Netspar Network for Studies on Pensions, Aging and Retirement

Theo Nijman

Bas Werker

Ralph Koijen

When Can Life-Cycle Investors Benefit from Time-Varying Bond Risk Premia?

Discussion Paper 2006 - 017

February, 2009 (revised version from January, 2006)

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

Ralph S.J. Koijen† University of Chicago

Theo E. Nijman‡ Tilburg University

Bas J.M. Werker§ Tilburg University

This version: February 2009

Abstract

We study the economic importance of time-varying bond risk premia in a life-cycle consumption and portfolio-choice problem for an investor facing short-sales and borrowing constraints. On average, the investor is able to time bond markets only as of age 45. Tilts in the optimal asset allocation in response to changes in bond risk premia exhibit pronounced life-cycle patterns. Taking as a point of reference an investor who conditions only on age and wealth, we compute the management fee this investor is willing to pay to account for either current risk premia or for both current and future risk premia. We find the fees to account for current risk premia to be economically sizeable, ranging up to 1% per annum, but this fee is comparable to the fee of the fully optimal strategy. To solve our model, we extend recently developed simulation-based techniques to life-cycle problems featuring multiple state variables and multiple risky assets.

*First version: September 2005. We are grateful to an anonymous referee, Jules van Binsbergen, Lans Bovenberg, Michael Brandt, Laurent Calvet, John Cochrane, Greg Duffee, Thijs van der Heijden, Frank de Jong, Martin Lettau, Hanno Lustig, Claus Munk, Hans Schumacher, Ken Singleton, Georgios Skoulakis, Carsten Sorensen, Raman Uppal, Otto Van Hemert, Stijn Van Nieuwerburgh, Luis Viceira, Rui Yao, and seminar participants at Duke university - Fuqua School of Business, University of Maastricht, Tilburg University, 2007 EFA Meetings (Ljubljana), 2007 WFA Meetings (Montana), the Imperial College Financial Econometrics Conference London, and the 5th Winterschool of Financial Mathematics for helpful comments and suggestions.

†University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, 60637 Chicago, Illinois. Phone: (773)834-4199. Email: ralph.koijen@chicagogsb.edu. Koijen is also associated with Netspar.

‡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.

Bond risk premia vary over time. For instance, Cochrane and Piazzesi (2005 2006), report that a single predictor variable, which is a linear combination of forward rates, explains up to 44% of the variation in long-term bond returns in excess of the one-year rate at an annual horizon. This suggests that investors can construct dynamic bond strategies that take advantage of this stylized fact. Indeed, Sangvinatsos and Wachter (2005) show that (unconstrained) long-term investors can realize large gains by exploiting time variation in bond risk premia. One might argue that such utility gains resonate with the stellar growth in the fixed-income mutual fund industry. As Figure 1 illustrates, total assets under management has grown tremendously during the last two decades. The same holds true for the number of fixed-income mutual funds.

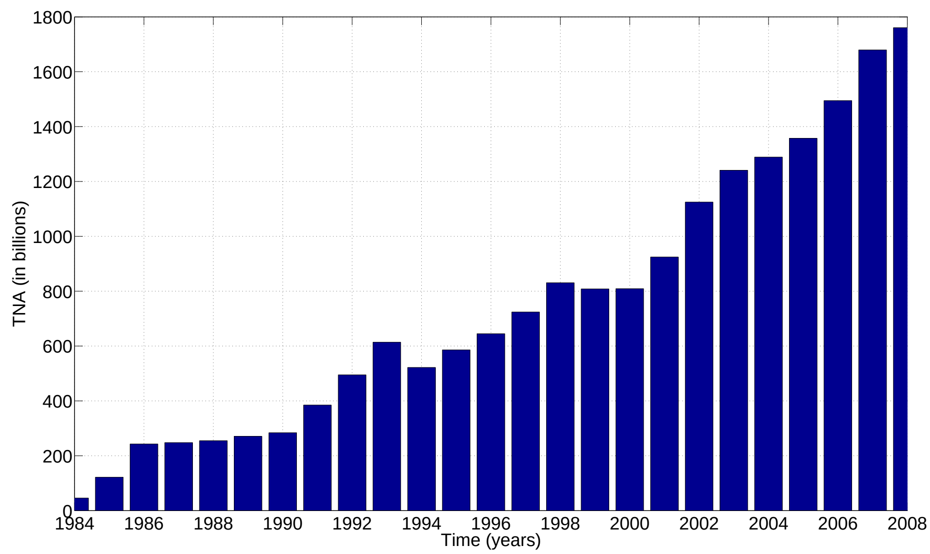

Figure 1: Growth in the fixed-income mutual fund industry from 1984 to May 2008
The figure displays the total net assets of all fixed-income mutual funds. Source: Trends In Mutual Fund Investing from the Investment Company Institute over the period 1984 to May 2008. The fixed-income mutual fund industry is taken to be the sum of "Taxable Bond Funds" and "Municipal Bond Funds."

Our main question is whether the gains an investor realizes by accounting for variation in bond risk premia carry over to a carefully calibrated life-cycle model in which households receive labor income and are restricted by borrowing and short-sales constraints. We take as a natural point of reference a strategy that conditions on households' wealth and age only, but not on the current term structure of interest rates.[1] Subsequently, we compute the maximum annual management fee a household would be willing to pay to access a fund implementing a strategy that incorporates information about current bond risk premia. However, we assume that the fund reverts to the strategy that conditions only on age and wealth after one period. We call this strategy the one-period conditional

[1]This problem has been studied by for instance in Viceira (2001) and Cocco, Gomes, and Maenhout (2005).

strategy. Likewise, we calculate the maximum fee that can be charged by a so-called life-cycle fund implementing the fully optimal dynamic strategy.[2] We find that the fees of the one-period conditional strategy can be sizeable and reach a maximum of around 1% per annum. Comparing such numbers to low-cost, actively-managed mutual funds, our results suggest that it may be possible to benefit from variation in bond risk premia. Second, we find that the fees of the one-period conditional strategy and the life-cycle strategy virtually coincide. It implies that the investor gains little by accounting for the fact that she can also time bond markets in future periods.

To understand the gains of incorporating information about variation in bond risk premia, we find that we can split the life-cycle roughly into four periods. The first period (age 25 to 35) is characterized by a large stock of non-tradable human capital. Borrowing constraints inhibit the household to capitalize on future labor income to increase today's consumption. The investor, therefore, consumes (almost) all income available and hardly participates in financial markets. During the second stage (age 35 to 45), the investor accumulates some financial wealth and allocates it almost exclusively to equity markets. Human capital resembles to a large extent a (non-tradable) position in inflation-linked bonds, which reduces the individual's effective risk aversion. Because the investor cannot borrow using its human capital as collateral, it would have to reduce the equity allocation to invest in long-term bonds. This leads to opportunity costs that are too high for an empirically plausible range of bond risk premia. In the third period (age 45 to 55), the individual holds substantial bond positions as its human capital diminished sufficiently. In addition, the individual optimally tilts the portfolio towards long-term nominal bonds in periods of high bond risk premia. These tilts are economically significant and the holdings range from -20% to $+45\%$ for a plausible range of bond risk premia. During the last period (age 55 to 65), the stock of human capital has largely been depleted and the individual acts more conservatively as a result. In addition to stocks and long-term bonds, the individual holds cash positions of up to 20% on average just before retirement. The opportunity costs of reducing the cash position to tilt the portfolio to long-term bonds are smaller than the opportunity costs of cutting back on the equity allocation. This implies that in periods of high bond risk premia, the investor first reduces the cash position and, only for relatively high bond risk premia, the equity allocation as well. This results in pronounced life-cycle patterns in the tilts caused by variation in bond risk premia.

Perhaps the most important economic question is how costly it is to ignore information about current bond risk premia in the consumption and investment policies. To address this question, we first compute a proportional management fee that makes the investor indifferent between investing in a fund that implements the one-period conditional strategy and the strategy that conditions only on age and wealth. Our first important

[2] An alternative interpretation of our approach is that we compute the maximum fee a financial planner can charge to implement dynamic fixed-income strategies on behalf of households.

finding is that the management fees can be sizeable. It suggests households can benefit from variation in bond risk premia using low-cost bond mutual funds. Next, we offer the investor a fund or strategy that is dynamically optimal over the life-cycle. This strategy accounts for the fact that the investor can also exploit variation in future bond risk premia. We again compare this strategy to the one that conditions on wealth and age only. Our second important finding is that the fee that the life-cycle fund can charge to make the household indifferent between this advanced strategy and a simpler strategy that conditions on wealth and age only virtually coincides with the fees of one-period conditional strategy. As such, there are only small gains possible when accounting for the fact that the investor can also time bond markets in future periods.[3] We also compare the portfolio weights of the one-period conditional strategy to the ones of the life-cycle strategy as an alternative way to illustrate this result. The difference between both strategies can be interpreted as a form of hedging demands. We find the difference between both strategies to be hump-shaped over the life-cycle, but they are small and amount to a position of at most 4%.

In the richest asset menu that we consider, the investor has access to two long-term bonds with different maturities, stocks, and a nominal cash account. By studying the optimal investment strategy over the life-cycle, we derive the optimal maturity structure of the bond portfolio for fixed-income mutual funds targeting a specific age group. We find that the duration of the overall portfolio is hump-shaped over the life-cycle, whereas the duration of the fixed-income portfolio is decreasing over the life-cycle. In the early stages of the life-cycle, the investor allocates almost all capital to stocks and, if at all, a small fraction to the longest-term bond. If the household ages, it becomes effectively more risk averse. This implies that the allocation to equity gradually declines and that the household allocates more capital to the bond with the shorter maturity. This shortens the duration of the fixed-income portfolio. Combining both effects leads to the hump-shaped pattern for the overall portfolio and the decreasing pattern if we restrict attention to the fixed-income portfolio.

We calibrate our model on the basis of US data over the period January 1959 to December 2005. Using this model, we show that the individual can exploit short-term variation in bond risk premia mostly during later stages of the life-cycle. Life-cycle constraints prevent investors to do so before that age. As the individual ages, two effects come into play. First, borrowing and short-sales constraints become less restrictive due to the decreased ratio of human capital to financial wealth. Second, the investor becomes more conservative due to the lower amount of human capital. The first effect by itself would lead to more sensitive allocations to risky assets over the life cycle. The second effect would, on the contrary, lead to less sensitive allocations. For persistent risk factors

[3]The investor still conditions on wealth and age, which implies that the investor does not act myopically.

that are priced, the first effect is more pronounced and we find the sensitivity of the bond allocation in response to changes in the term structure variables to increase over the life cycle.

To derive our results, we also make an important methodological contribution by extending the simulation-based approach by Brandt, Goyal, Santa-Clara, and Stroud (2005). Specifically, we improve upon the optimization over the optimal asset allocation and show how to optimize over consumption in a computationally efficient way by combining the simulation-based approach with the endogenous grid method introduced by Carroll (2006). We therefore show how simulation-based techniques are useful to solve complex life-cycle problems with multiple state variables. A separate appendix that is available online contains further details. In another application, Chapman and Xu (2007) use our approach to solve for the optimal consumption and investment problem of mutual fund managers.

Our model of the financial market accommodates time-varying interest rates, inflation rates, and bond risk premia. It is closely related to Brennan and Xia (2002) and Campbell and Viceira (2001), but both papers assume bond risk premia to be constant. These papers study the optimal demand for long-term bonds and show that it is optimal to hedge time variation in real interest rates, in particular for conservative investors (see also Wachter (2003)). Sangvinatsos and Wachter (2005) do allow for time variation in bond risk premia. They conclude that long-term investors that are not restricted by portfolio constraints and not endowed with non-tradable labor income can realize large economic gains by both timing bond markets and hedging time variation in bond risk premia. This is in line with the recent asset allocation literature, which emphasizes the importance of time-varying risk premia for both tactical, short-term investors and strategic, long-term investors, see Barberis (2000), Brandt (1999), and Campbell and Viceira (1999), Campbell, Chan, and Viceira (2003), Jurek and Viceira (2007), and Wachter (2002). However, the focus of these papers is not on life-cycle investors with its inherent constraints and labor income.

Our paper also relates to the life-cycle literature, see Cocco, Gomes, and Maenhout (2005), Gomes and Michaelides (2005), Gourinchas and Parker (2002), Heaton and Lucas (1997), and Viceira (2001). These papers focus predominantly on the impact of risky, non-tradable human capital on the consumption and portfolio choice decision. These studies find (i) that there are strong age effects in the optimal asset allocation as a result of changing human capital, (ii) binding liquidity constraints during early stages of the individual's life-cycle, (iii) a negative relation between income risk and the optimal equity allocation, and (iv) a high sensitivity of the optimal asset allocation to correlation between income risk and financial market risks. However, these papers restrict attention to financial markets with constant investment opportunities,[4] including constant interest

[4] Gourinchas and Parker (2002) focus on optimal consumption policies and wealth accumulation, and

and inflation rates, and bond risk premia.

Closest to our paper are presumably Munk and Sørensen (2005) and Van Hemert (2006). Both papers allow for risky, non-tradable labor income and impose standard constraints on the strategies implemented. Munk and Sørensen (2005) accommodate stochastic real rates, but assume inflation rates and bond risk premia to be constant. Van Hemert (2006) does allow for stochastic inflation rates and includes housing, but assumes risk premia to be constant. We allow for time variation in bond risk premia instead and analyze how individuals can benefit from such time variation over the life-cycle. We thus examine the interaction between exploiting time variation in investment opportunities and both realistic life-cycle constraints and changing labor income.

1 Financial market and the individual's problem

1.1 Financial market

Our financial market accommodates time variation in bond risk premia. The model we propose is closely related to Brennan and Xia (2002), Campbell and Viceira (2001), and Sangvinatsos and Wachter (2005). Brennan and Xia (2002) and Campbell and Viceira (2001) propose two-factor models of the term structure, where the factors are identified as the real interest rate and expected inflation. Both models assume that bond risk premia are constant. Sangvinatsos and Wachter (2005) use a three-factor term structure model with latent factors and accommodate time variation in bond risk premia, in line with Duffee (2002). We consider a model with a factor structure as in Brennan and Xia (2002) and Campbell and Viceira (2001), but generalize these models by allowing for time-varying bond risk premia.

The asset menu of the life-cycle investor includes a stock (index), long-term nominal bonds, and a nominal money market account. The dynamics of the term structure of interest rates is governed by two state variables X_{1t} and X_{2t}. To accommodate the first-order autocorrelation in the interest rate, we model $X_t = (X_{1t}, X_{2t})'$ to be mean-reverting around zero, i.e.,

$$\mathrm{d}X_t = -K_X X_t \mathrm{d}t + \begin{bmatrix} I_{2\times 2} & 0_{2\times 2} \end{bmatrix} \mathrm{d}Z_t, \tag{1}$$

where $I_{k\times k}$ is a $k-$dimensional identity matrix and $Z \in \mathbb{R}^{4\times 1}$ is a vector of independent Brownian motions driving the uncertainty in the financial market. Any correlation between the processes is captured by the volatility vectors. Following Dai and Singleton (2000), we normalize K_X to be lower triangular.

abstract from optimal life-cycle portfolio choice.

The instantaneous nominal interest rate, R^0, is assumed to be affine in both factors

$$R_t^0 \quad = \quad \delta_{0R} + \delta_{1R}' X_t,\ \ \delta_{0R} > 0. \tag{2}$$

We postulate a process for the (commodity) price index

$$\frac{\mathrm{d}\Pi_t}{\Pi_t} = \pi_t \mathrm{d}t + \sigma_\Pi' \mathrm{d}Z_t,\ \ \sigma_\Pi \in \mathbb{R}^4,\ \Pi_0 = 1, \tag{3}$$

where π_t denotes the instantaneous expected inflation. Instantaneous expected inflation is assumed to be affine in both factors

$$\pi_t = \delta_{0\pi} + \delta_{1\pi}' X_t,\ \ \delta_{0\pi} > 0. \tag{4}$$

Concerning the stock index, S, we postulate

$$\frac{\mathrm{d}S_t}{S_t} = \left(R_t^0 + \eta_S\right)\mathrm{d}t + \sigma_S' \mathrm{d}Z_t,\ \ \sigma_S \in \mathbb{R}^4, \tag{5}$$

where η_S the (constant) equity risk premium.

