investor consumes all wealth at time T the value function at that moment equals

$$J_T(w_T, X_T, y_T) = \frac{\phi w_T^{1-\gamma}}{1-\gamma}. \tag{8}$$

The Bellman equation for all other points in time reads

$$J_t(w_t, X_t, y_t) = \max_{(c_t, x_t) \in \mathcal{K}_t} \left(\frac{c_t^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}_t \left(J_{t+1}(w_{t+1}, X_{t+1}, y_{t+1}) \right) \right). \tag{9}$$

It is important to remark that the problem is homogenous in (c_t, w_t, y_t). The value function in (1) is homogenous of degree $(1 - \gamma)$, the budget dynamics and liquidity constraint in (2) and (5) are homogenous of degree 1, and the portfolio constraints in (6) are homogenous of degree 0. We can therefore reduce the number of state variables by one and normalize either by income or wealth. To this end, recall that Koijen, Nijman, and Werker (2006) model real income in any specific period as

$$y_t = \exp(g_t + \nu_t + \epsilon_t), \tag{10}$$

with $\nu_{t+1} = \nu_t + u_{t+1}$, where $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ and $u_t \sim N(0, \sigma_u^2)$. This representation resembles Cocco, Gomes, and Maenhout (2005) and allows for both transitory (ϵ) and permanent (u) shocks. In line with Cocco, Gomes, and Maenhout (2005), we normalize all variables with

$\exp(\nu_t)$, implying

$$w_t^\nu = w_t \exp(-\nu_t),\ y_t^\nu = \exp(g_t + \epsilon_t),\ c_t^\nu = c_t \exp(-\nu_t). \tag{11}$$

The dynamic budget constraint changes to

$$w_{t+1}^\nu = (w_t^\nu - c_t^\nu)\left(x_t'\left(r_{t+1} - \iota r_{t+1}^f\right) + r_{t+1}^f\right)\exp(-u_{t+1}) + y_{t+1}^\nu, \tag{12}$$

and the Bellman equation to (with a modified value function $J_t^\nu(w_t^\nu, X_t)$)

$$J_t^\nu(w_t^\nu, X_t) = \max_{(c_t^\nu, x_t)\in\mathcal{K}_t}\left(\frac{c_t^{\nu(1-\gamma)}}{1-\gamma} + \beta\mathbb{E}_t\left(\exp((1-\gamma)u_{t+1})J_{t+1}^\nu(w_{t+1}^\nu, X_{t+1})\right)\right), \tag{13}$$

together with the terminal condition

$$J_T^\nu(w_T^\nu, X_T) = \frac{\phi}{1-\gamma}\left(w_T^\nu\right)^{1-\gamma}. \tag{14}$$

This normalization reduces the state variables at time t to (X_t, w_t^ν). In order to recover the original state variables and decision variables, we multiply the variables by $\exp(\nu_t)$.

2 Method

We now discuss the numerical approach used in Koijen, Nijman, and Werker (2006) that combines, in essence, the methods of Brandt, Goyal, Santa-Clara, and Stroud (2005) and Carroll (2006). Brandt, Goyal, Santa-Clara, and Stroud (2005) propose to approximate the conditional expectations that we encounter in solving the dynamic program by polynomial expansions in the state variables. The regression coefficients of these expansions are estimated via cross-sectional regressions across a set of simulated trajectories of asset returns and state variables (see for more details below). This does require that we can actually simulate the state variables. This not the case however for one of the state variables in our problem that is endogenous,[2] namely financial wealth, w_t^ν. The way to deal with endogenous state variables is by constructing a grid. We follow Carroll (2006) and construct a grid in wealth after consumption ($a_t^\nu = w_t^\nu - c_t^\nu$) as opposed to financial wealth (w_t^ν). As we will discuss in detail below, this particular choice of the grid allows us to solve for consumption in

[2] A state variable is said to be endogenous if its value at time t depends on decisions taken before time t.

analytically rather than solving the Euler equation numerically. Next to combining these two methods, the following two sections will introduce two extensions that improve the portfolio and consumption optimization.

