Network for Studies on Pensions, Aging and Retirement

Netspar THESES

Maarten Meeuwis

The Impact of Stochastic Interest Rates in a Defined Contribution Pension Scheme

MSc Thesis 2013-019

The Impact of Stochastic Interest Rates in a Defined Contribution Pension Scheme

by
Maarten Meeuwis (798022)
BSc Tilburg University 2012

A thesis submitted in partial fulfillment of the requirements for the degree of
Master of Science in Quantitative Finance and Actuarial Science

Tilburg School of Economics and Management
Tilburg University

Supervisors:
Prof. dr. B. Melenberg (Tilburg University)
Dr. T.D. Markwat (Robeco)
Drs. R.D.J. Molenaar (Robeco)
Dr. ir. P. Beekhuizen (Robeco)

Second reader:
Prof. dr. R.J. Mahieu (Tilburg University)

July 29, 2013

Abstract

We find optimal allocation paths for a life-cycle investor with stochastic labor income and predetermined pension contributions, who faces borrowing and short-selling constraints and converts his wealth into an annuity at retirement. We specify an affine term structure model, allowing for stochastic interest rates, stochastic inflation, and bond return predictability. We calibrate the model parameters on the basis of data for a European investor. We develop a method based on numerical simulations to solve the allocation problem by means of dynamic programming. Compared to a market with constant interest rates, we find that the investor has additional hedging demands for bonds. When bond returns are predictable, the investor should time bond markets. These timing effects tend to dominate life-cycle effects.

KEYWORDS: LIFE-CYCLE INVESTING, HUMAN CAPITAL, AFFINE TERM STRUCTURE MODELS, TIME-VARYING BOND RISK PREMIA

Acknowledgments

This Master thesis was written during an internship at Robeco. The views expressed in this thesis are those of the author and do not necessarily reflect those of Robeco.

Writing this thesis would definitely not have been possible without the support I received from several people. First of all, I would like to thank my supervisor at Tilburg University, prof. dr. B. Melenberg. His feedback, suggestions, and warnings were very helpful in writing this thesis.

Second, I would like to thank Robeco for coming up with an interesting research question and providing me with the opportunity of writing my Master thesis at Robeco's Quantitative Strategies department. Special thanks go to my committed supervisors at Robeco, dr. T.D. Markwat, drs. R.D.J. Molenaar, and dr. ir. P. Beekhuizen. Their comments and assistance were very valuable.

Finally, I would like to thank my family and friends for their support during the process of writing my Master thesis.

Contents

1 Introduction

We will focus on asset allocation decisions in a defined contribution (DC) pension scheme. A DC pension scheme is a form of second pillar retirement provision. In contrast to defined benefit (DB) pension plans, only the joint contribution of the participant and the employer is predetermined, and there are no future benefits specified in advance. All investment risk is fully borne by the participants. Because of the limited risk for the employer and the ongoing changes in accounting rules and pension regulations, DC plans have recently been increasing in popularity.

The allocation of accumulated pension wealth to the assets that are available on the market is one of the most important financial planning decisions that people need to make. Due to the long investment horizon and the possibility to rebalance the long-term portfolio frequently, it can be very difficult to solve the problem of choosing the optimal allocation. At each point in time, the investor needs to take into account all possible future scenarios and the corresponding future (optimal) allocations. In general, this problem can only be solved by using complex dynamic programming techniques.

There are, however, some situations in which closed-form solutions to the long-term investment problem are available. These solutions rest upon the assumption that interest rates are constant over time. In this study, we relax the unrealistic assumption of a constant risk-free rate and determine the implications on the optimal asset allocations in a DC pension scheme. When interest rates are stochastic rather than fixed, not only stock returns, but also bond returns are stochastic. In that case, the participant in the DC scheme faces interest rate risk. The optimal exposure to financial assets now changes over time. Due to the long investment horizon, small changes in interest rates can have a huge impact on the investor's total wealth at the end of the investment horizon. Hedging interest rate risk will therefore play an important role in the asset allocation decision.

Under some very simplifying assumptions, the optimal allocation for a long-term investor is constant over time and is the same as the optimal allocation for a short-term investor, the so-called myopic solution. One of these assumptions is that there is a risk-free rate which is constant over time. Although this assumption is highly unrealistic, the resulting solution is still widely used in the pension industry. One of the reasons is the difficulty of solving a more general asset allocation problem.

The fact that the optimal allocation for a long-term investor is in this case constant over time does not imply that optimal portfolio allocations of accumulated pension savings in DC plans are constant over time. The reason is that labor income also plays an important role for long-term investors. Working individuals will have future contributions to the pension plan, and therefore own not only financial assets but also human capital. Human capital is defined as the present value of future labor income. Via this human capital, agents have already implicit investments in risky and risk-free assets. Since it is in this case optimal to allocate total wealth on the basis of the myopic portfolio, the allocation of financial wealth should be chosen such that the allocation of total wealth is equal to the myopic portfolio.

The evolution of human and financial wealth over time determines the optimal allocation of financial wealth to stocks and bonds, see Bodie, Merton, and Samuelson (1992) and

Campbell and Viceira (2002). The safer the salary payout from human capital is, the more it behaves like a bond. Thus, if human capital behaves mostly like a bond, a young investor has already a huge implicit investment in bonds via his human capital. Therefore, he should optimally invest all his financial wealth in risky assets. On the other hand, if human capital is stock-like, a young investor should invest in the risk-free asset because he has already invested in the risky asset via his human capital (see Benzoni, Collin-Dufresne, and Goldstein (2007)). Note that Benzoni et al. (2007) do not find that labor income is stock-like for young investors because of a high contemporaneous correlation between stock returns and labor income shocks, but because of cointegration between labor income and dividends.

This thesis relates to the aforementioned papers in the life-cycle literature by focusing on including human capital in the asset allocation decision. We solve the life-cycle optimization problem for an agent that works from the age of 25 to the age of 65 and saves a part of his income for retirement according to a predetermined scheme. For instance, in the Netherlands, people are required to save up to a certain percentage of their income for retirement without having to pay taxes over these savings. Since this age-dependent maximum is determined such that the participant will under regular market conditions end up with the desired pension benefits, actual retirement savings in a DC plan are almost always equal to the maximum tax-free savings. Hence, future savings as a percentage of labor income are in reality often predetermined and not part of the life-cycle optimization problem. We therefore deviate from the usual optimization problem in the life-cycle literature that takes into account consumption smoothing (see e.g. Bodie, Detemple, Otruba, and Walter (2004)).

We adopt a model specification of labor income that is based on the representation that is assumed by Cocco, Gomes, and Maenhout (2005) and others. Labor income is stochastic, and we allow for correlation between labor income shocks and stock returns. The part of labor income risk that can be explained by risk factors that are traded on the market is called systematic risk, the rest of labor income risk is idiosyncratic. The exposure of labor income shocks to stock returns determines the systematic labor income risk and plays a very important role in the asset allocation decision. Idiosyncratic risk cannot be diversified away and will therefore not lead to significantly different allocations.

In the life-cycle optimization problem, we include annuitization at retirement. We consider individuals that want to maximize real pension benefits, i.e., real annuity payments. Because annuities insure individuals against mortality risk, the benefits of annuities are considerable. Yaari (1965) shows that if annuities are fairly priced and agents have no bequest motives, it is optimal to invest all wealth in annuities. In the Netherlands, it is in fact obligatory to convert all accumulated second pillar pension savings into (nominal) annuities at retirement. When interest rates are stochastic, this means that individuals face considerable conversion risk, because the price of the annuity depends heavily on long-term yields. As a result, life-cycle investors that buy annuities at retirement will choose a different allocation at the years before retirement than investors that want to maximize their wealth at retirement.

As said, it is not straightforward to solve the life-cycle optimization problem of a general long-term investor. Solving this problem involves the computation of several conditional expectations. Because of time-varying investment opportunities and allocation constraints, the conditional expectations depend on the specific moment in time and on the current state of the economy. There are no analytical solutions available to solve for the optimal allocation. Therefore, we implement a modified version of the advanced simulation-based approach that is proposed by Koijen, Nijman, and Werker (2010). In this method, conditional expectations are approximated by performing regressions. The big advantage of this method is that it is very generally applicable. Its use is not constrained to specific distributional assumptions, utility functions, or a maximum number of state variables or financial assets. Since the ratio between human capital and financial wealth is endogenous, we make a grid for this variable.

Our research is most closely related to the paper of Koijen et al. (2010). They include a similar type of labor income in the asset allocation decision, and also allow for stochastic interest rates, stochastic inflation rates and time-varying bond risk premia. However, our research differs from their study in a number of aspects. First of all, as said, we exclude consumption decisions from the life-cycle optimization problem, which is not done by most studies such as Koijen et al. (2010). Further, we use the actual price of an annuity at retirement, leading to additional hedging demands, while Koijen et al. (2010) do not take into account conversion risk. In contrast to Koijen et al. (2010), we include real yields in the dataset and do not impose the restriction that the price of orthogonal inflation shocks is equal to zero. Finally, we value human capital explicitly, which has the advantage that we do not need to make an endogenous grid for the wealth level and can use a constant grid for the ratio between human capital and financial wealth instead.

In this thesis, we make the transition from a model with a constant risk-free rate and constant investment opportunities to a model with stochastic interest rates and time-varying investment opportunities. We do this in two steps. As a first step, we adopt a three-factor affine term structure model. In this model, bond yields are no longer constant, but instead depend on the state of the world. This means that investment returns fluctuate over time as a result of changing interest rates. However, we assume that investment opportunities in terms of excess returns of the risky assets on the market are constant, by imposing risk premia to be constant. In this way, we find the optimal allocation to be relatively insensitive to the state of the economy.

The absence of pronounced state effects allows us to compare the general allocation path over time to the “simple” allocation when interest rates and investment opportunities are assumed to be fixed. We find that both paths have similar shapes. The advice of financial planners turns out to be still valid for individuals with bond-like human capital when interest rates are stochastic: invest financial wealth fully in stocks at the beginning of the life-cycle, and reduce equity holdings over time as the fraction of financial wealth in total wealth grows. On the other hand, individuals with stock-like human capital should allocate a significantly larger part of their financial wealth to bonds when they are young. As compared to the basic model with constant interest rates, we find that the life-cycle investor replaces part of his stock holdings by bond holdings at each point in time. The reason for this is that he wants to be hedged against future decreases in investment returns caused by changing short rates. Due to the negative relation between bond returns and short rates, long-term bonds can be used to hedge the risk of disappointing future investment returns.

Although the three-factor model provides us with useful insights in the way how allocations should change when interest rates are stochastic rather than fixed, this restricted affine term structure model with constant risk premia is not able to replicate the properties of historically observed yields. Further, it does not allow for predictability in bond returns, which contradicts empirical observations (see for instance Sangvinatsos and Wachter (2005)). Therefore, as a second step, we adopt a more general four-factor model, that includes two yield factors (following e.g. Ang, Piazzesi, and Wei (2006)), and that allows for bond predictability. We find that this model provides a much better fit to historically observed yields.

Consistent with Sangvinatsos and Wachter (2005) and Koijen et al. (2010), we find that the life-cycle investor should exploit bond predictability by timing bond markets. The optimal allocation is therefore highly state-dependent. Under regular financial market conditions, we find a similar pattern of the optimal allocations over time as for the three-factor model without predictability. Individuals with bond-like human capital should allocate their financial wealth fully to stocks when they are young, and should quickly replace their stock holdings with bond holdings as they grow older. In comparison to the three-factor model, we find that under the same conditions investors now have an even higher demand for bonds. This is caused by the fact that bonds cannot only be used to hedge risk in the short rate, but also the risk in changing future bond risk premia. However, the optimal allocation is highly dependent on the state of the economy. When bond risk premia are relatively high, investors should react by increasing their allocation to bonds considerably, and vice versa. We find that state effects even dominate the life-cycle effects.

Our most general model of the financial market allows for stochastic interest rates, stochastic inflation, and time-varying bond risk premia. This model is related to the models of Campbell and Viceira (2001) and Brennan and Xia (2002). These studies show that there is a hedging demand for long-term bonds as a result of stochastic interest rates. However, they assume that risk premia are constant and they do not include labor income in the allocation decision. Sangvinatsos and Wachter (2005) include time-varying risk premia, but also do not take into account future labor income of a long-term investor. They find that bonds can be used to hedge time variation in bond risk premia. Munk and Sørensen (2010) include stochastic labor income and stochastic interest rates, but do not allow for stochastic inflation and time-varying bond risk premia.

We start with a review of the theory of long-term investing under the assumption of constant interest rates in Section 2. In Section 3, we will describe the general class of affine term structure models and define the two specific models that we will use. We will estimate the corresponding model parameters in Section 4. In Section 5, we will specify the life-cycle optimization problem that we will solve. Section 6 describes the optimization procedure that we use to solve this life-cycle optimization problem. The results for our three-factor model are presented in Section 7, the results for the four-factor model are discussed in Section 8. Section 9 concludes.

2 Myopic Portfolio Choice

In this section, we review the theory of long-term investing under the assumption of constant interest rates, and explain the situations for which long-term investors make the same choices as short-term investors. Subsequently, we will introduce the asset allocation problem of a so-called life-cycle investor.

2.1 Short-term investors

Before we consider the problem of a long-term investor, we first discuss short-term investing. This is called myopic portfolio choice. Afterwards, we will extend the analysis to long-term portfolio choice. We will find that under some conditions, the optimal portfolio for long-term investors is the same as the myopic portfolio that we will discuss here.

Suppose that a short-term investor has two assets available at time t: one risk-free asset with return $R_{f,t}$ in the period from time t to time $t+1$, and one risky asset with return R_{t+1} in the same period. Let $\mathbb{E}_t$ be the expectation operator that conditions on all information that is observed and known at time t. The return on the risky asset from period t to period $t+1$ has a conditional expectation of $\mathbb{E}_t R_{t+1}$ and a conditional variance of $\sigma_t^2 \equiv \mathrm{Var}_t R_{t+1}$. The investor puts a share of α_t in the risky asset. The return on his portfolio is given by $R_{p,t+1} = \alpha_t R_{t+1} + (1-\alpha_t) R_{f,t} = R_{f,t} + \alpha_t (R_{t+1} - R_{f,t})$, with corresponding conditional variance $\sigma_{p,t}^2 = \alpha_t^2 \sigma_t^2$.

Consider a conditional mean-variance optimizer with risk-aversion parameter $\gamma > 1$ that solves the following problem:

$$\max \left(\mathbb{E}_t R_{p,t+1} - \frac{\gamma}{2} \sigma_{p,t}^2 \right) = R_{f,t} + \max \left(\alpha_t (\mathbb{E}_t R_{t+1} - R_{f,t}) - \frac{\gamma}{2} \alpha_t^2 \sigma_t^2 \right). \tag{1}$$

From the first-order condition it follows that the solution to this maximization problem is

$$\alpha_t = \frac{\mathbb{E}_t R_{t+1} - R_{f,t}}{\gamma \sigma_t^2}. \tag{2}$$

Hence, the portfolio share in the risky asset is given by the risk premium divided by the conditional variance times the risk-aversion coefficient.

The result above was obtained by assuming that the investor is a mean-variance optimizer. However, we are interested in the optimal allocation for an investor that wants to maximize utility over his future wealth, since this is what we assume to be the investor's objective. We obtain similar results in case the investor does not only care about the mean and variance of his portfolio returns, but instead derives utility from his wealth at the end of the period. The utility that he derives from wealth W is described by the utility function $u(W)$. He solves the following optimization problem:

$$\begin{array}{ll} \max & \mathbb{E}_t u(W_{t+1}) \\ \text{subject to} & W_{t+1} = (1 + R_{p,t+1}) W_t. \end{array} \tag{3}$$

The utility function $u(W)$ plays an important role in the optimization problem (3). To solve the problem, we need to make an assumption about this utility function. We follow Campbell and Viceira (2002) by assuming that agents have the power utility function

$$u(W) = \frac{W^{1-\gamma} - 1}{1 - \gamma} \tag{4}$$

with risk-aversion parameter $\gamma > 1$. This utility function has the advantage that the coefficient of relative risk aversion, defined as

$$R(W) = -\frac{Wu''(W)}{u'(W)}, \tag{5}$$

is equal to γ for any wealth level W and thus constant. In this case, people with different wealth levels will make the same decision about the fraction of wealth that they are willing to pay to avoid a gamble of a given size relative to wealth. As Campbell and Viceira (2002) point out, this is empirically plausible. Recall that next period's wealth is the product of this period's wealth and the portfolio return, which means that financial risks are multiplicative. Consumption and wealth have increased greatly over the past two centuries, which makes financial risks much bigger in absolute sense, while the relative scale is still the same. Since there is no evidence of long-term trends in interest rates and risk premia due to this growth effect, investors are indeed willing to pay roughly the same relative amount to avoid a relative risk, independent of their absolute wealth level.

In the remainder of this thesis, we will make modeling assumptions about log returns instead of regular asset returns. Log returns are denoted by lowercases. The log return on the risk-free asset in the period from t to $t+1$ is defined by $r_{f,t} \equiv \log(1 + R_{f,t})$, and the log return on the stock in this same period is defined by $r_{t+1} \equiv \log(1+R_{t+1})$. Let $\sigma_t^2 \equiv \text{Var}_t(r_{t+1})$ be the conditional variance of the log stock return.

Theorem 1. *Assume that the return on the risky asset follows a lognormal distribution and the investor has a power utility function as described by (4). The optimal share of the risky asset is then approximately given by the equivalent of the mean-variance solution (2):*

$$\alpha_t = \frac{\mathbb{E}_t(r_{t+1}) - r_{f,t} + \frac{1}{2}\sigma_t^2}{\gamma \sigma_t^2}.$$

Proof. See Appendix A. □

2.2 Long-term investors

Early results of Samuelson (1969) and Merton (1969, 1971) prove that under two sets of conditions, a long-term investor chooses the same portfolio as a short-term investor: the myopic portfolio. It is assumed that the risk-free rate is constant; the conditions will not be valid when interest rates are stochastic. Therefore, this result is not very realistic.

First, when the investor has log utility (which is the limiting case for power utility with $\gamma \to 1$), the portfolio choice will always be myopic. The reason for this is that investors maximize expected log returns. As the log return over multiple periods is just given by the sum of the log returns over the single periods, the expected log return is maximized by maximizing the expected log return over each period separately. It follows that the long-term investor makes the same decisions as the short-term investor.

The other set of conditions is that the investor has a power utility function and returns are identically and independently distributed. We will illustrate this by considering a two-period investment problem. In principle, future optimal allocations (α_{t+1}) depend on the future state of the world, since they can be chosen in the future when new information is available. We have seen that the wealth level has no influence on the optimal allocation for a short-term investor with a power utility function. By replacing the one-period return by a multi-period return, it follows that this also holds for the long-term investor with a constant relative risk aversion. Hence, the optimal allocation does not depend on past returns. Since the returns are i.i.d., no additional information about future returns becomes available as time goes by. This means that the optimal portfolio share in each period is deterministic and can already be determined at the current point in time.

Let R_f be the constant risk-free rate. We denote the constant log return on the risk-free asset by r_f, the constant expected log return on the risky asset by μ, and the constant variance of the return on the risky asset by σ^2. Following the derivations in Appendix A, the two-period log return of the portfolio $r_{p,t,t+2} \equiv r_{p,t+1} + r_{p,t+2}$ with conditional variance $\sigma^2_{p,t,t+2}$ is approximated by

$$r_{p,t+1} + r_{p,t+2} \approx 2r_f + \alpha_t(r_{t+1} - r_f) + \frac{1}{2}\alpha_t(1-\alpha_t)\sigma^2 + \alpha_{t+1}(r_{t+2} - r_f) + \frac{1}{2}\alpha_{t+1}(1-\alpha_{t+1})\sigma^2.$$

This two-period return has a variance of $(\alpha_t^2 + \alpha_{t+1}^2)\sigma^2$ and a mean log return of

$$2r_f + (\alpha_t + \alpha_{t+1})(\mu - r_f) + \frac{1}{2}\alpha_t(1-\alpha_t)\sigma^2 + \frac{1}{2}\alpha_{t+1}(1-\alpha_{t+1})\sigma^2.$$

For an investor with power utility, this leads to the following optimization problem:

$$\begin{aligned} &\max\left\{\mathbb{E}_t r_{p,t,t+2} + \frac{1}{2}(1-\gamma)\sigma^2_{p,t,t+2}\right\} = \\ &\max\left\{(\alpha_t + \alpha_{t+1})\left(\mu - r_f + \frac{1}{2}\sigma^2\right) - \frac{\gamma}{2}\left(\alpha_t^2 + \alpha_{t+1}^2\right)\sigma^2\right\}. \end{aligned} \tag{6}$$

In the last period, the long-term investor has become a short-term investor. This means that the decision of the long-term investor is the same as the decision of the short-term investor. From the optimization problem it follows that for a given sum $\alpha_t + \alpha_{t+1}$, the investor will minimize total variance by choosing α_t and α_{t+1} to be equal. For a solution to be optimal, it should therefore hold that $\alpha_t = \alpha_{t+1}$. By induction, we can conclude that the long-term investor will in each period choose the same optimal portfolio share, namely the myopic portfolio choice.

In conclusion, we have seen that the optimal portfolio of a long-term investor with power utility is the same as the myopic portfolio under the condition of constant risk-free rates and i.i.d. returns of the risky asset. No matter what the investment horizon is, the optimal portfolio is always the same. Merton (1969) was one of the first to study the problem of long-term investing and to find conditions for the myopic strategy to be optimal. Therefore, we will refer to this solution for long-term investors as the Merton solution.

2.3 Including human capital

Financial planners often recommend a negative relation between equity holdings and age. Young individuals should invest a large fraction of their wealth in the risky asset, and should reduce their equity holdings as they grow older. They have formulated the well-known "100-minus-your-age"-rule. As a rule of thumb, you should subtract your age from 100 to find the percentage of wealth that is to be allocated to stocks. The remainder is to be invested in the risk-free asset. This goes totally against the Merton solution that prescribes a constant allocation over time. Are we missing something?

Yes, we are. Remember that we are interested in the optimal allocation of the savings that an individual has already accumulated for his retirement. However, the optimization problem for this individual is not the same as the problem that we have discussed before for the long-term investor. The long-term investor invests a given amount of money at the beginning of the period and does not add or subtract money before the end of the period. The working individual, however, saves every year a part of his income for retirement and therefore adds money to the accumulated savings during the period.

2.3.1 Solving the life-cycle optimization problem

Define the fraction of labor income that participants in the defined contribution pension plan save for retirement at time t by ρ_t. We assume that these values for ρ_t are exogenously determined. Hence, the consumption and savings at time t are predetermined and are not part of the optimization problem. To solve for the optimal allocation of the pension money, we should take into account future contributions to the pension scheme. In the optimization problem, we should explicitly model future pension savings at each point in time. This is called the life-cycle optimization problem.

At first sight, the new problem seems much more difficult to solve, since the optimization problem is now different at each point in time. Though, solving this difficult problem becomes straightforward by recognizing that future labor income is also a form of wealth. Although labor income cannot be traded since the abolishment of slavery, it can be treated like an asset. We define human capital to be the present value of future pension contributions. A person's total wealth thus consists of both financial wealth and human capital. Human capital should be taken into account in the investment decision. Since the yearly contribution rate is fixed and we therefore leave out consumption decisions, we are interested in the part of the human capital that is destined for the pension scheme. We denote this value by H_t.

Hence, the total wealth of the investor is equal to the sum of financial wealth and human capital, $W_t = F_t + H_t$.

If we model wealth as the sum of financial wealth and human capital, then all changes in the wealth level can be modeled as returns on assets in the portfolio of the life-cycle investor, be it financial assets or human capital. During the life-cycle, human capital becomes financial capital when labor income is paid and part of this income is saved for retirement, but over the whole period no additional wealth is added or subtracted from the total amount of wealth.

We assume for this moment that labor income can be replicated perfectly by the assets that are available on the financial market. That is, labor income is potentially stochastic, but there is only systematic labor income risk (risk traded on the financial market) and no idiosyncratic labor income risk (additional labor income risk that is not traded on the financial market). Since there is only one risky asset on the market, the stock, all labor income risk originates from an exposure to stock returns. In this case, human capital is equivalent to a portfolio of stocks and the risk-free asset. This assumption is not realistic, but it is useful to obtain an analytical solution to the life-cycle problem. Later on, we will include idiosyncratic labor income risk as well, but we will find that the effects of idiosyncratic labor income risk on the optimal allocations are negligible.

Because total wealth is the sum of financial wealth and human capital and both types of wealth are invested in a portfolio of stocks and the risk-free asset, total wealth itself is invested partly in stocks and partly in the risk-free asset. The problem of determining the optimal allocation of total wealth to stocks is equivalent to the problem of the long-term investor that was described before. Under the assumption of power utility and i.i.d. returns on the risky asset, we have seen that the optimal strategy of the investor is to invest a constant share α of his total wealth in the risky asset. This solution was found by Bodie et al. (1992).

Since the investor is not able to sell his human capital, he is constrained to holding an amount of H_t in human wealth. The allocation of human capital to the two available assets is predetermined and cannot be influenced. Let k be the exposure of human capital to stock returns. This means that human capital is equivalent to an investment of kH_t in stocks and $(1-k)H_t$ in the risk-free asset. Since we assumed that labor income shocks (if any) and stock returns are perfectly correlated and labor income shocks are typically less volatile than stock returns, we would generally expect to find $|k| < 1$. We should choose the allocation of financial wealth such that the allocation of total wealth to the risky asset is equal to the optimal value α. Therefore, the optimal share of the risky asset as a percentage of financial wealth is

$$\hat{\alpha}_t = \frac{\alpha(F_t + H_t) - kH_t}{F_t} = \alpha + (\alpha - k)\frac{H_t}{F_t}. \tag{7}$$

To illustrate solution (7), we consider a numerical example. Suppose that it is optimal for a life-cycle investor with a certain risk aversion to allocate 20% of his total wealth to the risky asset. This individual has a human capital that has an exposure of $k = 10\%$ to the risky asset. We focus on two moments in his life: a point in the beginning of his life-cycle, at which he has a lot of human capital and little financial wealth, and a point at the end of his life-cycle, at which he has accumulated a lot of financial wealth but at which he has little human capital left. Table 1 gives a stylized example for the allocations at both points in time. While the individual has the same amount of his total wealth invested in stocks when he is young as when he is old, his financial wealth is invested for as much as 60% in stocks at the younger age, while it is invested for only 22.5% in stocks at the older age.

Table 1: Allocations for the same individual at different ages.

(a) Allocations for a young person

Type of wealth	Stocks	Bonds
Human capital	80,000	720,000
Financial wealth	120,000	80,000
Total wealth	200,000	800,000

(b) Allocations for an old person

Type of wealth	Stocks	Bonds
Human capital	20,000	180,000
Financial wealth	180,000	620,000
Total wealth	200,000	800,000

2.3.2 Allocation restrictions

In practice, the individual cannot simply choose any allocation to the stock that he likes. Since most countries have laws against slavery, it is not possible to design a legally binding contract that allows one to sell his human capital, which is the reason that labor income cannot be traded. This also has the implication that it is not possible to borrow against future labor income. However, if equation (7) leads to an $\hat{\alpha}$ that is larger than one, it would mean that we would have to borrow against future labor income. Therefore, we impose the restriction that the allocation of financial wealth to stocks is not larger than one. Further, it is for individuals also often impossible to go short in stocks. We assume that short-selling is indeed not allowed and impose the restriction that the allocation of financial wealth to stocks is non-negative.

By imposing these restrictions on the allocation of financial wealth, the unconstrained optimization problem of finding the optimal allocation of total wealth becomes a constrained optimization problem. Before, we have seen that the myopic portfolio is optimal for the long-term investor when there are no allocation constraints. We ignore potential deviations in the optimal allocation between the short-term investor and the long-term investor as a result from including the allocation constraints and simply take the constrained myopic portfolio as a solution for life-cycle investors.[1] Due to the specific form of the objective function (a concave parabola, see Appendix A), the optimal solution to the constrained myopic problem is the allocation that is closest to the unconstrained optimal allocation α. Hence, the myopic

[1]For this problem, there is no longer a straightforward analytical solution available. We could use stochastic dynamic programming to find the optimal allocation at each point in time. Instead, we chose to focus on an approximate solution here.

optimal allocation of financial wealth at time t is given by:

$$\hat{\alpha}_t = \max\left[\min\left\{\alpha + (\alpha - k)\frac{H_t}{F_t}, 1\right\}, 0\right]. \tag{8}$$

2.3.3 Valuation of human capital

In order to find the value of H_t, we need to have a specific model for labor income. In this thesis, we will focus on an individual that starts working at the age of 25, works for $T = 40$ years, and retires at the age of 65. We assume for the moment that inflation is constant and let π be the constant inflation level. Besides its exposure k to stock return shocks, we assume that labor income grows with inflation and with a deterministic exponential growth factor that represents growth in labor productivity. Recall that r_{t+1} is the log return on the stock at time $t+1$. We assume that this log return is normally distributed with mean $r_f + \mu$ and variance σ^2: $r_{t+1} = r_f + \mu + \epsilon_{t+1}$, with $\epsilon_{t+1} \overset{i.i.d.}{\sim} N(0, \sigma^2)$. Define the labor income at time t by L_t. We normalize the labor income at time 0 (corresponding to the age of 25) to be equal to 1. We have the following model for labor income:

$$\begin{cases} L_0 = 1 \\ L_{t+1} = L_t \exp(g_{t+1} + \pi + k\epsilon_{t+1}) \quad \text{for } t = 0, \ldots, T-2 \\ L_T = 0 \end{cases}, \tag{9}$$

where g_t is the deterministic growth in labor productivity at time t. An explicit expression for L_t is given by

$$L_t = \exp\left(G_t + \pi t + k\sum_{s=1}^{t}\epsilon_s\right), \tag{10}$$

where $G_t = \sum_{s=1}^{t} g_s$.