To complete our model, we specify an affine model for the term structure of interest rates by assuming that the prices of risk are affine in X_t. More precisely, the nominal state price density $\phi^\$$ is given by

$$\frac{\mathrm{d}\phi_t^\$}{\phi_t^\$} = -R_t^0 \mathrm{d}t - \Lambda_t' \mathrm{d}Z_t, \tag{6}$$

with

$$\Lambda_t = \Lambda_0 + \Lambda_1 X_t. \tag{7}$$

We thus adopt the essentially affine model as proposed by Duffee (2002). In the terminology of Dai and Singleton (2000), the model is the maximal $\mathbb{A}_0(2)$ model. The conditions specified in Duffie and Kan (1996) to ensure that bond prices are exponentially affine in the state variables are satisfied. Hence, we find for the price of a nominal bond at time t with a maturity $t + \tau$,

$$P(t, t+\tau) = \exp(A(\tau) + B(\tau)' X_t), \tag{8}$$

so that the corresponding yield is $y_t^\tau = -A(\tau)/\tau - B(\tau)' X_t/\tau$.

This specification accommodates time variation in bond risk premia as advocated by, for instance, Dai and Singleton (2002) and Cochrane and Piazzesi (2005). As we assume the equity risk premium to be constant, we have

$$\sigma_S'\Lambda_t = \eta_S, \tag{9}$$

which restricts Λ_0 and Λ_1.

1.2 Individual's preferences, labor income, and constraints

We consider a life-cycle investor who starts working at age 0 and retires at age T. The individual derives utility from real consumption, C_t/Π_t, and real retirement capital, W_T/Π_T. The individual's preferences are summarized by a time-separable, constant relative risk aversion utility index. More formally, the individual solves

$$\max_{(C_t,x_t)\in\mathcal{K}_t} \mathbb{E}_0\left(\sum_{t=t_0}^{T-1}\frac{\beta^t}{1-\gamma}\left(\frac{C_t}{\Pi_t}\right)^{1-\gamma}+\frac{\varphi\beta^T}{1-\gamma}\left(\frac{W_T}{\Pi_T}\right)^{1-\gamma}\right), \tag{10}$$

where φ governs the utility value of terminal wealth relative to intermediate consumption, β denotes the subjective discount factor, and $\mathcal{K}_t$ summarizes the constraints that have to be satisfied by the consumption and investment strategy at time t. We discuss these constraints in detail below. The fraction of wealth allocated to the risky assets at time t is indicated by x_t. The remainder, $1-x_t'\iota$, is allocated to a nominal cash account.

The nominal, gross asset returns are denoted by R_t and the nominal, gross return on the single-period cash account is indicated by R_t^f. The dynamics of financial wealth, W_t, is then given by

$$W_{t+1} = (W_t - C_t)\left(x_t'\left(R_{t+1}-\iota R_t^f\right)+R_t^f\right)+Y_{t+1}, \tag{11}$$

in which Y_t denotes the income received at time t in nominal terms. The supply of labor is assumed to be exogenous.[5] For notational convenience, we formulate the problem in real terms, with small letters indicating real counterparts, i.e.,

$$c_t = \frac{C_t}{\Pi_t},\ w_t = \frac{W_t}{\Pi_t},\ r_t = \frac{R_t\Pi_t}{\Pi_{t-1}},\ r_t^f = \frac{R_{t-1}^f\Pi_{t-1}}{\Pi_t},\ y_t = \frac{Y_t}{\Pi_t}. \tag{12}$$

The resulting budget constraint in real terms reads

$$w_{t+1} = (w_t - c_t)\left(x_t'\left(r_{t+1}-\iota r_{t+1}^f\right)+r_{t+1}^f\right)+y_{t+1}. \tag{13}$$

The state variables 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

[5] Chan and Viceira (2000) relax this assumption and consider an individual who can supply labor income flexibly instead.

the consumption and investment policy. First, we assume that the investor is liquidity constrained, i.e.,

$$c_t \le w_t, \tag{14}$$

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

$$x_t \ge 0 \text{ and } \iota' x_t \le 1. \tag{15}$$

Formally, we have

$$\mathcal{K}(w_t) = \{(c, x) : c \le w_t,\ x \ge 0, \text{ and } \iota' x \le 1\}. \tag{16}$$

Note that the investor cannot default within the model as a result of these constraints.[6]

We model real income in any specific period as

$$y_t = \exp(g_t + \nu_t + \epsilon_t), \tag{17}$$

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 follows Cocco, Gomes, and Maenhout (2005) and allows for both transitory (ϵ) and permanent (u) shocks to labor income. We calibrate g_t consistently with Cocco, Gomes, and Maenhout (2005) to capture the familiar hump-shaped pattern in labor income over the life-cycle (see Section 1.4.3 for details). In our benchmark specification, both income shocks are uncorrelated with financial market risks. In Section 3, we also consider the case in which permanent income shocks, i.e., u_t, are correlated with financial market risks.

We solve for the individual's optimal consumption and investment policies by dynamic programming. The investor consumes all financial wealth in the final period, which implies that we exactly know the utility derived from terminal wealth w_T. More specifically, the time-T value function is given by

$$J_T(w_T, X_T, y_T) = \frac{\varphi w_T^{1-\gamma}}{1-\gamma}. \tag{18}$$

For all other time periods, we have the following Bellman equation

$$J_t(w_t, X_t, y_t) = \max_{(c_t, x_t) \in \mathcal{K}_t} \left(\frac{c_t^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}_t \left(J_{t+1}(w_{t+1}, X_{t+1}, y_{t+1}) \right) \right). \tag{19}$$

We provide further details on the solution approach in Section 1.5.

[6] Davis, Kubler, and Willen (2003) and Cocco, Gomes, and Maenhout (2005) accommodate costly borrowing and allow the investor to default (endogenously) within their model.

1.3 Computing management fees

We are the first to study the impact of time-varying bond risk premia on the optimal strategies of life-cycle investors. However, this also raises the question how costly it is to ignore variation in bond risk premia. To examine this question, we take as a point of reference a strategy, called the benchmark strategy, in which the investor conditions its decisions on wealth and age only. This strategy, therefore, ignores the variation in the term structure of interest rates. Next, we grant the household access to one of two mutual funds. The first fund implements the so-called one-period conditional strategy. This strategy incorporates the information in the term structure at time t for the optimal investment and consumption choice, but uses the benchmark strategy as of the next period. Hence, this fund simplifies the computation of the optimal strategy by assuming that after the current period, the strategy depends on wealth and age only. It therefore ignores the possibility to time bond markets in subsequent periods. The second fund, which we call the life-cycle fund, implements the dynamically-optimal investment and consumption strategy.

With these three strategies in hand, we compare the impact on the household's utility. As a way to compare the different strategies, we compute the management fee that makes the household indifferent between the more sophisticated strategy and the strategy that conditions on wealth and age only. The proportional management fee, which is denoted by χ, enters the budget constraint (13) in the following way

$$w_{t+1} = (w_t - c_t)(1-\chi)\left(x_t'\left(r_{t+1} - \iota r_{t+1}^f\right) + r_{t+1}^f\right) + y_{t+1}. \tag{20}$$

It is important to note that the optimal consumption and portfolio policy may be affected by the presence of fees. We account for this in our numerical approach. To compute the maximum management fee, we equate the value function induced by the benchmark strategy[7] ($J_{\text{Benchmark}}(W, X, t)$) to the value function induced by, for instance, the one-period conditional strategy with a fee χ ($J_{\text{One-period}}(W, X, \chi, t)$)

$$J_{\text{Benchmark}}(W, X, t) = J_{\text{One-period}}(W, X, \chi, t). \tag{21}$$

We then solve for χ to compute the fee that makes the household indifferent between the one-period conditional and the unconditional strategy. We compute the management fee that the life-cycle fund can charge along the same lines.

We use management fees as a way to compare strategies or funds instead of losses in certainty-equivalent consumption, which is more conventional (Cocco, Gomes, and Maenhout (2005)). We argue that annual management fees are easier to interpret and can

[7] Note that despite the fact that the policies do not depend on the current state X_t, the induced value function does depend on the state variables.

for instance be compared to the fees charged by actively-managed mutual funds relative to passive funds. This would correspond to the case in which the household requires a fund manager or financial planner[8] to implement the more advanced strategies on their behalf. The difference between the fees of the active and passive fund would correspond to the fees we compute in the context of our model (see also Gruber (1996) for a similar interpretation of management fees).

Finally, we would like to stress that management fees and certainty-equivalent consumption losses are closely connected. In Appendix B, we derive the exact link between both measures in a simple continuous-time model. We will discuss some of the results from that analysis below to understand the management fees that we compute.

1.4 Estimation of the model

We now estimate our specification of the financial market introduced in Section 1.1. Section 1.4.1 describes the data that we use in estimation and we report in Section 1.4.2 the estimation results. In Section 1.4.3 we provide the individual-specific parameters of the individual's preferences and income process.

1.4.1 Data

We use monthly US data as of January 1959 to December 2005 to estimate our specification of the financial market. We use six yields in estimation with 3-month, 6-month, 1-year, 2-year, 5-year, and 10-year maturities, respectively. The monthly US government yield data are the same as in Duffee (2002) and Sangvinatsos and Wachter (2005) to December 1998. These data are taken from McCulloch and Kwon up to February 1991 and extended using the data in Bliss (1997) to December 1998. We extend the time series of 1-year, 2-year, 5-year, and 10-year yields to December 2005 using data from the Federal Reserve bank of New York. The data on the 3-month and 6-month yield are extended to December 2005 using data from the Federal Reserve Bank of St. Louis.[9] Data on the price index have been obtained from the Bureau of Labor Statistics. We use

[8] It is motivated by the observation that households may lack the knowledge to implement dynamic fixed-income strategies themselves. This limited sophistication is central to the debate on household finance (Campbell (2006), Huberman and Jiang (2006), and Calvet, Campbell, and Sodini (2007)), and we take it seriously. We therefore take the perspective of a household and compute the maximum fee it would be willing to pay to access a fund that implements the one-period conditional strategy or the life-cycle strategy.

[9] The yield data for the period January 1999 to December 2005 are available at `http://www.federalreserve.gov/pubs/feds/2006` and `http://research.stlouisfed.org/fred2`. The data from the Federal Reserve Bank of New York are available for the cross-section of long-term yields (1-year, 2-year, 5-year, and 10-year) as of August 1971. The correlation over the period August 1971 to December 1998 of these yields with the data used in Duffee (2002) equals 99.95%, 99.97%, 99.94%, and 99.85%, respectively. The 3-month and 6-month yields are available as of January 1982. The correlation of these data over the period January 1982 to December 1998 with the data used in Duffee (2002) equals 99.96% and 99.95% for 3-month and 6-month yields.

the CPI-U index to represent the relevant price index for the investor. The CPI-U index represents the buying habits of the residents of urban and metropolitan areas in the US.[10] We use returns on the CRSP value-weighted NYSE/Amex/Nasdaq index data for stock returns.

1.4.2 Estimation

We use the Kalman filter with unobserved state variables X_{1t} and X_{2t} to estimate the model by maximum likelihood. We assume that all yields have been measured with error in line with Brennan and Xia (2002) and Campbell and Viceira (2001). Details on the estimation procedure are in Appendix C.

The relevant processes in estimation are $K_t = (X_t', \log \Pi_t, \log S_t)'$ for which the joint dynamics can be written as

$$\mathrm{d}K_t = \left(\begin{bmatrix} 0_{2\times 1} \\ \delta_{0\pi} - \frac{1}{2}\sigma_\Pi'\sigma_\Pi \\ \delta_{0R} + \eta_S - \frac{1}{2}\sigma_S'\sigma_S \end{bmatrix} + \begin{bmatrix} -K_X & 0_{2\times 2} \\ \delta_{1\pi}' & 0_{1\times 2} \\ \delta_{1R}' & 0_{1\times 2} \end{bmatrix} K_t \right) \mathrm{d}t + \Sigma_K \mathrm{d}Z_t, \tag{22}$$

with $\Sigma_K = (\Sigma_X', \sigma_\Pi, \sigma_S)'$ and $\Sigma_X = \begin{bmatrix} I_{2\times 2} & 0_{2\times 2} \end{bmatrix}$. An unrestricted volatility matrix, Σ_K, would be statistically unidentified and we therefore impose the volatility matrix to be lower triangular.

The price of unexpected inflation risk cannot be identified on the basis of data on the nominal side of the economy alone. We impose that the part of the price of unexpected inflation risk that cannot be identified using nominal bond data equals zero. Since inflation-linked bonds have been launched in the US only as of 1997, the data available is insufficient to estimate this price of risk accurately. This restriction is in line with the recent literature, see for instance Ang and Bekaert (2005), Campbell and Viceira (2001), and Sangvinatsos and Wachter (2005).

Formally, these constraints on the prices of risk imply

$$\Lambda_t \quad = \quad \Lambda_0 + \Lambda_1 X_t = \begin{bmatrix} \Lambda_{0(1)} \\ \Lambda_{0(2)} \\ 0 \\ \star \end{bmatrix} + \begin{bmatrix} \Lambda_{1(1,1)} & \Lambda_{1(1,2)} \\ \Lambda_{1(2,1)} & \Lambda_{1(2,2)} \\ 0 & 0 \\ \star & \star \end{bmatrix} X_t, \tag{23}$$

where the '$\star$' in the last row indicate that these parameters are chosen to satisfy the restriction that the equity risk premium is constant (i.e., $\sigma_S'\Lambda_0 = \eta_S$ and $\sigma_S'\Lambda_1 = 0$).

We report the estimation results in Table 1. The parameters are expressed in annual

[10] See `http://www.bls.gov` for further details.

terms. The standard errors are computed using the outer product gradient estimator. The parameters σ_u ($u = 0.25, 0.5, 1, 2, 5, 10$) correspond to the volatility of the measurement errors of the bond yields at the six maturities used in estimation.

We briefly summarize the relevant aspects of our estimation results. First, we find that both the instantaneous short rate and expected inflation are increasing in both X_1 and X_2. Second, X_2 is estimated to be more persistent than X_1. The first-order autocorrelation implied by our estimates on an annual frequency equals 0.503 and 0.861, respectively. It corresponds to a half-life of 1 year for X_1 and approximately 5 years for X_2. Third, we find that stock and bond returns are negatively correlated with inflation innovations.

We now turn to the prices of risk and implied risk premia. The equity risk premium (η_S) is estimated to be 5.4%, which reflects the historical equity risk premium. We further find that the unconditional price of risk of X_1 is more negative than the price of risk of X_2, i.e., $\left|\Lambda_{0(1)}\right| > \left|\Lambda_{0(2)}\right|$. Table 2 reports the risk premia on nominal bonds along with their volatilities. We set the factors equal to their unconditional expectation. Nominal bond risk premia range from almost 60bp for a 1-year bond somewhat over 2% for a 10-year bond. The Sharpe ratio of 5-year nominal bonds is somewhat higher than for 10-year bonds (0.24 versus 0.18). The impact of the term structure factors on bond risk premia, i.e., the time variation in prices of risk, is governed by Λ_1. Figure 3 presents the 5-year nominal bond risk premia for a realistic range of X_1 and X_2. First, we find that risk premia are decreasing in X_1 and increasing in X_2. Second, we find that bond risk premia are more sensitive to shifts in X_2 than X_1, which is caused by the high persistence of X_2 discussed earlier.