We start to simulate N trajectories of T periods of both asset returns, state variables, and income innovations. We indicate the realized state at time t in trajectory i by ω_{it}, $i = 1, \ldots, N$. The realized value of the exogenous state variables (the real rate and expected inflationin our model) is denoted by $X_t(\omega_{it})$ and likewise for returns and income innovations. We choose an $M-$dimensional grid for financial wealth *after* consumption, i.e. $a_t^\nu = w_t^\nu - c_t^\nu$. The wealth grid points are indicated by $a_t^\nu(j)$, $j = 1, \ldots, M$. In total, we have $(M \times N)$ grid points at every point in time.

The problem is solved by means of dynamic programming. We outline the general recursion. In what follows, we omit the superscripts ν for notational convenience.

- **Time T:** At retirement, it is optimal for the investor to consume all wealth available, i.e. $c_T = w_T$. Hence, the value function equals

$$J_T(w_t, X_T) = \phi \frac{w_T^{1-\gamma}}{1-\gamma}. \tag{15}$$

- **Time T-1:** The problem at time $T-1$ can be summarized as

$$\max_{(c_{T-1}, x_{T-1}) \in \mathcal{K}_{T-1}} \frac{c_{T-1}^{1-\gamma}}{1-\gamma} + \phi\beta\mathbb{E}\left(\frac{w_T^{1-\gamma}}{1-\gamma} e^{(1-\gamma)u_T} \middle| X_{T-1} = X_{T-1}(\omega_{i(T-1)}), a_{T-1} = a_{T-1}(j) \right). \tag{16}$$

We first optimize over the asset allocation and then, given the optimal asset allocation, optimize over consumption, following Carroll (2006). The optimal asset allocation follows from

$$\begin{aligned} & x_{T-1}(\omega_{i(T-1)}, j) = \\ & \operatorname{argmax}_{x_{T-1} \in \mathcal{K}^x} \mathbb{E}\left(\frac{w_T^{1-\gamma}}{1-\gamma} e^{(1-\gamma)u_T} \middle| X_{T-1} = X_{T-1}(\omega_{i(T-1)}), a_{T-1} = a_{T-1}(j) \right), \end{aligned} \tag{17}$$

with $\mathcal{K}^x = \{(x) : x \geq 0 \text{ and } \iota' x \leq 1\}$ summarizing the constraints that apply specifically to the asset allocation. The first order conditions for the asset allocation read in turn

$$\mathbb{E}\left(\phi\beta e^{(1-\gamma)u_T} w_T^{-\gamma} \left(r_{t+1} - r_{t+1}^f \right) \middle| X_{T-1}, a_{T-1} \right) + \lambda - \mu\iota = 0, \tag{18}$$

where λ and μ are (non-negative) Kuhn-Tucker multipliers corresponding to the portfolio constraints.[3] In particular, the multipliers λ correspond to the short-sales constraints on each of the risky assets and μ to the borrowing constraint. Next, we approximate the conditional expectations in (17) and (20). Brandt, Goyal, Santa-Clara, and Stroud (2005) propose to approximate the conditional expectations with a polynomial expansion in the state variables. We then approximate for each of the n assets other than the cash account

$$\mathbb{E}\left(\phi\beta e^{(1-\gamma)u_T} w_T^{-\gamma}\left(r_{s,t+1}-r_{t+1}^f\right)\Big| X_{T-1}, a_{T-1}\right) \approx \theta_s(x, a_{T-1})' f(X_{T-1}), \tag{19}$$

with $s = 1, \ldots, n$. The projection coefficients, $\theta_s(x, a_{T-1})$, are subsequently estimated via cross-sectional regressions across the simulated trajectories. It is important to note that the regression coefficients depend on the portfolio weights. Using this approximation, we can solve for the optimal asset allocation. Section 3 introduces an optimization method that turns out to be very accurate and fast in our problem.

The optimal consumption is then, given the optimal asset allocation, determined by (using the standard first order conditions)

$$c_{T-1}^{\star}(\omega_{i(T-1)}, j) = \left\{\phi\beta\mathbb{E}\left(e^{(1-\gamma)u_T} w_T^{\star-\gamma} r_T^p \mid X_{T-1} = X_{T-1}(\omega_{i(T-1)}), a_{T-1} = a_{T-1}(j)\right)\right\}^{-\frac{1}{\gamma}} \tag{20}$$

with $w_T^{\star}$ the wealth level resulting from the optimal investment strategy and r_T^p refers to the portfolio return corresponding to the optimal strategy, $x_{T-1}^{\star}$. The final step is to construct the *endogenous* grid for cash-on-hand, i.e.,