From the specification of L_t, it follows that the log return on human capital at time t has an exposure of k to shocks to the risky asset and that there is no idiosyncratic labor income risk. Note that the logarithm of labor income is modeled to be affine in shocks to the risky asset, rather than labor income itself. As a result, labor income cannot be replicated by a linear portfolio of stocks and bonds. However, as a first-order approximation, the exposure can be obtained by investing the fraction k of human capital in stocks and the remaining part $(1-k)$ in the risk-free asset. We ignore the small replication errors that arise due to this approximation. This means that indeed the amount kH_t is invested in the risky asset via the human capital, while the amount $(1-k)H_t$ is invested in the risk-free asset.

We assume that the financial market is free of arbitrage. This implies that asset with the same exposures to financial risks have the same returns. Thus, the return on human capital at time t should (approximately) be equal to $R_f + k(R_t - R_f)$. According to standard asset pricing theory, we can obtain the market-consistent value of a future payment by discounting the expected value of the future payment by the expected return (otherwise, the expected return would not be the expected return). Let $H_0^{(t)}$ be the market-consistent value at time 0 of the labor income contribution to the pension plan at time t. For the total human capital, we have $H_0 = \sum_{t=1}^{T-1} H_0^{(t)}$, where T is the retirement age. For $H_0^{(t)}$, we derive

$$\begin{aligned} H_0^{(t)} &= \frac{\mathbb{E}(\rho_t L_t)}{\text{Expected return on human capital up to period } t} \\ &\approx \frac{\mathbb{E}(\rho_t L_t)}{(1 + R_f + k(\mathbb{E}(R_t) - R_f))^t} \\ &= \frac{\rho_t \exp(G_t + \pi t + \frac{1}{2}k^2\sigma^2)}{\left(\exp(r_f) + k(\exp(r_f + \mu + \frac{1}{2}\sigma^2) - \exp(r_f))\right)^t}. \end{aligned} \tag{11}$$

2.3.4 Results for bond-like human capital

To illustrate, we show a typical allocation over time for a model in which the continuously compounded nominal risk-free rate is equal to 4% and the inflation is equal to 2%. Further, we assume that log stock returns in excess over the nominal risk-free rate have a mean of 3.0% and a standard deviation of 22.7%. These parameter values are consistent with empirical data (see Section 4).

Several studies have investigated the relation between stock returns and returns on human capital. For most people, human capital is mostly bond-like. For people that do not work in the financial industry, the estimated effects of stock returns on labor income are small. As the benchmark parameter, we therefore take $k = 2\%$, which means a small positive effect of stock returns on labor income. We will get back to this point and refer to relevant studies in Section 5.4.

As parameters for the yearly contribution rate and the deterministic growth in labor productivity, we take input data from the Netherlands. In the Netherlands, people that are enrolled in a defined contribution scheme save a part of their income for retirement without having to pay taxes on these savings. The maximum fraction of total income that can be saved free of taxes depends on the age of the individual. These maximums are announced by the Dutch tax authority and are called "maximum premiestaffel". This "premiestaffel", which is a life-cycle path of contributions to the pension scheme as a percentage of labor income, is determined such that under regular market conditions, the individual has an acceptable replacement rate (retirement income as a percentage of average income in the working life). Therefore, we assume that the contributions to the pension scheme are equal to the maximum tax-free contributions in the Netherlands. We use the "maximum premiestaffel" from 2009. The yearly contributions are given in Table 2. A new "premiestaffel" will be introduced as of 2014.

We use $N = 10,000$ sample paths for the returns on the risky asset. For each sample path and at each point in time, we calculate the "optimal" allocation according to (8). Since the ratio between financial wealth and human capital depends on past investment returns, the allocation at a certain point in time varies slightly over the N different sample paths. The mean allocation across all sample paths is given in Figure 1, for different values of the risk-aversion parameter γ.

Table 2: Benchmark parameters for the individual's labor income.

Age at time t	Labor income growth g_t	Contribution ρ_t	Age at time t	Labor income growth g_t	Contribution ρ_t
25-29	3%	7.3%	45-49	1%	16.3%
30-34	3%	8.9%	50-54	1%	20.0%
35-39	2%	10.9%	55-59	0%	24.8%
40-44	2%	13.3%	60-64	0%	31.1%

The results that we get in Figure 1 are consistent with the advice of the financial planner. For young people, human capital is their most valuable asset. Since exposure of labor income to stock returns is only 2%, human capital has the nature of a risk-free bond. This means that young people are already heavily invested in the risk-free asset. Since it is optimal to have an exposure to stocks which is considerably larger than 2%, depending on the risk-aversion parameter γ, they should invest all their financial wealth in stocks. Older people have more financial wealth and less human capital compared to young people. However, it is still optimal for them to have an exposure of α to stocks as a fraction of total wealth. Older people should therefore decrease their equity holdings and invest the difference in the risk-free asset. Just one year before retirement, the agent has no human capital left. It is optimal for him to follow the myopic strategy and invest a fraction α of his financial wealth in the risky asset.

Figure 1: Mean allocations for the life-cycle investor with human capital ($k = 2\%$), when there is a constant risk-free rate.

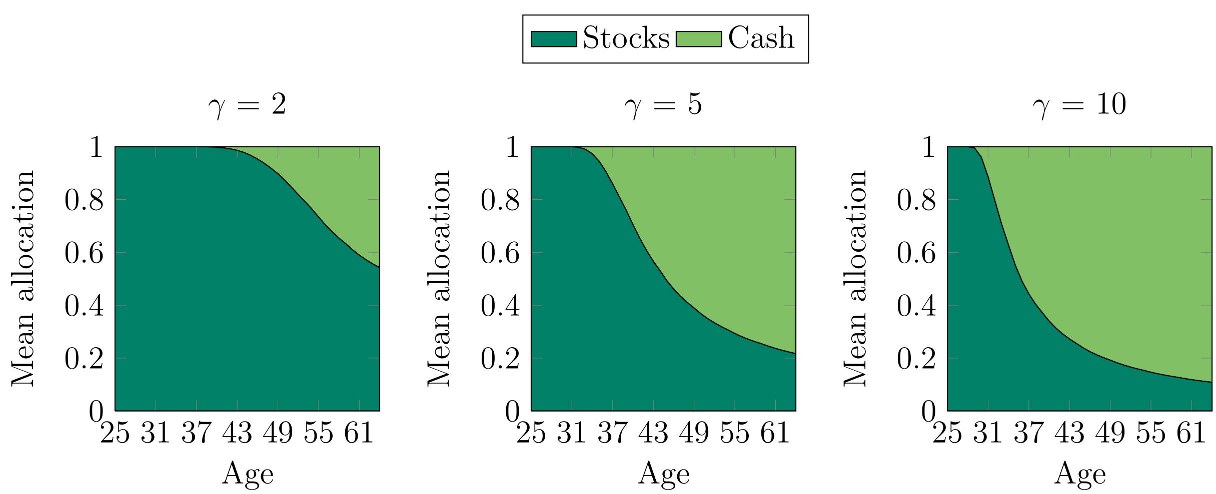

2.3.5 Results for different natures of labor income

It becomes clear from formula (7) that the optimal allocation of financial wealth depends heavily on the nature of labor income, i.e., the positions that the life-cycle investor has already in stocks and in bonds implicitly via his human capital. If human capital has little exposure to stock returns and is bond-like, the investor should follow the conventional advice to reduce equity holdings over time. However, if human capital has much exposure to stock returns and is stock-like, the investor should invest all his financial wealth in bonds when he is young. We will repeat the previous analysis for an investor with $\gamma = 5$ that has different exposures of his labor income to equity risk. Later on, we will find that the volatility of labor income shocks is about one fifth of the volatility of equity returns. If labor income shocks are perfectly correlated to equity shocks, this means that the exposure k is approximately equal to 20% (positive correlation) or -20% (negative correlation). We will compare the optimal allocations for these values of k to the allocation for $k = 2\%$. The results are presented in Figure 2.

Figure 2: Mean allocations for the life-cycle investor ($\gamma = 5$) with human capital, when there is a constant risk-free rate.

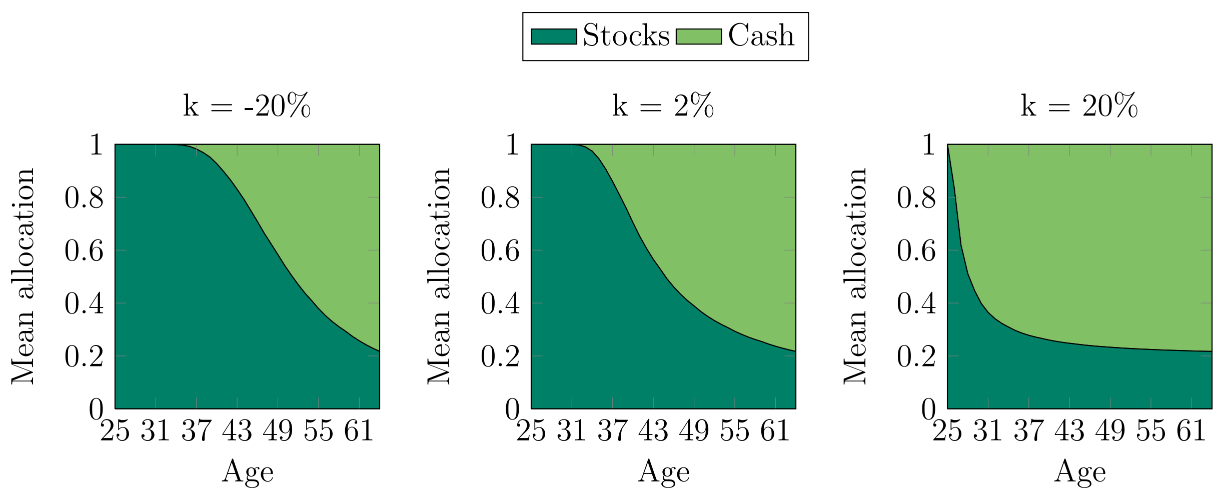

Since the myopic portfolio for $\gamma = 5$ allocates slightly more than 20% to the risky asset, we see a decreasing pattern in the allocation of stocks over time for all three values for k. Of course, this would change if we would consider values for k that would even be higher in absolute value. However, we can clearly see the impact that the nature of human capital has on the optimal allocations. The higher the implicit investments in stocks via human capital, the lower the optimal allocation to stocks for financial wealth. Since the investor has no human capital left just before his retirement, all three paths end at the same allocation: the myopic portfolio choice.

In the following chapters, we will relax the assumptions that were made in this chapter and find the resulting allocations. Most importantly, we will relax the assumption of a constant interest rate, by introducing stochastic interest rates. Further, we also include inflation risk, annuitization, idiosyncratic labor income risk, and return predictability in the analysis.

3 Stochastic interest rates

In the previous section, we have focused on the allocation decision of a life-cycle investor when there is a constant risk-free rate. The goal of this study is to determine what the optimal allocation to stocks and bonds over the life-cycle will be when interest rates are stochastic rather than fixed. In that case, both human capital and bond returns become stochastic, in addition to stochastic returns on the risky asset.

We will model the term structure of interest rates by introducing a VAR(1)-model and specifying an affine term structure model. Examples of such models are the Vasicek (1977) model and the Brennan and Xia (2002) model. First, we will introduce the general framework. Second, we will specify two affine term structure models. Finally, we will estimate the parameters of these affine models by using market data.

3.1 Modeling stochastic interest rates

In order to model the term structure of interest rates, we specify a pricing kernel and use the assumption of no arbitrage opportunities to derive the term structure of interest rates. Many influential papers on modeling interest rates rely on a specification of the pricing kernel, see for instance Campbell and Viceira (2001) and Piazzesi (2010). By assuming that prices of risk are affine in the state variables, we arrive at an essentially affine term structure model as introduced by Duffee (2002) and specified in discrete time by Ang and Piazzesi (2003).

An alternative to this approach is to compose a dataset of historically observed yields, perform a factor analysis of yield changes and express the covariance matrix of yields in terms of the few factors that drive their movements. Litterman and Scheinkman (1991) find three principal sources of variation in the yield curve: a "level" factor, a "slope" factor, and a "curvature" factor. However, the problem with this approach is that this method easily leads to a representation of the term structure that allows for arbitrage opportunities. Since we do not want to rely on a purely statistical model that allows for arbitrage opportunities, we follow the pricing kernel method instead.

3.1.1 Specifying a pricing kernel

Under the assumption of no arbitrage opportunities, the pricing kernel determines the prices of all assets in the economy. Using a pricing kernel to price assets in the market is equivalent to several other valuation methods, for instance risk-neutral valuation (see e.g. Cochrane (2005)). When markets are complete, all payoffs can be replicated by existing assets on the financial market. In that case, the pricing kernel is uniquely defined. For incomplete markets, there exist infinitely many kernels that lead to the same prices of the existing assets on the market. This means that we have to make specific modeling assumptions for the pricing kernel, in addition to the assumptions for the asset dynamics.

The use of a pricing kernel can be illustrated by considering a representative-agent framework. Suppose that we have a two-period economy, where the representative agent maximizes his expected utility, which can be written as

$$\mathbb{E}_t(U(C_t, C_{t+1})) = u(C_t) + \delta \mathbb{E}_t(u(C_{t+1})), \tag{12}$$

where C_t is the consumption at time t and δ the subjective discount factor. Now introduce a number of assets with price vector $\boldsymbol{P}_t$ and (random) payoff vector $\boldsymbol{X}_{t+1}$, where $\boldsymbol{X}_{t+1}$ consists of future prices $\boldsymbol{P}_{t+1}$ plus dividend or coupon payments. Suppose that the agent is endowed with an initial wealth W_t. At time t, the agent can choose to invest (part of) his wealth in financial assets for consumption at time $t+1$. Define $\boldsymbol{\alpha}$ to be the vector of allocations to the assets. The agent has the following two budget restrictions:

$$\begin{aligned} C_t + \boldsymbol{\alpha}' \boldsymbol{P}_t &= W_t \\ C_{t+1} &= \boldsymbol{\alpha}' \boldsymbol{X}_{t+1}. \end{aligned} \tag{13}$$

Substituting the budget constraints into the objective function leads to the following unconstrained maximization problem:

$$\max_{\boldsymbol{\alpha}} \{ u(W_t - \boldsymbol{\alpha}' \boldsymbol{P}_t) + \delta \mathbb{E}_t(u(\boldsymbol{\alpha}' \boldsymbol{X}_{t+1})) \}. \tag{14}$$

Under some regularity conditions, the differentiation operator and the expectation operator can be interchanged. This is an application of Leibniz's rule (see Flanders (1973)). We assume that the utility function u is sufficiently smooth to satisfy these conditions. The first-order condition then yields

$$\boldsymbol{P}_t = \mathbb{E}_t \left[\delta \frac{u'(C_{t+1})}{u'(C_t)} \boldsymbol{X}_{t+1} \right] \equiv \mathbb{E}_t[M_{t+1} \boldsymbol{X}_{t+1}]. \tag{15}$$

Hence, we have found a basic asset pricing formula for this economy. Given the payoff of an asset and the investor's utility function and consumption choices C_t and C_{t+1}, we can determine the market prices of the assets on the basis of the pricing kernel M_{t+1}. If we assume that the utility function is strictly increasing, the pricing kernel M_{t+1} is always positive.

We have illustrated the use of a pricing kernel in a simple two-period for a certain model of the behavior of economic agents. In fact, the pricing relation $\boldsymbol{P}_t = \mathbb{E}_t[M_{t+1} \boldsymbol{X}_{t+1}]$ is equally valid in a general multi-period setting; it follows directly from absence of arbitrage (see e.g. Cochrane (2005)). A different model leads to a different specification of the pricing kernel. For the purpose of asset pricing, it is useful to focus directly on the specification of the pricing kernel, without considering the underlying economic model explicitly.

We will follow many of the influential papers in the recent literature on modeling term structures of interest rates by specifying a pricing kernel which is affine in the risk factors (see for example Campbell and Viceira (2001) and Brennan and Xia (2002)). We choose our description of the financial market such that it allows for time-varying interest rates, inflation rates, and risk premia.

3.1.2 Affine term structure models

We assume that the prices of nominal bonds, real bonds, and the stock market index are determined on the basis of a single state vector $\boldsymbol{x}_t$. Many term structure models assume that the state vector follows a vector autoregressive (VAR) model, that the short rate and the inflation rate are affine functions of the state vector, that the pricing kernel is exponentially affine in the shocks that drive the state vector, and that prices of risk are affine in the state vector. As we will see later, this model leads to an affine term structure of interest rates. The existence of an explicit solution for the bond price function (Piazzesi, 2010) makes affine models very attractive. Another advantage is that our approach is very flexible, since we do not need to rely on a general economic model. By introducing time-varying risk premia, we are able to model time variation in expected bond returns. This is consistent with the stylized fact that bond returns are predictable. For instance, Cochrane and Piazzesi (2005) report that a single variable, constructed by a linear combination of forwards rates, explains up to 44% of the variance of excess bond returns .

In general, we have a state vector $\boldsymbol{x}_t \in \mathbb{R}^N$ that obeys the following process:

$$\boldsymbol{x}_{t+1} = \boldsymbol{\mu} + \Gamma(\boldsymbol{x}_t - \boldsymbol{\mu}) + \Sigma\boldsymbol{\epsilon}_{t+1}, \tag{16}$$

where $\boldsymbol{\epsilon}_{t+1} \overset{i.i.d.}{\sim} N(0, I_N)$, with I_N the identity matrix of size $N \times N$. We impose the volatility matrix Σ to be lower triangular, since a full matrix Σ would not be statistically identified. Note that $\text{Var}_t(\boldsymbol{x}_{t+1}) = \Sigma\Sigma'$, so we can use a Cholesky decomposition to identify the lower triangular matrix Σ. Thus, we allow the error terms to be contemporaneously correlated, but assume that they are homoskedastic and independently distributed over time. This assumption of homoskedasticity is of course restrictive, since we do not allow for the possibility that state variables can predict changes in risk levels. As Campbell and Viceira (2002) point out, however, several authors such as Campbell (1987), Harvey (1989, 1991), and Glosten, Jagannathan, and Runkle (1993) have found only modest effects on optimal asset allocations. Since we only have a limited sample period available, we therefore abstain from including time-varying variances of the error terms.

Let $Y_t^{(n)}$ be the annually compounded nominal yield with maturity n at time t, and let Π_t be the level of the price index at time t, with the normalization $\Pi_0 = 1$. Define $y_t^{(n)}$ to be the continuously compounded nominal yield at time t with maturity n. The continuously compounded nominal short rate $y_t^{(1)}$ and the log (price) inflation rate π_t will be modeled as linear functions of the state vector $\boldsymbol{x}_t$. They are given by

$$\begin{aligned} y_t^{(1)} &\equiv \log(1 + Y_t^{(1)}) = \boldsymbol{\delta}_y'\boldsymbol{x}_t \\ \pi_t &\equiv \log(\Pi_t) - \log(\Pi_{t-1}) = \boldsymbol{\delta}_\pi'\boldsymbol{x}_t \end{aligned} \tag{17}$$

where $\boldsymbol{\delta}_y, \boldsymbol{\delta}_\pi \in \mathbb{R}^N$ are parameters.

Finally, we assume that the pricing kernel M_t is given by

$$-\log M_{t+1} = y_t^{(1)} + \frac{1}{2}\boldsymbol{\lambda}_t'\boldsymbol{\lambda}_t + \boldsymbol{\lambda}_t'\boldsymbol{\epsilon}_{t+1}, \tag{18}$$

where the prices of risk $\boldsymbol{\lambda}_t$ are affine in the state variables $\boldsymbol{x}_t$:

$$\boldsymbol{\lambda}_t = \boldsymbol{\Lambda}_0 + \Lambda_1 \boldsymbol{x}_t. \tag{19}$$

We denote the price of an n-year zero-coupon bond at time t by $P_t^{(n)}$. By definition, the continuously compounded nominal yield is given by

$$y_t^{(n)} \equiv -\frac{\log P_t^{(n)}}{n}. \tag{20}$$

One year from now, the n-year bond has become a bond with maturity $n-1$ and price $P_{t+1}^{(n-1)}$. As before, absence of arbitrage implies that the pricing relation

$$P_t^{(n)} = \mathbb{E}_t \left(M_{t+1} P_{t+1}^{(n-1)} \right) \tag{21}$$

is valid. Since the value of an immediate payoff of 1 is by definition 1, the price of a bond with maturity 0 is $P_t^{(0)} \equiv 1$. Note that the pricing kernel M_t is specified such that

$$\begin{aligned} P_t^{(1)} &= \mathbb{E}_t \left(M_{t+1} P_{t+1}^{(0)} \right) = \mathbb{E}_t (M_{t+1}) \\ &= \mathbb{E}_t \left[\exp \left\{ -y_t^{(1)} - \frac{1}{2} \boldsymbol{\lambda}_t' \boldsymbol{\lambda}_t - \boldsymbol{\lambda}_t' \boldsymbol{\epsilon}_{t+1} \right\} \right] \\ &= \exp \left\{ -y_t^{(1)} - \frac{1}{2} \boldsymbol{\lambda}_t' \boldsymbol{\lambda}_t + \frac{1}{2} \boldsymbol{\lambda}_t' \boldsymbol{\lambda}_t \right\} = \left(1 + Y_t^{(1)} \right)^{-1}. \end{aligned} \tag{22}$$

By imposing no-arbitrage conditions, this model specification leads to a term structure of interest rates. In this model, the price at time t of a nominal bond with maturity n can be shown to satisfy

$$P_t^{(n)} = \exp(-A(n) - \boldsymbol{B}(n)' \boldsymbol{x}_t), \tag{23}$$

where

$$\begin{aligned} A(n) &= A(n-1) + \boldsymbol{B}(n-1)'(I - \Gamma)\boldsymbol{\mu} - \frac{1}{2} \boldsymbol{B}(n-1)' \Sigma \Sigma' \boldsymbol{B}(n-1) - \boldsymbol{B}(n-1)' \Sigma \boldsymbol{\Lambda}_0 \\ \boldsymbol{B}(n) &= (\Gamma - \Sigma \Lambda_1)' \boldsymbol{B}(n-1) + \boldsymbol{\delta}_y, \end{aligned} \tag{24}$$

with initial conditions $A(1) = 0$ and $B(1) = \boldsymbol{\delta}_y$. This leads to the n-period yield

$$y_t^{(n)} = a(n) + \boldsymbol{b}(n)' \boldsymbol{x}_t, \tag{25}$$

with $a(n) \equiv A(n)/n$ and $\boldsymbol{b}(n) \equiv \boldsymbol{B}(n)/n$. For a derivation of this result, we refer to Appendix B.

The formula for the price of a nominal bond directly leads to an expression for the return on a nominal zero-coupon bond, by recognizing that a bond which has maturity n at time t has become a bond with maturity $n-1$ at time $t+1$. From applying the bond pricing equation (23), it follows that the log return on an n-year nominal bond is given by

$$r_{b,t+1}^{(n)} = y_t^{(1)} - \frac{1}{2} \boldsymbol{B}(n-1)' \Sigma \Sigma' \boldsymbol{B}(n-1) - \boldsymbol{B}(n-1)' \Sigma \boldsymbol{\lambda}_t - \boldsymbol{B}(n-1)' \boldsymbol{\epsilon}_{t+1}. \tag{26}$$

For a derivation, see Appendix B.

Equation (25) allows us to model the whole yield curve. According to the Expectation Hypothesis, bond yields are expected values of average future short rates (see Piazzesi, 2010):

$$\text{EH: } Y_t^{(n)} = \frac{1}{n} \sum_{i=0}^{n-1} Y_{t+i}^{(1)}. \tag{27}$$

If we set Λ_1 equal to the zero matrix, writing out the recurrence relations in (24) yields:

$$\begin{aligned} A(n) &= \boldsymbol{\delta}_y'(I-\Gamma)\boldsymbol{\mu} + \boldsymbol{\delta}_y'(I-\Gamma^2)\boldsymbol{\mu} + \ldots + \boldsymbol{\delta}_y'(I-\Gamma^{n-1})\boldsymbol{\mu} + \phi^{(n)} \\ \boldsymbol{B}(n) &= \boldsymbol{\delta}_y + \Gamma'\boldsymbol{\delta}_y + \ldots + (\Gamma')^{n-1}\boldsymbol{\delta}_y \end{aligned} \tag{28}$$

where $\phi^{(n)}$ is the collection of Jensen's inequality terms and constant risk premium terms. The corresponding yields are

$$\begin{aligned} y_t^{(n)} &= \frac{A(n)}{n} + \frac{\boldsymbol{B}(n)'}{n}\boldsymbol{x}_t \\ &= \frac{1}{n}\left(\boldsymbol{\delta}_y'\boldsymbol{x}_t + \boldsymbol{\delta}_y'(\boldsymbol{\mu} + \Gamma(\boldsymbol{x}_t - \boldsymbol{\mu})) + \ldots + \boldsymbol{\delta}_y'(\boldsymbol{\mu} + \Gamma^{n-1}(\boldsymbol{x}_t - \boldsymbol{\mu})) + \phi^{(n)}\right) \\ &= \frac{1}{n}\sum_{i=0}^{n-1} \mathbb{E}_t y_{t+i}^{(1)} + \frac{1}{n}\phi^{(n)}. \end{aligned} \tag{29}$$

From (29), we infer that if $\boldsymbol{\Lambda}_0$ and Λ_1 are equal to the zero vector and the zero matrix, respectively, then long-term yields are the average of future expected short yields (ignoring the small terms resulting from Jensen's inequality). This means that a local version of the Expectation Hypothesis holds (Piazzesi, 2010). When $\boldsymbol{\Lambda}_0$ is nonzero but Λ_1 is equal to zero, the difference between long rates and future expected short rates is no longer equal to zero, but still does not depend on the state vector $\boldsymbol{x}_t$. In that case, there is no time variation in expected bond returns. In contrast, a nonzero Λ_1 implies that risk premia vary over time. In that case, $\boldsymbol{\lambda}_t$ is not constant over time and expected bond returns depend on the state vector. This means that there is predictability in bond returns. Of course, this has a large impact on the allocation decision of a life-cycle investor. In order to benefit from high bond risk premia and to avoid low risk premia, we will find that he should time bond markets when there is predictability.

As will become clear later on, we want to know the price of a real annuity, which consists of a portfolio of real bonds. Furthermore, since labor income grows (partially) with price inflation, it also has (some of) the characteristics of a real bond. Therefore, we are not only interested in nominal yields but also in real yields. We can derive the real term structure in a similar way as we did for the nominal term structure. If we hold at time t an n-year zero-coupon real bond that pays 1 unit in real terms, then at time $t+1$ this bond has become a zero-coupon real bond with maturity $n-1$. Hence, the nominal return on an n-year real bond is $\frac{P_{t+1}^{R(n-1)}\Pi_{t+1}}{P_t^{R(n)}\Pi_t}$.[2] It follows that the price at time t of a real bond with maturity n satisfies

$$P_t^{R(n)} = \exp(-A^R(n) - \boldsymbol{B}^R(n)'\boldsymbol{x}_t), \tag{30}$$

where

$$\begin{aligned} A^R(n) &= A^R(n-1) + (\boldsymbol{B}^R(n-1) - \boldsymbol{\delta}_\pi)'(I-\Gamma)\boldsymbol{\mu} \\ &\quad -\frac{1}{2}(\boldsymbol{B}^R(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\Sigma'(\boldsymbol{B}^R(n-1) - \boldsymbol{\delta}_\pi) \\ &\quad -(\boldsymbol{B}^R(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\boldsymbol{\Lambda}_0 \\ \boldsymbol{B}^R(n) &= (\Gamma - \Sigma\Lambda_1)'(\boldsymbol{B}^R(n-1) - \boldsymbol{\delta}_\pi) + \boldsymbol{\delta}_y. \end{aligned} \tag{31}$$

The initial conditions are given by

$$\begin{aligned} A^R(1) &= -\boldsymbol{\delta}_\pi'(I-\Gamma)\boldsymbol{\mu} - \frac{1}{2}\boldsymbol{\delta}_\pi'\Sigma\Sigma'\boldsymbol{\delta}_\pi + \boldsymbol{\delta}_\pi'\Sigma\boldsymbol{\Lambda}_0 \\ \boldsymbol{B}^R(1) &= \boldsymbol{\delta}_y + (\Sigma\Lambda_1 - \Gamma)'\boldsymbol{\delta}_\pi. \end{aligned} \tag{32}$$

This leads to the n-period real yield

$$y_t^{R(n)} = a^R(n) + \boldsymbol{b}^R(n)'\boldsymbol{x}_t, \tag{33}$$

with $a^R(n) \equiv A^R(n)/n$ and $\boldsymbol{b}^R(n) \equiv \boldsymbol{B}^R(n)/n$. For a derivation of this result, we refer to Appendix C.