Panel A of Table 3 presents the correlations between the assets that are possibly included in the asset menu, while Panel B of Table 3 reports the correlation between the risk premia on 5-year and 10-year nominal bonds and the same asset returns. These correlations are important as they drive the hedging demands formed by the investor to hedge against future changes in investment opportunities. Stock returns and nominal bond returns are positively correlated, consistent with Sangvinatsos and Wachter (2005). The correlations in Panel B indicate that long-term bond returns are negatively correlated with bond risk premia. It implies that a long position in these bonds can be used to hedge adverse changes in bond risk premia. After all, a decrease in the risk premium on long-term nominal bonds is likely to occur simultaneously with a positive return on these bonds. Consequently, the optimal allocation to long-term bonds of an unconstrained, long-term investor that is not endowed with a stream of labor income is positive and increasing in the investment horizon, see also Sangvinatsos and Wachter (2005).

Next, we compare the fit of the model and examine the average yield curve, the volatility of bond returns, and the predictability of bond returns. We summarize the results below and conclude that our tractable two-factor model provides a reasonably good fit to the data, in particular for the questions we are interested in. We follow Dai and Singleton (2002) and Sangvinatsos and Wachter (2005). That is, we simulate 5,000 sample paths of the same length as our sample. For each of the samples, we first compute the average yield and we then plot the average and the 95%-confidence interval of the sample means. We also compute the sample values using the data directly. The results are presented in Figure 4.

We find that all means are comfortably in the 95%-confidence interval. As such, the model does a good job reproducing the average yields. We repeat this exercise for the volatility of yields for the maturities we have in our sample. The results are presented in Figure 5. Also in this case, we find that the standard deviation of yields that we compute directly from the data lies well within the confidence bounds that we generate from our model. We conclude that the model is able to match the cross-sectional moments. As Sangvinatsos and Wachter (2005) remark, this is non-trivial because the model needs to fit a set of time-series and cross-sectional moments.

We analyze the bond return predictability implied by the data and our model. To this end, compute the Campbell and Shiller (1991) long-horizon regressions. We consider the following regression ($n > m$)

$$y_{t+m}^{n-m} - y_t^n = \beta_0 + \beta_1 \frac{m\left(y_t^n - y_t^m\right)}{n-m} + \varepsilon_{t+m}. \tag{24}$$

Under the expectations hypothesis $\beta_1 = 1$. Empirically, the regressions coefficients are negative and decreasing for longer maturities (Campbell and Shiller (1991) and Dai and Singleton (2002)). The predictive regression is estimated using overlapping monthly data, which we replicate within our model by simulating monthly data. Figure 6 shows that our model reproduces the empirically-observed pattern in the regression coefficients and the estimates based on our sample fall within the 95%-confidence bounds.

1.4.3 Individual-specific parameters

We now specify the parameters that govern the individual's preferences and labor income process. In our benchmark specification, we set the coefficient of relative risk aversion to $\gamma = 5$ and the subjective discount factor to $\beta = 0.96$. The investor consumes and invests from age 25 to age 65. The income process is calibrated to the model of Cocco, Gomes, and Maenhout (2005). In the benchmark specification, we focus on an individual with high school education, but without a college degree, that is, the "High School" individual in Cocco, Gomes, and Maenhout (2005). The variance of the transient shocks then equals $\sigma_u^2 = 0.0738$ and of the permanent shocks $\sigma_\epsilon^2 = 0.0106$. The function g_t, $t \in [25, 65]$, in

(17) is modeled by a third order polynomial in age

$$g_t = \alpha_0 + \alpha_1 t + \alpha_2 t^2/10 + \alpha_3 t^3/100, \tag{25}$$

and captures the hump-shaped pattern in labor income. The parameters are set to $\alpha_1 = 0.1682$, $\alpha_2 = -0.0323$, and $\alpha_3 = 0.0020$. The parameters α_1, α_2, and α_3 follow from Cocco, Gomes, and Maenhout (2005) and the constant is chosen so that the income level at age 25 equals \$20,000.

At retirement, we assume that all wealth in converted into an inflation-linked annuity that is priced on the basis of the unconditional expectation of the "real interest rate."[11] Koijen, Nijman, and Werker (2007b) study the asset allocation problem for an individual who allocates her retirement capital to various annuity products. They show that the hedging demands before retirement induced by this retirement choice are negligible. We therefore abstract from conversion risk caused by the annuitization decision. We simplify the retirement problem further by assuming that the individual dies with probability one at age 80. This allows us to determine φ in (10) as the utility derived from annuitizing retirement wealth.[12,13] We have an annual decision frequency in our model.

We further analyze in Section 3 the impact of individual-specific characteristics, like risk preferences, education level, correlation of human capital with asset returns. We also modify the asset menu of the investor. We assume in our benchmark specification that the individual has access to the stock index, 5-year nominal bonds, and cash. Section 3 studies alternative asset menus, in which the 5-year bond is replaced by a 10-year bond and a menu in which the investor can trade both a 3-year and a 10-year bond.

[11] The "real rate" is computed under the assumption that the price of unexpected inflation risk equals zero: $r^0 = \delta_{0R} - \delta_{0\pi}$.

[12] Specifically, denote the price of a real annuity at retirement by A_T. The time-T value function induced by annuitization is then given by

$$\sum_{t=T}^{80} \beta^{t-T} \frac{(W_T/A_T)^{1-\gamma}}{1-\gamma},$$

i.e.,

$$\varphi = \frac{A_T^{\gamma-1}}{1-\gamma} \sum_{t=T}^{80} \beta^{t-T}.$$

[13] We thus focus on the life-cycle investment and consumption problem in the pre-retirement period, consistent with for instance Benzoni, Collin-Dufresne, and Goldstein (2006). The investor's preference to save for retirement consumption is captured by the assumption that the investor derives utility from annuitized wealth up to age 80.

1.5 Solution technique

Life-cycle problems generally do not allow for analytical solutions, and we use numerical techniques instead. Numerical dynamic programming using Gaussian quadrature is the leading solution technique in life-cycle models, see for instance Cocco, Gomes, and Maenhout (2005). This approach becomes infeasible given our number of state variables and we therefore adopt the simulation-based approach developed recently by Brandt, Goyal, Santa-Clara, and Stroud (2005). Simulation-based techniques are well suited to deal with multiple exogenous state variables that can be simulated. However, life-cycle problems are usually characterized by (at least) one endogenous state variable, which is in our case financial wealth normalized by current income, which depends on previous choices. This variable cannot be simulated and we therefore need to specify a grid. However, instead of specifying a grid for (normalized) wealth before consumption, as is typically done, we construct a grid for wealth after consumption. By specifying the grid in this way, it is possible to compute the consumption policy analytically as has been shown by Carroll (2006).[14]

In addition to combining the endogenous grid method with simulation-based techniques, we also extend them in two important ways. First, we solve a problem with three risky assets, namely two long-term bonds and stocks, and a cash account. This requires us to develop a fast and accurate solution method with several risky assets. If we study the first-order condition that we use to solve for the optimal portfolio,[15] then it takes the following form

$$0 = \mathbb{E}_t\left(\beta e^{-\gamma u_{t+1}} c_{t+1}^{\star -\gamma}\left(r_{t+1} - r_{t+1}^f \iota\right)\right) + \lambda - \mu\iota. \tag{26}$$

In this first-order condition, λ denotes a vector of Lagrange multiplier corresponding to the short-sales constraints and μ is a Lagrange multiplier corresponding to the borrowing constraint. Because we work backwards in dynamic programming, we know the optimal consumption policy, $c_{t+1}^{\star}$, at time $t+1$ at all states. As in Brandt, Goyal, Santa-Clara, and Stroud (2005), we first simulate N paths for T periods, so that we have a cross-section of N paths at time t and $t+1$. We then construct $z_{t+1} = \beta e^{-\gamma u_{t+1}} c_{t+1}^{\star -\gamma}\left(r_{t+1} - r_{t+1}^f \iota\right)$ and regress it on a polynomial in the state variables at time t, $f(X_t)$. This leads to the approximation, in this case for asset s

$$\mathbb{E}_t\left(\beta e^{-\gamma u_{t+1}} c_{t+1}^{\star -\gamma}\left(r_{s,t+1} - r_{t+1}^f\right)\right) \simeq \theta_s(x, a_t)' f(X_t), \tag{27}$$

[14] Barillas and Fernández-Villaverde (2006) extend the endogenous grid approach of Carroll (2006) to solve problems with multiple endogenous state variables and illustrate the efficiency gains realized by this method.

[15] Koijen, Nijman, and Werker (2007a) contains further details on the derivation of these first-order conditions.

in which a_t denotes (normalized) wealth after consumption. We emphasize that the regression coefficients, $\theta_s(x, a_t)$, depend on the portfolio that the investor chooses, x. Without any further approximation, we would substitute (27) into (26) and solve for the optimal portfolio making sure that the Lagrange multipliers are non-negative. However, this would imply that for each state X_t and level of wealth after consumption a_t we need to compute the cross-sectional regressions in evaluating a new guess for the optimal portfolio. This is computationally infeasible. Our main technical insight is that the projection coefficients, $\theta_s(x, a_t)$, are smooth functions of the portfolio weights, x. We therefore consider a set of test portfolios on a rather coarse grid. For each of the test portfolios, we compute the projection coefficients and subsequently use the approximation

$$\theta_s(x, a_t) \simeq \Psi_s(a_t) g(x), \tag{28}$$

which implies

$$\mathbb{E}_t \left(\beta e^{-\gamma u_{t+1}} c_{t+1}^{\star -\gamma} \left(r_{s,t+1} - r_{t+1}^f \right) \right) \simeq g(x)' \Psi(a_t)' f(X_t). \tag{29}$$

The second-stage projection coefficients, $\Psi(a_t)$, are computed using the test portfolios. Using this approximation, we can compute the optimal portfolio relatively fast and dramatically reduce the number of regressions. Koijen, Nijman, and Werker (2007a) explains the method in more detail.

The second extension is the way in which we compute the optimal consumption policy. Brandt, Goyal, Santa-Clara, and Stroud (2005) also addresses the intermediate consumption problem, but the resulting optimal consumption strategy is not ensured to be strictly positive. In our model, the optimal consumption strategy is given by

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

in which r_{t+1}^p denotes the real return on the optimal portfolio. Clearly, it is important to ensure that the conditional expectation remains strictly positive. To this end, we approximate

$$\mathbb{E}_t \left(\beta e^{-\gamma u_{t+1}} c_{t+1}^{\star -\gamma} r_{t+1}^p \right) \simeq \exp(\bar{\theta}_0 + \bar{\theta}' f(X_t)). \tag{31}$$

Hence, we modify the simulation-based approach to ensure that the optimal consumption strategy remains strictly positive, thereby showing how simulation-based methods can be used to solve life-cycle problems with several exogenous states variables and multiple risky assets. An in-depth discussion of the numerical method is provided in the technical appendix Koijen, Nijman, and Werker (2007a).

2 Life-cycle investors and bond risk premia

We present the optimal policies for the benchmark specification concerning preference parameters, income process, and asset menu as discussed in Section 1.4.3. Section 2.1 analyzes the optimal life-cycle strategy. In Section 2.2, we compare the optimal life-cycle strategy to the one-period conditional strategy. In Section 2.3, we compute the management fees that would make the investor indifferent between a strategy that conditions on age and wealth only and either the one-period conditional or the life-cycle strategy.

2.1 Optimal life-cycle portfolio choice

Figure 7 presents the optimal average allocation over the life-cycle to stocks, 5-year nominal bonds, and cash. The vertical axis displays the average portfolio choice alongside the investor's age on the horizontal axis.

Between age 25 and 35, the individual optimally allocates all financial wealth, which is little to begin with,[16] to equity. In our benchmark specification, labor income can be viewed as a portfolio of bonds, perturbed with an idiosyncratic risk factor. This implies that the investor's total wealth, which is the sum of financial and human wealth, has a large exposure to interest rate risk. The investor, therefore, prefers to expose its financial wealth to other priced factors. Because the investor cares about the exposure of total wealth to the risk factors instead of financial wealth only, it is optimal to tilt the portfolio to equity. The individual starts to accumulate financial wealth between age 35 and 45. Nevertheless, the stock of human capital is sufficiently large for the individual to hold predominantly equity. The investor holds significant positions in long-term nominal bonds (that is, larger than 20% on average) only as of age 45, as human capital has depleted sufficiently by then to reduce its effect on the portfolio choice. Between age 50 and 55, the investor has a positive demand for stocks, long-term nominal bonds, and cash. The reduction in human capital is equivalent to an increase in the individual's effective risk aversion coefficient (Bodie, Merton, and Samuelson (1992)). Campbell and Viceira (2001) and Brennan and Xia (2002) show in addition that more conservative investors prefer to invest in inflation-linked bonds to hedge inflation risk. If inflation-linked bonds are not part of the asset menu, the investor allocates its capital to cash instead (Campbell and Viceira (2001)). Prior to retirement, the investor allocates on average 40% to stocks, 40% to 5-year nominal bonds, and the remaining 20% to cash.

[16] Cocco, Gomes, and Maenhout (2005) also find that the borrowing constraint binds during the first decade of the individual's life-cycle. An unconstrained life-cycle investor optimally capitalizes future labor income to increase today's consumption as a result of the hump-shaped pattern in labor income. This is, however, prohibited by the borrowing constraint and the investor consumes (almost) all income available.

We now illustrate how the optimal conditional allocation to the three assets responds to changes in bond risk premia. To this end, we present tilts in the optimal portfolio caused by variation in bond risk premia over the life-cycle. Figure 8 displays the optimal life-cycle allocation to stocks (Panel A), 5-year nominal bonds (Panel B), and cash (Panel C) for an empirically plausible range of either X_1 (left panels) or X_2 (right panels). These figures are constructed by first regressing the optimal asset allocations along all trajectories of the simulation-based method at a certain point in time on a second-order polynomial expansion (including cross-terms) in the prevailing state variables. The axes are different across figures for expository reasons.

Figure 8 shows that tilts in the optimal allocation to any of the three assets in response to changes in bond risk premia exhibit pronounced life-cycle patterns. Up to age 35, financial wealth is allocated almost exclusively to equity, regardless of the prevailing bond risk premia. Between age 35 and 45, the individual's allocation starts to tilt to long-term nominal bonds mostly if X_2 is high, which is the more persistent risk factor and has a bigger effect on risk premia (see Figure 3). Also, bond risk premia are increasing in X_2 and decreasing in X_1, explaining the opposite response in the optimal allocation to long-term bonds. As the investor reduces the equity allocation as of age 45, tilts in the optimal portfolio are large and can easily range from -20% to $+45\%$ in the allocation to long-term bonds in response to changes in X_2, and from -20% to $+20\%$ in case of X_1. Since the investor optimally holds no cash during this stage of the life-cycle, the equity allocation experiences exactly the opposite tilts once compared to the long-term bond allocation. The borrowing constraint prohibits the investor to borrow cash to take further advantage of high bond risk premia. The investor either has to reduce the equity allocation or forfeit high bond risk premia. We find that the optimal stock-bond mix is very sensitive to changes in X_2 during this period. Qualitatively, the results are the same for X_1 in this period, although in the opposite direction as bond risk premia decrease in X_1, but the quantitative impact is much smaller as suggested already by Figure 3. As of age 55, the optimal investment portfolio also contains cash positions due to the reduced stock of human capital. This impacts both the investor's willingness and ability to time bond risk premia. First, the investor acts more conservatively and timing risk premia adds less value as a result. Second, portfolio constraints are no longer binding, which may actually induce a larger value of timing risk premia. We find that the second effect dominates. Tilts in the bond allocation steadily increase as the investor ages. In addition, to tilt the optimal portfolio towards long-term bonds, the investor first reduces the cash allocation and only for high bond risk premia, the equity allocation as well. This implies that the opportunity costs of reducing the equity allocation exceed the costs induced by reducing the cash allocation. In sum, we find that tilts in the equity allocation in response to changes in risk premia are hump-shaped, but tilts in the cash allocation increase as the investor ages, consistent with the tilts in the allocation to long-term bonds.