$$w_{T-1}(\omega_{i(T-1)}, j) = c_{T-1}^{\star}(\omega_{i(T-1)}, j) + a_{T-1}(j), \tag{21}$$

which results in a different endogenous grid for every trajectory. The major advantage of the endogenous grid method is that we do not have to optimize over consumption numerically since we can solve the first-order condition in closed-form instead. Barillas and Fernández-Villaverde (2006) illustrate in a stochastic neoclassical growth model the computational gains that can be achieved by the endogenous grid method of Carroll (2006). The conditional expectations required for the optimal consumption policy are approximated along similar lines. Section 4 discusses the approximation for the optimal

[3] Together with the complementary slackness conditions, we have a set of necessary, and due to concavity of the value function, sufficient conditions of optimality.

consumption policy that we use in more detail.

For each time point, we collect the optimal investment and consumption policy, as well as the (endogenous) wealth grid.

- **Time $\mathbf{t = T - 2, \ldots, 1}$:** We now discuss the recursion at all other points in time. Suppose we have optimized the policies as of time $t+1$ onwards and we have available (i) the endogenous cash-on-hand grid at time $t+1$ (w_{t+1}) and (ii) the optimal consumption strategy in the next period ($c^\star_{t+1}$). This information suffices to determine the optimal asset allocation at time t and (given the optimal asset allocation) the optimal consumption strategy in turn. The value function at time t reads as

$$J_t(w_t, X_t, t) = \max_{(c_t, x_t) \in \mathcal{K}_t} \left(\frac{c_t^{1-\gamma}}{1-\gamma} + \mathbb{E}_t \left(\beta e^{(1-\gamma)u_{t+1}} J_{t+1}(w_{t+1}, X_{t+1}, t+1) \right) \right), \tag{22}$$

consistently with (9) before the normalization.

The first order conditions are given by

$$\begin{aligned} 0 &= \mathbb{E}_t \left(\beta e^{(1-\gamma)u_{t+1}} \frac{\partial J(w_{t+1}, X_{t+1}, t+1)}{\partial w_t} \left(r_{t+1} - r^f_{t+1} \iota \right) \right) + \lambda - \mu \iota, & (23) \\ c_t^{\star -\gamma} &= \mathbb{E}_t \left(\beta e^{(1-\gamma)u_{t+1}} \frac{\partial J(w_{t+1}, X_{t+1}, t+1)}{\partial w_t} r^p_{t+1} \right). & (24) \end{aligned}$$

for respectively the asset allocation and consumption policy. r^p_{t+1} refers to the portfolio return resulting from the optimal investment strategy, $x^\star_t$. The Kuhn-Tucker multipliers, λ and μ, are non-negative and satisfy the complementary slackness conditions

$$\lambda_s x_{st} = 0 \; (s = 1, \ldots, n), \text{ and } \mu(x_t' \iota - 1) = 0. \tag{25}$$

We next take the total derivative of (22) with respect to w_t, which results in

$$\frac{\partial J_t(w_t, X_t, t)}{\partial w_t} = \mathbb{E}_t \left(\beta e^{(1-\gamma)u_{t+1}} \frac{\partial J(w_{t+1}, X_{t+1}, t+1)}{\partial w_t} r^p_{t+1} \right), \tag{26}$$

i.e.,

$$c_t^{\star -\gamma} = \frac{\partial J_t(w_t, X_t, t)}{\partial w_t}. \tag{27}$$

The first order conditions then read

$$0 = \mathbb{E}_t\left(\beta e^{(1-\gamma)u_{t+1}} c_{t+1}^{\star-\gamma}\left(r_{t+1} - r_{t+1}^f \iota\right)\right) + \lambda - \mu\iota, \tag{28}$$

$$c_t^{\star-\gamma} = \mathbb{E}_t\left(\beta e^{(1-\gamma)u_{t+1}} c_{t+1}^{\star-\gamma} r_{t+1}^p\right). \tag{29}$$

Note that the objective function is concave in both the portfolio allocation and consumption. The Kuhn-Tucker conditions are as a result also sufficient.