We are also interested in the nominal return on a real zero-coupon bond. This return follows from applying the bond pricing equation (30). The nominal log return on an n-year real bond is given by:

$$\begin{aligned} r_{b,t+1}^{R(n)} &= y_t^{(1)} - \frac{1}{2}(\boldsymbol{B}^R(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\Sigma'(\boldsymbol{B}^R(n-1) - \boldsymbol{\delta}_\pi) \\ &\quad -(\boldsymbol{B}^R(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\boldsymbol{\lambda}_t - (\boldsymbol{B}^R(n-1) - \boldsymbol{\delta}_\pi)'\boldsymbol{\epsilon}_{t+1}. \end{aligned} \tag{34}$$

For a derivation, we again refer to Appendix C.

3.2 Two affine term structure models

We will now specify two models for the financial market in the class of affine term structure models. We will follow Brennan and Xia (2002) and Ang et al. (2006) by specifying the factors explicitly. In contrast, the financial market models that are proposed by Koijen et al. (2011) and Sangvinatsos and Wachter (2005), among others, contain term structures with latent factors. These models are estimated by making use of the Kalman filter. While a

[2] In reality, payments of real bonds are linked to a lagged value of the price index, rather than to the current value of the price index. This is the so-called "indexation lag". For simplicity, we assume that there is no indexation lag.

latent factor model approach allows for somewhat more flexibility, it leads to identification issues and greatly complicates the estimation procedure. Therefore, we assume that the term structure factors are observable.

We start with a simplified model in which only one yield factor is present and in which both stock and bond returns are not predictable. Afterwards, we will introduce a more general model, which allows for predictability and has a much better fit to historical term structures.

3.2.1 A three-factor model

We start with a model that is very similar to the model that Koijen and Nijman (2006) use. In this model, there is only one yield in the state vector. This allows for a no-arbitrage specification in which excess bond returns are not predictable, i.e., the price-of-risk parameter Λ_1 can be set equal to zero. We therefore assume that there is no predictability, neither in stock returns nor in bond returns. Later on, this will provide us with useful insights in the general outcomes, without having pronounced state effects due to return predictability. We define the following state vector:

$$\boldsymbol{x}_t = \begin{pmatrix} y_t^{(1)} \\ \pi_t \\ r_{s,t} - y_{t-1}^{(1)} \end{pmatrix}, \tag{35}$$

where $r_{s,t} = \log(S_t) - \log(S_{t-1})$ is the log return on the stock index. Note that in this model, the nominal short rate and the inflation rate are by definition linear functions of the state variable in (17):

$$\begin{aligned} y_t^{(1)} &= \boldsymbol{e}_1'\boldsymbol{x}_t \\ \pi_t &= \boldsymbol{e}_2'\boldsymbol{x}_t, \end{aligned}$$

where $\boldsymbol{e}_i$ denotes the i-th unit vector. We assume that the dynamics of the state variable satisfy (16). As said, for this simplified model, we exclude predictability in asset returns. For stock returns, this means that the last row of Γ is set equal to zero: $\boldsymbol{e}_3'\Gamma = \mathbf{0}$. For consistency, we also let the other state variables not depend on lagged stock returns: $\Gamma\boldsymbol{e}_3 = \mathbf{0}$. In order to have no predictability for bond returns, we set Λ_1 to be equal to the zero matrix, which implies $\boldsymbol{\lambda}_t = \boldsymbol{\Lambda}_0$.

3.2.2 A four-factor model

We now introduce our more general model, in which we allow for predictability in bond returns and in which we include an additional yield factor to improve the model fit. We define the state vector to be:

$$\boldsymbol{x}_t = \begin{pmatrix} y_t^{(1)} \\ y_t^{(10)} - y_t^{(1)} \\ \pi_t \\ r_{s,t} - y_{t-1}^{(1)} \end{pmatrix}. \tag{36}$$

Note that we follow Ang et al. (2006) by including two yield factors to model the dynamics of yields: the short rate is closely related to the first principal component of yields ("level") and the spread on the 10-year bond is closely related to the second principal component of yields ("slope"). Ang et al. (2006) argue that it is not necessary to include a third factor ("curvature") to capture variation of yields. They find that the third principal component accounts for less than 0.3%, using quarterly data. Therefore, it is sufficient to include two yield factors. Further, we include the macroeconomic quantity inflation as a factor. Finally, we also include the log return on the stock index as a factor.

We again assume that the dynamics of the state variables satisfy (16). As we have included the nominal short rate $y_t^{(1)}$ and the inflation rate π_t as state variables, (17) is automatically satisfied with $\boldsymbol{\delta}_y = \boldsymbol{e}_1$ and $\boldsymbol{\delta}_\pi = \boldsymbol{e}_3$. The vector of risk premia $\boldsymbol{\lambda}_t$ is now not constant but an affine function of the state variable $\boldsymbol{x}_t$, satisfying (18).

While we allow for bond return predictability, we do not allow for stock return predictability on the basis of yields, inflation or an additional factor such as dividend yield. Empirical findings on predictability are highly ambiguous. For instance, by looking at out-of-sample performance, a healthy skepticism appears to be appropriate when it comes to predicting the equity premium (see e.g. Goyal and Welch (2008)). We therefore choose to abstract from stock return predictability and impose the restriction that the last row of Γ is equal to zero. We further abstract from predictability with respect to stock returns by assuming that the last column of Γ is equal to zero, which implies that stock returns are not used to predict future values of the other state variables.

3.3 Estimation procedure

For the class of affine models that we have discussed before, we have several parameters: the VAR parameters $\boldsymbol{\mu}, \Gamma$, and Σ, and the price-of-risk parameters $\boldsymbol{\Lambda}_0$ and Λ_1. To obtain efficient estimates for these parameters, we should in principle estimate these parameters simultaneously. We could do so by using maximum likelihood. This method is computationally intensive and difficult to apply when we have limited data on long-term nominal yields and on real yields, which makes it for instance hard to estimate covariances.

Ang et al. (2006) compare estimating a model with latent factors in one step via maximum likelihood to the estimates from a two-step procedure with explicitly specified yield curve factors, in which first the VAR parameters are estimated and then the price-of-risk parameters. While a one-step maximum likelihood procedure is more efficient, the two-step procedure leads to almost the same estimates. Moreover, including explicitly specified yield factors in the state variable makes the model more tractable. Therefore, we will adopt this estimation procedure. First, we will estimate the VAR parameters by using maximum likelihood. Second, we will estimate the price-of-risk parameters by minimizing a weighted average of squared estimation errors.

3.3.1 Estimation of VAR parameters

Thus, we start with estimating the VAR parameters $\boldsymbol{\mu}, \Gamma$, and Σ. For both the three-factor model and the four-factor model, the last row and the last column of Γ (corresponding to stock returns) are restricted to be zero. Hence, we have a restricted VAR model. While for an unrestricted VAR model estimating the regression equations line-by-line by using standard OLS is efficient[3], this is not the case for a restricted VAR model. We use maximum likelihood to estimate the VAR parameters.

While OLS is not efficient in this case, it is a consistent method to estimate the VAR parameters. To illustrate a potential pitfall of the estimation procedure, we suppose for the moment that we would be using OLS to estimate the VAR parameters. As noted by Campbell and Viceira (2002), standard unconstrained least-squares fits exactly the mean of the variables in the VAR excluding the first observation. In this case, the population unconditional moment condition

$$\mathbb{E}(\boldsymbol{x_t} - \Gamma \boldsymbol{x}_{t-1}) = (I - \Gamma)\boldsymbol{\mu}$$

is replaced by its sample equivalent

$$\frac{1}{T}\sum_{t=1}^{T}\boldsymbol{x_t} - \Gamma\frac{1}{T}\sum_{t=1}^{T}\boldsymbol{x}_{t-1} = (I - \Gamma)\hat{\boldsymbol{\mu}},$$

where T denotes in this section the length of the sample data for the state vector.

Due to empirically high persistence in interest rates, a small difference in the sample mean of the sample without the first observation and the sample without the last observation can lead to a high difference between the full-sample mean and the estimated value for $\boldsymbol{\mu}$. Therefore, Campbell and Viceira (2002) impose the restriction that the unconditional means of the variables implied by the VAR coefficient estimates equal the full-sample means:

$$(I - \Gamma)\frac{1}{T+1}\sum_{t=0}^{T}\boldsymbol{x_t} = (I - \Gamma)\hat{\boldsymbol{\mu}} \Leftrightarrow \frac{1}{T+1}\sum_{t=0}^{T}\boldsymbol{x_t} = \hat{\boldsymbol{\mu}}.$$

Although we will use maximum likelihood and not OLS to estimate the parameters, we still want to avoid the situation that estimated long-term means μ are significantly different from the sample averages of the state variables. Therefore, we not only impose restrictions on Γ, but we also restrict the estimates for μ to be equal to the sample averages of the state variables.

Hence, we have restrictions on both $\boldsymbol{\mu}$ and Γ. We will estimate the free parameters of the VAR model by means of maximum likelihood estimation. Recall that the error terms

[3]Unrestricted VAR models form a subclass of Seemingly Unrelated Regression (SUR) models. A SUR model is a model that consists of a number of classical regression equations, with regressors and coefficients that vary across the different equations. The error terms are contemporaneously correlated, but have no intertemporal correlation. Line-by-line OLS is efficient (see Zellner (1962)) if the regressors are exactly the same across the different regression equations, which is the case for a VAR model.

$\Sigma\boldsymbol{\epsilon}_t$ $(= \boldsymbol{x}_t - \boldsymbol{\mu} - \Gamma(\boldsymbol{x}_{t-1} - \boldsymbol{\mu}))$ are i.i.d. and follow a multivariate normal distribution. The likelihood function is given by

$$\mathcal{L}(\boldsymbol{\mu}, \Gamma, \Sigma | \boldsymbol{x}_0, \ldots, \boldsymbol{x}_T) = \prod_{t=1}^{T} f(\boldsymbol{x}_{t-1}, \boldsymbol{x}_t | \boldsymbol{\mu}, \Gamma, \Sigma) \tag{37}$$

where

$$\begin{aligned} & f(\boldsymbol{x}_{t-1}, \boldsymbol{x}_t | \boldsymbol{\mu}, \Gamma, \Sigma) \\ &= \frac{1}{\sqrt{(2\pi)^N |\Sigma\Sigma'|}} \exp\left(-\frac{1}{2} (\boldsymbol{x}_t - \boldsymbol{\mu} - \Gamma(\boldsymbol{x}_{t-1} - \boldsymbol{\mu}))' (\Sigma\Sigma')^{-1} (\boldsymbol{x}_t - \boldsymbol{\mu} - \Gamma(\boldsymbol{x}_{t-1} - \boldsymbol{\mu})) \right). \end{aligned} \tag{38}$$

For $\boldsymbol{\mu}$, we plug in the restricted value $\hat{\boldsymbol{\mu}} = \frac{1}{T+1} \sum_{t=0}^{T} \boldsymbol{x}_t$. The free parameters in Γ and Σ maximize the likelihood function $\mathcal{L}(\hat{\boldsymbol{\mu}}, \Gamma, \Sigma | \boldsymbol{x}_0, \ldots, \boldsymbol{x}_T)$. Multiplication by a constant and transformation by an increasing function does not change the location of the maximum, which makes it more convenient to maximize $\frac{1}{T} \log(\mathcal{L}(\hat{\boldsymbol{\mu}}, \Gamma, \Sigma | \boldsymbol{x}_0, \ldots, \boldsymbol{x}_T))$. We will rely on numerical optimization techniques to find the argument that maximizes this expression.

We also determine the standard errors for the parameter estimates in the VAR model. In Appendix D, we describe the method for finding these standard errors.

3.3.2 Estimation of price-of-risk parameters

After we have estimated the VAR parameters at the first stage of the estimation procedure, we estimate the price-of-risk parameters $\boldsymbol{\Lambda}_0$ and Λ_1 in the second stage. We take the parameters that we have found at the first stage as given. We consider data on N_1 nominal yields and N_2 real yields over time and find $\boldsymbol{\Lambda}_0$ and Λ_1 by matching model yields to observed yields. The model yields follow from (25): $\hat{y}_t^{(n)} = \frac{A(n)}{n} + \frac{\boldsymbol{B}(n)}{n}' \boldsymbol{x}_t$. Since we have already estimates for $\boldsymbol{\mu}, \Gamma$, and Σ, only $\boldsymbol{\Lambda}_0$ and Λ_1 are left as unknown parameters for the functions $A(n)$ and $\boldsymbol{B}(n)$. We assume that all yields that are not included as factors are measured with an observation error that has mean zero.

An important remark is that we cannot simply choose any combination $\{\boldsymbol{\Lambda}_0, \Lambda_1\}$ that fits best. Note that this model specification leads to two expressions for the 10-year yield $y_t^{(10)}$ in the four-factor model. First of all, this yield follows directly from the state vector, as it is the sum of the first two components of the state vector: $y_t^{(10)} = (\boldsymbol{e}_1 + \boldsymbol{e}_2)' \boldsymbol{x}_t$. Second, from the term structure formula that followed from no-arbitrage restrictions, we have $y_t^{(10)} = \frac{A(10)}{10} + \frac{\boldsymbol{B}(10)}{10}' \boldsymbol{x}_t$. Hence, we need to choose $\boldsymbol{\Lambda}_0$ and Λ_1 in the second stage such that the two expressions for the 10-year yield that we have are equal to each other: $A(10) = 0$ and $\boldsymbol{B}(10) = 10(\boldsymbol{e}_1 + \boldsymbol{e}_2)$. We only impose these restrictions for the four-factor model. In the three-factor model, the 10-year yield does not follow directly from the state vector, and there we assume that it is measured with an observation error like the other yields.

Remember that we have derived bond yields such that bond returns do not allow for arbitrage opportunities. Stock returns, however, follow directly from the VAR model. If $\boldsymbol{\Lambda}_0$ and Λ_1 can take arbitrary values, it could be the case that expected stock returns do not match to the compensation that is obtained from the exposure of stock returns to the risk factors as prescribed by the price-of-risk parameters. To find the restrictions on $\boldsymbol{\Lambda}_0$ and Λ_1 as a result of including stock returns in the state vector, we use the no-arbitrage pricing equation $\boldsymbol{P}_t = \mathbb{E}_t[M_{t+1}\boldsymbol{P}_{t+1}]$. Let $\boldsymbol{e}_s$ be the unit vector that corresponds to the position of the excess stock return in the state vector. We obtain

$$\begin{aligned} \boldsymbol{e}_s'\Sigma\boldsymbol{\Lambda}_0 &= \boldsymbol{e}_s'(I-\Gamma)\boldsymbol{\mu} + \frac{1}{2}\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{e}_s \\ \boldsymbol{e}_s'\Sigma\Lambda_1 &= \boldsymbol{e}_s'\Gamma. \end{aligned} \tag{39}$$

The derivation is given in Appendix E. Note that we have for both the three-factor model and the four-factor model imposed the restriction that $\boldsymbol{e}_s'\Gamma$ is equal to zero (no stock return predictability). In combination with the no-arbitrage restriction, this implies that $\boldsymbol{e}_s'\Sigma\Lambda_1 = \mathbf{0}$.

As a final constraint, we include the constraint that all eigenvalues of the matrix $\Gamma - \Sigma\Lambda_1$ should be smaller than one in absolute value. This last constraint guarantees that yields with a very long maturity will not 'explode'.

Let T_n be the number of observations for the nominal yield with maturity n, and T_n^R the number of observations for the real yield with maturity n. We solve for $\boldsymbol{\Lambda}_0$ and Λ_1 by minimizing the sum of squared fitting errors of the model, scaled by the number of observations per maturity:

$$\begin{aligned} \min \quad & \sum_{n=1}^{N_1}\frac{1}{T_n}\sum_{t=T-T_n+1}^{T}\left(\hat{y}_t^{(n)} - y_t^{(n)}\right)^2 + \sum_{n=1}^{N_2}\frac{1}{T_n^R}\sum_{t=T-T_n^R+1}^{T}\left(\hat{y}_t^{R(n)} - y_t^{R(n)}\right)^2 \\ \text{subject to} \quad & \{\boldsymbol{\Lambda}_0, \Lambda_1\} \in \Lambda, \end{aligned} \tag{40}$$

where Λ denotes the space of possible parameter values $\{\boldsymbol{\Lambda}_0, \Lambda_1\}$ that satisfy all constraints that are described above. We again rely on numerical optimization techniques to find the optimum to (40).

4 Estimation results

We will now estimate the parameters of both the three-factor and the four-factor model that we have described in Section 3.2, by applying the estimation procedure that we have discussed in Section 3.3.

4.1 Data

To estimate the parameters for our specifications of the financial markets, we use yearly data on nominal bonds, real bonds, excess stock returns, and inflation. We are particularly interested in implications for European pension savers. Since the Deutsche Mark is seen as the most important predecessor of the euro, we use German data. For the first stage of the estimation procedure, we need data on the nominal short rate, the 10-year nominal yield (only for the four-factor model), inflation, and stock returns.

The nominal yield data are taken from the Deutsche Bundesbank, the central bank of the Federal Republic of Germany. We use the term structure of interest rates on listed Federal securities. Nominal 1-year rates are available from December 1972 onwards. The same holds for 10-year rates. For realized inflation and stock returns, we obtain data from Datastream. To calculate price inflation, we use the German Consumer Price Index (unadjusted CPI). Finally, for stock returns, we assume that the investor invests in the MSCI Europe Index. This index is quoted in dollars, but obviously the investor is interested in returns in the local currency (Deutsche Mark/euro). Therefore, we use exchange rates to calculate the realized returns of the position in stocks (not hedged for currency risk). Since the excess stock return is defined as the stock return minus the 1-year rate one year earlier, we have data on the state vector from December 1973 to December 2012, a period of 40 years.

In the second stage of the estimation procedure, we match model yields to observed yields. For this purpose, we include nominal and real yields with maturities of 1-year, 2-year, 5-year, 8-year, 10-year, 15-year, 20-year, and 30-year in our sample. Recall that for the three-factor model, we considered the 1-year rate to be measured without observation error. For the four-factor model, this holds for the 1-year rate and the 10-year rate, which both follow directly from the state variable.

All nominal yields are taken from the Deutsche Bundesbank. All yields with maturities up to 10 years are available from the beginning of our sample period (1973). The 15-year and 20-year nominal bonds have been traded since 1986, the 30-year bond is quoted since 2000. Real interest rates could be obtained from inflation-linked bonds. However, due to the limited availability of German inflation-linked bonds, there are no zero-coupon yields on inflation-linked bonds available. Another possibility to determine real interest rates is by taking inflation swap rates as a measure for break-even inflation (the difference between nominal and real yields). As Haubrich, Pennacchi, and Ritchken (2012) show, the difference between the nominal yield and the zero-coupon swap rate is a valid measure for the corresponding real yield. Since no German inflation swaps are available, we use European inflation swaps as a proxy. These swap rates are available as of the end of 2004. We determine the real rates by subtracting the log (zero-coupon) swap rate from the nominal log rates.

Table 3 presents statistics for the state variables in our dataset. At the end of 2011 and 2012, the nominal 1-year yield was slightly negative. This goes against some influential term structure models in the literature that only allow for positive interest rates, see e.g. Cox, Ingersoll and Ross (1985). In our model specification, negative interest rates are not excluded.

The average term structure spread has been positive, which means that the term structure is normally upward sloping. There have been seven periods in which the term spread was negative, the last one was in 1992. With the exception of 1986, yearly inflation has always been positive.

Table 3: Sample statistics for the yearly data (in %).

Variable	Sample period	Average	Std. dev.	Minimum	Maximum
$y_t^{(1)}$	1973-2012	4.74	2.65	-0.04	9.92
$y_t^{(10)} - y_t^{(1)}$	1973-2012	1.22	1.11	-1.01	2.86
π_t	1973-2012	2.60	1.78	-1.01	7.45
$r_{s,t} - y_{t-1}^{(1)}$	1973-2012	3.00	22.65	-60.70	36.53
$y_t^{(2)}$	1973-2012	4.94	2.52	-0.04	9.63
$y_t^{(5)}$	1973-2012	5.49	2.31	0.39	9.34
$y_t^{(8)}$	1973-2012	5.82	2.15	1.00	9.42
$y_t^{(10)}$	1973-2012	5.96	2.06	1.37	9.57
$y_t^{(15)}$	1986-2012	5.33	1.65	2.00	8.70
$y_t^{(20)}$	1986-2012	5.47	1.61	2.24	8.71
$y_t^{(30)}$	2000-2012	4.17	1.08	2.18	5.62
$y_t^{R(1)}$	2004-2012	0.00	1.38	-1.90	1.53
$y_t^{R(2)}$	2004-2012	0.13	1.33	-1.75	1.65
$y_t^{R(5)}$	2004-2012	0.49	1.10	-1.34	1.75
$y_t^{R(8)}$	2004-2012	0.83	0.92	-0.85	1.89
$y_t^{R(10)}$	2004-2012	1.01	0.85	-0.60	1.97
$y_t^{R(15)}$	2004-2012	1.30	0.74	-0.12	2.12
$y_t^{R(20)}$	2004-2012	1.40	0.73	0.01	2.21
$y_t^{R(30)}$	2004-2012	1.31	0.86	-0.17	2.17

4.2 Parameters three-factor model

Table 4 presents our estimates for the three-factor model. The estimated long-term means $\boldsymbol{\mu}$ are equal to the full-sample means, which was the restriction that we imposed upon estimation. The 1-year yield and inflation are found to be significantly dependent on their own past, where in particular the short-term yield shows high persistency. Only the diagonal values of Γ are statistically significantly different from zero. Most of the variance in both state variables is explained, even though the estimation technique (maximum likelihood) does not necessarily focus on maximizing the model's R^2.

Table 4: Parameter estimates (standard errors in brackets) and model fit for the three-factor model.

Variable	$\boldsymbol{\mu}$ (in %)	Γ			Σ (in %)			R^2	$\boldsymbol{\Lambda}_0$
$y_t^{(1)}$	4.74	0.93	-0.14	0	1.32	0	0	0.72	-0.19
	(0.42)	(0.13)	(0.19)		(0.14)				
π_t	2.60	0.14	0.55	0	0.49	0.90	0	0.59	-0.21
	(0.28)	(0.10)	(0.14)		(0.15)	(0.10)			
$r_{s,t}$	3.00	0	0	0	2.85	1.41	21.73		0.29
$-y_{t-1}^{(1)}$	(3.58)				(4.09)	(3.65)	(2.46)		

The volatility of the shock to the nominal short rate is 1.32%, which is considerably smaller than the unconditional volatility of the realized nominal short rate in the sample (2.65%). The annual volatility of the orthogonal component corresponding to inflation shocks is equal to 0.90%, which is also smaller than the realized sample standard deviation of 1.78%. Inflation has a significant positive exposure to short rate shocks. The loading of stock returns on the orthogonal component is by far the largest, and accounts for almost the whole volatility (21.73% versus 22.65%). The exposures of stock returns to the other risk factors are insignificant.

The estimated risk premia for the orthogonal components of short rate shocks and inflation shocks are negative. Table 5 shows the exposure of nominal and real bonds to the risk factors in $\boldsymbol{\epsilon}_t$ and gives the resulting risk premia on nominal and real bonds with maturities of 5 and 10 years. Nominal bond returns depend negatively on shocks to the nominal short rate, but positively on inflation shocks. For nominal bond returns, the negative dependency prevails in determining the risk premium, which is positive. Real bonds are indexed with the price level and therefore have a nominal return that has a strong positive dependence on inflation. Due to the relatively short horizon, the 5-year real bond has relatively little exposure to the nominal short rate and has a negative risk premium. The risk premium on the 10-year real bond is slightly positive due to the increased (negative) exposure to the nominal short rate.

Table 5: Bond return characteristics for the three-factor model.

Maturity	Exposures ϵ_1	ϵ_2	ϵ_3	Risk premium	Volatility
5-year nominal bond	-4.37%	0.52%	0	0.74%	4.40%
10-year nominal bond	-7.27%	1.42%	0	1.11%	7.41%
5-year real bond	-2.32%	2.30%	0	-0.04%	3.27%
10-year real bond	-4.13%	3.01%	0	0.16%	5.11%

Table 6 illustrates the resulting correlations between asset returns. These correlations are important, since they form the basis of hedging demands of the investor that wants to be hedged against adverse shocks in future investment opportunities. The correlation between bond returns and stock returns is negative, but small in absolute value. The correlation between the return on a 10-year nominal bond and its real counterpart is very high: more than 90%. This is a result of the availability of real yields in the data. We only have data on real yields for the past decade, in which nominal yields and real yields have moved simultaneously. However, it is important to remark that nominal yields and real yields do not necessarily move closely at periods in which inflation rates are less stable.

We see that the impact of inflation risk on nominal bonds is very small. Because all variation in bond returns is in this model explained by the exposure to risk in the short rate and inflation risk, this means that nominal bond returns are almost perfectly (negatively) correlated to shocks to the short rate. The correlation between next period's short rate and the nominal return on a 10-year real bond is about -80%.

Table 6: Correlations between asset returns in the three-factor model.

	Stock return	10-yr nom. bond	10-yr real bond
Stock return	1		
10-year nominal bond return	-0.115	1	
10-year real bond return	-0.067	0.906	1
Short rate	0.130	-0.981	-0.808

Finally, we compare the fit of the modeled yields to the observed yields in our data set. We do this by comparing the average yields and volatility of the yields from our model to the average and standard deviation of the observed yields in our dataset. We follow Sangvinatsos and Wachter (2005) and Koijen et al. (2010) in performing this comparison. We simulate 10,000 sample paths of our state variables for a period of 40 years, which is equal to the sample period of our data. At each point in time and for all sample paths, we calculate the yields for a number of maturities. Since only nominal yields with a maturity up to 10 years have been observed for the whole sample period, we focus on these maturities. Afterwards, we calculate for all maturities and for all sample paths the realized average yield and realized volatility over the sample path. From these 10,000 Monte Carlo observations, we determine the average and a 95% confidence interval. We also compute the historical realized average yields and the volatility of the yields. The results are presented in Figure 3 and Figure 4.

Figure 3: Average yields, historically and in the three-factor model.

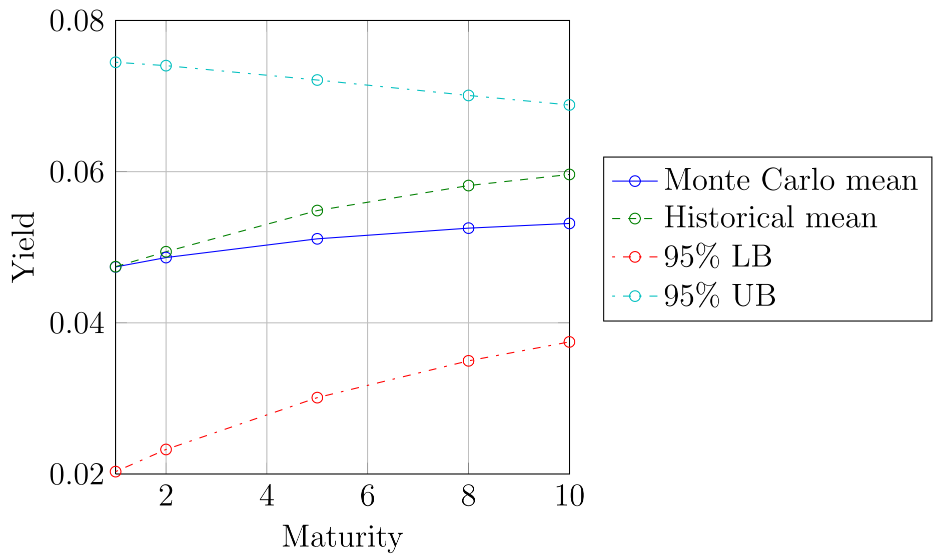

We see that our restricted three-factor model fails to replicate the shape of the average yield in the data as a function of the maturity. While the average yields from the model are almost flat as a function of the maturity, the historical mean is clearly upward sloping. The confidence interval has a large width, which reflect the well-known result that means of return series are generally hard to estimate precisely. For that reason, the observed means are comfortably in the 95% confidence interval. Standard deviations are known to be estimated more precisely, which is confirmed by the smaller width of the confidence interval for the standard deviations of the yields. For maturities above 8 years, the historically observed volatility of the yields exceeds the upper bound of the 95% confidence interval. This gives us the indication that the simple three-factor model is not able to match the cross-sectional moments of the data. We try to improve the model fit by looking at the four-factor model with time-varying risk premia that was described in Section 3.2.2.