Figure 9 summarizes the key insights of Figure 8 in a more compact way. We measure tilts in the optimal portfolio as the difference in the optimal allocation when X_1 and X_2 range from -1.65 unconditional standard deviations to plus 1.65 unconditional standard deviations around their unconditional means. More precisely, we indicate the difference in the allocation to asset i, for an investor of age t, caused by a change in X_1 by $\Delta^i_{X_1,t}$ and for changes in X_2 by $\Delta^i_{X_2,t}$, $i =$ stocks, bonds, or cash. All other state variables equal their unconditional expectation. For instance, to summarize the impact of X_1 on the allocation to stocks, we define

$$\Delta^{\text{stocks}}_{X_1,t} = x^{\text{stocks}}_t(X_{1,\max}) - x^{\text{stocks}}_t(X_{1,\min}), \tag{32}$$

where $X_{1,\max}$ equals the unconditional mean of X_1 plus 1.65 standard deviations and $X_{1,\min}$ is set to the unconditional mean of X_1 minus 1.65 standard deviations. Hence, a positive $\Delta^i_{X_1,t}$ implies an increase in the age-t allocation to asset i if X_1 increases.[17]

Figure 9 presents the asset allocation sensitivities, $\Delta^i_{j,t}$, for changes in either X_1 (left panel) or X_2 (right panel) over the life-cycle for all three assets. By construction, the sum of the differences at any moment of the life-cycle equals zero. In line with Figure 8, tilts in the optimal asset allocation induced by variation in bond risk premia exhibit pronounced life-cycle patterns. Especially during the second part of the life-cycle, these tilts are economically significant, in particular for the more persistent risk factor X_2. Tilts in the allocation to long-term bonds and cash are monotonically increasing over the life-cycle, whereas tilts in the equity allocation are hump-shaped over the life-cycle.

In addition to changes in bond risk premia, the optimal policies depend on the realizations of asset returns and labor income innovations to date. As argued before, it is the ratio of human capital to financial wealth that determines the investor's effective risk aversion and, as a result, portfolio choice. Figure 10 presents tilts in the optimal allocation to stocks, 5-year nominal bonds, and cash due to an increase in financial wealth from its 25%-quantile to its 75%-quantile.[18] In other words, we subtract the optimal asset allocation at the 25%-quantile from the optimal allocation at the 75%-quantile.

We find that the optimal allocation to equity is reduced for higher levels of financial wealth, while the allocation to long-term bonds and cash increases. Higher than average levels of financial wealth reduce the impact of human capital on total wealth and increase the individual's effective risk aversion. Therefore, the individual selects a more conservative strategy. An alternative interpretation is that a series of good returns essentially shifts the individual forward in the life-cycle. Figure 7 shows that the optimal

[17]Because all tilts are monotonically increasing in the state variables, this measure indeed summarizes the first-order effect of changes in the state variables presented in Figure 8.

[18]In contrast to both term structure variables, financial wealth is not stationary. To obtain comparable tilts over the life-cycle, we compare the difference in asset allocation at different quantiles of the financial wealth distribution at a particular age.

allocation to equity is decreasing in age, whereas the optimal allocations to long-term bonds and cash increase over the life-cycle. In contrast, lower than average levels of financial wealth level lead to a more prevalent role of human capital in the composition of total wealth. This resembles an individual who acts as in an earlier stage of her life-cycle. High asset returns or, equivalently, low income innovations tilt the optimal allocation away from stocks to long-term bonds and cash at any moment in the life-cycle. Negative asset returns, or positive labor income innovations, result in exactly the opposite effect. The effects are quantitatively most pronounced around age 50 and amount to at most 15%.

2.2 Comparing the life-cycle and the one-period conditional strategy

In this section, we compare the optimal life-cycle strategy to the one-period conditional strategy. The latter strategy assumes that the investor reverts to the strategy that conditions on age and wealth only as of the next period. The difference between both strategies can be interpreted as a form of hedging demands. When calculating the differences in the portfolio holdings, we average over both term structure variables. We cannot average over financial wealth, since both strategies will result in different levels of average wealth over the life-cycle. Figure 10 shows that this affects the investment strategy. We therefore condition on the average financial wealth level following from the optimal life-cycle effects to avoid any wealth effects. However, because we consider a life-cycle problem, one might argue that it is not as pure a measure of hedging demands as in asset-only investment problems, because the investor does account for the fact that she has a long horizon. The difference in allocations is the result of the investor's ability to respond to bond risk premia in subsequent periods. This definition is, in a life-cycle framework as we have, closest to the dynamic asset allocation literature, see Campbell, Chan, and Viceira (2003) and Sangvinatsos and Wachter (2005).

Figure 11 displays the resulting hedging demands. The optimal hedging demands turn out to be long in 5-year nominal bonds and are financed by a reduction in the allocation to stocks and, in particular, cash. Long-term bond returns are negatively correlated with future bond risk premia (Table 3, Panel B) so that a long position in bonds position pays off exactly in those states of the economy where bond risk premia are low. However, the hedging demands are strikingly small. The axes of Figure 11 range only from -3% to 4%. This result can be explained by considering how the investor actually finances the hedging demand. The investor has to reduce the equity allocation up to, say, age 50-55 to hold hedging demands that are long 5-year nominal bonds as a result of borrowing constraints. The opportunity costs induced by cutting back on the equity allocation turn out to be too high. The individual, therefore, reduces the hedging demands. As soon as the optimal investment strategy also includes a non-zero cash position, the investor uses this cash position to construct hedging demands. However, the individual is then already at age 55 and hedging adds little value for the remaining period. After all, Wachter (2002) shows in a model with a time-varying equity risk premium that the effective duration is substantially shortened for intermediate consumption problems as opposed to terminal wealth problems.

In sum, we find that tilts in the optimal portfolio in response to variation in bond risk premia exhibit pronounced life-cycle patterns, which are different for X_1 and X_2. We further show that hedging demands are long bonds, and short equity and cash, but quantitatively rather small. Since small hedging demands can in principle still lead to large utility costs once ignored, we compute in the next section the management fees that would make the investor indifferent between the various strategies.

2.3 Computing management fees

In this section, we compute the annual management fee an investor is willing to pay to follow either the one-period conditional strategy or the life-cycle strategy. We compute the fees over the life-cycle. That is, we ask for instance the question: "Suppose a 47-year old investor gains access to the one-period conditional strategy. What is the annual management fee that this investor would be willing to pay to switch from the strategy that conditions only on age and wealth to this one-period conditional strategy?"

The results are presented in the top panel of Figure 12. First, we find that the fees of both strategies amount to at most 43 basis points (bp) per annum. It implies that exploiting variation in bond risk premia may be within reach for households if they use low-cost mutual funds to implement their strategies. Second, if we compare the fees that the life-cycle fund and the fund implementing the one-period conditional strategy can charge, we find that they virtually coincide. In the context of our model, the optimal strategy virtually coincides with a strategy that assumes that as of next period, the investor reverts to the strategy conditioning on age and wealth only.

In the next sections, we analyze whether these results are robust to changes in the asset menu (Section 3) and different individual characteristics (Section 4).

3 Alternative and multiple long-term bonds

In this section, we analyze different asset menus. We consider the case with two long-term bonds with maturities equal to 3 years and 10 years, and compare the results to a model with just one long-term bond with a maturity equal to either 5 years or 10 years. We also considered the case in which the household can trade 5-year and 10-year bonds, but the returns on these two securities are even more highly correlated (97%) than 3-year and 10-year bond returns (90%). Also, this asset menu of long-term bonds has been used in Sangvinatsos and Wachter (2005). However, the exact maturity choice of long-term bonds is inconsequential, both qualitatively and quantitatively, for the point we are making. We focus on the fees of the one-period conditional and the life-cycle fund, which are the key statistics of our model. Table 4 and Figure 12 summarize the main result for the benchmark household we consider.

We find that the household prefers the 10-year bond over the 5-year bond to be included in the asset menu. The reason is that 10-year bonds have a higher exposure to X_2, which implies that it requires a smaller reduction in the allocation to stocks to obtain a particular expected inflation exposure. If we replace 5-year bonds by 10-year bonds, the management fees calculated in Section 2.3 increase to about 106bp. If we include both the 3-year and 10-year long-term bond in the investor's asset menu, we find that the management fees that the household is willing to pay increases slightly from 106bp to at most 108bp. Taken together, we conclude that the main conclusions of the paper are not affected by including multiple long-term bonds in the asset menu. Using low-cost mutual funds, households can benefit from variation in bond risk premia, primarily in later stages of the life-cycle. By comparing the fees on the one-period conditional and life-cycle strategy, we conclude that the additional gains of the life-cycle strategy are small relative to the one-period conditional strategy.

However, it is also interesting to look at the composition of the portfolio over the life-cycle. This answers how to optimally choose the maturity structure of the one-period conditional and life-cycle fund. We define two duration measures that summarize the properties of the fixed-income portfolio. Recall that we now have three fixed-income securities (cash, 3-year nominal bond, and the 10-year nominal bond). We can compute the overall duration of the portfolio by computing

$$Dur_{P,t} \equiv x_{Cash,t} + 3 \times x_{3Y,t} + 10 \times x_{10Y,t}, \tag{33}$$

and the duration of the fixed-income portfolio

$$Dur_{FI,t} \equiv \frac{x_{Cash,t} + 3 \times x_{3Y,t} + 10 \times x_{10Y,t}}{x_{Cash,t} + x_{3Y,t} + x_{10Y,t}}. \tag{34}$$

The portfolio duration is affected by the fraction invested in equity, whereas the latter focusses on the composition of the fixed-income portfolio. Both metrics are of interest. The duration measures are depicted in Figure 13, in which we average across states. We find that the portfolio duration is hump-shaped over the life-cycle, whereas the duration of the fixed-income portfolio is steadily declining. The hump-shaped pattern in the portfolio duration is the result of borrowing constraints for young households that gradually relax if the household ages. However, the value of human capital drops as well, which makes the household effectively more risk-averse, leading to the decline towards the end of the life-cycle.

4 Alternative individual characteristics

We also compute the optimal strategies for investors with different individual characteristics, like (i) risk preferences, (ii) education level, and (iii) different correlations between the permanent income shocks and stock returns. We briefly summarize the main effects, because none of our main results is materially affected. Detailed results for average portfolio holdings, tilts in allocations, hedging demands, and management fees are available upon request.

Risk preferences First, we solve our model for more aggressive, $\gamma = 3$, and more conservative, $\gamma = 7$, individuals. We find that the more aggressive individual ($\gamma = 3$) allocates on average more to equity than the benchmark individual. As a result, the borrowing constraint binds for a longer period and the individual's optimal investment strategy is less sensitive to variation in bond risk premia. In contrast, more conservative investors ($\gamma = 7$) shift earlier into long-term nominal bonds and cash. We find that the tilts in the optimal portfolio are, therefore, larger than for the benchmark investor. This is the result of two effects. First, the more conservative investor is not constrained and does not allocate all financial wealth to equity. As such, the opportunity costs amount to reducing the cash position which are smaller than reducing the equity position. Second, very conservative investors do not care about variation in risk premia and are only concerned with selection the portfolio to optimally smooth consumption over time. For our specification, the first effect still dominates for the conservative investor and tilts are more pronounced for the conservative individual.

The management fees of the one-period conditional and the life-cycle strategy coincide for $\gamma = 3$ and peak at 36bp. For $\gamma = 7$, the maximum fee of the one-period conditional strategy equals 49bp and for the life-cycle strategy 52bp.

Education level The education level of the benchmark individual we consider is "High School" according to the classification of Cocco, Gomes, and Maenhout (2005). We now consider different income processes as well that correspond to the "No High School" individual in Cocco, Gomes, and Maenhout (2005), for which we have $\alpha_1 = 0.1684$, $\alpha_2 = -0.0353$, and $\alpha_3 = 0.0023$ in (25), and the "College" individual that is characterized by $\alpha_1 = 0.3194$, $\alpha_2 = -0.0577$, and $\alpha_3 = 0.0033$ in (25). Higher education levels correspond to higher expected income growth. This implies that the borrowing constraints during early stages of the life-cycle bind for a longer period as the individual prefers to capitalize future income for the purpose of consumption smoothing. We find that tilts in the optimal portfolio are larger for individuals in the lowest education group, in particular up to age 45. The borrowing and short-sales constraints during the first two decades are less restrictive for individuals with low education levels and they can therefore benefit from variation in bond risk premia. Individuals with high education levels allocate most financial wealth to equity as their income pattern is very steep and thus more back-loaded. This effect is particularly pronounced up to age 45-50. From that age onwards, tilts in the optimal portfolios are very similar across all education groups. The management fees virtually coincide with the benchmark case we consider.

Correlation between income and financial market risks Finally, we examine the impact of correlation between labor income risk and asset returns. Specifically, we consider the case where permanent income innovations (u) are correlated with equity returns. This correlation likely depends on the individual's occupation, education level, age, and gender. Cocco, Gomes, and Maenhout (2005) estimate the correlation between permanent labor income shocks and stock returns between -1% and 2%. Heaton and Lucas (2000) report estimates between -7% and 14%. Munk and Sørensen (2005) provide an estimate for this correlation of 17%. Finally, Davis and Willen (2000) report estimates between -25% and 30% for the correlations between a broad equity index and labor income innovations. Regarding the correlation between labor income risk and industry-specific equity risk, the correlation ranges between -10% and 40% in their estimates, depending on the individual's education level, age, and gender. We follow Viceira (2001) instead and consider an individual whose labor permanent labor income innovations exhibit a correlation of 25% with the stock index returns.[19]

The positive correlation of permanent income shocks with equity returns has two effects, namely the substitution effect and the value effect. The substitution effect implies that non-tradable labor income substitutes for investments in particular assets in the investment portfolio. In the benchmark specification, human capital is essentially a position in inflation-indexed bonds, if we abstract from idiosyncratic risk. This tilts the portfolio towards equity and away from long-term nominal bonds. We now consider a case in which labor income innovations and equity returns are positively correlated. This implies that the investor will reduce the equity allocation (see also Viceira (2001)) and possibly allows for a larger impact of variation in bond risk premia. The value effect refers to the change in the value of human capital for different correlations of income

[19]Alternative labor income dynamics have been proposed in the literature. Benzoni, Collin-Dufresne, and Goldstein (2006) show that even though labor income and stock markets have a low correlation in the short run, this correlation increases at longer horizons. Benzoni, Collin-Dufresne, and Goldstein (2006) subsequently analyze a model in which labor income and stock market prices are co-integrated. Storesletten, Telmer, and Yaron (2004) and Lynch and Tan (2006) show that idiosyncratic labor income risk varies considerably over the business cycle. Lynch and Tan (2006) provide furthermore evidence that income growth is higher during economically good times.

innovations with financial risks that are prices. For instance, if labor income is perfectly correlated with equity returns, future income is effectively discounted at the expected real return on equity instead of the yield on an inflation-linked bond. This reduces the ratio of human capital to financial wealth and increases the effective risk aversion in turn. We indeed find larger tilts in the beginning of the life-cycle. The positive correlation of income innovations crowds out the equity investment and allows the investor to time bond markets instead. By the same token, the hedging demands increase slightly, which again reflects the more important role of bond market timing when labor income and stock returns are positively correlated. Management fees are at most 49bp (44bp) per annum for the life-cycle strategy (one-period conditional strategy) if the investor can trade stocks, 5-year long-term bonds, and cash. As such, our main conclusions are robust to such correlations between labor income risk and stock returns.