We first solve for the optimal investment strategy via (28). The optimal consumption strategy is then given by

$$c_t^{\star} = \left(\mathbb{E}_t\left(\beta e^{(1-\gamma)u_{t+1}} c_{t+1}^{\star-\gamma} r_{t+1}^p\right)\right)^{-\frac{1}{\gamma}}. \tag{30}$$

Using the optimal consumption strategy, we construct the endogenous wealth grid as before

$$w_t(\omega_{it}, j) = c_t^{\star}(\omega_{it}, j) + a_t(j). \tag{31}$$

Both the investment and consumption policy only require the optimal consumption strategy at time $t+1$. We again approximate all conditional expectations we encounter by polynomial expansions in the basis functions and estimate the coefficients via cross-sectional regressions like in Brandt, Goyal, Santa-Clara, and Stroud (2005).

The optimization over the asset allocation boils down to solving for the root of (28), taking into the borrowing and short-sales constraints. Since we know the optimal consumption policy at time $t+1$ only at the (endogenous) grid points, we interpolate the consumption policy linearly for intermediate values. It is well-known that the optimal consumption policy is almost linear in wealth, except for very low wealth levels. We therefore select a grid with a triple exponential growth rate between the grid points, see Carroll (2006) for further details.

In interpolating the consumption policy, we have to take into account the liquidity constraint in (5). To this end, we always include the grid point $a_t(1) = 0$ as lowest grid point, i.e., the wealth after consumption equals zero. Since optimal consumption is increasing in wealth, we know that the liquidity constraint binds for any wealth level (w_{t+1}) smaller than the endogenous wealth level corresponding to this lowest grid point,

$w_{t+1}(\omega_{it}, 1) = a_{t+1}(1) + c^{\star}_{t+1}(\omega_{it}, 1)$. Numerically, this constraint can be implemented by adding the point $(c^{\star}_{t+1}, w_{t+1}) = (0, 0)$ to the endogenous grid as suggested by Carroll (2006). This implements exactly the binding liquidity constraint in the region $w_{t+1}(\omega_{it}, \cdot) \in [0, a_{t+1}(1) + c^{\star}_{t+1}(\omega_{it}, 1)]$ if we interpolate the next period's consumption policy linearly.

For each time point, we collect the optimal investment and consumption policy, as well as the (endogenous) wealth grid.

This method results in the optimal policies at all $(N \times M)$ grid points at each point in time. The last step is to start from the initial state vector, X_0, and to simulate forward. Depending on the realized wealth levels, we implement the optimal investment and consumption strategies. This generates the optimal policies along the N trajectories, which we either average or project on the state variables to illustrate how the policies respond to changes in the state variables. This results in the figures presented in Koijen, Nijman, and Werker (2006) to summarize the optimal strategies over the life-cycle.

The next two sections improve upon the numerical procedure described above. First, we develop an accurate and fast optimization technique under portfolio constraints. Second, we show how to solve for the intermediate consumption problem efficiently using the simulation-based method. The main complication is that when we regress marginal utility derived from consumption in the next period on a polynomial expansion in the state variables, it is not ensured that the resulting fitted values will be strictly positive. We use an alternative approximation to the conditional expectations that are required to determine the consumption policies instead, which is ensured to be strictly positive. This is moreover without significantly increasing the computational effort required.

3 Enhanced portfolio optimization

The constrained portfolio optimization is the most time-consuming step in our numerical procedure. After all, we have to optimize the portfolio for each of the $(N \times M)$ grid points at a particular moment in time. Every different evaluation of a portfolio requires a cross-sectional regression across the N paths, see (19). This implies that we have to perform for each of the $(N \times M)$ grid points several cross-sectional regressions across N paths, where N is typically large (we take $N = 30,000$). This turns out to be very time-consuming and is therefore infeasible.

To circumvent this problem, Brandt, Goyal, Santa-Clara, and Stroud (2005) propose a recursive solution method that is based on a fourth-order expansion of the utility index. The starting values for this recursion are determined using a second-order expansion of the utility index, which can be solved for analytically. However, as remarked by DeTemple, Garcia, and Rindisbacher (2005), such a recursion is not ensured to converge to the optimum. In addition, it is non-trivial to include portfolio constraints in this recursion.

One possible alternative to circumvent this recursion is to construct a grid of possible portfolio policies and perform grid search instead. This is the approach taken in van Binsbergen and Brandt (2006) and Koijen, Rodriguez, and Sbuelz (2006). Yet, again, grid search becomes computationally intensive for fine grids and multiple risky assets.