Figure 4: Volatility of yields, historically and in the three-factor model.

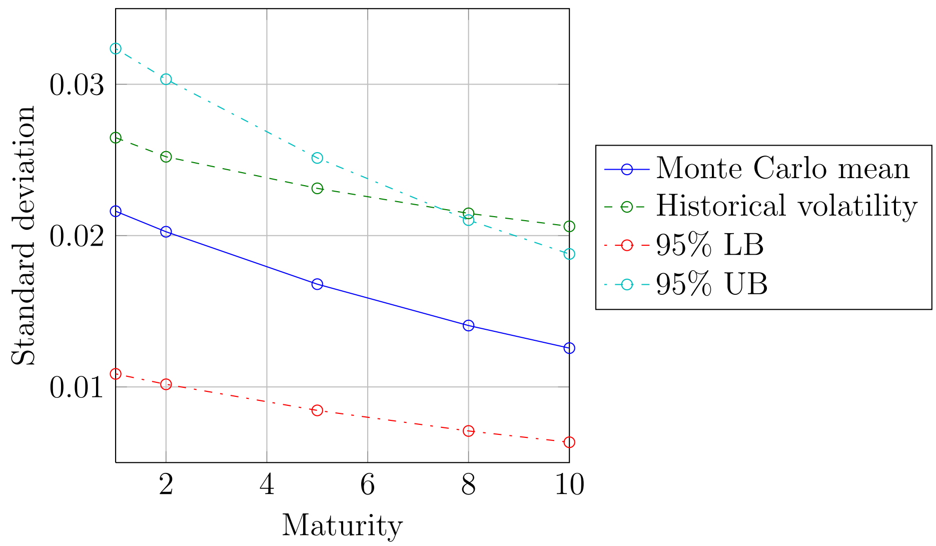

4.3 Parameters four-factor model

By including an extra state variable (term spread) and allowing for time-varying risk premia, we expect to improve the fit of the model to historically observed yields. Table 7 presents our parameter estimates for the four-factor model. Obviously, the restricted values of $\boldsymbol{\mu}$ are the same as before. Again, only the diagonal elements of Γ are statistically significantly different from zero. The nominal short rate is found to be somewhat less persistent when the term spread is included as a state variable, with a cumulative coefficient of 0.87 on its lagged value. The newly introduced state variable "term spread" has a coefficient of 0.58 on its own past.

We find that the term spread has significant negative exposure to shocks to the nominal short rate. We indeed would expect a negative relation, since long-term rates are expected to be less sensitive to shocks in the short rate than the short rate itself. Recall that according to the Expectations Hypothesis, long-term rates are the average of future expected short rates. All risk in equity returns is still almost exclusively determined by the orthogonal equity shock.

In Table 8, the exposures of bond returns to the four shocks to the economy are given. They have a similar magnitude as those of the three-factor model. Note that although stock returns do not predict future values of the state variables in our restricted model, they do influence future prices of risk. Therefore, bond returns have a small exposure to stock returns.

The prices of risk in $\boldsymbol{\lambda}_t$ depend in this model strongly on the state of the world, as given by the state vector $\boldsymbol{x}_t$. First of all, we set the state vector $\boldsymbol{x}_t$ equal to its unconditional expectation $\boldsymbol{\mu}$. In Table 8, the resulting risk premia on long-term bonds are denoted. When the state vector is equal to its long-term mean, the risk premia on the four included bonds

Table 7: Parameter estimates (standard errors in brackets) and model fit for the four-factor model

Variable	$\boldsymbol{\mu}$ (in %)	Γ				Σ (in %)				R^2
$y_t^{(1)}$	4.74	0.96	0.09	-0.15	0	1.32	0	0	0	0.72
	(0.42)	(0.17)	(0.30)	(0.19)		(0.14)				
$y_t^{(10)} - y_t^{(1)}$	1.22	-0.06	0.58	0.11	0	-0.67	0.57	0	0	0.33
	(0.17)	(0.10)	(0.19)	(0.12)		(0.12)	(0.06)			
π_t	2.60	0.12	-0.07	0.55	0	0.49	-0.14	0.88	0	0.59
	(0.28)	(0.12)	(0.22)	(0.14)		(0.15)	(0.15)	(0.10)		
$r_{s,t} - y_{t-1}^{(1)}$	3.00	0	0	0	0	2.56	-2.79	1.47	21.59	
	(3.58)					(4.22)	(3.66)	(3.74)	(2.48)	
Variable	$\boldsymbol{\Lambda}_0$	Λ_1								
$y_t^{(1)}$	0.42	-3.82	-31.32	-10.10	-0.20					
$y_t^{(10)} - y_t^{(1)}$	1.15	-13.11	-46.57	4.69	0.25					
π_t	-0.22	6.64	0.05	6.08	-1.51					
$r_{s,t} - y_{t-1}^{(1)}$	0.36	-1.69	-2.30	1.39	0.16					

are all positive. The Sharpe ratios (risk premium divided by volatility) for the 5-year bonds are somewhat higher than those of the 10-year bonds. Real bonds have a slightly higher Sharpe ratio than nominal bonds.

Table 8: Bond return characteristics for the four-factor model.

Maturity	Exposures ϵ_1	ϵ_2	ϵ_3	ϵ_4	Risk premium	Volatility
5-year nom. bond	-3.73%	-1.33%	0.03%	-0.09%	1.41%	3.96%
10-year nom. bond	-6.11%	-4.91%	0.01%	-0.02%	2.11%	7.84%
5-year real bond	-2.28%	-1.56%	1.95%	0.56%	1.39%	3.43%
10-year real bond	-4.26%	-4.72%	2.03%	0.69%	2.00%	6.71%

It is difficult to interpret the estimates of the price-of-risk parameters $\mathbf{\Lambda}_0$ and Λ_1 directly. Therefore, we consider the exposures of bond risk premia to the value of the state vector. The risk premium on an n-year nominal bond is defined by $-\boldsymbol{B}(n-1)'\Sigma(\mathbf{\Lambda}_0 + \Lambda_1 \boldsymbol{x}_t)$. In Table 9, we present the coefficients on the state vector, $-\boldsymbol{B}(n-1)'\Sigma\Lambda_1$, for two different nominal bonds and two different real bonds.

Table 9: Exposures of bond risk premia to the elements in the state vector.

Maturity	$y_t^{(1)}$	$y_t^{(10)} - y_t^{(1)}$	π_t	$r_{s,t} - y_{t-1}^{(1)}$
5-year nom. bond	0.32	1.79	0.31	0.00
10-year nom. bond	0.88	4.20	0.39	0.00
5-year real bond	0.41	1.43	0.28	-0.03
10-year real bond	0.90	3.52	0.34	-0.03

We see that bond risk premia depend positively on the first three state variables. Risk premia on nominal bonds show almost no dependence on past stock returns, wile risk premia on real bonds show a small negative exposure to equity returns. To illustrate the time variation in bond risk premia, we show the effects of individual changes in the state variables while keeping all other variables equal to their unconditional means. We express these changes in terms of the unconditional standard deviations of the state variables. The results for the risk premium on a nominal 10-year bond are presented in Figure 5. Since past stock returns have no impact on the risk premium on a 10-year nominal bond, we only included changes in the first three state variables.

In the figure, we see that changes in the state variables have a considerable impact on

Figure 5: Variation in the risk-premium on a 10-year nominal bond.

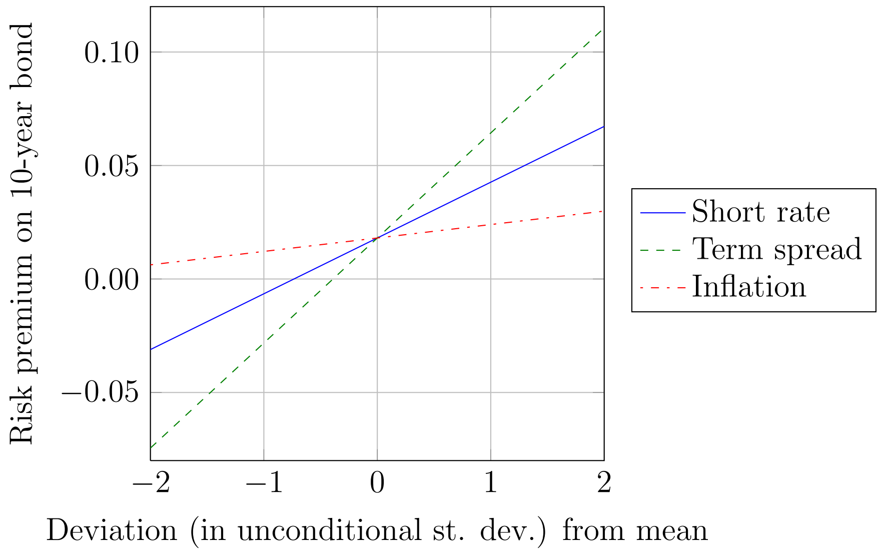

the expected excess return on bonds. This means that investment opportunities are largely affected by the state of the economy. The long-term investor should take this into account in his allocation decision. The order of magnitude of the deviations in risk premia is consistent with Sangvinatsos and Wachter (2005) and Koijen et al. (2010). Changes in the term spread have the biggest impact on bond risk premia, but also the value of the short rate has quite some impact on bond risk premia. Expected excess returns on bonds are not very sensitive to changes in the level of inflation.

As noted before, the correlations between asset returns are an important determinant of the hedging demand of a risk-averse investor. The model-implied correlations are similar to those that we found for the three-factor model. They are given in Table 10. The correlation between bond returns and stock returns is still very close to zero. Nominal bond returns and real bond returns are almost perfectly correlated. The relation between shocks to the nominal short rate and bond returns is still negative and very strong.

Since bond returns are now predictable, investment opportunities are heavily dependent on the specific state. This brings in new risks for the long-term investor. A risk-averse agent wants to be hedged against sudden adverse changes in investment opportunities. As we have seen for the three-factor model, bonds can be used to hedge the risk of disappointing future short yields. In the table, we see that there is also a negative correlation between bond returns and bond risk premia. This holds for both nominal bonds and real bonds. This means that bonds can not only be used to hedge against risk in short rates, but can also be used to hedge the risk of falling bond return expectations in excess of the short rate. This will play an important aspect in the optimal allocation for a risk-averse agent, because it increases the hedging demand for bonds.

Finally, we again compare the yields from our model to the yields in our dataset. As we

Table 10: Correlations between asset returns and risk premia in the four-factor model.

	Stock return	10-yr nom. bond	10-yr real bond
Stock return	1		
10-year nominal bond return	-0.014	1	
10-year real bond return	0.004	0.936	1
Short rate next period	0.117	-0.780	-0.635
Risk premium 10-yr nom. bond	0.160	-0.116	-0.220
Risk premium 10-yr real bond	0.438	-0.182	-0.309

did for the three-factor model, we create a confidence interval for the time-averaged yields and the volatility of yields generated by the model on the basis of 10,000 simulated sample paths. We simulate these paths of our state variables over a period of 40 years, which is equal to the sample period of our data. At each point in time and for all sample paths, we calculate the yields for a number of maturities. From the 10,000 Monte Carlo observations of the time-averaged yields, we determine the average and a 95% confidence interval. We also compute the historical realized average yields and volatility of the yields. The results are presented in Figure 6 and Figure 7.

Figure 6: Average yields, historically and in the four-factor model.

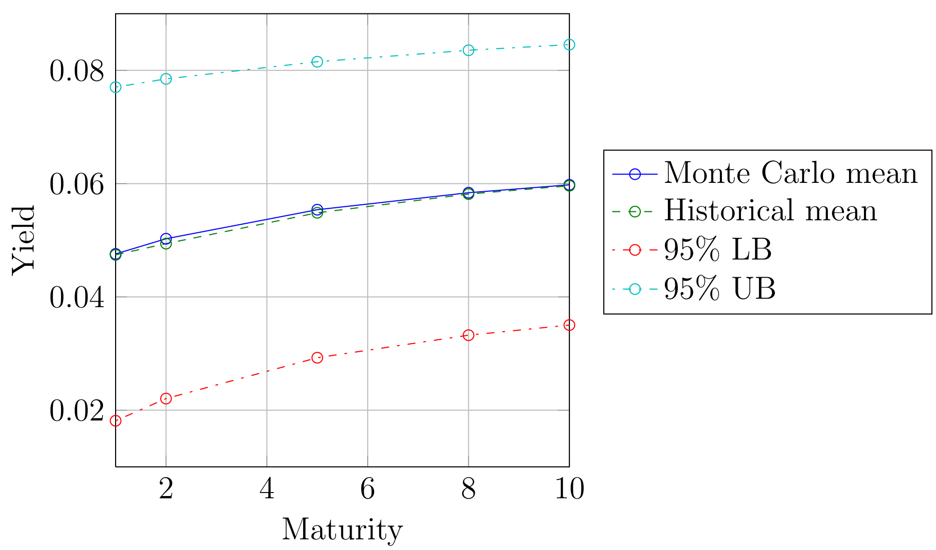

We see that our four-factor model matches the observed average yields in the data

Figure 7: Volatility of yields, historically and in the four-factor model.

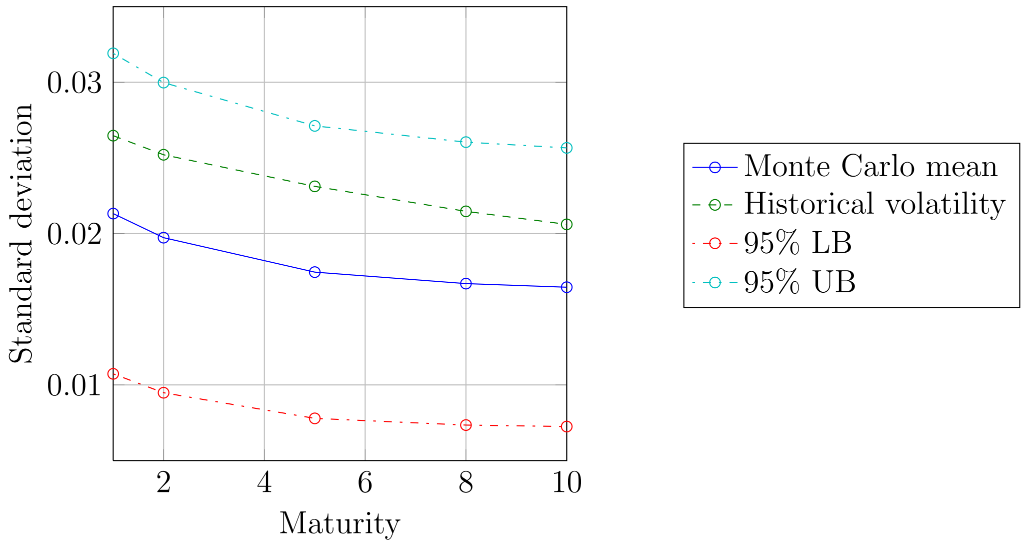

almost perfectly. Hence, the observed means are comfortably in the 95% confidence interval. Further, the standard deviation of the yields that we computed directly from the data lies comfortably in the 95% confidence interval that we obtained from the Monte Carlo simulations. We can conclude that the model is successful in matching the cross-sectional moments of the yields to the data. As we have seen for the three-factor model, this is not trivial, since there are numerous cross-sectional as well as time-series moments. We conclude that this four-factor model provides a reasonably good fit to the data, in contrast to the simple three-factor model without predictability.

5 Life-cycle investing

Our goal is to find optimal asset allocations for a defined contribution pension plan, by adopting either one of the models for the financial market that were introduced in the previous section. As we have seen before, the value of future contributions to the pension scheme has an important impact on the optimal allocation of financial wealth. Therefore, we are particularly interested in how stochastic interest rates influence human capital and consequently the optimal allocation decision based on human and financial capital. Before we introduce the life-cycle optimization problem, we will first introduce our model for labor income and derive a no-arbitrage valuation method for human capital.

5.1 Human capital

In Section 2.3, we have already seen a model for future labor income contributions to the pension scheme. This model allowed for labor income exposure to stock returns. Since equity risk is traded on the financial market, this is called systematic labor income risk. Here, we will describe the somewhat more general model for labor income that we use for solving the life-cycle optimization problem. This model also allows for idiosyncratic labor income risk. Further, we will describe how we value human capital in a market-consistent way by using the pricing kernel that we have specified.

We assume that labor income is paid at the beginning of each period, from the beginning of the working life at time $t = 0$ until time $T - 1$, the last year before retirement. As before in Section 2.3, let L_t be the income at time t, and let ρ_t be the fraction of this income that is contributed to the pension scheme. Labor income L_t is modeled as:

$$\begin{cases} L_0 = 1 \\ L_{t+1} = L_t \exp(g_{t+1} + \pi_{t+1} + u_{t+1}) \quad \text{for } t = 0, \ldots, T-2 \\ L_T = 0 \end{cases}, \tag{41}$$

where g_t is a deterministic growth factor and u_t is a labor income shock: $u_t \overset{i.i.d.}{\sim} N(0, \sigma_u^2)$. Since T is the retirement date, the individual receives his last income at time $T - 1$. An explicit expression for L_t is given by

$$L_t = \Pi_t \exp(G_t + \nu_t), \tag{42}$$

where $\Pi_0 = 1$, $G_t = \sum_{s=1}^{t} g_s$, and $\nu_t = \sum_{s=1}^{t} u_s$.

We allow for correlation between stock returns and labor income. Let $\eta_t \overset{i.i.d.}{\sim} N(0, 1)$ be independent of the shocks $\boldsymbol{\epsilon}_t$. We define u_t as follows:

$$u_t = k \cdot \boldsymbol{e}_s' \Sigma \boldsymbol{\epsilon}_t + \sqrt{\sigma_u^2 - k^2 \boldsymbol{e}_s' \Sigma \Sigma' \boldsymbol{e}_s} \cdot \eta_t. \tag{43}$$

In this way, the variance of u_t is indeed equal to σ_u^2. As before, k is the exposure of labor income to stock returns, describing the systematic labor income risk. In other words, it determines the extent to which labor income is stock-like. The idiosyncratic risk factor in this labor income model is η_t. Since labor income is generally less volatile than stock returns, we expect k to be smaller than one (in absolute value). In the literature, several authors try to specify the correlation between labor income and stock returns. Define this correlation by ρ. We have

$$\rho = \frac{\mathrm{Cov}_t(\pi_{t+1} + u_{t+1}, r_{s,t+1})}{\sqrt{\mathrm{Var}_t(\pi_{t+1} + u_{t+1})}\sqrt{\mathrm{Var}_t(r_{s,t+1})}} = \frac{(\boldsymbol{\delta}_\pi + k\boldsymbol{e}_s)'\Sigma\Sigma'\boldsymbol{e}_s}{\sqrt{\boldsymbol{\delta}_\pi'\Sigma\Sigma'\boldsymbol{\delta}_\pi + 2k\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{\delta}_\pi + \sigma_u^2}\sqrt{\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{e}_s}} \tag{44}$$

We want to determine the value of future contributions, human capital, in a market-consistent way. As discussed before, we can use the stochastic discount factor to find an arbitrage-free value for any payoff on the financial market. When labor income is deterministic in real terms ($\sigma_u = 0$), implying that it grows with the deterministic growth rate g_t and with price inflation, it has the same payoff as an inflation-linked bond. This means that labor income can be perfectly replicated on the financial market. By using the real term structure of interest rates, we can determine the value of human capital. If labor income is stochastic but there is only systematic labor income risk, the loading on the idiosyncratic risk factor η_t is equal to zero ($\sigma_u^2 = k^2\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{e}_s$). In that case, labor income can still be perfectly replicated by payoffs on the financial market. By imposing no-arbitrage restrictions, we can determine the value of human capital on the basis of the pricing kernel in our model.

However, when $\sigma_u^2 - k^2\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{e}_s > 0$, labor income cannot be replicated on the financial market due to its exposure to the non-traded income risk η_t. In that case, we are in the situation of incomplete markets. We can no longer apply no-arbitrage restrictions to find a unique price for human capital. No-arbitrage restrictions only lead to a range of prices. De Jong (2008) provides an overview of valuation methods for pension liabilities that depend on non-traded wage inflation. For simplicity, we assume that idiosyncratic labor income risk has a price of zero. In that case, we can still use the stochastic discount factor to value human capital.

The derivation of the discount factor for future contributions is similar to the derivation of nominal and real yields in our model. When labor income is correlated to stock returns, this discount factor includes an equity risk premium. Let $H_t^{(n)}$ be the value at time t of the contribution $\rho_{t+n}L_{t+n}$ at time $t+n$. In Appendix F, $H_t^{(n)}$ is shown to satisfy:

$$H_t^{(n)} = \rho_{t+n}\Pi_t \exp(G_{t+n} + \nu_t - A^L(n) - \boldsymbol{B}^L(n)'\boldsymbol{x}_t), \tag{45}$$

where

$$\begin{aligned} A^L(n) &= A^L(n-1) + (\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi)'(I - \Gamma)\boldsymbol{\mu} - (\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\boldsymbol{\Lambda}_0 + k\boldsymbol{e}_s'\boldsymbol{\Lambda}_0 \\ &\quad - \frac{1}{2}(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\Sigma'(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi) - \frac{1}{2}\sigma_u^2 + k\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{B}^L(n-1) \\ \boldsymbol{B}^L(n) &= (\Gamma - \Sigma\Lambda_1)'(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi) + k\Lambda_1'\Sigma'\boldsymbol{e}_s + \boldsymbol{\delta}_y. \end{aligned} \tag{46}$$

The initial conditions are given by

$$\begin{aligned} A^L(1) &= A^R(1) + k\boldsymbol{e}_s'\Sigma\boldsymbol{\Lambda}_0 - \frac{1}{2}\sigma_u^2 \\ \boldsymbol{B}^L(1) &= \boldsymbol{B}^R(1) + k\Lambda_1'\Sigma'\boldsymbol{e}_s. \end{aligned} \tag{47}$$

5.2 Life-cycle optimization problem

We will now discuss the optimization problem of a life-cycle investor. We generalize the problem that was discussed in Section 2.3 by including annuities and by focusing on real pension benefits instead of nominal benefits.

As before in Section 2.3, the investor has an initial total wealth assigned to his pension provision of $W_0 = F_0 + H_0$. The initial financial wealth F_0 equals the first contribution to the pension scheme of $\rho_0 L_0$. The initial value of the human capital that is destined for his pension provision can be obtained by discounting future contributions ($\{\rho_n L_n\}_{n=1,\ldots,T}$) by the appropriate discount rate at time $t = 0$:

$$H_0 = \sum_{n=1}^{T} \rho_n \exp(G_n - A^L(n) - \boldsymbol{B}^L(n)'\boldsymbol{x}_0). \tag{48}$$

Again, we assume that the investor has the power utility function (4). At time T, the investor annuitizes all of his wealth. The price of an annuity that pays 1 unit in real terms every year the individual is still alive is denoted by A_T. This means that at time T, this person can buy $\frac{W_T}{A_T}$ of these annuities, where W_T is defined as the terminal wealth in real terms. The individual wants to maximize his expected utility over the real pension benefits $\frac{W_T}{A_T}$.

Note that wealth consists of both financial wealth and human capital. As we have seen before, the allocation of human capital cannot be chosen. It follows directly from the nature of labor income. For instance, when labor income is deterministic in real terms, the allocation of human capital consists of a portfolio of real bonds. We can choose the allocation of financial wealth ourselves, but only under some restrictions. In order to be able to model these restrictions explicitly, we define $\boldsymbol{\alpha_t}$ to be the vector of allocations of financial wealth to the different assets that are available on the financial market. We impose the restrictions

$$\boldsymbol{\alpha_t} \geq \boldsymbol{0}, \qquad \boldsymbol{\iota}'\boldsymbol{\alpha_t} \leq 1, \tag{49}$$

where ι is a vector of ones. These constraints imply that no borrowing and no short selling are allowed.

Hence, we have the following optimization problem:

$$
\begin{aligned}
\max \quad & \mathbb{E}\left[\frac{1}{1-\gamma}\left(\frac{W_T}{A_T}\right)^{1-\gamma}\right] \\
\text{subject to} \quad & W_{t+1} = W_t(1+R_{p,t+1})\frac{\Pi_t}{\Pi_{t+1}},\ t=0,\dots,T-1 \\
& \boldsymbol{\alpha}_{T-1} \geq \mathbf{0} \\
& \boldsymbol{\iota}'\boldsymbol{\alpha}_{T-1} \leq 1.
\end{aligned} \tag{50}
$$

Note that $R_{p,t+1}$ is the nominal return on total wealth, which is the sum of financial wealth and human capital. Since the life-cycle investor wants to maximizes his expected utility over real pension benefits, wealth is expressed in real terms. By dividing by the realized inflation $\frac{\Pi_{t+1}}{\Pi_t}$, the nominal return on the portfolio is transformed into a real return.

There is a distinction between the asset menu for financial wealth and the asset menu for human capital. We assume that in each period, the individual can invest his financial wealth in two different risky assets: a stock index and a nominal bond with a maturity of 10 years. We denote the (nominal) returns on these assets by the vector $\boldsymbol{R}^{fin}_{t+1} \equiv \left(R_{s,t+1}, R^{(10)}_{b,t+1}\right)'$. The asset menu for human capital follows from the nature of labor income. We denote the (nominal) returns on human capital assets by the vector $\boldsymbol{R}^{hc}_{t+1} \equiv \left(R^{(1)}_{hc,t+1}, \dots, R^{(T-1)}_{hc,t+1}\right)'$. The (given) allocation that corresponds to human capital is denoted by $\boldsymbol{\beta}_t$. This allocation depends on both the time and the current state of the world. For instance, at time $T-4$, the individual has three labor income payments left (at $T-3$, $T-2$, and $T-1$). This implies that only the first three elements of $\boldsymbol{\beta}_{T-4}$ are nonzero. The specific weights depend on the value of the state vector.

For the optimization problem in (50), we need to have an expression for the return on total wealth. Total wealth is the sum of financial wealth and human capital, and therefore the return on total wealth is the weighted average of the return on financial wealth and the return on human capital. The return on financial wealth is given by $Y^{(1)}_t + \boldsymbol{\alpha}'_t(\boldsymbol{R}^{fin}_{t+1} - Y^{(1)}_t\boldsymbol{\iota})$. For the return on human capital, we have $Y^{(1)}_t + \boldsymbol{\beta}'_t(\boldsymbol{R}^{hc}_{t+1} - Y^{(1)}_t\boldsymbol{\iota})$. We let the variable $f_t \equiv \frac{F_t}{W_t}$ denote the fraction of financial wealth in total wealth. The return on the portfolio of total wealth is then given by:

$$R_{p,t+1} = Y^{(1)}_t + \boldsymbol{\alpha}'_t(\boldsymbol{R}^{fin}_{t+1} - Y^{(1)}_t\boldsymbol{\iota})f_t + \boldsymbol{\beta}'_t(\boldsymbol{R}^{hc}_{t+1} - Y^{(1)}_t\boldsymbol{\iota})(1-f_t). \tag{51}$$

As we have seen, the total amount of wealth after one period is given by the current wealth multiplied by the real return of the total portfolio. However, the individual values of financial wealth and human capital after one period are not simply given by their initial values multiplied by the real returns. Due to the fact that labor income is paid out, there is an outflow of human capital and an inflow of financial wealth. Financial wealth in the next period is given by

$$F_{t+1} = (1 + Y^{(1)}_t + \boldsymbol{\alpha}'_t(\boldsymbol{R}^{fin}_{t+1} - Y^{(1)}_t\boldsymbol{\iota}))F_t + \rho_{t+1}L_{t+1}, \tag{52}$$

i.e., the initial financial wealth multiplied by the return on financial wealth, plus the contribution in the next period. This leads to the following dynamics for f_t:

$$f_{t+1} = f_t \frac{1 + Y_t^{(1)} + \boldsymbol{\alpha}_t'(\boldsymbol{R}_{t+1}^{fin} - Y_t^{(1)}\boldsymbol{\iota})}{1 + R_{p,t+1}} + \frac{\rho_{t+1} L_{t+1}}{W_{t+1}}. \tag{53}$$

5.3 Portfolio returns

In the model that we use, the individual returns of the assets in the portfolios of financial wealth and human capital follow a lognormal distribution. Therefore, it will be useful to derive an expression for the log return on a portfolio of assets with lognormally distributed returns. The disadvantage of using log returns is that the log return of a portfolio of different assets is not the sum of the log returns of the individual assets multiplied by their portfolio weight. Therefore, we will first focus on deriving a general expression for the log return of a portfolio of assets, in line with the derivations for the Merton model in Appendix A.