5 Conclusions

We solve a realistic life-cycle consumption and investment problem for an investor who has access to stocks, long-term nominal bonds, and cash. The investor has to comply with borrowing and short-sales constraints and receives a stream of stochastic labor income. Consistent with recent empirical evidence, we accommodate time variation in bond risk premia. We take an investor who conditions on age and wealth only as our point of reference. We subsequently offer this investor either a one-period conditional fund or a life-cycle fund. The former fund assumes that the investor reverts to the strategy that conditions on wealth and age only after one period. The life-cycle fund implements a fully dynamically optimal strategy. We compute the maximum annual (proportional) management fee a household would be willing to pay to have access to either the one-period conditional or the life-cycle fund. We find that there are economically significant gains to account for variation in bond risk premia. The second finding is that the fees of the one-period conditional and the life-cycle strategy virtually coincide.

In this paper, we find that the term structure of interest rates is driven by two factors that have different levels of persistence. We find bond risk premia to be more sensitive to the persistent factor. We show that tilts in the optimal asset allocation as a result of changes in the persistent factor are larger. Tilts in the optimal allocation exhibit pronounced life-cycle patterns. Tilts in the allocation to long-term nominal bonds and cash are monotonically increasing in age, but hump-shaped for equity.

We solve our model with multiple long-term bonds and compute the optimal duration of the overall portfolio and the fixed-income portfolio over the life-cycle. We find that the former is hump-shaped, while the latter is monotonically declining over the life-cycle. To derive these results, we extend recently developed simulation-based methods to solve for the optimal consumption and portfolio choice. This allows us to solve life-cycle problems with a large number of state variables and multiple assets as is the case in our model. We check robustness of our results to risk preferences, education level, and different correlations of income risk and asset returns.

This paper can be extended in various directions. First, we abstract from housing in our life-cycle problem. If the investor finances the house via a mortgage, there are essentially two types, namely fixed-rate mortgages (FRM) and adjustable-rate mortgages (ARM).[20] FRMs can be interpreted as short positions in long-term bonds and ARMs as short positions in cash. This additional flexibility may allow the investor to improve upon the exposures to the term structure factors. A second extension would be to allow for an endogenous retirement decision and analyze how this interacts with the prevailing investment opportunities at bond markets (see for instance Farhi and Panageas (2007)).

[20]Koijen, Van Hemert, and Van Nieuwerburgh (2008) study the choice between an ARM and an FRM in a simple two-period model.

References

Ang, A., and G. Bekaert (2005): "The Term Structure of Real Rates and Expected Inflation," Working Paper Columbia University.

Barberis, N. (2000): "Investing for the Long Run When Returns Are Predictable," *The Journal of Finance*, 55, 225–264.

Barillas, J., and J. Fernández-Villaverde (2006): "A Generalization of the Endogenous Grid Method," *Journal of Economic Dynamics and Control*, Forthcoming.

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein (2006): "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets are Cointegrated," *The Journal of Finance*, Forthcoming.

Bergstrom, A. R. (1984): "Continuous Time Stochastic Models and Issues of Aggregation over Time," in *Handbook of Econometrics*.

Bliss, R. R. (1997): "Testing Term Structure Estimation Methods," in *Vol. 9 of Advances in Futures and Options Research*, ed. by P. Boyle, G. Pennacchi, and P. Ritchken, pp. 197–231. JAI Press, Greenwich, Connecticut.

Bodie, Z., R. C. Merton, and W. F. Samuelson (1992): "Labor Supply Flexibility and Portfolio Choice in a Life-Cycle Model," *Journal of Economic Dynamics and Control*, 16, 427–449.

Brandt, M. W. (1999): "Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach," *The Journal of Finance*, 54, 1609–1645.

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," *The Review of Financial Studies*, 18, 831–873.

Brennan, M. J., and Y. Xia (2002): "Dynamic Asset Allocation under Inflation," *The Journal of Finance*, 57, 1201–1238.

Calvet, L., J. Y. Campbell, and P. Sodini (2007): "Down or Out: Assessing the Welfare Costs of Household Investment Mistakes," *Journal of Political Economy*, 115, 707–747.

Campbell, J. Y. (2006): "Household Finance," *The Journal of Finance*, 61, 1553–1604.

Campbell, J. Y., Y. L. Chan, and L. Viceira (2003): "A Multivariate Model of Strategic Asset Allocation," *Journal of Financial Economics*, 67, 41–80.

Campbell, J. Y., and R. J. Shiller (1991): "Yield Spreads and Interest Rate Movements: A Bird's Eye View," *Review of Economic Studies*, 58, 495–514.

Campbell, J. Y., and L. Viceira (1999): "Consumption and Portfolio Decisions When Expected Returns Are Time Varying," *Quarterly Journal of Economics*, 114, 433–496.

——— (2001): "Who Should Buy Long-Term Bonds?," *American Economic Review*, 91, 99–127.

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

Chan, Y. L., and L. M. Viceira (2000): "Asset Allocation with Endogenous Labor Income: The Case of Incomplete Markets," Working Paper, Hong Kong University of Science and Technology and Harvard Business School.

Chapman, D. A., and Z. Xu (2007): "Career Concerns and the Active Fund Managers Problem," Working Paper, Boston College.

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

Cochrane, J. H., and M. Piazzesi (2005): "Bond Risk Premia," *American Economic Review*, 95, 138–160.

——— (2006): "Decomposing the Yield Curve," Working Paper, University of Chicago.

Dai, Q., and K. J. Singleton (2000): "Specification Analysis of Affine Term Structure Models," *The Journal of Finance*, 50, 1943–1978.

——— (2002): "Expectation Puzzles, Time-varying Risk Premia, and Affine Models of the Term Structure," *Journal of Financial Economics*, 63, 415–441.

Davis, S. J., F. Kubler, and P. Willen (2003): "Borrowing Costs and the Demand for Equity over the Life Cycle," Working Paper.

Davis, S. J., and P. Willen (2000): "Using Financial Assets to Hedge Labor Income Risk: Estimating the Benefits," Working Paper.

de Jong, F. (2000): "Time-Series and Cross-Section Information in Affine Term Structure Models," *Journal of Business & Economic Statistics*, 18, 300–314.

Duffee, G. R. (2002): "Term Premia and Interest Rate Forecasts in Affine Models," *The Journal of Finance*, 57, 405–443.

Duffie, D., and R. Kan (1996): "A Yield Factor Model of Interest Rates," *Mathematical Finance*, 6, 379–406.

Farhi, E., and S. Panageas (2007): "Saving and Investing for Early Retirement: A Theoretical Analysis," *Journal of Financial Economics*, 83(1), 87–122.

Gomes, F., and A. Michaelides (2005): "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," *The Journal of Finance*, 60, 869–904.

Gourinchas, P.-O., and J. A. Parker (2002): "Consumption over the Life-Cycle," *Econometrica*, 70, 47–89.

Gruber, M. (1996): "Another Puzzle: The Growth in Actively Managed Mutual Funds," *Journal of Finance*, 51, 783–810.

Harvey, A. C. (1989): *Forecasting, Structural Time Series Models and the Kalman Filter.* Cambridge University Press.

Heaton, J., and D. Lucas (1997): "Market Frictions, Savings Bahavior, and Portfolio Choice," *Macroeconomic Dynamics*, 1, 76–101.

——— (2000): "Portfolio Choice in the Presence of Background Risk," *The Economic Journal*, 110, 1–26.

Huberman, G., and W. Jiang (2006): "Offering vs. Choice in 401(k) Plans: Equity Exposure and Number of Funds," *Journal of Finance*, 61, 763–801.

Jurek, J., and L. M. Viceira (2007): "Optimal Value and Growth Tilts in Long-Horizon Portfolios," Working Paper, Harvard University.

Koijen, R. S. J., T. E. Nijman, and B. J. M. Werker (2007a): "Appendix Describing the Numerical Methods Used in 'When Can Life-cycle Investors Benefit from Time-varying Bond Risk Premia?'," Working Paper, Tilburg University and University of Chicago.

——— (2007b): "Optimal Portfolio Choice with Annuitization," Working Paper, Tilburg University and University of Chicago.

Koijen, R. S. J., O. Van Hemert, and S. Van Nieuwerburgh (2008): "Mortgage Timing," Forthcoming *Journal of Financial Economics*.

Lynch, A., and S. Tan (2006): "Labor Income Dynamics at Business-cycle Frequencies: Implications for Portfolio Choice," Working Paper, NYU Stern School of Business and Fordham University.

Munk, C., and C. Sørensen (2005): "Dynamic Asset Allocation with Stochastic Income and Interest Rates," Working Paper.

Sangvinatsos, A., and J. A. Wachter (2005): "Does the Failure of the Expectations Hypothesis Matter for Long-Term Investors?," *The Journal of Finance*, 60, 179–230.

Storesletten, K., C. Telmer, and A. Yaron (2004): "Cyclical Dynamics in Idiosyncratic Labor-market Risk," *Journal of Political Economy*, 112, 695–717.

Van Hemert, O. (2006): "Life-Cycle Housing and Portfolio Choice with Bond Markets," Working Paper, NYU Stern School of Business.

Viceira, L. M. (2001): "Optimal Portfolio Choice for Long-Horizon Investors with Non-Tradable Labor Income," *The Journal of Finance*, 56, 433–470.

Wachter, J. A. (2002): "Portfolio and Consumption Decisions Under Mean-Reverting Returns: An Exact Solution for Complete Markets," *Journal of Financial and Quantitative Analysis*, 37, 63–92.

——— (2003): "Risk Aversion and Allocation to Long-term Bonds," *Journal of Economic Theory*, 112, 325–333.

A Pricing nominal bonds

We derive prices of nominal bonds in the financial market described in Section 1, following the results on affine term structure models in, for instance, Duffie and Kan (1996) and Sangvinatsos and Wachter (2005).

To that extent, we assume that bond prices are smooth functions of time and the term structure factors X. Denote the price of a nominal bond at time t that matures at time T by $P(X_t, t, T)$. Since nominal bonds are traded assets, we must have that $\phi_t^{\$} P(X_t, t, T)$ is a martingale, where $\phi^{\$}$ is given in (6). This implies

$$-P_X K_X X + P_t + \frac{1}{2} tr\left(P_{XX}\right) - RP - P_X' \Lambda = 0, \tag{A.1}$$

where the subscripts of P denote partial derivatives with respect to the different arguments. Subsequently, Duffie and Kan (1996) have shown that in this case we obtain nominal bond prices that are exponentially affine in the state variables

$$P(X, t, t+\tau) = \exp\left(A(\tau) + B(\tau)' X\right). \tag{A.2}$$

Substituting this expression in (A.1) and matching the coefficients on the constant and the state variables X, we obtain the following set of ordinary differential equations

$$A'(\tau) = -B(\tau)'\Lambda_0 + \frac{1}{2} B(\tau)' B(\tau) - \delta_{0R}, \tag{A.3}$$

$$B'(\tau) = -\left(K_X' + \Lambda_1'\right) B(\tau) - \delta_{1R}, \tag{A.4}$$

where ι_2 denotes a two dimensional vector of ones. We also have the boundary conditions

$$A(0) = 0, \;\; B(0) = 0. \tag{A.5}$$

The ODEs can be solved in closed form, see for instance Dai and Singleton (2002). This leads to

$$B(\tau) = \left(K_X' + \Lambda_1'\right)^{-1} \left[\exp\left(-\left(K_X' + \Lambda_1'\right)\tau\right) - I_{2\times 2}\right] \delta_{1R}, \tag{A.6}$$

$$A(\tau) = \int_0^\tau A'(s) ds, \tag{A.7}$$

where $I_{2\times 2}$ denotes the two by two identity matrix.

B Linking welfare metrics

In this appendix we discuss the relation between two welfare metrics: the loss in certainty-equivalent consumption and management fees. The former metric has been used widely in the literature to quantify utility losses. For the purpose of our paper, however, management fees are more natural. We illustrate in a simple life-cycle model the link between certainty-equivalent consumption losses and management fees.

The model We consider a continuous-time model in which the household receives labor income at a constant rate Y. Labor income is assumed to be deterministic. The household can only transfer wealth intertemporally by investing in a cash account that earns a riskless interest rate r. We make these simplifying assumptions just to make the link as transparent as possible. The household's preferences are given by

$$J_t = \max_{(C_s)_{s\in[t,T]}} \int_t^T e^{-\beta s} \frac{C_s^{1-\gamma}}{1-\gamma},$$

and the optimization is subject to the budget constraint

$$\mathrm{d}W_t = (r - \chi) W_t \mathrm{d}t + \left(Y - C_t\right) \mathrm{d}t, \tag{B.1}$$

in which χ denotes the proportional management fee. Note that in this simple model, management fees can be interpreted as a reduction in the risk-free rate.

We compare the optimal strategy to a sub-optimal strategy that we can leave unspecified. This alternative strategy induces a value function $\tilde{J}_t$ that is lower than J_t. We subsequently use this value function to link the loss in certainty-equivalent consumption to the management fees that we report in this paper.

Management fees We first derive the optimal strategy for a general χ. We subsequently choose χ to ensure that $J_t(\chi)$ equals $\tilde{J}_t$. The value function J follows from the Bellman equation

$$\beta J = \max_{C_t} \left(\frac{C_t^{1-\gamma}}{1-\gamma} + J_W \left[(r-\chi)W_t + (Y - C_t) \right] + J_t \right),$$

which leads implies that the optimal rate of consumption equals: $C_t = J_W^{-\frac{1}{\gamma}}$. The value function takes the form

$$J_t(\chi) = \frac{(W_t + H_t(\chi))^{1-\gamma}}{1-\gamma} g_t(\chi)^{\gamma},$$

with

$$\begin{aligned} g_t(\chi) &= \frac{1}{A(\chi)} \left(1 - e^{-A(\chi)(T-t)} \right), \\ H_t(\chi) &= \int_t^T e^{-(r-\chi)(s-t)} Y \mathrm{d}s, \end{aligned}$$

in which H_t denotes the household's human capital. To solve for A, we substitute the optimal policy in the Bellman equation

$$\begin{aligned} \beta J &= \frac{J_W^{1-\frac{1}{\gamma}}}{1-\gamma} + J_W \left((r-\chi)W_t + \left(Y_t - J_W^{-\frac{1}{\gamma}} \right) \right) + J_W \frac{\partial H_t}{\partial t} + \frac{J\gamma}{g_t} \frac{\partial g_t}{\partial t} \\ &= \frac{\gamma J_W^{1-\frac{1}{\gamma}}}{1-\gamma} + J_W (r-\chi) (W_t + H_t) + J\gamma \left(A - \frac{1}{g_t} \right), \end{aligned}$$

in which

$$\frac{\partial H_t}{\partial t} = (r-\chi)H_t - Y_t, \; \frac{\partial g_t}{\partial t} = A g_t - 1, \text{ and } J_W = \frac{J(1-\gamma)}{W_t + H_t}.$$

This implies

$$\beta = \frac{\gamma}{g_t} + (1-\gamma)(r-\chi) + \gamma \left(A - \frac{1}{g_t} \right),$$

which leads to the expression for $A(\chi) = (\beta - (1-\gamma)(r-\chi))/\gamma$. The management fee solves

$$J_t(\chi) = \tilde{J}_t. \tag{B.2}$$

Certainty-equivalent consumption loss In this case, we set $\chi = 0$ and compute the certainty-equivalent consumption as the fraction of consumption the household is willing to sacrifice

$$\max_{(C_s)_{s\in[t,T]}} \int_t^T e^{-\beta s} \frac{((1+\theta) C_t)^{1-\gamma}}{1-\gamma} = (1+\theta)^{1-\gamma} J_t(0) = \tilde{J}_t.$$

The certainty-equivalent consumption loss is now given by

$$\theta = \left(\frac{\tilde{J}_t}{J_t(0)} \right)^{\frac{1}{1-\gamma}} - 1, \tag{B.3}$$

as in Cocco, Gomes, and Maenhout (2005).