We propose an alternative method that is motivated by the observation that the projection coefficients, $\theta_s(x, a_t)$, in (19) are smooth functions in the portfolio weights, x. To illustrate, suppose that we construct a grid for the possible portfolio policies and calculate the projection coefficients for each of these portfolio and all points in the wealth grid (a_t) at a particular point in time. Then the average R^2, averaged across all points in the wealth grid, of a regression of the projection coefficients on a constant and the portfolio weights (i.e., a first order approximation) is in general comfortably above 99%. We will exploit exactly this property in the method described below.

We fix the state in the wealth grid for the moment, so that the projection coefficients depend only on the portfolio weights, $\theta(x)$. We first estimate the projection coefficients on a coarse grid of H possible portfolios in the set of admissible portfolios, $\mathcal{K}^x$. We will refer to these portfolios as *test portfolios*. These H test portfolios, x^k, $k = 1, \ldots, H$, result in H projection coefficients, $\hat{\theta}_s^k = \hat{\theta}_s\left(x^k\right)$, $k = 1, \ldots, H$ and $s = 1, \ldots, n$. The next step is to parameterize the regression coefficients in the asset weights

$$\theta_s(x) \approx \Psi_s g(x), \tag{32}$$

with $g(x)$ a polynomial expansion in the portfolio weights. This implies in turn that the conditional expectation of marginal utility is approximated via, $s = 1, \ldots, n$

$$\mathbb{E}_t\left(\beta e^{(1-\gamma)u_{t+1}} c_{t+1}^{\star-\gamma}\left(r_{s,t+1} - r_{t+1}^f\right)\right) \approx g(x)'\Psi_s' f(X_t). \tag{33}$$

The efficiency gain is attainable due to the fact that we reduced N times the number of regressions per state required to find the optimum to H regressions, where H will be relatively small. Even though it is theoretically possible to estimate Ψ at once, it requires a very large regression, which makes it computationally unattractive. The final step is to solve for the root of the polynomial, taking into account the portfolio constraints for which fast numerical techniques are at our disposal.

In optimizing the portfolio weights under constraints, we employ that the value function is strictly concave, and the constraints linear. This implies that the Kuhn-Tucker conditions are also sufficient for a global maximum.

We consider an affine expansion of the projection coefficients, θ_s, $s = 1, \ldots, n$. If for the projection coefficients turn out to be non-linear, we can easily increase the order of the polynomial. However, the affine approximation turns out to be sufficient for the life-cycle problem in Koijen, Nijman, and Werker (2006). Hence, we solve the following system for each trajectory

$$0 = \begin{pmatrix} 1 \\ x \end{pmatrix}' \Psi_s f(X_t(\omega_i)) + \lambda_s - \mu \ (s = 1, \ldots, n), \tag{34}$$

$$\lambda_s x_s = 0 \ (s = 1, \ldots, n) \text{ and } \mu(x'\iota - 1) = 0, \tag{35}$$

$$\lambda \geq 0 \text{ and } \mu \geq 0. \tag{36}$$

The system can be written in more compact form if we define, for $p = 0, \ldots, n$,

$$\Theta_p^i = \begin{pmatrix} \Psi_{1(p+1,:)} \\ \cdots \\ \Psi_{n(p+1,:)} \end{pmatrix} f(X_t(\omega_i)), \tag{37}$$

so that we can rewrite (34) for all $s = 1, \ldots, n$ as

$$0_{n\times 1} = \Theta_0^i + \sum_{s=1}^{n} \Theta_s^i x_s + \lambda - \mu\iota. \tag{38}$$

This system can efficiently be solved for using standard numerical techniques.

To summarize, rather than using an iterative procedure or grid search, we propose to approximate the projection coefficients of the polynomial expansions. A low order approximation turns out to be very accurate and reduces the number of cross-sectional regressions required dramatically.

In the numerical part of the paper, we select test portfolios on a 20% grid. Refining this grid does not change our results at the reported precision. For two risky assets, we then have $H = 21$, which is is a considerably lower number of regressions than the $(N \times M)$ regressions multiplied by the number of function evaluations required to optimize the optimal portfolio. We use $N = 30,000$ simulated trajectories and $M = 128$ points to discretize the wealth space. Increasing the number of grid points for this endogenous state variables does not affect our results at the reported precision.