In general, we have a portfolio that consists of J risky assets. At time t, we invest a fraction $\xi_{j,t}$ of total wealth in risky asset j ($j = 1, \ldots, J$). The vector of all weights is defined by $\boldsymbol{\xi}_t$. The corresponding return vector at time $t+1$ is denoted by $\boldsymbol{R}_{t+1}$. The return of the portfolio is given by

$$1 + R_{p,t+1} = 1 + Y_t^{(1)} + \boldsymbol{\xi}_t'(\boldsymbol{R}_{t+1} - Y_t^{(1)}\boldsymbol{\iota}). \tag{54}$$

Therefore, for the log portfolio return $r_{p,t+1}$, we have:

$$\begin{aligned} r_{p,t+1} - y_t^{(1)} &= \log\left(1 + \boldsymbol{\xi}_t'\left(\frac{1 + \boldsymbol{R}_{t+1}}{1 + Y_t^{(1)}} - \boldsymbol{\iota}\right)\right) = \log\left(1 + \sum_{j=1}^{J} \xi_{j,t}(\exp(r_{j,t+1} - y_t^{(1)}) - 1)\right) \\ &\equiv f\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right). \end{aligned} \tag{55}$$

In a similar way as before, we use a second-order Taylor expansion to approximate this expression:

$$f\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right) \approx f(0) + \boldsymbol{\nabla f}(\boldsymbol{0})'\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right) + \frac{1}{2}\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right)' H_f(\boldsymbol{0})\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right),$$

where

$$\begin{aligned} \boldsymbol{\nabla f}(\boldsymbol{0}) &= \boldsymbol{\xi}_t \\ H_f(\boldsymbol{0}) &= \mathrm{diag}(\boldsymbol{\xi}_t) - \boldsymbol{\xi}_t\boldsymbol{\xi}_t'. \end{aligned}$$

Hence, we find as approximation

$$\begin{aligned} r_{p,t+1} - y_t^{(1)} \approx\ & \boldsymbol{\xi}_t'\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right) + \frac{1}{2}\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right)' \mathrm{diag}(\boldsymbol{\xi}_t)\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right) \\ & - \frac{1}{2}\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right)' \boldsymbol{\xi}_t\boldsymbol{\xi}_t'\left(\boldsymbol{r}_{t+1} - y_t^{(1)}\boldsymbol{\iota}\right). \end{aligned} \tag{56}$$

If we further approximate the second-order terms by their conditional expectations, we obtain

$$r_{p,t+1} - y_t^{(1)} \approx \boldsymbol{\xi}_t' \left(\boldsymbol{r}_{t+1} - y_t^{(1)} \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\xi}_t' \boldsymbol{\sigma}_r^2 - \frac{1}{2} \boldsymbol{\xi}_t' \Sigma_r \boldsymbol{\xi}_t. \tag{57}$$

Note that we are dealing here with the covariance matrix of the return vector $\boldsymbol{r}_{t+1}$, Σ_r, instead of the covariance matrix of the error term $\boldsymbol{\epsilon}_{t+1}$ that belongs to the state variables, which is Σ. Of course, Σ_r follows directly from Σ via the exposures of the returns to the different factors.

5.4 Individual-specific parameters

In Section 2.3, we have already discussed some properties of the life-cycle investor in which we are interested. In Section 4, we have discussed our estimates for the term structure parameters that we will use to obtain numerical solutions to the life-cycle optimization problem. In this section, we will specify the remaining individual-specific parameters that we will use as benchmark for the numerical results.

Recall from Section 2.3 that we consider an individual that receives his first labor income at the age of 25. Thus, at $t = 0$ the individual is 25 years old. After $T = 40$ years, the individual retires at the age of 65. As we will see, the risk-aversion coefficient γ will have a very high impact on the outcomes. Since risk-aversion can vary greatly among individuals and is difficult to estimate, we consider different values of γ. In the literature, risk aversion usually varies between $\gamma = 2$ (an aggressive investor) and $\gamma = 10$ (a conservative investor). Therefore, we do the analysis for $\gamma = 2$, $\gamma = 5$, and $\gamma = 10$.

For simplicity, we abstract from mortality and longevity risk and assume that all people in the annuity pool die precisely in the year after they have received their twentieth annuity payment at the age of 84. This means that the (fair) price of the annuity at time T is equal to the sum of the real discount factors for maturities 0 (immediate payoff at the retirement age) to 19 (last payment nineteen years later).

Labor income is stochastic and is potentially correlated to stock returns. The characteristics of labor income vary widely across different individuals, and depend among others on the industry sector, age, education level, gender, and race. Several studies have quantified the deterministic growth in labor productivity, the volatility of labor income, and the correlation between labor income and equity returns. As discussed before, we use the growth path that is assumed by the Dutch tax authority as input for the deterministic labor productivity growth, g_t. We also already discussed our values for the parameter ρ_t, the yearly contribution to the pension plan. The values for g_t and ρ_t are given in Table 2.

Palacios-Huerta (2003) estimates the mean return and volatility of human capital for American citizens. For white individuals with some college experience (middle group) and with 6 to 15 years of work experience, he finds a volatility of approximately 5%. This holds for both males (5.4%) and females (5.1%). The estimated volatilities for (white) people with a different education level or a different experience are similar. We therefore take $\sigma_u = 5\%$.

Also the correlation between human capital returns and stock returns depends heavily on individual characteristics. Cocco, Gomes, and Maenhout (2005) report correlations between stock returns and labor income shocks between -1% and 2%, depending on the education level. Davis and Willen (2000) estimate the correlation between income innovations and equity shocks in a range from -25% to 30%. Munk and Sørensen (2010) estimate the correlation to be 17%. We take a moderate correlation of 10% as benchmark. This is the correlation between stock returns and the whole innovation in labor income (including a correction for price inflation). For the parameters in our model, this boils (approximately) down to an exposure of $k = 2\%$ to stock returns. Hence, as benchmark parameter, we take $k = 2\%$.

6 Dynamic optimization procedure

In this section, we discuss our methodology to find the optimal allocations in a defined contribution pension scheme over time. Recall from Section 5.2 that $\boldsymbol{\alpha}_t$ denotes the allocation of financial wealth to the J risky assets that are available on the financial market. The resulting allocation to the one-period risk-free asset ("cash") at time t is given by $1 - \iota' \boldsymbol{\alpha}_t$.

We solve the optimization problem described in Section 5 by making use of stochastic dynamic optimization procedures. Examples of these multi-period optimizations are Blake, Wright, and Zhang (2011a, 2011b). However, these papers do not value human capital in a market-consistent way.[4] We start with the discussion of the optimal allocation at the last point in time. We are able to derive an analytic expression for the optimal allocation under our model assumptions. Afterwards, we discuss our method to find the optimal allocations at earlier points in time by making use of numerical simulation techniques.

We will apply the numerical approach that is applied in Koijen, Nijman, and Werker (2010) and explained in Koijen, Nijman, and Werker (2007). This method combines the methods of Brandt, Goyal, Santa-Clara, and Stroud (2005) and Carroll (2006). Brandt et al. (2005) propose to use functional approximations to evaluate the conditional expectations that we encounter when solving the problem dynamically, inspired by the approach of Longstaff and Schwarz (2001) for pricing American options. However, there is one state variable that we cannot simulate, namely the endogenous wealth level. Wealth is endogenous because the return of the total portfolio depends on the chosen portfolio weights. We can deal with this by constructing a grid in wealth (Carroll, 2006). Since for our problem not the wealth itself, but only the ratio of financial wealth to human capital plays a role, we instead construct a grid for f_t, financial wealth as a fraction of total wealth.

6.1 Final period

As is common in dynamic optimization, we start at the last period in time. The last decision that we need to make is the asset allocation decision at time $T-1$. Let $V_{T-1}(W_{T-1}, f_{T-1}, \boldsymbol{x}_{T-1})$ be the optimal value at time $T-1$, conditional on the information available at that moment in time. The problem at time $T-1$ is as follows:

$$
\begin{aligned}
V_{T-1}(W_{T-1}, f_{T-1}, \boldsymbol{x}_{T-1}) \equiv \quad & \max \quad && \mathbb{E}_{T-1}\left[\frac{1}{1-\gamma}\left(\frac{W_T}{A_T}\right)^{1-\gamma}\right] \\
& \text{subject to} \quad && W_T = W_{T-1}(1+R_{p,T})\frac{\Pi_{T-1}}{\Pi_T} \qquad (58) \\
& && \boldsymbol{\alpha}_{T-1} \geq \mathbf{0} \\
& && \boldsymbol{\iota}'\boldsymbol{\alpha}_{T-1} \leq 1,
\end{aligned}
$$

[4]Blake et al. (2011a, 2011b) have a model of the financial market with a risk-free bond and a risky stock, but discount future labor income with the yield on AA-grade corporate bonds that is determined on the basis of a single data point.

where the portfolio return $R_{p,T}$ is defined by (51):

$$R_{p,T} = Y_{T-1}^{(1)} + \boldsymbol{\alpha}'_{T-1}(\boldsymbol{R}_T^{fin} - Y_{T-1}^{(1)}\boldsymbol{\iota})f_{T-1} + \boldsymbol{\beta}'_{T-1}(\boldsymbol{R}_T^{hc} - Y_{T-1}^{(1)}\boldsymbol{\iota})(1 - f_{T-1}).$$

At time $T-1$, the individual receives his last income and hence the value of his remaining human capital, H_{T-1}, equals zero. This means that all wealth is financial wealth: $f_{T-1} = 1$. It follows that we have for the portfolio return:

$$R_{p,T} = Y_{T-1}^{(1)} + \boldsymbol{\alpha}'_{T-1}(\boldsymbol{R}_T^{fin} - Y_{T-1}^{(1)}\boldsymbol{\iota}).$$

Recall that the risk-aversion parameter γ is larger than one, which means that the fraction $\frac{1}{1-\gamma}$ is negative. Again assuming that borrowing and short selling is not possible, the final wealth of the investor is always non-negative. Thus, the objective is equivalent to minimizing the expected value of the real pension benefits to the power $1-\gamma$.

When we rewrite the optimization problem in logarithmic terms, we are able to derive an analytical expression for the allocation decision in the last period. Since the logarithmic function is increasing on its whole domain, minimizing the log of the (positive) expectation is equivalent to minimizing the expectation itself. We find the equivalent problem

$$\begin{array}{ll} \min & \log \mathbb{E}_{T-1}[\exp\{(1-\gamma)(w_T - a_T)\}] \\ \text{subject to} & w_T = w_{T-1} + r_{p,T} - \pi_T \\ & \boldsymbol{\alpha}_{T-1} \geq \mathbf{0} \\ & \boldsymbol{\iota}'\boldsymbol{\alpha}_{T-1} \leq 1, \end{array} \tag{59}$$

where $w_t \equiv \log(W_t), r_{p,T} \equiv \log(1 + R_{p,T})$ and $a_T \equiv \log(A_T)$.

From Section 5.3, we know that portfolio returns are approximately lognormally distributed. Further, we can replicate the annuity at retirement by buying a deferred annuity now. Remember that we assumed that the individual dies exactly twenty years after his retirement (see Section 5.4). Hence, replicating at time $T-1$ the annuity payments from time T onwards boils down to buying 20 zero-coupon inflation-linked bonds with maturities $1, \ldots, 20$. We can write

$$a_T = a_{T-1} + r_{a,T} - \pi_T, \tag{60}$$

where a_{T-1} is the value of a deferred real annuity that pays off from time T onwards, and $r_{a,T}$ is the nominal return on the portfolio of 20 zero-coupon inflation-linked bonds with maturities $1, \ldots, 20$. The nominal return on a real bond is also lognormally distributed in this model, which means that the whole expression $\exp\{(1-\gamma)(w_T - a_T)\}$ follows a lognormal distribution. By applying that $\mathbb{E}e^X = e^{\mu + 0.5\sigma^2}$ when $X \sim N(\mu, \sigma^2)$ and deleting terms that are irrelevant for the optimization, the optimization problem becomes

$$\begin{array}{ll} \min & \left[(1-\gamma)\mathbb{E}_{T-1}r_{p,T} + \dfrac{1}{2}(1-\gamma)^2(\text{Var}_{T-1}r_{p,T} - 2\text{Cov}_{T-1}(r_{p,T}, r_{A,T}))\right] \\ \text{subject to} & \boldsymbol{\alpha}_{T-1} \geq \mathbf{0} \\ & \boldsymbol{\iota}'\boldsymbol{\alpha}_{T-1} \leq 1. \end{array} \tag{61}$$

Note that from this reformulation, it becomes clear that the optimal allocation at time $T-1$ does not depend on the current wealth level W_{T-1}.

Finally, we use the expression that we derived in the previous subsection (equation (57)), which yields

$$\begin{aligned} \mathbb{E}_{T-1} r_{p,T} + \frac{1}{2}(1-\gamma)\mathrm{Var}_{T-1} r_{p,T} &= y_{T-1}^{(1)} + \boldsymbol{\alpha}_{T-1}' \left(\mathbb{E}_{T-1} \boldsymbol{r}_T^{fin} - y_{T-1}^{(1)} \boldsymbol{\iota} \right) \\ &\quad + \frac{1}{2} \boldsymbol{\alpha}_{T-1}' \boldsymbol{\sigma}_{r,fin}^2 - \frac{1}{2} \gamma \boldsymbol{\alpha}_{T-1}' \Sigma_{r,fin} \boldsymbol{\alpha}_{T-1}, \end{aligned} \tag{62}$$

where $\Sigma_{r,fin}$ is the covariance matrix that corresponds to the vector of log returns $\boldsymbol{r}_T^{fin}$ and $\boldsymbol{\sigma}_{r,fin}^2$ is the vector with the corresponding individual variances, i.e., the diagonal elements of $\Sigma_{r,fin}$.

After canceling the (negative) common factor $1-\gamma$, we obtain the reduced problem

$$\begin{aligned} \max \quad & \boldsymbol{\alpha}_{T-1}' \mathbb{E}_{T-1} (\boldsymbol{r}_T^{fin} - y_{T-1}^{(1)} \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\alpha}_{T-1}' \boldsymbol{\sigma}_{r,fin}^2 \\ & - \frac{1}{2} \gamma \boldsymbol{\alpha}_{T-1}' \Sigma_{r,fin} \boldsymbol{\alpha}_{T-1} - (1-\gamma) \mathrm{Cov}_{T-1} (\boldsymbol{\alpha}_{T-1}' \boldsymbol{r}_T^{fin}, r_{A,T}) \\ \text{s.t.} \quad & \boldsymbol{\alpha}_{T-1} \geq \mathbf{0} \\ & \boldsymbol{\iota}' \boldsymbol{\alpha}_{T-1} \leq 1. \end{aligned} \tag{63}$$

Let $\boldsymbol{\lambda}_{T-1}$ and μ_{T-1} be the Lagrange multipliers corresponding to the restrictions $\boldsymbol{\alpha}_{T-1} \geq \mathbf{0}$ and $\boldsymbol{\iota}' \boldsymbol{\alpha}_{T-1} \leq 1$, respectively. The corresponding Lagrangian is given by:

$$\begin{aligned} L(\boldsymbol{\alpha}_{T-1}, \boldsymbol{\lambda}_{T-1}, \mu_{T-1}) &= \boldsymbol{\alpha}_{T-1}' \mathbb{E}_{T-1} (\boldsymbol{r}_T^{fin} - y_{T-1}^{(1)} \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\alpha}_{T-1}' \boldsymbol{\sigma}_{r,fin}^2 - \frac{1}{2} \gamma \boldsymbol{\alpha}_{T-1}' \Sigma_{r,fin} \boldsymbol{\alpha}_{T-1} \\ & \quad - (1-\gamma) \mathrm{Cov}_{T-1} (\boldsymbol{\alpha}_{T-1}' \boldsymbol{r}_T^{fin}, r_{A,T}) + \boldsymbol{\lambda}_{T-1}' \boldsymbol{\alpha}_{T-1} - \mu_{T-1} (\boldsymbol{\iota}' \boldsymbol{\alpha}_{T-1} - 1). \end{aligned} \tag{64}$$

To solve the optimization problem, we make use of the fact that we have a concave objective function[5] and linear constraints. This means that solving the KKT conditions leads to a global maximum. The KKT conditions can be classified into four categories:

1. Primal feasibility: $\boldsymbol{\alpha}_{T-1} \geq 0, \boldsymbol{\iota}' \boldsymbol{\alpha}_{T-1} \leq 1$.

2. Dual feasibility: $\boldsymbol{\lambda}_{T-1} \geq 0, \mu_{T-1} \geq 0$.

3. Complementary slackness: $\lambda_{j,T-1} \alpha_{j,T-1} = 0 \quad (j = 1, \ldots, J)$, $\mu_{T-1} (\boldsymbol{\iota}' \boldsymbol{\alpha}_{T-1} - 1) = 0$.

4. Optimality: $\frac{\partial}{\partial \boldsymbol{\alpha}_{T-1}} L(\boldsymbol{\alpha}_{T-1}, \boldsymbol{\lambda}_{T-1}, \mu_{T-1}) = 0$. This gives:

$$\mathbb{E}_{T-1} (\boldsymbol{r}_T^{fin} - y_{T-1}^{(1)} \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_{r,fin}^2 - \gamma \Sigma_{r,fin} \boldsymbol{\alpha}_{T-1} - (1-\gamma) \mathrm{Cov}_{T-1} (\boldsymbol{r}_T^{fin}, r_{A,T}) + \boldsymbol{\lambda}_{T-1} - \mu_{T-1} \boldsymbol{\iota} = 0. \tag{65}$$

[5] The Hessian of the objective function is $-\gamma \Sigma_{r,fin}$, which is negative definite.

Solving this set of conditions is straightforward. When the allocation constraints are not binding, we get the following myopic portfolio:

$$\boldsymbol{\alpha}_{T-1} = (\gamma\Sigma_{r,fin})^{-1}\{\underbrace{\mathbb{E}_{T-1}(\boldsymbol{r}_T^{fin} - y_{T-1}^{(1)}\boldsymbol{\iota}) + \frac{1}{2}\boldsymbol{\sigma}_{r,fin}^2}_{\text{Return component}}\underbrace{-(1-\gamma)\text{Cov}_{T-1}(\boldsymbol{r}_T^{fin}, r_{a,T})}_{\text{Hedging component}}\}. \quad (66)$$

Note that we have split up the myopic portfolio allocation into a return component and a hedging component. The return component is the multi-asset equivalent of the myopic portfolio that we found in Section 2.1. However, we have changed the optimization problem from maximizing utility over nominal wealth to maximizing utility over annuitized real wealth. The additional hedging term accounts for correlation between asset returns and the price of a real annuity.

In case the wealth at maturity would not be annuitized, it can easily be verified that the covariance term in (65) is replaced by

$$-(1-\gamma)\text{Cov}_{T-1}(\boldsymbol{r}_T^{fin}, \pi_T).$$

In that case, when the allocation constraints are not binding, we get the myopic choice

$$\boldsymbol{\alpha}_{T-1} = (\gamma\Sigma_{r,fin})^{-1}\{\underbrace{\mathbb{E}_{T-1}(\boldsymbol{r}_T^{fin} - y_{T-1}^{(1)}\boldsymbol{\iota}) + \frac{1}{2}\boldsymbol{\sigma}_{r,fin}^2}_{\text{Return component}}\underbrace{-(1-\gamma)\text{Cov}_{T-1}(\boldsymbol{r}_T^{fin}, \pi_T)}_{\text{Hedging component}}\}. \quad (67)$$

6.2 Value at final period

Let α_{T-1}^* be the optimal asset allocation for a given value of the state vector $\boldsymbol{x}_{T-1}$, let $r_{p,T}^*$ be the corresponding portfolio return, and let W_T^* be the (random) value at time T that results from applying this optimal strategy. As we will see later on, the conditional expectation

$$Z_{T-1}(f_{T-1}, \boldsymbol{x}_{T-1}) \equiv \mathbb{E}_{T-1}\left[\frac{1}{A_T^{1-\gamma}}\left(1 + Y_{T-1}^{(1)} + \boldsymbol{\alpha}_{T-1}^{*\prime} R_T^{fin}\right)^{1-\gamma}\left(\frac{\Pi_{T-1}}{\Pi_T}\right)^{1-\gamma}\right]$$

will play an important role in the optimization procedure. Therefore, we will derive an analytic expression for this quantity. Note that this quantity is the optimal value of the optimization problem at the last period in time, divided by the initial wealth component

$W_{T-1}^{1-\gamma}$. We obtain:

$$\begin{aligned}
&\mathbb{E}_{T-1}\left[\frac{1}{A_T^{1-\gamma}}\left(1+Y_{T-1}^{(1)}+\boldsymbol{\alpha}_{T-1}^{*\prime}R_T^{fin}\right)^{1-\gamma}\left(\frac{\Pi_{T-1}}{\Pi_T}\right)^{1-\gamma}\right]\\
&=\mathbb{E}_{T-1}[\exp\{(1-\gamma)(-a_T+r_{p,T}^{*}-\pi_T)\}]\\
&=\mathbb{E}_{T-1}[\exp\{(\gamma-1)a_{T-1}+(1-\gamma)(r_{p,T}^{*}-r_{A,T})\}]\\
&=A_{T-1}^{\gamma-1}\exp\left\{(1-\gamma)\mathbb{E}_{T-1}(r_{p,T}^{*}-r_{A,T})+\frac{1}{2}(1-\gamma)^2\mathrm{Var}_{T-1}(r_{p,T}^{*}-r_{A,T})\right\}\\
&=A_{T-1}^{\gamma-1}\exp\left\{(1-\gamma)\left(\boldsymbol{\alpha}_{T-1}^{*\prime}(\mathbb{E}_{T-1}\boldsymbol{r}_T^{fin}-y_{T-1}^{(1)}\boldsymbol{\iota})+\frac{1}{2}\boldsymbol{\alpha}_{T-1}^{*\prime}\boldsymbol{\sigma}_{r,fin}^2-\frac{1}{2}\gamma\boldsymbol{\alpha}_{T-1}^{*\prime}\Sigma_{r,fin}\boldsymbol{\alpha}_{T-1}^{*}\right.\right.\\
&\left.\left.+y_{T-1}^{(1)}-\mathbb{E}_{T-1}r_{A,T}+\frac{1}{2}(1-\gamma)\mathrm{Var}_{T-1}r_{A,T}-(1-\gamma)\mathrm{Cov}_{T-1}(r_{p,T}^{*},r_{A,T})\right)\right\}. \qquad (68)
\end{aligned}$$

6.3 Earlier periods

We are now able to determine the optimal decision at earlier moments, by working backwards in time from time $T-2$ to time 0. At time t, we make use of the information that we have about the optimal allocation from time $t+1$ onwards.

At time $T-2$, we have the following objective function:

$$\begin{aligned}
&\mathbb{E}_{T-2}\left[\frac{1}{1-\gamma}\left(\frac{W_T}{A_T}\right)^{1-\gamma}\right]\\
&=\mathbb{E}_{T-2}\left[\frac{1}{1-\gamma}\frac{W_{T-2}^{1-\gamma}}{A_T^{1-\gamma}}\left(1+R_{p,T-1}\right)^{1-\gamma}\left(1+R_{p,T}\right)^{1-\gamma}\left(\frac{\Pi_{T-2}}{\Pi_T}\right)^{1-\gamma}\right]\\
&=\mathbb{E}_{T-2}\left(\mathbb{E}_{T-1}\left[\frac{1}{1-\gamma}\left\{\frac{W_{T-2}}{A_T}\left(1+R_{p,T-1}\right)\left(1+R_{p,T}\right)\frac{\Pi_{T-2}}{\Pi_T}\right\}^{1-\gamma}\right]\right), \qquad (69)
\end{aligned}$$

where $\boldsymbol{\alpha}_{T-2}$ is deterministic at time $T-2$ and $\boldsymbol{\alpha}_{T-1}(\boldsymbol{x}_{T-1})$ is chosen conditional on the information available one period ahead, at time $T-1$. Since the optimal allocation strategy at time $T-1$ is the same for all wealth levels W_{T-1}, this strategy will also be optimal at time $T-1$ in the two-period setting. We can already see that again the current wealth level (W_{T-2}) does not influence the optimal allocation strategy.

By applying the law of iterated expectations in the last step in (69), we arrive at the Bellman equation for this dynamic problem. A simple application of the Bellman equation implies that the problem at time t ($t=0,\ldots,T-2$) is as follows:

$$\begin{aligned}
V_t(W_t,f_t,\boldsymbol{x}_t)= \quad &\max \quad \mathbb{E}_t\left[V_{t+1}(W_{t+1},f_{t+1},\boldsymbol{x}_{t+1})\right]\\
&\text{subject to} \quad W_{t+1}=W_t(1+R_{p,t+1})\frac{\Pi_t}{\Pi_{t+1}}\\
&\boldsymbol{\alpha}_t\geq\mathbf{0}\\
&\boldsymbol{\iota}'\boldsymbol{\alpha}_t\leq 1,
\end{aligned} \qquad (70)$$

where $V_{T-1}(W_{T-1}, f_{T-1}, \boldsymbol{x}_{T-1})$ is the optimal value of the optimization problem at the final period, given by (58).

Suppose that we know the optimal investment strategy $\boldsymbol{\alpha}_s^*(\boldsymbol{x}_s)$ for all periods $s \geq t$. Let $R_{p,s}^*$ again be the portfolio return at time s that results from following the optimal asset allocation strategy. Define the auxiliary variable Z_t by

$$Z_t(f_t, \boldsymbol{x}_t) \equiv \mathbb{E}_t \left[\frac{1}{A_T^{1-\gamma}} \prod_{s=t}^{T-1} (1 + R_{p,s+1}^*)^{1-\gamma} \left(\frac{\Pi_s}{\Pi_{s+1}} \right)^{1-\gamma} \right] \tag{71}$$

By definition of the value function, it holds that

$$V_{t+1}(W_{t+1}, f_{t+1}, \boldsymbol{x}_{t+1}) = \frac{1}{1-\gamma} W_{t+1}^{1-\gamma} Z_{t+1}(f_{t+1}, \boldsymbol{x}_{t+1}), \tag{72}$$

where $W_{t+1} = W_t(1 + R_{p,t+1}) \frac{\Pi_t}{\Pi_{t+1}}$.