Linking both metrics Equations (B.2) and (B.3) show that certainty-equivalent consumption and management fees are one-to-one related via

$$(1+\theta)^{1-\gamma} J_t(0) = J_t(\chi),$$

which implies

$$\begin{aligned}\theta &= \left(\frac{J_t(\chi)}{J_t(0)}\right)^{\frac{1}{1-\gamma}} - 1 \\ &= \frac{W_t + H_t(\chi)}{W_t + H_t(0)} \left(\frac{g_t(\chi)}{g_t(0)}\right)^{\frac{\gamma}{1-\gamma}} - 1.\end{aligned}$$

Numerical example We consider a simple example with the following parameter values

$$Y_t = 1, r = 8\%, \beta = 0.04, T = 60, \gamma = 5,$$

which are consistent with the values in our benchmark case in the full model.[21] We consider a policy (which we do not have to specify for the purpose of this analysis) that implies that the household is willing to pay a constant management fee of 50bp over the life-cycle to switch to the optimal strategy. We subsequently compute the corresponding pattern of certainty-equivalent consumption losses and depict it in Figure 2.

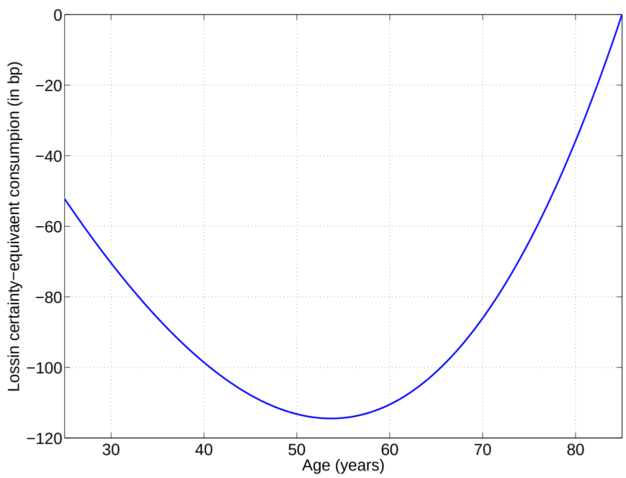

Figure 2: Linking welfare metrics
We compute the loss in certainty-equivalent consumption over the life-cycle if there is a strategy that induces a management fee of 50bp to make the household indifferent between that particular strategy and the optimal strategy.

There are two key insights. First, losses in certainty-equivalent consumption tend to be larger than management fees. Second, even if management fees are constant over the life-cycle, the loss in certainty-equivalent consumption is hump-shaped over the life-cycle. We verified whether this property carries over to the full model and we conclude that this is indeed the case.[22]

C Estimation procedure

Our estimation procedure in closely related to Sangvinatsos and Wachter (2005). The main difference is that we allow all yields to be measured with error, following de Jong (2000), Brennan and Xia (2002), and Campbell and Viceira (2001). However, we assume that the measurement errors are independent, both sequentially and cross-sectionally. The continuous time equations underlying the financial market

[21]We select a relatively high interest rate to proxy for the better investment opportunities that the investor has in the full model in which it can invest in stocks and long-term bonds.

[22]These results are available upon request.

in Section 1 can be written as

$$
\begin{aligned}
\mathrm{d}K_t &= \left(\begin{bmatrix} 0_{2\times 1} \\ \delta_{0\pi} - \frac{1}{2}\sigma_\Pi'\sigma_\Pi \\ \delta_{0R} + \eta_S - \frac{1}{2}\sigma_S'\sigma_S \end{bmatrix} + \begin{bmatrix} -K_X & 0_{2\times 2} \\ \delta_{1\pi}' & 0_{1\times 2} \\ \delta_{1R}' & 0_{1\times 2} \end{bmatrix} K_t \right) \mathrm{d}t + \Sigma_K \mathrm{d}Z_t, \\
&= \left(\Theta_0 + \Theta_1 K_t\right) \mathrm{d}t + \Sigma_K \mathrm{d}Z_t. \qquad \text{(C.1)}
\end{aligned}
$$

with

$$
K_t = \begin{bmatrix} X_t \\ \log \Pi_t \\ \log S_t \end{bmatrix}. \qquad \text{(C.2)}
$$

As K_t follow a standard multivariate Ornstein-Uhlenbeck process, we may write the exact discretization (see, e.g., Bergstrom (1984) and Sangvinatsos and Wachter (2005))

$$
K_{t+h} = \mu^{(h)} + \Gamma^{(h)} K_t + \varepsilon_{t+h}, \qquad \text{(C.3)}
$$

where $\varepsilon_{t+h} \overset{i.i.d.}{\sim} N\left(0_{4\times 1}, \Sigma^{(h)}\right)$ for appropriate $\mu^{(h)}$, $\Gamma^{(h)}$, and $\Sigma^{(h)}$ which we derive below. To derive the discrete time parameters, we consider the eigenvalue decomposition[23] $\Theta_1 = UDU^{-1}$. The parameters in the VAR(1) - model relate to the structural parameters via

$$
\begin{aligned}
\Gamma^{(h)} &= \exp\left(\Theta_1 h\right) = U \exp\left(Dh\right) U^{-1}; \qquad \text{(C.4)} \\
\mu^{(h)} &= \left[\int_t^{t+h} \exp\left(\Theta_1 \left[t+h-s\right]\right) ds \right] \Theta_0 \\
&= UFU^{-1}\Theta_0,
\end{aligned}
$$

where F is a diagonal matrix with elements $F_{ii} = h\alpha\left(D_{ii}h\right)$, with

$$
\alpha(x) = \frac{\exp(x) - 1}{x},
$$

and $\alpha(0) = 1$. The derivation of $\Sigma^{(h)}$ is a bit more involved. We have

$$
\begin{aligned}
\Sigma^{(h)} &= \int_t^{t+h} \exp\left(\Theta_1 \left[t+h-s\right]\right) \Sigma_K \Sigma_K' \exp\left(\Theta_1 \left[t+h-s\right]\right) \mathrm{d}s \qquad \text{(C.5)} \\
&= UVU',
\end{aligned}
$$

where V is a matrix with elements

$$
\begin{aligned}
V_{ij} &= \left[\int_t^{t+h} \exp\left(D\left[t+h-s\right]\right) U^{-1} \Sigma_K \Sigma_K' \left(U^{-1}\right)' \exp\left(D\left[t+h-s\right]\right) ds \right]_{ij} \qquad \text{(C.6)} \\
&= \left[U^{-1} \Sigma_K \Sigma_K' \left(U^{-1}\right)' \right]_{ij} \int_t^{t+h} \exp\left(\left[D_{ii} + D_{jj}\right]\left[t+h-s\right]\right) \mathrm{d}s \\
&= \left[U^{-1} \Sigma_K \Sigma_K' \left(U^{-1}\right)' \right]_{ij} h\alpha\left(\left[D_{ii} + D_{jj}\right] h\right).
\end{aligned}
$$

Using data on six yields, stock returns, and inflation, we estimate the model using the Kalman filter. The transition equation is given by (C.3). We assume that all yields are measured with measurement error, in line with de Jong (2000), Brennan and Xia (2002), and Campbell and Viceira (2001). On the other hand, Duffee (2002) and Sangvinatsos and Wachter (2005) select certain maturities and fit

[23]Note that, since K_X is a diagonal matrix, the eigenvalues of Θ_1 are given by κ_1, κ_2, and 0 (with multiplicity two). Recall that a square matrix is diagonalizable if and only if the dimension of the eigenspace of every eigenvalue equals the multiplicity of the eigenvalue. This condition is satisfied for Θ_1.

these exactly, which is tantamount to identifying the factors. In line with these papers,[24] we assume the measurement to be Gaussian and independent of the innovations in the transition equation. The likelihood can subsequently be constructed using the error-prediction decomposition, see for instance Harvey (1989).

[24]Notably, de Jong (2000) allows for cross-sectional correlation between the measurement errors.

D Tables and figures

Parameter	Estimate	Standard error	Parameter	Estimate	Standard error
Expected inflation: $\pi_t = \delta_{0\pi} + \delta'_{1\pi} X_t$					
$\delta_{0\pi}$	4.20%	1.03%			
$\delta_{1\pi(1)}$	1.69%	0.19%	$\delta_{1\pi(2)}$	0.50%	0.24%
Nominal interest rate: $R_t = \delta_{0R} + \delta'_{1R} X_t$					
δ_{0R}	5.89%	1.66%			
$\delta_{1R(1)}$	1.92%	0.14%	$\delta_{1R(2)}$	1.03%	0.26%
Dynamics term structure factors: $\mathrm{d}X_t = -K_X X_t \mathrm{d}t + \Sigma_X \mathrm{d}Z_t$					
$K_{X(1,1)}$	0.687	0.177	$K_{X(2,1)}$	-0.350	0.119
$K_{X(2,2)}$	0.172	0.035			
Dynamics inflation: $\mathrm{d}\Pi_t/\Pi_t = \pi_t \mathrm{d}t + \sigma'_\Pi \mathrm{d}Z_t$					
$\sigma_{\Pi(1)}$	0.02%	0.05%	$\sigma_{\Pi(3)}$	0.98%	0.03%
$\sigma_{\Pi(2)}$	0.11%	0.04%			
Dynamics equity index: $\mathrm{d}S_t/S_t = (R_t + \eta_S)\mathrm{d}t + \sigma'_S \mathrm{d}Z_t$					
η_S	5.38%	2.48%			
$\sigma_{S(1)}$	-1.98%	0.57%	$\sigma_{S(3)}$	-1.74%	0.68%
$\sigma_{S(2)}$	-1.79%	0.69%	$\sigma_{S(4)}$	14.82%	0.32%
Prices of risk: $\Lambda_t = \Lambda_0 + \Lambda_1 X_t$					
$\Lambda_{0(1)}$	-0.293	0.127	$\Lambda_{1(1,1)}$	-0.103	0.182
$\Lambda_{0(2)}$	-0.158	0.071	$\Lambda_{1(2,2)}$	-0.168	0.038
			$\Lambda_{1(1,2)}$	-0.101	0.078
			$\Lambda_{1(2,1)}$	0.504	0.132
Volatility measurement error: σ_u $(u = 0.25, 0.5, 1, 2, 5, 10)$					
$\sigma_{0.25}$	0.47%		σ_2	0.11%	
$\sigma_{0.5}$	0.22%		σ_5	0.00%	
σ_1	0.00%		σ_{10}	0.22%	

Table 1: Estimation results for the financial market model
The financial market model is estimated by Maximum Likelihood using monthly US data on six bond yields, inflation, and stock returns over the period from January 1959 up to December 2005. The bond maturities used in estimation are 3-month, 6-month, 1-year, 2-year, 5-year, and 10-year. All yields are assumed to be measured with error. The CPI-U index is used to represent the relevant price index for the investor. Stock returns are based on the CRSP value-weighted NYSE/Amex/Nasdaq index. The parameters are expressed in annual terms. The standard errors are determined using the outer product gradient estimator.

Maturities	1-year	5-year	10-year
Risk premia	0.58%	1.55%	2.06%
Volatilities	1.77%	6.36%	11.75%

Table 2: Risk premia and volatilities

The table presents the (instantaneous) risk premia and return volatility on 1-year, 5-year, and 10-year bonds using the estimation results in Table 1. The term structure factors equal their unconditional expectation ($X_t = 0_{2\times 1}$). The risk premia and volatilities are expressed in annual terms.

Panel A: Correlation of asset returns			
	Stock return	5-year nom. bond return	10-year nom. bond return
Stock return	1		
5-year nom. bond return	0.153	1	
10-year nom. bond return	0.125	0.964	1
Panel B: Correlation of asset returns with risk premia			
	Stock return	5-year nom. bond return	10-year nom. bond return
Risk premium on 5-year nom. bond	0.073	-0.081	-0.343
Risk premium on 10-year nom. bond	0.086	-0.010	-0.275

Table 3: Correlations between asset returns and risk premia

Panel A presents the correlations between stock returns, 5-year nominal bonds, and 10-year nominal bonds using the estimates reported in Table 1. Panel B depicts the correlation between the same asset returns and either 5-year or 10-year nominal bond risk.

		Household's age						
		30	35	40	45	50	55	60
5-year bonds only	One-period conditional strategy	28	30	31	35	38	38	42
	Life-cycle strategy	31	33	35	39	41	41	43
10-year bonds	One-period conditional strategy	83	86	88	91	93	93	93
	Life-cycle strategy	99	101	103	106	105	102	97
Both 3-year and 10-year bonds	One-period conditional strategy	85	88	89	93	95	96	97
	Life-cycle strategy	101	103	105	108	108	106	101

Table 4: Management fees for different long-term bonds available to the household

The table reports the management fees, expressed in basis points per annum, of the one-period conditional and the life-cycle funds for different asset menus. The different columns are computed by varying the household's age at which we equate the value functions. The first two rows correspond to an asset menu that includes stocks, 5-year nominal bonds, and a cash account (that is, 1-year bond). For the third and fourth row, we replace the 5-year bond with the 10-year bond. The last two rows correspond to an asset menu that includes stocks, both 3-year and 10-year nominal bonds, and cash.

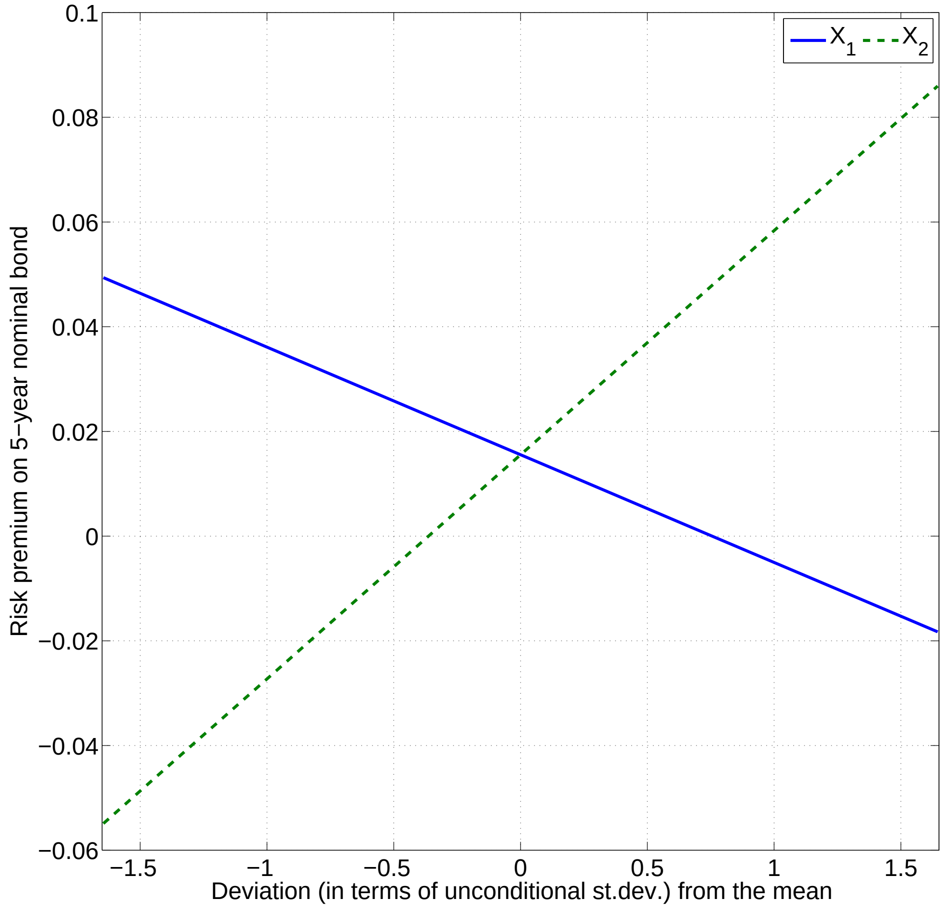

Figure 3: Time-variation in risk premia

The figure presents the 5-year nominal bond risk premium for different values of the term structure variables X_1 and X_2. The horizontal axis depicts the value of the term structure variables expressed in unconditional standard deviations around their means. All risk premia are expressed in annual terms.