4 Optimal consumption policies

The optimal consumption policy is determined in (30). We approximated the conditional expectations so far by a polynomial expansion in the state variables. Yet, this approximation is inappropriate for the optimal consumption strategy. As can be seen from (30), we need to make sure that the conditional expectation remains strictly positive, otherwise the implied optimal consumption will be negative. This section develops an approximation that ensures that the conditional expectation remains strictly positive, and does remain computationally tractable. To this end, we approximate the logarithm of the conditional expectation with a polynomial expansion in the state variables

$$\mathbb{E}_t(Q_{t+1}) \approx \exp(\bar{\theta}_0 + \bar{\theta}' f(X_t)), \tag{39}$$

with

$$Q_{t+1} = \beta e^{(1-\gamma)u_{t+1}} c_{t+1}^{-\gamma\star} r_{t+1}^{p}, \tag{40}$$

where the projection coefficients $(\bar{\theta})$ are obviously different from the ones in the linear approximation (θ).

This approximation ensures indeed that the optimal consumption strategy remains strictly positive. However, we emphasize that this approximation comes with little additional computational cost. Even though the coefficients have to be estimated via non-linear least squares, this regression has to be performed only once for the M grid points for the endogenous state variable. This stems from the fact that optimal consumption is known analytically once the conditional expectation are determined. Otherwise, we would have to perform this regression for iteration of the optimization, which makes it infeasible. Thus by adopting the endogenous grid method, we are able to approximate the conditional expectations that lead to the optimal consumption strategy so that they remain strictly positive. This allows us to solve for the intermediate consumption problem using the simulation-based approach.

5 Optimization for different types of investors

To understand the value of myopically timing bond markets as well as hedging time variation in investment opportunities, Koijen, Nijman, and Werker (2006) introduce three types of investors. The Strategic Investor implements the optimal strategies derived so far. The Conditionally Myopic Investor does time bond markets myopically, but abstract from any time-dependence in the state variables governing the investment opportunity set. As a result, the Conditionally Myopic Investor will not hold hedging demands. The third and final investor is the Unconditionally Myopic Investor. This investor ignores time variation in investment opportunities all together. We summarize how we can derive the optimal strategies, and thus the value function, for these three types of investors.

The Strategic Investor follows the strategies that have been derived so far. The Conditionally Myopic Investor does take the conditional one-period ahead ditribution of returns, $\mathcal{L}(R_{t+1} \mid X_t)$, into account. However, the Conditionally Myopic Investor does not take into account that future returns possibly correlate with future investment opportunities, which induces hedging demands. This can be implemented in the simulation-based approach by randomizing the pairs of returns at time $t + 1$ and the state variables at time t cross-sectionally. This preserves the one-period ahead conditional return distribution, but state variables and returns are otherwise independent over time.

The Unconditionally Myopic Investor ignores any time variation in investment opportunities. This is easily accommodated by only including a constant in the cross-sectional regressions that lead in turn to the optimal investment and consumption strategy. Since we solve this problem for each point in the wealth grid, this investor does take into account the ratio of financial wealth relative to human capital, but ignores time-variation in investment opportunities.

References

Barillas, J., and J. Fernández-Villaverde (2006): "A Generalization of the Endogenous Grid Method," Working Paper, New York University and the University of Pennsylvania.

Brandt, M. W., A. Goyal, P. Santa-Clara, and J. R. Stroud (2005): "A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Return Predictability," *Review of Financial Studies*, 18, 831–873.

Carroll, C. D. (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," Forthcoming.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005): "Consumption and Portfolio Choice over the Life-Cycle," *Review of Financial Studies*, 18, 491–533.

DeTemple, J. B., R. Garcia, and M. Rindisbacher (2005): "Intertemporal Asset Allocation: a Comparison of Methods," *Journal of Banking and Finance*, 29, 2821–2848.

Koijen, R. S. J., T. E. Nijman, and B. J. M. Werker (2006): "When Can Life-cycle Investors Benefit from Time-varying Bond Risk Premia?," Working Paper, Tilburg University.

Koijen, R. S. J., J.-C. Rodriguez, and A. Sbuelz (2006): "Momentum and Mean-Reversion in Strategic Asset Allocation," Working Paper.

van Binsbergen, J. H., and M. W. Brandt (2006): "Optimal Asset Allocation in Asset Liability Management," Working Paper, Duke University.