Recall from Section 5.2 (equation (51)) that we have for the portfolio return:

$$R_{p,t+1} = Y_t^{(1)} + \boldsymbol{\alpha}_t'(\boldsymbol{R}_{t+1}^{fin} - Y_t^{(1)}\boldsymbol{\iota}) f_t + \boldsymbol{\beta}_t'(\boldsymbol{R}_{t+1}^{hc} - Y_t^{(1)}\boldsymbol{\iota})(1 - f_t).$$

At time $t < T - 1$, the individual possesses both financial wealth and human capital. The fraction f_t of total wealth that is financial wealth changes over time both as a result of asset returns and as a result of human capital that becomes financial capital. The dynamics for f_t are given by (53):

$$f_{t+1} = f_t \frac{1 + Y_t^{(1)} + \boldsymbol{\alpha}_t'(\boldsymbol{R}_{t+1}^{fin} - Y_t^{(1)}\boldsymbol{\iota})}{1 + R_{p,t+1}} + \frac{\rho_{t+1} L_{t+1}}{W_{t+1}}.$$

The corresponding Lagrangian for the optimization problem (70) is given by:

$$L(\boldsymbol{\alpha}_t, \boldsymbol{\lambda}_t, \mu_t) = \mathbb{E}_t \left[V_{t+1}(W_{t+1}, f_{t+1}, \boldsymbol{x}_{t+1}) \right] + \boldsymbol{\lambda}_t' \boldsymbol{\alpha}_t - \mu_t(\boldsymbol{\iota}'\boldsymbol{\alpha}_t - 1). \tag{73}$$

We now get the following set of KKT conditions for this problem:

$$\begin{cases} \boldsymbol{\alpha}_t \geq & \mathbf{0} \\ \boldsymbol{\iota}'\boldsymbol{\alpha}_t \leq & 1 \\ \boldsymbol{\lambda}_t \geq & \mathbf{0} \\ \mu_t \geq & 0 \\ \lambda_{j,t}\alpha_{j,t} = & 0 \quad (j = 1, \ldots, J) \\ \mu_t(\boldsymbol{\iota}'\boldsymbol{\alpha}_t - 1) = & 0 \\ \mathbf{0} = & \frac{\partial}{\partial \boldsymbol{\alpha}_t} \mathbb{E}_t \left[V_{t+1}(W_{t+1}, f_{t+1}, \boldsymbol{x}_{t+1}) \right] + \boldsymbol{\lambda}_t - \mu_t \boldsymbol{\iota} \end{cases}, \tag{74}$$

where, following from applying equation (72):

$$
\begin{aligned}
&\frac{\partial}{\partial \boldsymbol{\alpha}_t}\mathbb{E}_t\left[V_{t+1}(W_{t+1}, f_{t+1}, \boldsymbol{x}_{t+1})\right] \\
&= \mathbb{E}_t\left[\frac{\partial V_{t+1}(W_{t+1}, f_{t+1}, \boldsymbol{x}_{t+1})}{\partial W_{t+1}}\frac{\partial W_{t+1}}{\partial \boldsymbol{\alpha}_t} + \frac{\partial V_{t+1}(W_{t+1}, f_{t+1}, \boldsymbol{x}_{t+1})}{\partial f_{t+1}}\frac{\partial f_{t+1}}{\partial \boldsymbol{\alpha}_t}\right] \\
&= \mathbb{E}_t\left[f_t W_t^{1-\gamma}(1+R_{p,t+1})^{-\gamma}Z_{t+1}(f_{t+1}, \boldsymbol{x}_{t+1})(\boldsymbol{R}_{t+1}^{fin} - Y_t^{(1)}\boldsymbol{\iota})\left(\frac{\Pi_t}{\Pi_{t+1}}\right)^{1-\gamma}\right] \\
&\quad + \mathbb{E}_t\left[W_t^{1-\gamma}\frac{1}{1-\gamma}(1+R_{p,t+1})^{1-\gamma}\left(\frac{\Pi_t}{\Pi_{t+1}}\right)^{1-\gamma}\frac{\partial Z_{t+1}(f_{t+1}, \boldsymbol{x}_{t+1})}{\partial f_{t+1}}\frac{\partial f_{t+1}}{\partial \boldsymbol{\alpha}_t}\right] \\
&= f_t W_t^{1-\gamma}\mathbb{E}_t\left[\boldsymbol{\Omega}_{t+1}(\boldsymbol{\alpha}_t, f_{t+1}, \boldsymbol{x}_{t+1})\right], \qquad (75)
\end{aligned}
$$

with

$$
\begin{aligned}
\boldsymbol{\Omega}_{t+1}(\boldsymbol{\alpha}_t, f_{t+1}, \boldsymbol{x}_{t+1}) \;\equiv\;& (1+R_{p,t+1})^{-\gamma}Z_{t+1}(f_{t+1}, \boldsymbol{x}_{t+1})(\boldsymbol{R}_{t+1}^{fin} - Y_t^{(1)}\boldsymbol{\iota})\left(\frac{\Pi_t}{\Pi_{t+1}}\right)^{1-\gamma} \qquad (76) \\
&+\frac{1}{1-\gamma}(1+R_{p,t+1})^{1-\gamma}\left(\frac{\Pi_t}{\Pi_{t+1}}\right)^{1-\gamma}\frac{\partial Z_{t+1}(f_{t+1}, \boldsymbol{x}_{t+1})}{\partial f_{t+1}}\frac{1}{f_t}\frac{\partial f_{t+1}}{\partial \boldsymbol{\alpha}_t}.
\end{aligned}
$$

Recall that we work backwards in time to determine the optimal allocation at each point in time, depending on the state of the economy. We were able to derive an analytic expression for the optimal allocation at the last point in time. In Section 6.2, we also derived an expression for $Z_{T-1}(f_{T-1}, \boldsymbol{x}_{T-1})$. This means that we have all information needed for the optimization problem at time $T-2$. In order to find the optimal allocation at $T-2$, we should solve the KKT conditions (74). Solving these conditions implies the evaluation of the expectation in (75). Unfortunately, we are not able to derive an analytical expression for this complicated expectation. Therefore, we need to rely on approximation techniques to solve the first-order conditions. In the next subsection, we will propose a method that uses numerical simulations to find the optimal allocations over time.

6.4 Numerical simulation techniques

In order to approximate the conditional expectation in (75), we make use of numerical simulations. We simulate N trajectories of the state variables over the whole period of T years. We take the long-term mean $\boldsymbol{\mu}$ as starting point for all these sample paths. Once we know the value of the state variables, we also know the returns on stocks and bonds. For the sample paths, the corresponding wealth levels that result from applying the optimal allocation strategy are endogenous, as a result of endogenous returns on financial wealth. Since the term $f_t W_t^{1-\gamma}$ in (75) is a known positive constant at time t, it can be ignored in the first-order condition. Thus, as stated before, the initial wealth does not play a role for the optimal allocation. Hence, we do not need to make a wealth grid for this problem, but we do need a grid for f_t (which enters the equation through f_{t+1}). Therefore, we make an

M-dimensional grid grid for f_t. This means that at every point in time, we have in total $M \times N$ grid points. At each grid point, we want to determine the optimal asset allocation vector $\boldsymbol{\alpha}_t$.

The key step in our approach is to replace the conditional expectation with a polynomial expansion in the state variables, as proposed by Brandt et al. (2005). For each asset j ($j = 1, \ldots, J$) we employ the following approximation:

$$\mathbb{E}_t \left[\Omega_{j,t+1}(\boldsymbol{\alpha}_t, f_{t+1}, \boldsymbol{x}_{t+1}) \right] \approx \boldsymbol{\theta}_j(\boldsymbol{\alpha}_t, f_t)' \boldsymbol{f}(\boldsymbol{x}_t), \tag{77}$$

where $\Omega_{j,t+1}$ is the j-th element of the vector $\boldsymbol{\Omega}_{t+1}$.

Hence, we perform a cross-sectional regression across the simulated sample paths. Note that the regression coefficients depend on both the chosen allocation strategy and the initial financial wealth fraction f_t. For each possible combination of asset j, allocation $\boldsymbol{\alpha}_t$, and financial wealth fraction f_t, we have different regression coefficients $\boldsymbol{\theta}_j(\boldsymbol{\alpha}_t, f_t)$. In order to find the optimal allocation, we need to perform a regression for each possible allocation. This means that for each of the M different values of f_t, we need to perform many regressions to find the optimal asset allocation, since the regression coefficients depend on the specific allocation.

This method is not computationally feasible, since it requires too many regressions. Therefore, we make use of another approximation, which results from the observation that the regression coefficients $\boldsymbol{\theta}_j(\boldsymbol{\alpha}_{T-1}, f_{T-1})$ are smooth functions of the portfolio weights $\boldsymbol{\alpha}_{T-1}$ (see Koijen et al. (2010)). For that reason, we do not evaluate the regression coefficients for a large number of portfolio weights, but instead we parameterize the regression coefficients $\boldsymbol{\theta}_j(\boldsymbol{\alpha}_t, f_t)$ as follows:

$$\boldsymbol{\theta}_j(\boldsymbol{\alpha}_t, f_t) = \Psi_j(f_t) \boldsymbol{g}(\boldsymbol{\alpha}_t), \tag{78}$$

where Ψ_j is a matrix with the number of rows equal to the length of the vector $\boldsymbol{f}(\boldsymbol{x}_t)$ and with the number of columns equal to the length of the vector $\boldsymbol{g}(\boldsymbol{\alpha}_t)$. This implies that we can write the conditional expectation as

$$\mathbb{E}_t \left[\Omega_{j,t+1}(\boldsymbol{\alpha}_t, f_{t+1}, \boldsymbol{x}_{t+1}) \right] = \boldsymbol{g}(\boldsymbol{\alpha}_t)' \Psi_j(f_t)' \boldsymbol{f}(\boldsymbol{x}_t). \tag{79}$$

Suppose that we have at time t (an estimate of) the realization of the value of $Z_{t+1}(f_{t+1}, \boldsymbol{x}_{t+1})$ for the N sample paths and M realizations of f_{t+1}. At time $T-2$, we can use equation (68) to determine these future realizations. We proceed by calculating for any initial fraction f_t on the grid the approximated realization of $\Omega_{j,t+1}(\boldsymbol{\alpha}_t, f_{t+1}, \boldsymbol{x}_{t+1})$ for all sample paths and for H different portfolio strategies $\boldsymbol{\alpha}_t$. These portfolios are called "test portfolios". We then estimate the coefficients $\Psi_j(f_t)$ by performing a regression of the realizations on the regressors that are obtained by combining the functions $\boldsymbol{f}(\boldsymbol{x}_t)$ and $\boldsymbol{g}(\boldsymbol{\alpha}_t)$.

After we have determined all regression coefficients, we then solve for all sample paths i

$(i = 1, \ldots, N)$ the system of equations

$$\begin{cases} \boldsymbol{\alpha}_t \geq & \mathbf{0} \\ \boldsymbol{\iota}'\boldsymbol{\alpha}_t \leq & 1 \\ \boldsymbol{\lambda}_t \geq & \mathbf{0} \\ \mu_t \geq & 0 \\ \lambda_{j,t}\alpha_{j,t} = & 0 \quad (j = 1, \ldots, J) \\ \mu_t(\boldsymbol{\iota}'\boldsymbol{\alpha}_t - 1) = & 0 \\ 0 = & \boldsymbol{g}(\boldsymbol{\alpha}_t)'\Psi_j(f_t)'\boldsymbol{f}(\boldsymbol{x}_t^{(i)}) + \lambda_{j,t} - \mu_t \quad (j = 1, \ldots, J) \end{cases} . \tag{80}$$

This gives us the optimal asset allocation $\boldsymbol{\alpha}_t^*$ at time t.

In order to solve for the optimal allocations, we supposed that we knew the future realizations of $Z_{t+1}(f_{t+1}, \boldsymbol{x}_{t+1})$ for all grid points. While this is true at time $t = T - 2$, we have not yet determined the realizations of the values $Z_t(f_t, \boldsymbol{x}_t)$ for $t < T - 1$. When we have found the optimal asset allocations $\boldsymbol{\alpha}_t^*$ for all sample paths, we therefore need to calculate the realizations of

$$\begin{aligned} Z_t(f_t, \boldsymbol{x}_t) &= \mathbb{E}_t\left[\frac{1}{A_T^{1-\gamma}}\prod_{s=t}^{T-1}(1 + R_{p,s+1}^*)^{1-\gamma}\left(\frac{\Pi_s}{\Pi_{s+1}}\right)^{1-\gamma}\right] \\ &= \mathbb{E}_t\left[(1 + R_{p,t+1}^*)^{1-\gamma}\left(\frac{\Pi_t}{\Pi_{t+1}}\right)^{1-\gamma} Z_{t+1}(f_{t+1}, \boldsymbol{x}_{t+1})\right], \end{aligned} \tag{81}$$

which we need for the optimization problem in period $t - 1$. We again approximate the conditional expectation by performing a regression. An important remark here is that the conditional expectation should always be positive, since an increase in the level of wealth should never lead to a lower value. After all, for the CRRA utility function (4), more wealth is always better. Therefore, we follow the approach of Koijen et al. (2010) by approximating the logarithm of the conditional expectation with a polynomial expansion in the state variables, rather than the conditional expectation itself. This means that we employ the approximation

$$\mathbb{E}_t\left[(1 + R_{p,t+1}^*)^{1-\gamma}\left(\frac{\Pi_t}{\Pi_{t+1}}\right)^{1-\gamma} Z_{t+1}(f_{t+1}, \boldsymbol{x}_{t+1})\right] = \exp\{\bar{\boldsymbol{\theta}}(f_t)'\boldsymbol{f}(\boldsymbol{x}_t)\}. \tag{82}$$

Hence, we perform a nonlinear regression and compute the fitted values of the regression to estimate the realizations of the conditional expectation. We then go back one period in time and repeat this procedure until we arrive at time $t = 0$.

6.5 Numerical optimization settings

We end this section on the optimization procedure by describing the settings that we used to obtain the numerical results in the following chapters. We used $N = 10,000$ simulations of sample paths for the state variables. Further, we took a linear grid for f_t between 0.01 and 0.99 with step size of 0.02, so $M = 50$. We took both functions $\boldsymbol{f}(\cdot)$ and $\boldsymbol{g}(\cdot)$ to be affine functions: $\boldsymbol{f}(\boldsymbol{x}_t) \equiv (1, \boldsymbol{x}_t')'$ and $\boldsymbol{g}(\boldsymbol{\alpha}_t) \equiv (1, \boldsymbol{\alpha}_t')'$. For the test portfolios, we specified allocations to stocks and to bonds between 0% and 100%, with steps of 20%. The set of test portfolios then consists of all combinations that leave a non-negative allocation to cash. In total, there are $H = 21$ test portfolios.

We checked our settings for robustness, and we can conclude that the results are robust to all optimization settings. For instance, we could easily take more complicated function specifications for $\boldsymbol{f}(\cdot)$ and $\boldsymbol{g}(\cdot)$, but this did not lead to significantly different outcomes. Further, an increase in the value of N, M, or H did not lead to significantly different results. Using our settings, it took us about three hours to obtain the resulting allocation paths for the $N = 10,000$ samples. Potential gains in computational time could be made when the grid for f_t is chosen flexibly instead of considering all grid points from 0 to 1 at each point in time, while in the beginning of the life-cycle the fraction is always relatively low and at the end of the life-cycle the fraction is always relatively high.

7 Results for three-factor model without predictability

In this section, we will present the numerical results for the three-factor model. As parameters for the financial market, we take the estimated values from Section 4.2. Ultimately, we are interested in the optimal allocations over time for a life-cycle investor with the properties that were discussed in Section 5. Before we discuss the results for this problem, we first consider simplified cases in which we discuss the effects of annuitization and in which we compare allocations when interest rates are stochastic to the allocations for the Merton model with human capital (see Section 2.3).

7.1 Long-term investing without labor income

We start the analysis with the simplified case in which there is no labor income. At time $t = 0$, we have a lump sum of money that we invest until retirement at $T = 40$, rebalancing each year. There is no additional accumulation or decumulation over the life-cycle. This means that $L_t \equiv 0$ and hence $H_t \equiv 0$ for all t. In other words, we have the optimization problem that was described in the previous section, with $f_t \equiv 1$ for $t = 0, \ldots, T-1$. Hence, no grid is needed, and the optimal allocation depends only on the value of the state vector. We simulate $N = 10,000$ sample paths for the state vector and use the method that was discussed in the previous section to determine the optimal allocations. We consider two different cases: the case in which wealth is converted into a real annuity at retirement and the case in which only cash (in real terms) determines the utility at time T.

7.1.1 No annuitization at retirement

When wealth is not converted into an annuity at retirement, the utility of the end-of-period real wealth W_T is $u(W_T) = \frac{W_T^{1-\gamma}-1}{1-\gamma}$. In the formulas from the previous section, we plug in $A_T \equiv 1$ and $r_{A,T} \equiv \pi_T$ to arrive at this situation.

The expected return and the variance of the stock are constant over time. Static mean-variance analysis with a one-period risk-free asset ('cash') suggests that all investors should hold a single portfolio of risky assets, and that only the ratio of cash to the risky portfolio depends on the risk aversion of individuals. This would mean that the ratio between stocks and bonds should be constant over time. However, this is not the outcome when we are dealing with interest rate risk. As Campbell and Viceira (2001) point out, static portfolio analysis can be seriously misleading when interest rates fluctuate over time and investors have long horizons.

When the risk-free rate varies over time, long-term investors should invest a greater percentage of their portfolio in long-term bonds than short-term investors (see e.g. Brennan and Xia (2000) and Wachter (2003)). The reason is that long-term bonds can be used to stabilize long-run payoffs, since bond returns are negatively correlated to interest rates. As Sangvinatsos and Wachter (2005) note, we should be careful with this argument here. The agent wants to hedge the real risk-free rate, while he can only invest in nominal bonds.

Hence, the argument is only valid when nominal bonds are negatively correlated to real interest rates.

For the parameters that we have estimated (see Section 4.2), we find that nominal yields and real yields are highly correlated. This means that nominal 10-year bonds can indeed be used to hedge the real risk-free rate. As a result, long-term investors do not only have a myopic demand (given by (67)) for nominal 10-year bonds, which decreases when risk-aversion increases, but they also have a hedging demand for these long-term bonds. This hedging demand increases with the value of γ. For extremely risk-averse agents, therefore, the demand for bonds is almost exclusively determined by this hedging demand. Of course, as time progresses, long-term investors become short-term investors and only the myopic demand is left. For that reason, there is a sharp decrease in the optimal allocation to bonds for risk-averse agents.

Figure 8: Mean allocations for the three-factor model without human capital and without annuitization at retirement.

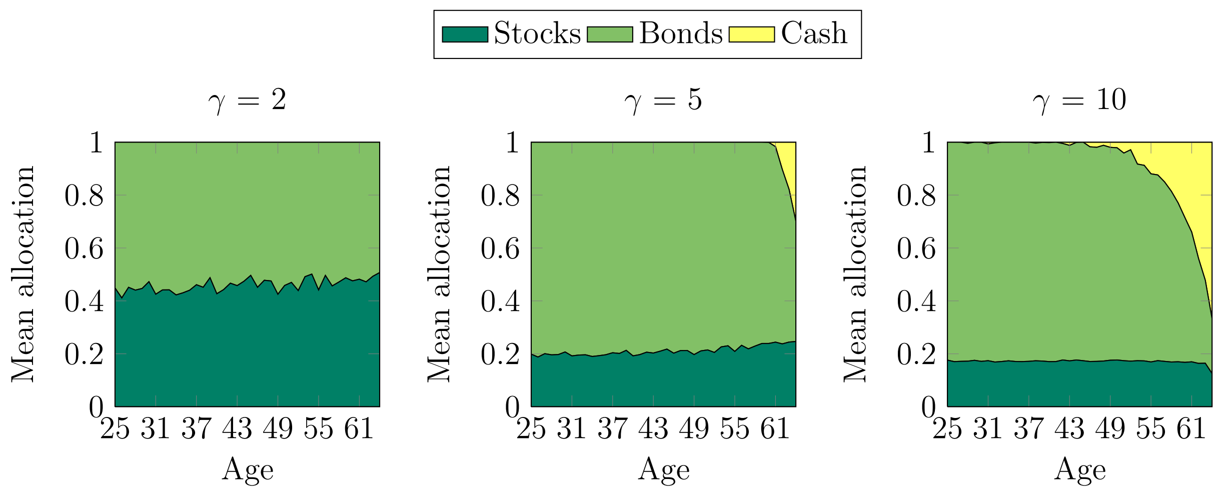

Figure 8 illustrates the optimal allocations that we find for our specific model of the financial market and for different values of the risk-aversion parameter γ. By comparing the allocation at the beginning of the period to the allocation at the end of the period, we can split up the demand for 10-year bonds in a hedging demand and a myopic demand. For agents with high risk aversion, the hedging demand for bonds is high, while the myopic demand for stocks and bonds shrinks when γ becomes large. For $\gamma = 2$ and $\gamma = 5$, the demand for bonds is mainly determined by the myopic demand. In contrast, for the investor with $\gamma = 10$, most of the demand for bonds comes from a hedging demand. For all investors, it is optimal to hold a full portfolio of risky assets in the first half of the period. The results that we get for this part of the life-cycle are close to those of Campbell and Viceira (2001), who treat the problem of an infinitely-lived agent. At the beginning of the period, the retirement age is still very far away, which implies that the agents act without focusing on the specific target date yet. After twenty years, the most risk-averse agent with $\gamma = 10$ starts to replace long-term bonds by cash, because he has a small myopic demand for bonds and a large myopic demand for cash. The agent with $\gamma = 5$ switches to cash only five years before retirement, and the most aggressive investor with $\gamma = 2$ keeps investing all his wealth in stocks and long-term bonds, ending up with a myopic portfolio of about 50% in stocks and 50% in bonds at the final period. Of course, this myopic demand strongly depends on the estimates of the return characteristics of stocks and bonds.

In Figure 8, the mean allocation is given for each point in time. This mean is calculated on the basis of the simulated sample paths. Recall that excess log returns do not depend on $\boldsymbol{x}_t$ in our model, since there is no predictability. Further, the state variable does not play a role in determining the current wealth, since the investor has only financial wealth and no human capital. Therefore, the state variable has hardly any impact on the optimal allocation in this problem. This means that the optimal allocations are approximately the same for all sample paths. The quantiles of the resulting allocations are therefore very close to the mean allocation.

7.1.2 Annuitization at retirement

We now include annuities in the optimization problem, implying that the utility function is $u(\frac{W_T}{A_T}) = \frac{(W_T/A_T)^{1-\gamma}-1}{1-\gamma}$. We perform the same analysis as for the problem without annuitization. The results are presented in Figure 9.

The demand for stocks is almost the same for the problem with annuitization as it was for the problem without annuitization. However, at the end of the life-cycle, the investor replaces cash from the myopic portfolio by long-term bonds, since he wants to be hedged against conversion risk. At the end of the period, a small change in interest rates has a huge impact on the resulting utility, since the price of the real annuity depends heavily on the term structure of interest rates at $T = 40$. Nominal and real interest rates are closely related, which means that nominal 10-year bonds provide a good hedge against the interest rate risk of a real annuity. The duration of the annuity gives an approximation for this change in the value of the annuity. The duration of the annuity at retirement date T is about 9, depending on the state of the economy.

The hedging demand for bonds increases with risk-aversion, and only plays an important role for the conservative investor with $\gamma = 10$. While the demand for bonds at the end of the period is for this conservative agent much larger than in the previous situation without annuitization, we still see that he increases his allocation to cash at the very end. This can be explained by the fact that the amount of interest rate risk in the value of the annuity becomes lower as the investor approaches his retirement date, since the annuity payments are closer to that point in time. This means that the hedging demand for long-term bonds also drops.

When real interest rates are relatively high, the duration is relatively low. This means that the sensitivity to interest rate changes is lower. Therefore, the hedging demand for long-term bonds will be smaller compared to the hedging demand when interest rates are relatively low. The opposite is true when real interest rates are relatively low and the

Figure 9: Mean allocations for the three-factor model without human capital, but with annuitization at retirement.

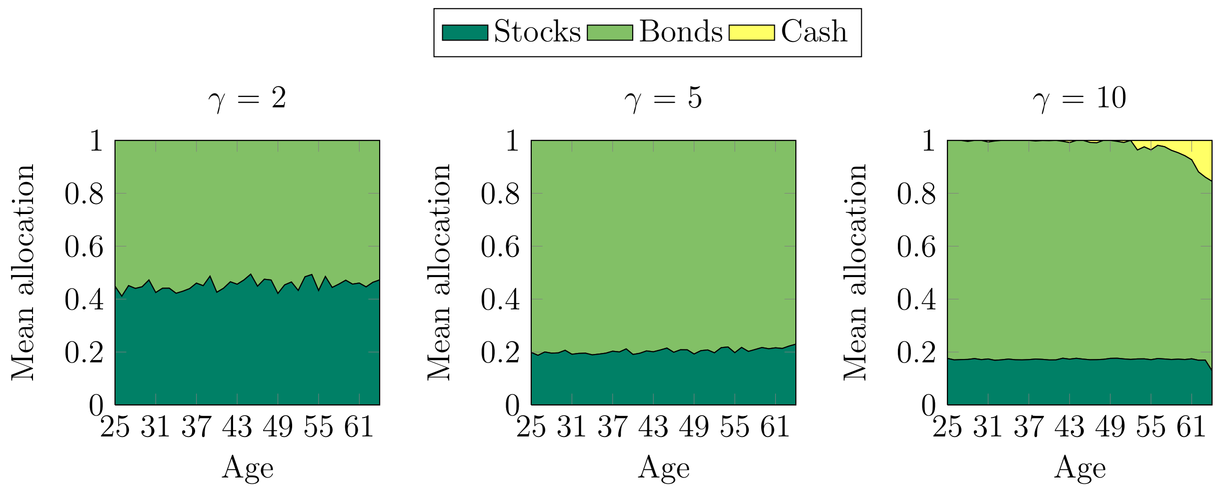

duration is relatively high. Hence, we have some state dependency, but this state dependency is negligible. The quantiles of the resulting allocations are still close to the mean allocation.

7.2 Including stochastic labor income

We now come to the results of our main optimization problem, in which agents contribute every year a part of their income to the defined contribution pension scheme. Hence, we extend on the work of Bodie et al. (1992) and related contributions that define wealth as the sum of financial wealth and human capital. The results for the model without labor income have shown that the optimal allocation is significantly different for a conservative investor when he annuitizes his wealth at retirement instead of just maximizing his real pension savings. Since we are interested in the decision for an agent that buys annuities at retirement, we explicitly include annuitization in the allocation decision problem from now on.

At each point in time, we determine for each sample path and for each point on the grid for f_t the optimal allocation to stocks and to 10-year nominal zero-coupon bonds. We use the method that was discussed in Section 6. Once we have found these $T \times N \times M$ allocations, we start at time 0 and determine the N paths for $f_t, t = 0, \ldots, T-1$, by applying the optimal strategy that we have found. Note that for all paths, the values $f_0 = \frac{\rho_0 L_0}{W_0}$ and $f_{T-1} = 1$ are the same. We analyze the optimal allocations for different specifications of labor income.

7.2.1 Systematic labor income risk only

Before we discuss the results for the most general specification of labor income in which the individual faces systematic labor income risk as well as idiosyncratic labor income risk (see Section 5.1), we first investigate the situation in which there is no idiosyncratic labor income risk. This is the situation in which the exposure to the idiosyncratic labor income risk factor η_t is equal to zero. This holds when $\sigma_u^2 = k^2 \boldsymbol{e}_s' \Sigma \Sigma' \boldsymbol{e}_s$. In that case, labor income can be perfectly replicated by the assets that are available on the financial market. This allows us to compare the resulting allocations to the results from Section 2.3, in which the risk-free rate was assumed to be constant and in which there was also no idiosyncratic labor income risk.

The resulting mean allocations to stocks, bonds, and cash at each point in time for labor income with an exposure of $k = 2\%$ to stock returns are given in Figure 10, for different values of the risk-aversion parameter γ. Due to different returns on financial wealth and human capital, optimal allocations vary over the different sample paths. However, because of constant investment opportunities in terms of expected excess returns over time, the resulting allocations are similar for all sample paths. Therefore, it suffices to show the mean allocations.

Figure 10: Mean allocations for the three-factor model with human capital ($k = 2\%$, systematic labor income risk only) and with annuitization at retirement.

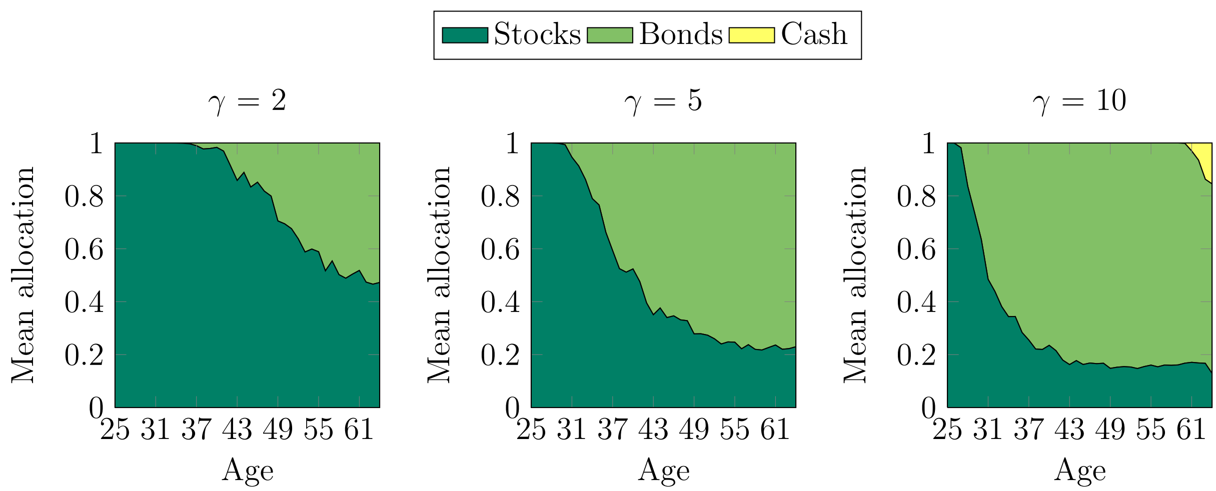

At the age of 25, the wealth of the individual consists almost exclusively of human capital. His financial wealth is the first contribution to the pension scheme, and accounts for only 1.2% of his total wealth including human capital. In the benchmark specification, labor income is much more (real) bond-like than it is stock-like. Therefore, the individual has already implicitly invested heavily in long-term bonds via his human capital. Because of diversification benefits and the relatively high expected return on stocks, it is optimal to invest financial wealth fully in stocks, even for agents with a high risk-aversion.