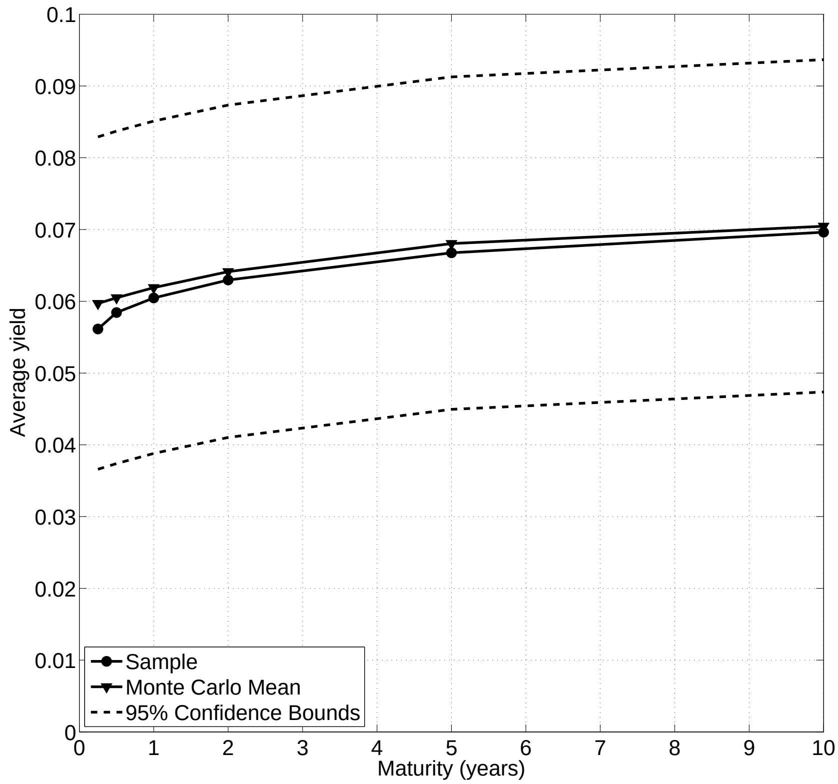

Figure 4: Average yields in the model and in the data

We simulate 5,000 sample paths of the same length as our sample. For each of the samples, we calculate the average yield and we then plot this average and a 95%-confidence interval. We also compute the sample averages using the data directly. The dots (triangles) correspond to the maturities that we use in estimation.

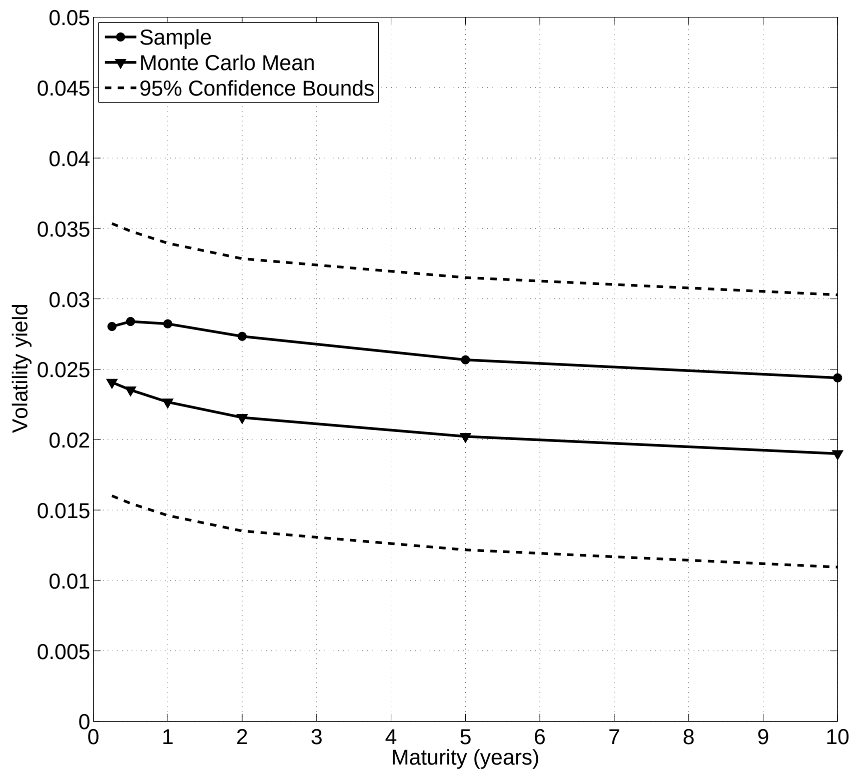

Figure 5: Volatility of yields in the model and in the data

We simulate 5,000 sample paths of the same length as our sample. For each of the samples, we calculate the volatility of yields and we then plot the average and a 95%-confidence interval. We also compute the sample volatility of yields using the data directly. The dots (triangles, boxes) correspond to the maturities that we use in estimation.

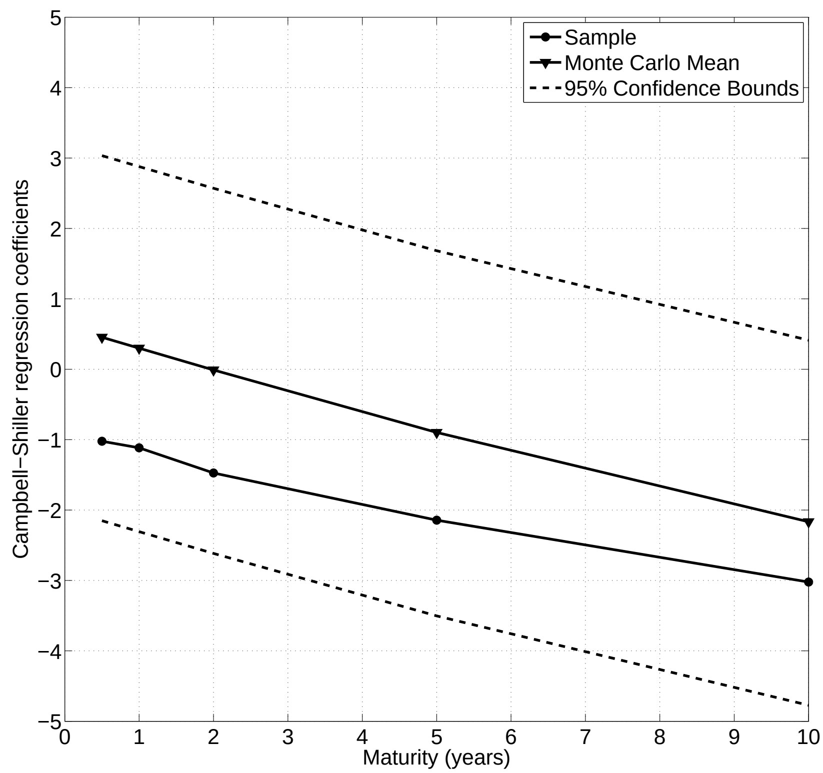

Figure 6: Return predictability in the model and in the data
We simulate 5,000 sample paths of the same length as our sample and run the predictive regression as in (24). For each of the samples, we compute the regression coefficient and we then plot the average and a 95%-confidence interval. The dots (triangles, boxes) correspond to the maturities that we use in estimation.

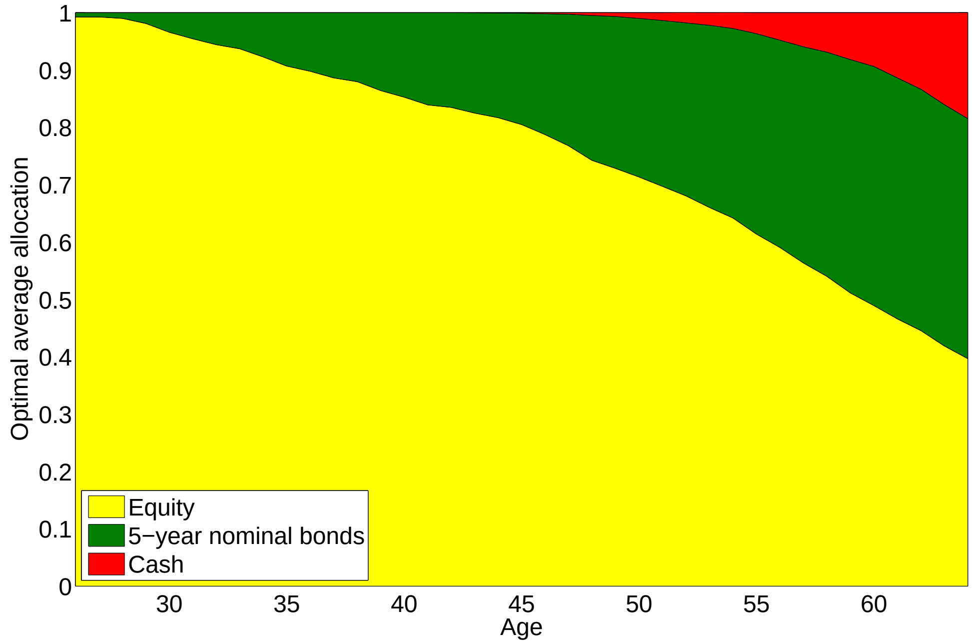

Figure 7: Optimal average asset allocation to stocks, 5-year nominal bonds, and cash
Optimal asset allocation over the individual's life-cycle, averaged over all state variables. The individual allocates capital to stocks, 5-year nominal bonds, and cash. The individual's coefficient of relative risk aversion equals $\gamma = 5$ and the time preference parameter $\beta = 0.96$. The vertical axis displays the average allocation and the horizontal axis indicates the individual's age.

Panel A: Optimal conditional equity allocation

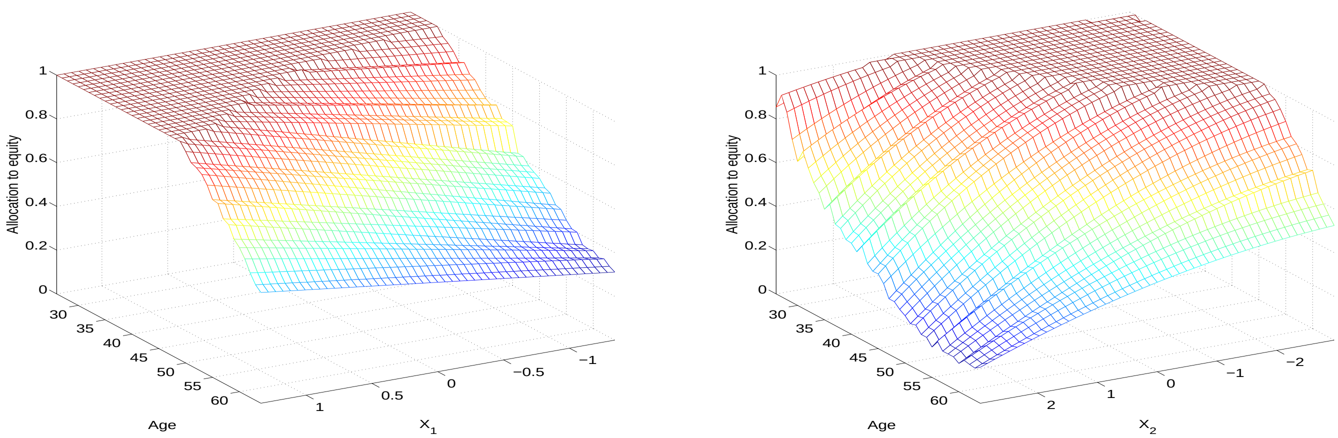

Panel B: Optimal conditional 5-year nominal bond allocation

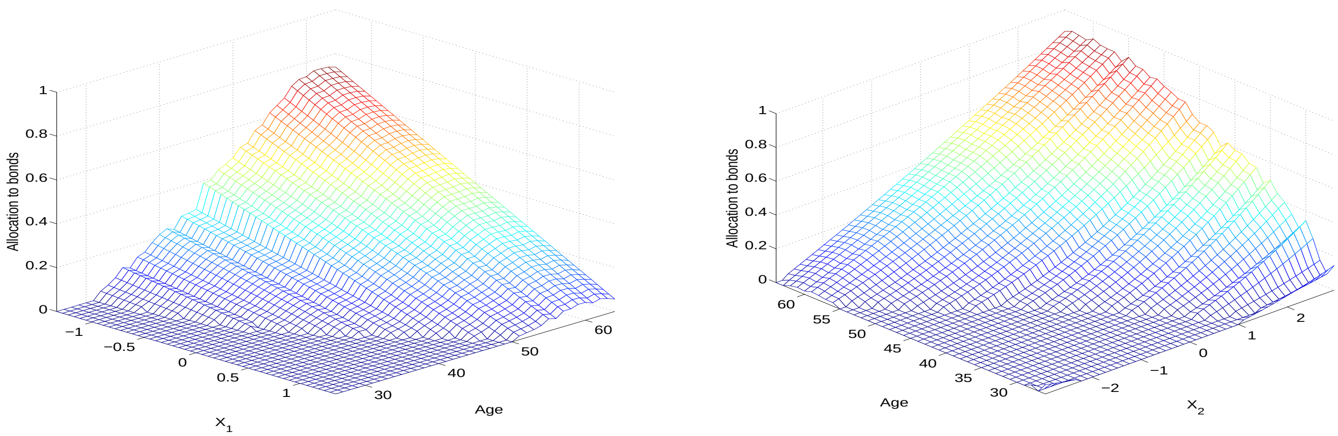

Panel C: Optimal conditional cash allocation

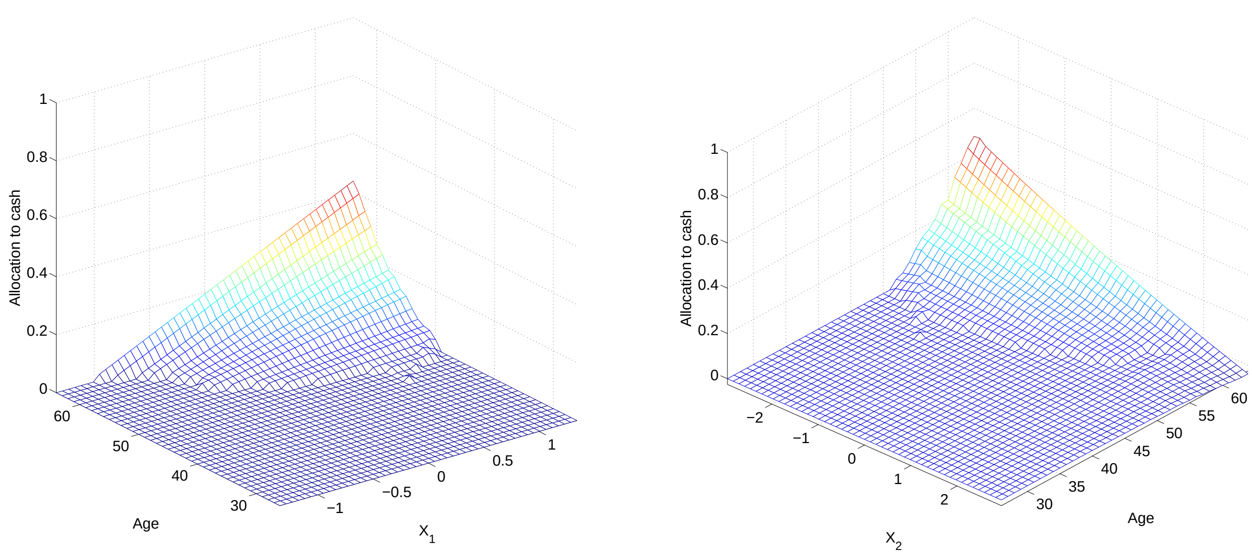

Figure 8: Optimal conditional asset allocation to stocks, 5-year nominal bonds, and cash

Optimal conditional asset allocation over the individual's life-cycle for an empirically plausible range of either X_1 (left panels) or X_2 (right panels). The individual allocates financial wealth to stocks (Panel A), 5-year nominal bonds (Panel B), and cash (Panel C). The individual's coefficient of relative risk aversion equals $\gamma = 5$ and the time preference parameter $\beta = 0.96$. The vertical axes displays the conditional allocation and the horizontal axes the individual's age and either X_1 or X_2. The axes are different across figures for expository reasons.