As we have seen for the long-term investor without human capital, the optimal allocation of total wealth is more or less constant over time (see Figure 9). As time goes by, the share of financial wealth in total wealth increases. This means that the allocation of financial wealth towards stocks should decrease and the allocation of financial wealth towards bonds should increase. In the beginning, human capital still dominates, implying that the investor mainly holds stocks. At the age of 64, the agent has no human capital left, which means that the allocation in the last period is exactly the same as the allocation of the long-term investor without human capital. In this last period, he has become a short-term investor. The higher the risk-aversion γ, the faster the investor increases his allocation of financial wealth to a portfolio mainly consisting of bonds.

The patterns that we find in Figure 10 are comparable to the patterns in Figure 1 (in Section 2.3), that shows the allocations that we found for the model in which interest rates and inflation are constant. However, the stock weights are always lower for the new model in which interest rates, inflation, and labor income are stochastic. The main reason is that stocks are replaced by long-term bonds to hedge real interest rate risk in future investment opportunities and in the price of the annuity at retirement. Hence, taking into account interest rate risk leads to a more prudent optimal allocation.

Similar to what we did in Section 2.3, we compare the optimal allocations for different values of k. In Section 5.4, we specified a labor income volatility of $\sigma_u = 5\%$. By assuming that $k = 2\%$ and that there is no idiosyncratic labor income risk, labor income volatility is negligible. We now still assume that there is no idiosyncratic labor income risk, but choose k such that the labor income volatility is close to 5%. By taking $k = 20\%$ or $k = -20\%$, we arrive at this situation. The resulting allocations are shown in Figure 11.

Figure 11: Mean allocations for the three-factor model with human capital (systematic labor income risk only) and with annuitization at retirement, for $\gamma = 5$.

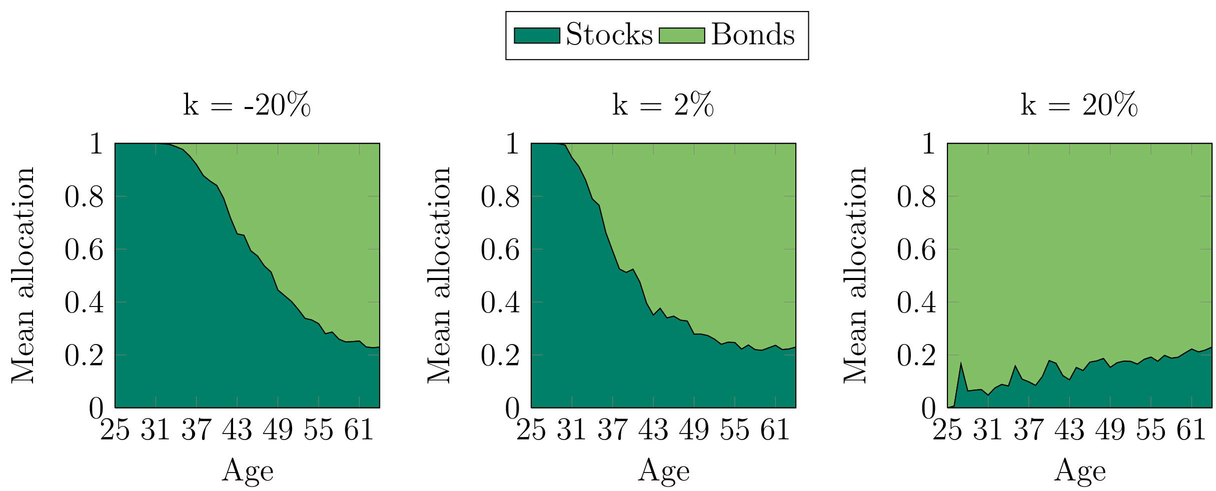

The effects of changes in the exposure of labor income to equity returns is in line with what would be expected. If labor income is relatively more stock-like, than the allocations of financial wealth in the beginning of the period are more loaded on bonds and less on stocks. Note that although the final allocation has an investment in stocks which is larger than 20%, a young life-cycle investor with $k = 20\%$ invests his wealth almost completely in bonds. This is remarkably different from the solution that we found when interest rates were assumed to be constant. The reason is that the life-cycle investor wants to be hedged against interest rate risk. The agent with $k = -20\%$ has an implicit short position in stocks, and therefore allocates a greater percentage of his financial wealth to stocks than in the benchmark case.

From the results in this section, we can conclude that the main takeaway from the analysis of the life-cycle allocation problem with constant interest rates is still valid: the nature of human capital decides the pattern of the optimal allocation over time. When labor income is relatively riskless, a young investor has already a large position in bonds. He optimally invests all his wealth in stocks, even though he now faces interest rate risk. As the individual grows older, he reduces his position in equity. This is consistent with conventional financial planning advice. However, when labor income has a large exposure to stock returns and is therefore stock-like, a young investor should invest in bonds only, and should increase his holdings in the stock as he ages. Compared to the solution when interest rates are constant, risk-averse investors of all ages should increase there bond holdings and decrease there equity holdings in order to hedge for interest rate risk.

7.2.2 Idiosyncratic labor income risk

The results for systematic labor income risk have provided us with valuable insights in how the optimal allocation changes when interest rates are stochastic rather than fixed. However, in reality individuals do not only face systematic risk. The largest uncertainty in labor income is idiosyncratic risk, i.e., exposure to an additional shock that is not related to shocks on the financial market.

Figure 12: Mean allocations for the three-factor model with human capital ($k = 2\%, \sigma_u = 5\%$) and with annuitization at retirement.

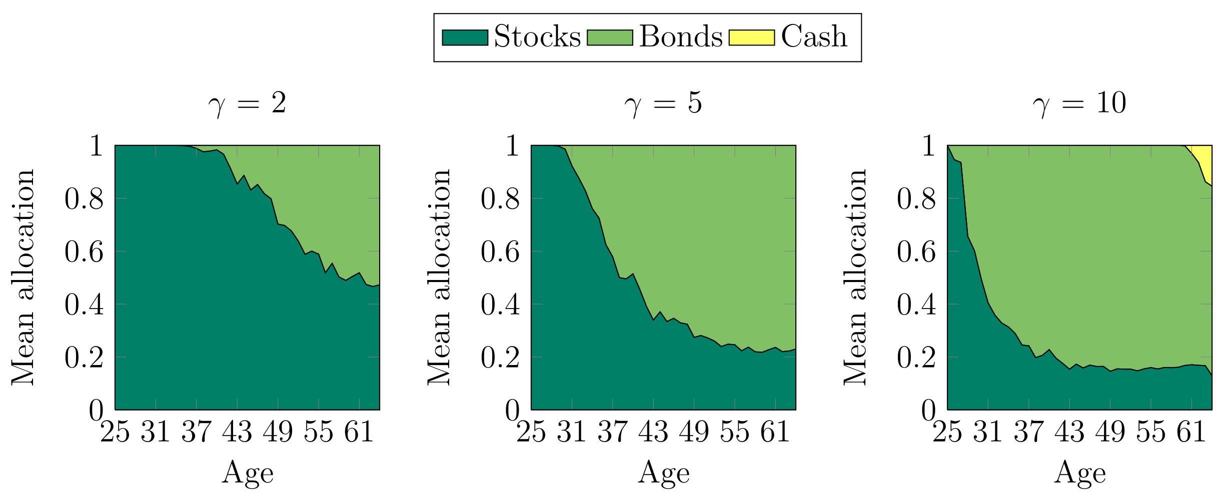

We now focus the optimal strategy for the benchmark specification of labor income as discussed in Section 5.4, where idiosyncratic labor income risk is present and dominates the systematic labor income risk. In this specification, labor income shocks have a standard deviation of 5% and labor income shocks have an expose of $k = 2\%$ to equity returns. This means that labor income is only moderately correlated to equity returns with a correlation of about 10%. The results are presented in Figure 12.

It turns out that idiosyncratic labor income risk has hardly any impact on the optimal allocation to financial assets. The only difference between the results that were presented in Figure 10 and the results that are presented here is the additional idiosyncratic labor income risk which is now present. We see that there is almost no difference in the resulting life-cycle allocation paths.

Because individuals face additional income risk that cannot be diversified away, they become somewhat more risk-averse to shocks in their financial assets. For the most risk-averse agent with $\gamma = 10$, there is a small decrease in the optimal allocation to stocks and a small increase in the allocation to bonds. As said, the results are almost negligible. Also the quantiles of the allocations for the different sample paths are almost completely the same as for situation with only systematic labor income risk. We can conclude that for the parameter values that we used, the only thing that really matters is the exposure of labor income to stock returns, and not the riskiness of labor income itself.

8 Results for four-factor model with predictability

In the previous section, we have seen how allocations over the life-cycle change if interest rates are stochastic rather than fixed. We did the analysis for a three-factor model, in which investment opportunities are constant over time. For this model, the conditional distribution of stock returns and bond returns is always the same and does not depend on the state of the world. However, in Section 4, we have seen that this simple three-factor model does not explain historical data sufficiently. To match historically observed yields more closely, we should include time-varying risk premia. This means that investment opportunities vary over time and depend on the state of the world.

We repeat the analysis of the previous section with the earlier described four-factor model that allows for time-varying risk premia. We again do this by using the optimization procedure described in Section 6. We simulate $N = 10,000$ sample paths and determine for each sample path and for each point on the M-dimensional grid for f_t the optimal allocation to stocks and to 10-year nominal zero-coupon bonds. Once we have found these $T \times N \times M$ allocations, we start at time 0 and determine the N paths for $f_t, t = 0, \ldots, T-1$, by applying the optimal strategy that we have found.

In the previous section, we have seen the results for a variety of settings of the model, including long-term investing with human capital versus investing without human capital, annuitization at retirement versus no annuitization, different exposures of labor income shocks to stock returns, different volatilities of labor income shocks, and idiosyncratic labor income risk versus only systematic labor income risk. The implications of each of these settings will remain valid for the four-factor model. We will therefore focus immediately on the model in which we are particularly invested, which is the model that uses the benchmark parameters discussed in Section 5.4 and for which the implications of the three-factor model were discussed in Section 7.2.2.

8.1 State dependency

For the three-factor model, it was sufficient to illustrate the mean allocation over all sample paths at each point in time. Although the ratio between financial wealth and human capital depends on yields and investment returns that are path-specific, the optimal allocations were very similar across the different sample paths. The reason for this was that investment opportunities (excess returns of financial assets) were constant and not path-specific. This no longer holds for the four-factor model with predictability. Here, the optimal allocations vary greatly cross-sectionally and over time. This is because expected bond returns depend heavily on the state of the world, as we have seen for instance in Figure 5 that shows bond risk premia as a function of the state variables.

The fact that bond risk premia have such a large influence on the optimal allocation makes it hard to compare allocations across different sample paths. To illustrate the variability of allocations when there is predictability and to show the effects of different bond risk premia, Figure 13 describes the optimal allocation over time for a specific sample path and also gives bond risk premia over time. Although the life-cycle effects discussed in the previous section certainly play a role, we clearly see a very strong relation between bond risk premia and the allocation to bonds. Investors should time bond markets, by allocating a large portion to bonds if expected returns are high and by avoiding bond investments if expected returns are low. In fact, from inspecting the resulting allocation paths such as the path in Figure 13, we conclude that these timing effects dominate the life-cycle effects.

Figure 13: The left graph shows a specific allocation path for the four-factor model with human capital ($k = 2\%, \sigma_u = 5\%$) and with annuitization at retirement, for an individual with $\gamma = 5$. In the right graph the allocations to bonds are compared to the corresponding bond risk premia.

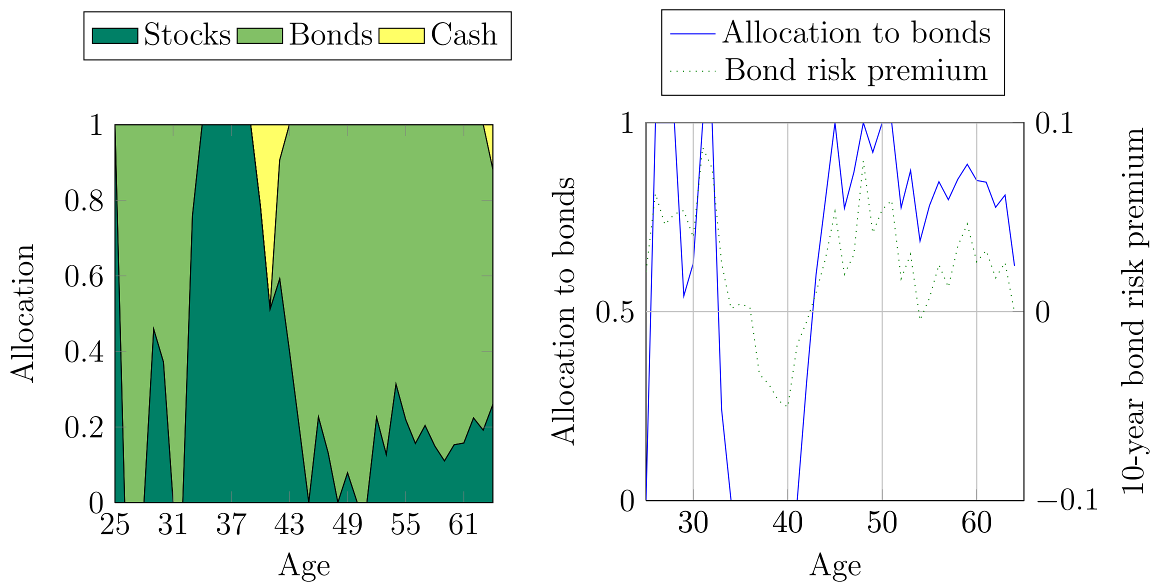

8.2 Life-cycle effects

To make a distinction between investment opportunity effects and life-cycle effects, we give the optimal allocation over time for some very special sample paths. So far, we have calculated the optimal allocations for the sample paths that we used as input for the approximation of the conditional expectations in our dynamic optimization problem. Thus, the input (sample paths to approximate conditional expectations) and output (resulting allocations for sample paths) were based on the same simulated sample paths. However, once we have approximated the first-order conditions as a function of the state variables and the ratio between human capital and financial wealth, we are able to determine the optimal allocation for any path of the state variables that we can think of. Instead of presenting the results for the sample paths that we also used as input for the optimization procedure, we present the results for three specific sample paths in which we are interested.

First of all, we consider a path in which the state variable is always equal to its long-term mean $\boldsymbol{\mu}$. This means that the error term $\boldsymbol{\epsilon}_t$ is equal to the zero vector for all t. For different values of the risk-aversion parameter γ, the resulting allocations over time are presented in Figure 14.

Figure 14: Allocations for the four-factor model with human capital ($k = 2\%, \sigma_u = 5\%$) and with annuitization at retirement for a path in which the state vector is always equal to its long-term mean $\boldsymbol{\mu}$.

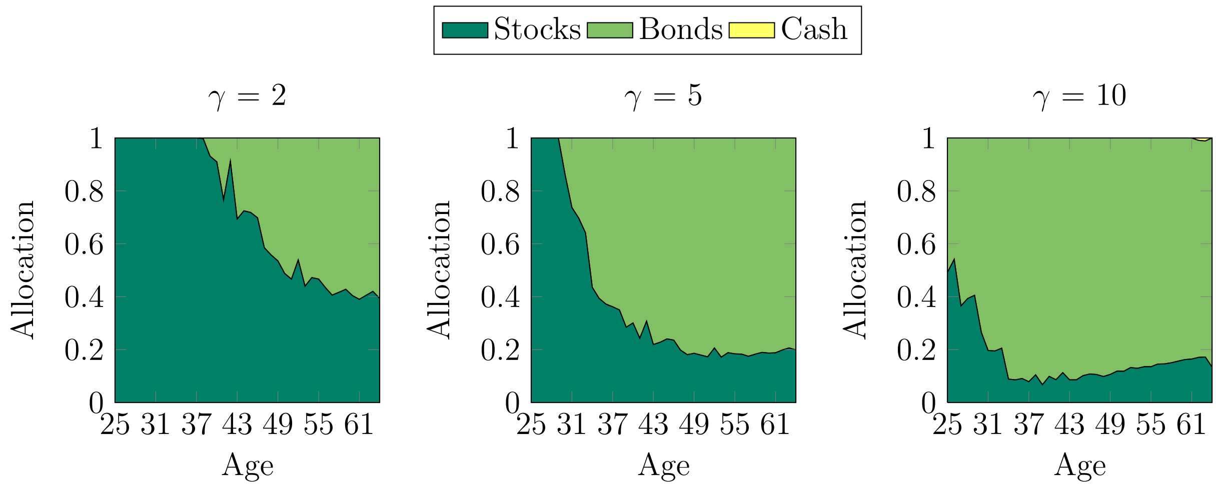

The allocations for $\boldsymbol{x}_t = \boldsymbol{\mu}$ represent the allocations under "normal" conditions. We find paths that are similar to the paths that we have seen before. Young agents with a labor income that has a 2% exposure to stock returns invest their financial wealth fully in stocks (unless they are very risk-averse), and older people reduce their equity holdings until they arrive at the myopic portfolio (which includes conversion risk hedging because of the annuitization) at the last year before retirement. However, compared to Figure 12, the allocation to stocks is lower and the allocation to bonds is higher. The more risk-averse the agent is, the larger are the differences between the two figures.

Recall that the demand for bonds consists of a myopic demand and a hedging demand. By looking at the allocation at the very end of the period, we can infer what the myopic demand is. Because we used the same data to estimate the parameters of the three-factor model and the parameters of the four-factor model, we would expect that under "normal" conditions the myopic demands for bonds in both models are comparable. By comparing again Figure 14 to Figure 12, we see that this is indeed the case. Hence, the reason for an increased demand for bonds must be an increase in the hedging demand for bonds.

In Section 4, we have already seen the reason why life-cycle investors have even more incentives to buy long-term bonds instead of stocks in this four-factor model: long-term bonds cannot only be used to hedge against adverse changes in nominal yields, but can also be used to hedge against adverse changes in bond risk premia. This is due to the negative correlation between bond returns and future nominal and real bond risk premia, see Table 10. As a result of this predictability effect, bonds are safer assets to long-term investors than to short-term investors. Since human capital is bond-like (at least, in the benchmark specification) and the investor cannot sell his human capital, he is constrained to holding a significant position in bonds via his human capital, even if bond risk premia are very low. This makes the need for a hedge against low bond risk premia particularly high. Thus, there are two ways in which long-term bonds provide a hedge against changes in future investment opportunities, and which does not apply to stocks. In Figure 14, we see that for a young conservative investor this implies that he does not offset his investment in bonds via his human capital by investing his financial wealth fully in stock, but instead invests his financial wealth partially in bonds to increase his bond holdings even further.

8.3 Different investment opportunities

It is also interesting to look at changes to the optimal allocation when bond risk premia change. In Figure 13, we have already seen that not only life-cycle effects, but also bond risk premia have a great impact on the optimal allocation. Bond risk premia are determined by the state variables. In Section 4.3, we have discussed the impact of the different state variables on bond risk premia. We have seen that past stock returns have no impact on the expected return on a 10-year nominal bond in our model. The three other state variables influence bond risk premia positively.

Figure 15: Allocations for the four-factor model with human capital ($k = 2\%, \sigma_u = 5\%$) and with annuitization at retirement for a path in which the short rate is always equal to its long-term mean plus one unconditional standard deviation, and all other state variables are equal to their long-term means.

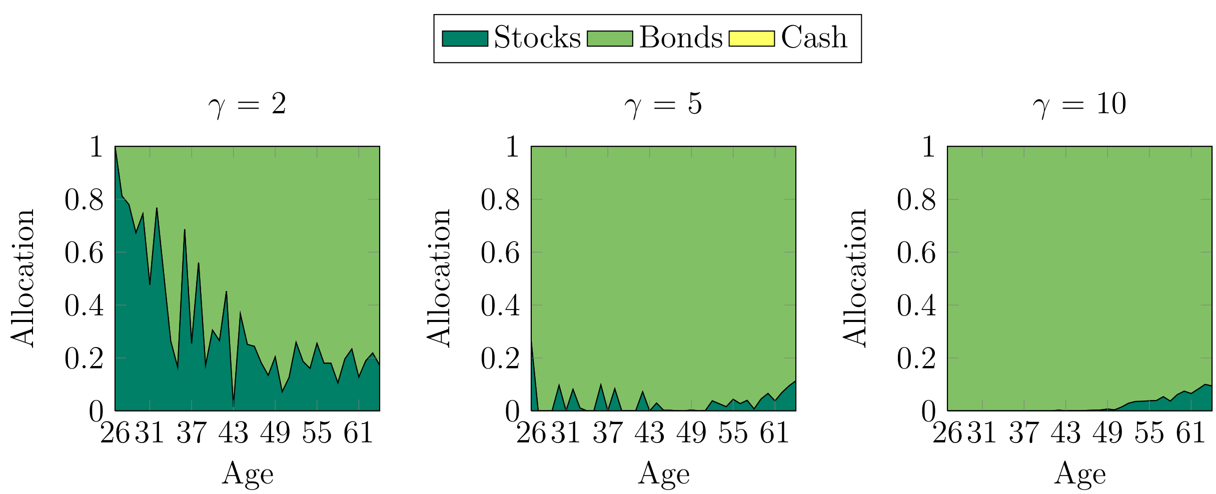

We consider a path in which the nominal short rate is always equal to 7.54%, its long-term mean (4.74%) plus one unconditional standard deviation (2.80%). All other state variables are equal to their long-term means. This means that the risk premium on bonds is relatively high. Since 10-year nominal bonds are in this case more attractive than under regular conditions, the myopic demand for bonds is higher than before. We would therefore expect increased holdings in the 10-year nominal bond. This is indeed the case, as can be seen in Figure 15.

The differences between Figure 14 and Figure 15 are huge. Under these market conditions, bonds provide a very good Sharpe ratio. Combined with the hedging benefits, bonds are very attractive. Conservative investors almost exclusively invest in the 10-year nominal bond, even when they are young and already have a huge exposure towards interest rate risk.

In contrast, when the nominal short rate is an unconditional standard deviation below its mean, bond risk premia are relatively low. This makes bonds a bad investment for a short-term investor without hedging demands. The life-cycle investor has a significant hedging demand for bonds and will therefore still invest in nominal 10-year bonds. Figure 16 illustrates the resulting allocations.

Figure 16: Allocations for the four-factor model with human capital ($k = 2\%, \sigma_u = 5\%$) and with annuitization at retirement for a path in which the short rate is always equal to its long-term mean minus one unconditional standard deviation, and all other state variables are equal to their long-term means.

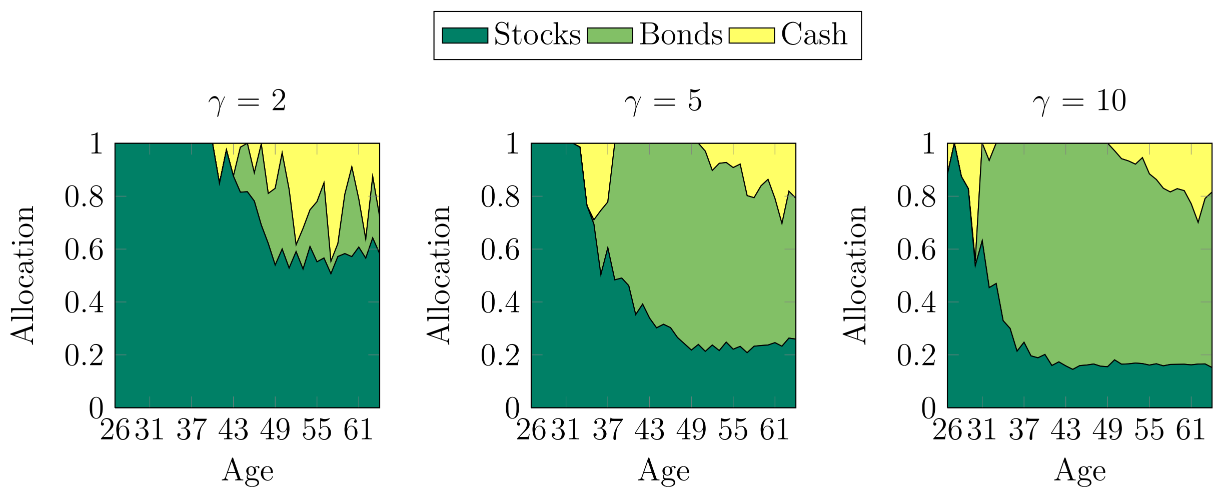

Roughly speaking, we can split up the allocation path for low bond risk premia in three different periods. In the first period, the life-cycle investor has already a huge exposure to bonds via his human capital. His hedging demands are already met via this exposure, and he has no additional myopic demand for bonds that needs to be fulfilled via his financial wealth. As time goes by, human capital shrinks and does not longer provide a sufficiently large investment in bonds to meet the hedging demands. At this middle period, the risk-averse agent invests a considerable amount in bonds, even though expected bond returns are low. At the final years before retirement, the hedging demand for bonds decreases, because the life-cycle investor is no longer exposed to the risk that future investment opportunities are bad. However, since he annuitizes his wealth at retirement, he still has a hedging demand for bonds. As a result, he replaces part of his portfolio in bonds with an investment in cash, but not the whole portfolio. This illustrates the importance of taking into account the conversion risk at retirement, because otherwise a much larger part of the investment in bonds would have been replaced by cash at the end of the period.

As said, the effects of changes in the other state variables on bond risk premia have the same sign. Therefore, they affect the optimal allocations similarly. From Figure 5 in Section 4, we infer that changes in the term spread have the largest impact on bond risk premia. Changes in the level of inflation have the least impact. Thus, changes in the term spread affect the optimal allocation in a same way as changes in the nominal short rate do, but the effects of changes in the term spread are larger. On the other hand, changes in the inflation level have similar but lower effects on the resulting allocations.

9 Conclusions

In this thesis, we have solved a realistic life-cycle optimization problem with borrowing and short-selling constraints for a participant in a defined contribution pension scheme that annuitizes his wealth at retirement. In the asset allocation decision of long-term investors, future contributions to the scheme should be taken into account. When human capital is bond-like, the optimal allocation is consistent with conventional financial planning advice. Young agents should invest a large part of their savings in stocks, since they already have a huge implicit investment in bonds. Annuitization at retirement induces extra hedging demand for bonds at the end of the period.

Compared to the solution for a model in which interest rates are constant, we found a larger allocation to bonds and a smaller allocation to stocks over the whole life-cycle. The reason is that investors want to be hedged against drops in the nominal short rate. When bond returns are predictable, investors should time bond markets. In our model, bond timing effects dominated life-cycle effects. In terms of life-cycle effects, we found an increased hedging demand for bonds, because bonds can not only be used to hedge short rate risk but also the risk of a future decrease in expected excess bond returns.

To make a gradual shift from a model with constant interest rates and constant investment opportunities to a model with stochastic interest rates and return predictability, we adopted two different models in the class of affine term structure models. Both models allowed for time-varying yields and inflation. The three-factor model was restricted such that there was no predictability in returns. While this model provided us with valuable insights in general life-cycle patterns, it did not do well in terms of explaining historically observed yields. The four-factor model extended the three-factor model by including a term spread factor and allowing for bond return predictability. This model did much better in explaining past yields.

Our model for labor income accounted for systematic as well as idiosyncratic labor income risk. If labor income is correlated to stock returns, agents have an implicit investment in stocks via their human capital. We have seen that the exposure of labor income to shocks on the equity market is an important determinant of life-cycle allocation paths. If this exposure is high, human capital is more stock-like than bond-like, and young agents should invest their financial wealth fully in bonds rather than in stocks. Including (some) idiosyncratic labor income risk in our model had hardly any impact on the resulting optimal allocations over time.

Since in practice pension contributions are often predetermined, we fixed future consumption and savings. This changes optimal allocations, since it is no longer possible to smooth consumption over time. In order to solve the optimization problem with dynamic programming, we adopted the simulation-based method of Koijen et al. (2010). We chose to value human capital explicitly and to make a grid for the endogenous fraction of financial wealth in total wealth. Since this fraction is always between zero and one, this made the optimization procedure more tractable. Our optimization framework is not restricted to the specific optimization problem and is very generally applicable. It can for instance also readily be used to solve problems with an increased number of financial assets, a different type of labor income, and an increased number of state variables.

There are several ways in which this research can be extended. One could aim for a more general model of the economy, that allows for instance for heteroskedasticity in the state variables or time-varying means because of changing inflation targets. Further, the optimization problem might be extended, by including the housing decision into the problem or by modeling an endogenous choice of annuitization. One could also allow for stock return predictability by including possible predictor variables such as dividend yield in the model. Since this is a heavily debated subject, we abstracted from stock predictability in this thesis. Finally, the model for labor income might be improved by allowing for cointegration between labor income and dividend yield (see Benzoni et al. (2007)), or by including the risk of becoming unemployed and having no labor income at all.