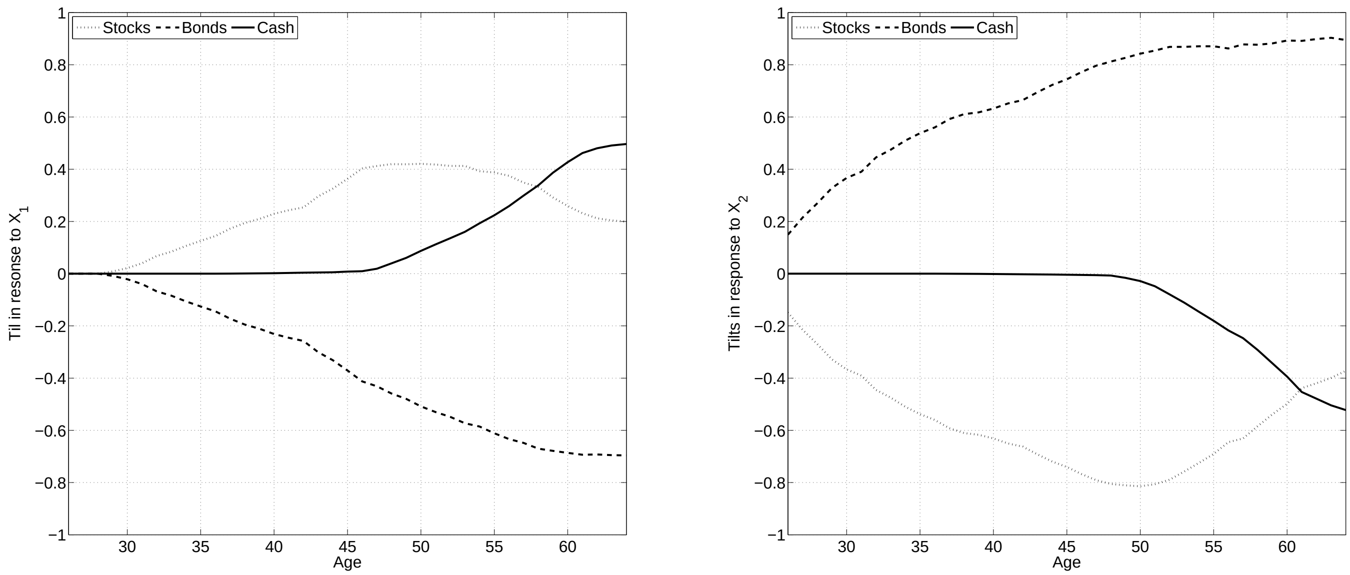

Figure 9: Tilts in the optimal asset allocation induced by bond risk premia
Tilts in the optimal allocation to stocks, 5-year nominal bonds, and cash in response to changes in either X_1 (left panel) or X_2 (right panel). We measure tilts in the optimal portfolio as the difference in the optimal allocation when X_1 and X_2 range from -1.65 unconditional standard deviations to plus 1.65 unconditional standard deviations around their unconditional means. A positive tilt implies an increase in the age-t allocation to a particular asset if the state variable increases. The individual's coefficient of relative risk aversion equals $\gamma = 5$ and the time preference parameter $\beta = 0.96$. The vertical axes display the difference in the optimal allocation and the horizontal axes the individual's age.

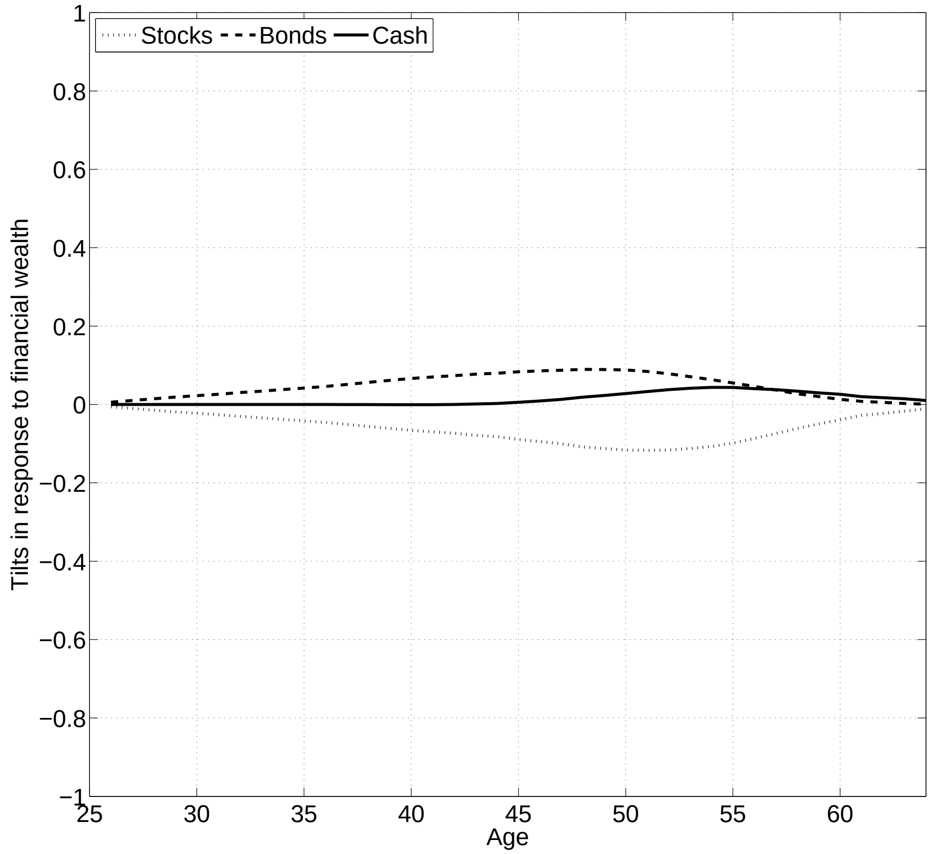

Figure 10: Tilts in the asset allocation induced by financial wealth

Tilts in the optimal allocation to stocks, 5-year nominal bonds, and cash in response to changes in financial wealth. We measure tilts in the optimal portfolio as the difference in the optimal allocation when financial wealth ranges from 25% to 75% of the distribution at a particular age. A positive tilt implies an increase in the age-t allocation to a particular asset if financial wealth increases. The individual's coefficient of relative risk aversion equals $\gamma = 5$ and the time preference parameter $\beta = 0.96$. The vertical axis displays the difference in the optimal allocation and the horizontal axes the individual's age.

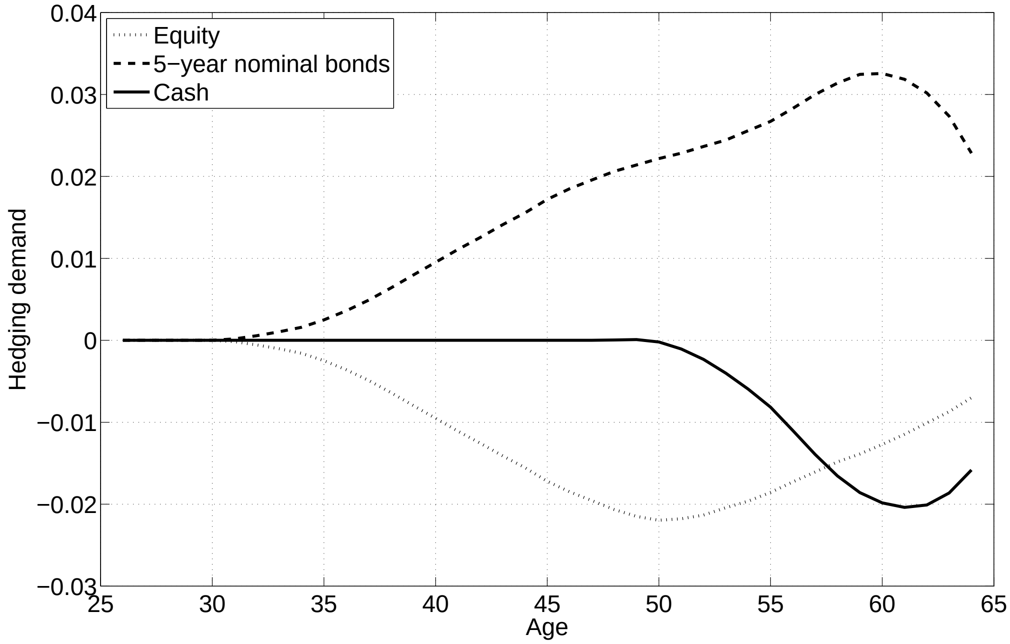

Figure 11: Hedging demands induced by time-varying bond risk premia over the life-cycle

Hedging demands using stocks, 5-year nominal bonds, and cash induced by time-varying bond risk premia. The hedging demands are calculated at each point in the life cycle by comparing the optimal strategies of the life-cycle fund and the one-period conditional fund. The main text provides further details. We refer in particular to Section 2.2 for the definition and interpretation of hedging demands in the context of our model. The individual's coefficient of relative risk aversion equals $\gamma = 5$ and the time preference parameter $\beta = 0.96$. The vertical axis displays the hedging demands and the horizontal axis the individual's age.

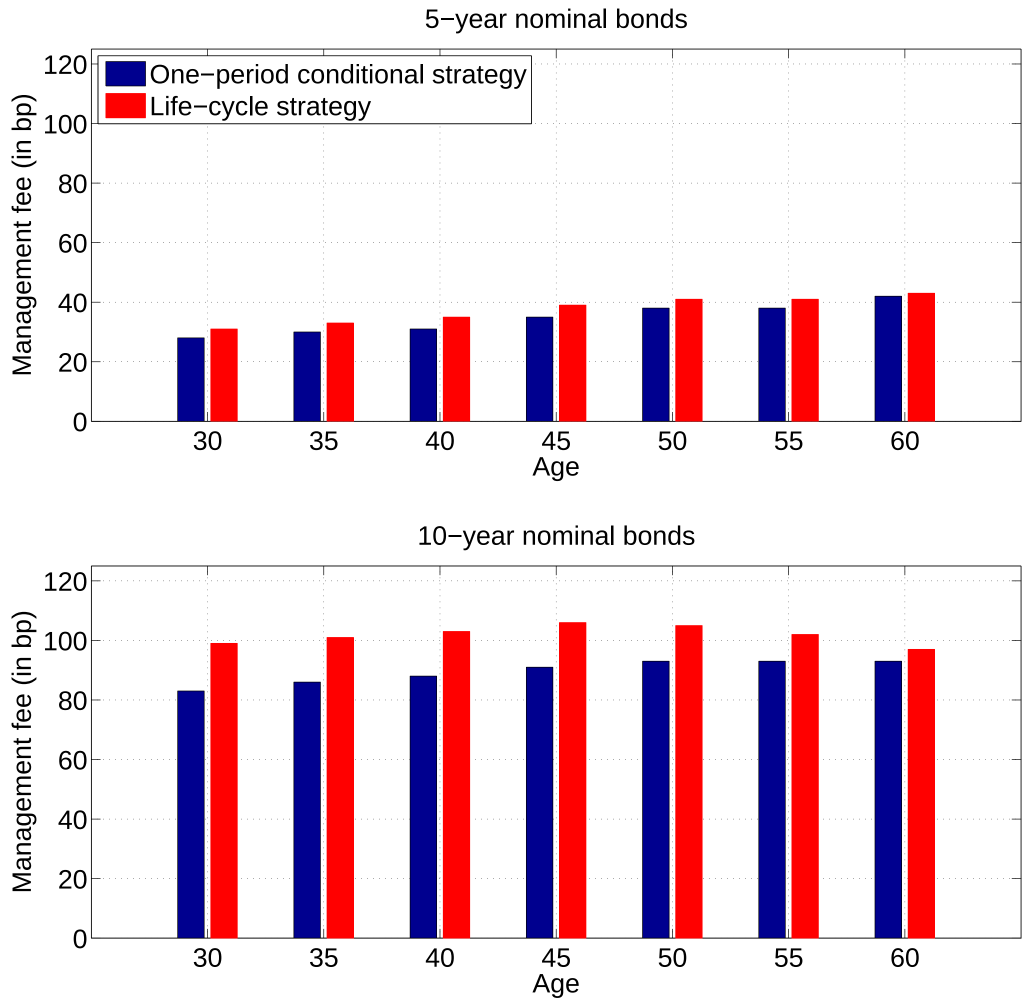

Figure 12: Maximum management fees expressed in basis points over the life-cycle
The figure depicts the maximum management fee (expressed in basis points per annum) of the one-period conditional fund and the life-cycle fund. The fee is computed by equating the value function induced by the strategy that conditions only on age and wealth to the value function induced by either the one-period conditional strategy or the life-cycle strategy and a particular fee. The top panel displays the results for an asset menu including stocks, 5-year nominal bonds, and cash. The bottom panel is for an asset menu in which the 5-year bond is replaced by a 10-year bond.

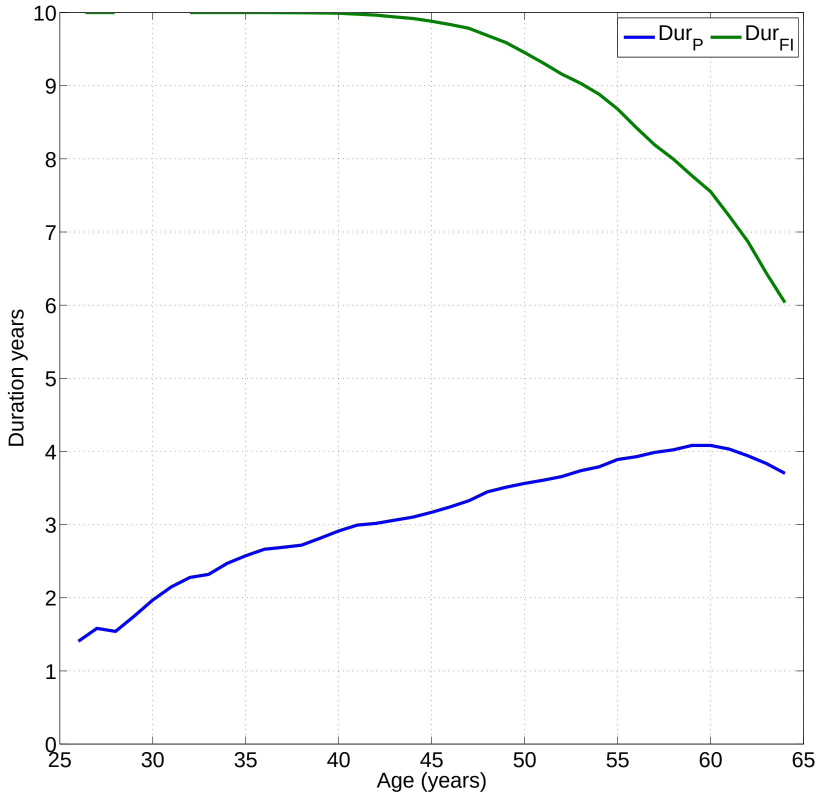

Figure 13: Duration of the overall and the fixed-income portfolio over the life-cycle
The blue (bottom) line depicts the optimal duration expressed in years of the overall portfolio over the life-cycle as defined in (33). The green (top) line corresponds to the optimal duration of the fixed-income portfolio over the life-cycle that is defined in (34).