References

Ang, A., and Piazzesi, M., 2003, "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables", Journal of Monetary Economics, 50(4) 745-787.

Ang, A., Piazzesi, M., and Wei, M., 2006, "What Does The Yield Curve Tell Us About GDP Growth?", Journal of Econometrics, 131(1-2) 359-403.

Benzoni, L., Collin-Dufresne, P., and Goldstein, R.S., 2007, "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets are Cointegrated", Journal of Finance, 62(5) 2123-2167.

Blake, D., Wright, D., and Zhang, Y., 2011a, "Target-Driven Investing: Optimal Investment Strategies in Defined Contribution Pension Plans under Loss Aversion", Pension Institute Working Paper.

Blake, D., Wright, D., and Zhang, Y., 2011b, "Age-Dependent Investing: Optimal Funding and Investment strategies in Defined Contribution Pension Plans when Members are Rational Life Cycle Financial Planners", Pension Institute Working Paper.

Bodie, Z., Detemple, J.B., Otruba, S., and Walter, S., 2004, "Optimal Consumption-Portfolio Choices and Retirement Planning", Journal of Economic Dynamics and Control, 28(6) 1115-1148.

Bodie, Z., Merton, R.C., and Samuelson, W.F., 1992, "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model", Journal of Economic Dynamics and Control, 16(3-4) 427-449.

Brandt, M.W., Goyal, A., Santa-Clara, P., and Stroud, J.R., 2005, "A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Return Predictability", The Review of Financial Studies, 18(3) 831-873.

Brennan, M.J., and Xia, Y., 2002, "Dynamic Asset Allocation under Inflation", The Journal of Finance, 57(3) 1201-1238.

Campbell, J.Y., 1987, "Stock Returns and the Term Structure", Journal of Financial Economics, 18(2) 373-399.

Campbell, J.Y., and Viceira, L.M., 2001, "Who Should Buy Long-Term Bonds?", American Economic Review, 87(1) 181-191.

Campbell, J.Y., and Viceira, L.M., 2002, "Strategic Asset Allocation: Portfolio Choice for Long Term Investors", Oxford University Press.

Campbell, J.Y., and Viceira, L.M., 2003, "A Multivariate Model of Strategic Asset Allocation", Journal of Financial Economics, 67(1) 41-80.

Carroll, C.D., 2006, "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems", Economics Letters, 91(3) 312-320.

Cocco, J., Gomes, F., and Maenhout, P., 2005, "Consumption and Portfolio Choice over the Life Cycle", Review of Financial Studies, 18(2) 491-533.

Cochrane, J.H., 2005, "Asset Pricing", revised edition, Princeton University Press.

Cochrane, J.H., and Piazzesi, M., 2005, "Bond Risk Premia", American Economic Review, 95(1) 138-160.

Cox, J.C., Ingersoll, J.E., and Ross, S.A., 1985, "A Theory of the Term Structure of Interest Rates", Econometrica, 53(2) 385-408.

de Jong, F., 2008, "Valuation of Pension Liabilities in Incomplete Markets", Journal of Pension Economics and Finance, 7(3) 277-294.

Duffee, G.R., 2002, "Term Premia and Interest Rate Forecasts in Affine Models", Journal of Finance, 57(1) 405-443.

Flanders, H., 1973, "Differentiation under the Integral Sign", The American Mathematical Monthly, 80(6) 615-627.

Glosten, L.R., Jagannathan, R., and Runkle, D., 1993, "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks", Journal of Finance, 48(2) 663-679.

Goyal, A., and Welch, I., 2008, "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction", Review of Financial Studies, 21(4) 1455-1509.

Harvey, C., 1989, "Time-Varying Conditional Covariances in Tests of Asset Pricing Models", Journal of Financial Economics, 22(2) 305-334.

Harvey, C., 1991, "The World Price of Covariance Risk", Journal of Finance, 46(1) 111-157.

Haubrich, J., Pennacchi, G., and Ritchken, P., 2012, "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps", Review of Financial Studies, 25(5) 1588-1629.

Koijen, R.S.J., and Nijman, Th.E., 2006, "Valuation and Risk Management of Inflation Sensitive Pension Rights", in: N. Kortleve, T.E. Nijman, and E. Ponds (Eds.), "Fair Value and Pension Fund Management", Amsterdam: Elsevier Science.

Koijen, R.S.J., Nijman, Th.E., and Werker, B.J.M., 2007, "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.

Koijen, R.S.J., Nijman, Th.E., and Werker, B.J.M., 2010, "When Can Life-cycle Investors Benefit from Time-varying Bond Risk Premia?", Review of Financial Studies, 23(2) 741-780.

Koijen, R.S.J., Nijman, Th.E., and Werker, B.J.M., 2011, "Optimal Annuity Risk Management", Review of Finance, 15(4) 799-833.

Litterman, R., and Scheinkman, J., 1991, "Common Factors Affecting Bond Returns", Journal of Fixed Income, June 54-61.

Longstaff, F.A., and Schwartz, E.S., 2001, "Valuing American Options by Simulation: A Simple Least-Squares Approach", Review of Financial Studies, 14(1) 113-148.

Merton, R.C., 1969, "Lifetime Portfolio Selection Under Uncertainty: The Continuous Time Case", Review of Economics and Statistics, 51(3) 247-257.

Merton, R.C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model", Journal of Economic Theory, 3(4) 373-413.

Munk, C., and Sørensen, C., 2010, "Dynamic Asset Allocation with Stochastic Income and Interest Rates", Journal of Financial Economics, 96(3) 433-462.

Palacios-Huerta, I., 2003, "An Empirical Analysis of the Risk Properties of Human Capital Returns", The American Economic Review, 93(3) 948-964.

Piazzesi, M., 2010, "Affine Term Structure Models", Handbook of Financial Econometrics, Elsevier, chapter 12.

Samuelson, P., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming", Review of Economics and Statistics, 51(3) 239-246.

Sangvinatsos, A., and Wachter, J.A., 2005, "Does the Failure of the Expectation Hypothesis Matter for Long-Term Investors?", Journal of Finance, 60(1) 179-230.

Vasicek, O., 1977, "An Equilibrium Characterization of the Term Structure of Interest Rates", Journal of Financial Economics, 5(2) 177-188.

Wachter, J.A., 2003, "Risk Aversion and Allocation to Long-Term Bonds", Journal of Economic Theory, 112(2) 325-333.

Yaari, M.E., 1965, "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer", Review of Economic Studies, 32(2) 137-150.

Zellner, A., 1962, "An Efficient Method of Estimating Seemingly Unrelated Regression Equations and Tests for Aggregation Bias", Journal of the American Statistical Association, 57(298) 348-368.

A Myopic portfolio

Proof of Theorem 1.

For the short-term investor that wants to maximize his expected utility at the end of the period, we have the maximization problem (3):

$$\begin{array}{ll} \max & \mathbb{E}_t u(W_{t+1}) \\ \text{subject to} & W_{t+1} = (1 + R_{p,t+1})W_t. \end{array}$$

We assume that investors have a power utility function (4). Further, we assume that the return on the risky asset follows a lognormal distribution. In that case, next-period wealth W_{t+1} is also lognormal. This means that $w_{t+1} \equiv \log(W_{t+1})$ has a normal distribution. We know that if $X \sim N(\mu, \sigma^2)$, it holds that $\mathbb{E}e^X = e^{\mu+0.5\sigma^2}$. We rewrite the optimization problem as

$$\begin{aligned} \max \log \mathbb{E}_t \frac{W_{t+1}^{1-\gamma} - 1}{1-\gamma} &= -\log(1-\gamma) + \max \log \mathbb{E}_t \exp\{(1-\gamma)w_{t+1}\} \\ &= -\log(1-\gamma) + \max \left\{ (1-\gamma)\mathbb{E}_t w_{t+1} + \frac{1}{2}(1-\gamma)^2 \sigma_{w,t}^2 \right\}, \quad (83) \end{aligned}$$

where $w_{t+1} = w_t + r_{p,t+1}$, with $r_{p,t+1} \equiv \log(1 + R_{p,t+1})$. Let $\sigma_{p,t}^2 \equiv \mathrm{Var}_t(r_{p,t+1})$ be the conditional variance of the log portfolio return.

Plugging in the budget constraint and canceling out common constants leads to the reduced problem

$$\max \left\{ \mathbb{E}_t r_{p,t+1} + \frac{1}{2}(1-\gamma)\sigma_{p,t}^2 \right\}. \quad (84)$$

The advantage of log returns is that they can be aggregated over time. Unfortunately, this comes at the cost of being unable to find portfolio returns as a linear combination of the returns on the single assets. Therefore, we will approximate the log portfolio return, which is given as follows:

$$1 + R_{p,t+1} = 1 + R_{f,t} + \alpha_t(R_{t+1} - R_{f,t}) \Leftrightarrow \frac{1 + R_{p,t+1}}{1 + R_{f,t}} = 1 + \alpha_t \left(\frac{1 + R_{t+1}}{1 + R_{f,t}} - 1 \right). \quad (85)$$

By taking logarithms again, we find

$$r_{p,t+1} - r_{f,t} = \log(1 + \alpha_t(\exp(r_{t+1} - r_{f,t}) - 1)) \equiv f(r_{t+1} - r_{f,t}). \quad (86)$$

We use a second-order Taylor expansion to approximate this expression:

$$f(r_{t+1} - r_{f,t}) \approx f(0) + f'(0)(r_{t+1} - r_{f,t}) + \frac{f''(0)}{2}(r_{t+1} - r_{f,t})^2,$$

where

$$\begin{aligned} f'(x) &= \frac{1}{1 + \alpha_t(\exp(x) - 1)}\alpha_t \exp(x), \\ f''(x) &= \frac{1}{1 + \alpha_t(\exp(x) - 1)}\alpha_t \exp(x) - \frac{1}{(1 + \alpha_t(\exp(x) - 1))^2}\alpha_t^2 \exp(2x). \end{aligned}$$

Hence, we find

$$r_{p,t+1} - r_{f,t} = \alpha_t(r_{t+1} - r_{f,t}) + \frac{1}{2}\alpha_t(1-\alpha_t)(r_{t+1} - r_{f,t})^2. \tag{87}$$

If we further approximate the second-order term $(r_{t+1} - r_{f,t})^2$ by its conditional expectation σ_t^2, we obtain

$$r_{p,t+1} - r_{f,t} = \alpha_t(r_{t+1} - r_{f,t}) + \frac{1}{2}\alpha_t(1-\alpha_t)\sigma_t^2. \tag{88}$$

This leads to an expected excess return of $\mathbb{E}_t(r_{p,t+1} - r_{f,t}) = \alpha_t(\mathbb{E}_t(r_{t+1}) - r_{f,t}) + \frac{1}{2}\alpha_t(1-\alpha_t)\sigma_t^2$ and a variance of $\sigma_{p,t}^2 = \alpha_t^2\sigma_t^2$. The optimization problem becomes

$$\max\left\{\alpha_t(\mathbb{E}_t(r_{t+1}) - r_{f,t}) + \frac{1}{2}\alpha_t(1-\alpha_t)\sigma_t^2 + \frac{1}{2}(1-\gamma)\alpha_t^2\sigma_t^2\right\}. \tag{89}$$

The first-order condition now yields

$$\begin{aligned} &\mathbb{E}_t(r_{t+1}) - r_{f,t} + \frac{1}{2}(1-\alpha_t)\sigma_t^2 - \frac{1}{2}\alpha_t\sigma_t^2 + (1-\gamma)\alpha_t\sigma_t^2 = 0 \Rightarrow \\ &\mathbb{E}_t(r_{t+1}) - r_{f,t} + \frac{1}{2}\sigma_t^2 - \gamma\alpha_t\sigma_t^2 = 0 \Leftrightarrow \alpha_t = \frac{\mathbb{E}_t(r_{t+1}) - r_{f,t} + \frac{1}{2}\sigma_t^2}{\gamma\sigma_t^2}. \end{aligned} \tag{90}$$

Campbell and Viceira (2001) show that the approximation error converges to zero in the continuous time limit.

B Derivation of nominal term structure

By induction, we can show that the price at time t of a nominal bond with maturity n satisfies (23):

$$P_t^{(n)} = \exp(-A(n) - \boldsymbol{B}(n)'\boldsymbol{x}_t).$$

First of all, recall that $P_t^{(1)} = \exp\left(-y_t^{(1)}\right)$, so the equation is satisfied for maturity 1 ($A(1) = 0$ and $\boldsymbol{B}(1) = \boldsymbol{\delta}_y$). Now, suppose that the equation is valid up to maturity $n-1$. Since the distribution of M_{t+1} and $P_{t+1}^{(n-1)}$ conditional on the information at time t is lognormal, the product $M_{t+1}P_{t+1}^{(n-1)}$ is also lognormally distributed. Again, we use the result

$$\log(\mathbb{E}_t(X_{t+1})) = \mathbb{E}_t(\log(X_{t+1})) + \frac{1}{2}\mathrm{Var}_t(\log(X_{t+1}))$$

for a lognormally distributed random variable X_{t+1}. We obtain

$$\begin{aligned}
\log\left(P_t^{(n)}\right) &= \log\left(\mathbb{E}_t\left(M_{t+1}P_{t+1}^{(n-1)}\right)\right) \\
&= \mathbb{E}_t\left(\log\left(P_{t+1}^{(n-1)}\right) + \log M_{t+1}\right) + \frac{1}{2}\mathrm{Var}_t\left(\log\left(P_{t+1}^{(n-1)}\right) + \log(M_{t+1})\right) \\
&= -A(n-1) - \boldsymbol{B}(n-1)'(\boldsymbol{\mu} + \Gamma(\boldsymbol{x}_t - \boldsymbol{\mu})) - \boldsymbol{\delta}_y'\boldsymbol{x}_t \\
&\quad + \frac{1}{2}\boldsymbol{B}(n-1)'\Sigma\Sigma'\boldsymbol{B}(n-1) + \boldsymbol{B}(n-1)'\Sigma\boldsymbol{\lambda}_t \\
&= -A(n-1) - \boldsymbol{B}(n-1)'(\boldsymbol{\mu} + \Gamma(\boldsymbol{x}_t - \boldsymbol{\mu})) - \boldsymbol{\delta}_y'\boldsymbol{x}_t \\
&\quad + \frac{1}{2}\boldsymbol{B}(n-1)'\Sigma\Sigma'\boldsymbol{B}(n-1) + \boldsymbol{B}(n-1)'\Sigma(\boldsymbol{\Lambda}_0 + \Lambda_1\boldsymbol{x}_t) \\
&= -A(n) - \boldsymbol{B}(n)'\boldsymbol{x}_t,
\end{aligned} \tag{91}$$

where

$$\begin{aligned}
A(n) &= A(n-1) + \boldsymbol{B}(n-1)'(I-\Gamma)\boldsymbol{\mu} - \frac{1}{2}\boldsymbol{B}(n-1)'\Sigma\Sigma'\boldsymbol{B}(n-1) - \boldsymbol{B}(n-1)'\Sigma\boldsymbol{\Lambda}_0 \\
\boldsymbol{B}(n) &= (\Gamma - \Sigma\Lambda_1)'\boldsymbol{B}(n-1) + \boldsymbol{\delta}_y.
\end{aligned} \tag{92}$$

The initial conditions are given by $A(1) = 0$ and $\boldsymbol{B}(1) = \boldsymbol{\delta}_y$. An explicit solution to the recurrence relation for $\boldsymbol{B}(n)$ is given by $\boldsymbol{B}(n) = (I - \Gamma' + \Lambda_1'\Sigma')^{-1}(I - (\Gamma' - \Lambda_1'\Sigma')^n)\boldsymbol{\delta}_y$.

Now that we have found an expression for the term structure of nominal bonds, we are able to determine holding returns of nominal bonds. We do this again by recognizing that a bond with maturity n at time t has become a bond with maturity $n-1$ at time $t+1$.

Hence, we obtain:

$$\begin{aligned} r_{b,t+1}^{(n)} &= \log\left(P_{t+1}^{(n-1)}\right) - \log\left(P_t^{(n)}\right) \\ &= A(n) - A(n-1) + \boldsymbol{B}(n)'\boldsymbol{x}_t - \boldsymbol{B}(n-1)'\boldsymbol{x}_{t+1} \\ &= y_t^{(1)} - \frac{1}{2}\boldsymbol{B}(n-1)'\Sigma\Sigma'\boldsymbol{B}(n-1) - \boldsymbol{B}(n-1)'\Sigma\boldsymbol{\Lambda}_0 \\ &\quad +(\boldsymbol{B}(n)' - \boldsymbol{B}(n-1)'\Gamma - \boldsymbol{e}_1')\boldsymbol{x}_t - \boldsymbol{B}(n-1)'\boldsymbol{\epsilon}_{t+1} \\ &= y_t^{(1)} - \frac{1}{2}\boldsymbol{B}(n-1)'\Sigma\Sigma'\boldsymbol{B}(n-1) - \boldsymbol{B}(n-1)'\Sigma(\boldsymbol{\Lambda}_0 + \Lambda_1\boldsymbol{x}_t) - \boldsymbol{B}(n-1)'\boldsymbol{\epsilon}_{t+1} \\ &= y_t^{(1)} - \frac{1}{2}\boldsymbol{B}(n-1)'\Sigma\Sigma'\boldsymbol{B}(n-1) - \boldsymbol{B}(n-1)'\Sigma\boldsymbol{\lambda}_t - \boldsymbol{B}(n-1)'\boldsymbol{\epsilon}_{t+1}. \end{aligned} \tag{93}$$

C Derivation of real term structure

We will derive the real term structure in a similar way as we did for the nominal term structure. This means that we will again apply an induction argument. We denote the price of an n-year real bond by $P_t^{R(n)}$. We show that the price at time t of a real bond with maturity n satisfies (30):

$$P_t^{R(n)} = \exp(-A^R(n) - \boldsymbol{B}^R(n)'\boldsymbol{x}_t).$$

First of all, we determine the value of an inflation-linked bond with a maturity of one year. The nominal payoff of an n-year real bond one year from now is $P_{t+1}^{R(n-1)} \cdot \frac{\Pi_{t+1}}{\Pi_t}$. Since we have modeled inflation as an affine function of the state variables, the whole product $M_{t+1}P_{t+1}^{R(n-1)}\frac{\Pi_{t+1}}{\Pi_t}$ is lognormally distributed. For the real bond with $n = 1$, we have a payoff of $P_{t+1}^{R(0)}\frac{\Pi_{t+1}}{\Pi_t} = 1 \cdot \frac{\Pi_{t+1}}{\Pi_t} = \exp(\pi_{t+1})$. We determine the value of this bond as follows:

$$\begin{aligned}
\log\left(P_t^{R(1)}\right) &= \log\left(\mathbb{E}_t\left[M_{t+1}P_{t+1}^{R(0)}\frac{\Pi_{t+1}}{\Pi_t}\right]\right) = \log(\mathbb{E}_t[M_{t+1}\exp(\pi_{t+1})]) \\
&= \mathbb{E}_t[\log(M_{t+1}) + \pi_{t+1}] + \frac{1}{2}\mathrm{Var}_t[\log(M_{t+1}) + \pi_{t+1}] \\
&= -\boldsymbol{\delta}_y'\boldsymbol{x}_t - \frac{1}{2}\boldsymbol{\lambda}_t'\boldsymbol{\lambda}_t + \boldsymbol{\delta}_\pi'(\boldsymbol{\mu} + \Gamma(\boldsymbol{x}_t - \boldsymbol{\mu})) \\
&\quad + \frac{1}{2}\boldsymbol{\lambda}_t'\boldsymbol{\lambda}_t + \frac{1}{2}\boldsymbol{\delta}_\pi'\Sigma\Sigma'\boldsymbol{\delta}_\pi - \boldsymbol{\delta}_\pi'\Sigma\boldsymbol{\lambda}_t \\
&= \boldsymbol{\delta}_\pi'(I - \Gamma)\boldsymbol{\mu} + \frac{1}{2}\boldsymbol{\delta}_\pi'\Sigma\Sigma'\boldsymbol{\delta}_{\pi 1} - \boldsymbol{\delta}_\pi'\Sigma\boldsymbol{\lambda}_t - (\boldsymbol{\delta}_y - \Gamma'\boldsymbol{\delta}_\pi)'\boldsymbol{x}_t \\
&= \boldsymbol{\delta}_\pi'(I - \Gamma)\boldsymbol{\mu} + \frac{1}{2}\boldsymbol{\delta}_\pi'\Sigma\Sigma'\boldsymbol{\delta}_\pi \\
&\quad - \boldsymbol{\delta}_\pi'\Sigma(\boldsymbol{\Lambda}_0 + \Lambda_1\boldsymbol{x}_t) - (\boldsymbol{\delta}_y - \Gamma'\boldsymbol{\delta}_\pi)'\boldsymbol{x}_t \\
&= -A^R(1) - \boldsymbol{B}^R(1)'\boldsymbol{x}_t,
\end{aligned} \tag{94}$$

where

$$\begin{aligned}
A^R(1) &= -\boldsymbol{\delta}_\pi'(I - \Gamma)\boldsymbol{\mu} - \frac{1}{2}\boldsymbol{\delta}_\pi'\Sigma\Sigma'\boldsymbol{\delta}_\pi + \boldsymbol{\delta}_\pi'\Sigma\boldsymbol{\Lambda}_0 \\
\boldsymbol{B}^R(1) &= \boldsymbol{\delta}_y + (\Sigma\Lambda_1 - \Gamma)'\boldsymbol{\delta}_\pi.
\end{aligned} \tag{95}$$

Hence, the equation $P_t^{R(1)} = \exp(-A^R(1) - \boldsymbol{B}^R(1)'\boldsymbol{x}_t)$ is satisfied. Now suppose that the pricing equation $P_t^{R(n)} = \exp(-A^R(n) - \boldsymbol{B}^R(n)'\boldsymbol{x}_t)$ is satisfied up to maturity $n - 1$. We subsequently obtain

$$
\begin{aligned}
\log\left(P_t^{R(n)}\right) &= \log\left(\mathbb{E}_t\left(M_{t+1}P_{t+1}^{R(n-1)}\frac{\Pi_{t+1}}{\Pi_t}\right)\right) \\
&= \mathbb{E}_t\left(\log\left(P_{t+1}^{R(n-1)}\right)+\pi_{t+1}+\log(M_{t+1})\right) \\
&\quad +\frac{1}{2}\mathrm{Var}_t\left(\log\left(P_{t+1}^{R(n-1)}\right)+\pi_{t+1}+\log(M_{t+1})\right) \\
&= -A^R(n-1)-(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'(\boldsymbol{\mu}+\Gamma(\boldsymbol{x}_t-\boldsymbol{\mu}))-\boldsymbol{\delta}_y'\boldsymbol{x}_t \\
&\quad -\frac{1}{2}\boldsymbol{\lambda}_t'\boldsymbol{\lambda}_t+\frac{1}{2}\boldsymbol{B}^R(n-1)'\Sigma\Sigma'\boldsymbol{B}^R(n-1)+\frac{1}{2}\boldsymbol{\delta}_\pi'\Sigma\Sigma'\boldsymbol{\delta}_\pi+\frac{1}{2}\boldsymbol{\lambda}_t'\boldsymbol{\lambda}_t \\
&\quad -\boldsymbol{B}^R(n-1)'\Sigma\Sigma'\boldsymbol{\delta}_\pi+\boldsymbol{B}^R(n-1)'\Sigma\boldsymbol{\lambda}_t-\boldsymbol{\delta}_\pi'\Sigma\boldsymbol{\lambda}_t \\
&= -A^R(n-1)-(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'(\boldsymbol{\mu}+\Gamma(\boldsymbol{x}_t-\boldsymbol{\mu}))-\boldsymbol{\delta}_y'\boldsymbol{x}_t \\
&\quad +\frac{1}{2}(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\Sigma\Sigma'(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi) \\
&\quad +(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\Sigma(\boldsymbol{\Lambda}_0+\Lambda_1\boldsymbol{x}_t) \\
&= -A^R(n)-\boldsymbol{B}^R(n)'\boldsymbol{x}_t, \qquad (96)
\end{aligned}
$$

where

$$
\begin{aligned}
A^R(n) &= A^R(n-1)+(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'(I-\Gamma)\boldsymbol{\mu} \\
&\quad -\frac{1}{2}(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\Sigma\Sigma'(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi) \\
&\quad -(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\Sigma\boldsymbol{\Lambda}_0 \\
\boldsymbol{B}^R(n) &= (\Gamma-\Sigma\Lambda_1)'(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)+\boldsymbol{\delta}_y. \qquad (97)
\end{aligned}
$$

The initial conditions are given by (95). An explicit solution to the recurrence relation for $\boldsymbol{B}^R(n)$ is given by $\boldsymbol{B}^R(n)=(I-\Gamma'+\Lambda_1'\Sigma')^{-1}(I-(\Gamma'-\Lambda_1'\Sigma')^n)(\boldsymbol{\delta}_y+(\Lambda_1'\Sigma'-\Gamma')\boldsymbol{\delta}_\pi)$.

We now determine the nominal return on an n-year zero-coupon real bond in the period from t to $t+1$. At time $t+1$, the n-year real bond has a nominal return of

$$
\frac{P_{t+1}^{R(n-1)}\Pi_{t+1}}{P_t^{R(n)}\Pi_t}.
$$

This means that the nominal log return on an n-year real bond is given by

$$
\begin{aligned}
r_{b,t+1}^{R(n)} &= \log\left(P_{t+1}^{R(n-1)}\right)+\pi_{t+1}-\log\left(P_t^{R(n)}\right) \\
&= A^R(n)-A^R(n-1)+\boldsymbol{B}^R(n)'\boldsymbol{x}_t-\boldsymbol{B}^R(n-1)'\boldsymbol{x}_{t+1}+\boldsymbol{\delta}_\pi'\boldsymbol{x}_{t+1} \\
&= y_t^{(1)}-\frac{1}{2}(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\Sigma\Sigma'(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)-(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\Sigma\boldsymbol{\Lambda}_0 \\
&\quad +(\boldsymbol{B}^R(n)'-(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\Gamma-\boldsymbol{\delta}_y)\boldsymbol{x}_t-(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\boldsymbol{\epsilon}_{t+1} \\
&= y_t^{(1)}-\frac{1}{2}(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\Sigma\Sigma'(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi) \\
&\quad -(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\Sigma(\boldsymbol{\Lambda}_0+\Lambda_1\boldsymbol{x}_t)-(\boldsymbol{B}^R(n-1)-\boldsymbol{\delta}_\pi)'\boldsymbol{\epsilon}_{t+1}. \qquad (98)
\end{aligned}
$$

D Standard errors of VAR parameters

In this appendix, we describe the method for finding the standard errors of the parameter estimates for the VAR model. Recall from Section 3.3.1 that the estimation procedure for these VAR parameters consisted of two steps. Since we imposed the restriction that the long-term means of the state variables should be equal to the full-sample means, we first estimate the parameter $\boldsymbol{\mu}$. Afterwards, we calculate the other parameters Γ and Σ by using maximum likelihood.

Step 1
In the first step, we estimate the long-term mean $\boldsymbol{\mu}$ with its sample equivalent:

$$\hat{\boldsymbol{\mu}} = \frac{1}{T+1}\sum_{t=0}^{T} \boldsymbol{x_t}.$$

From the Central Limit Theorem, it follows that

$$\sqrt{T+1}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}) \to N(0, S), \tag{99}$$

where S denotes the unconditional covariance matrix of the state vector $\boldsymbol{x}_t$. This already gives us the standard errors for $\hat{\boldsymbol{\mu}}$.

Step 2
In the second step, we estimate the parameters Γ and Σ on the basis of maximum likelihood. Without focusing on technical details, we assume that the maximum likelihood estimator satisfies all regularity conditions that are needed to derive its well-known asymptotic properties. Let $\boldsymbol{\theta} \equiv \{\Gamma, \Sigma\}$ denote the vector of all VAR parameters that are estimated in this second step. We have

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \frac{1}{T}\sum_{t=1}^{T} \log f(\boldsymbol{x}_{t-1}, \boldsymbol{x}_t | \hat{\boldsymbol{\mu}}, \boldsymbol{\theta}). \tag{100}$$

The first-order condition yields

$$\frac{1}{T}\sum_{t=1}^{T} \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\boldsymbol{x}_{t-1}, \boldsymbol{x}_t | \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\theta}}) = 0. \tag{101}$$

We expand the function on the left side of the equality sign by using a first-order Taylor approximation around the true parameter values $\boldsymbol{\mu}$ and $\boldsymbol{\theta}$:

$$\begin{aligned} \frac{1}{T}\sum_{t=1}^{T} \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\boldsymbol{x}_{t-1}, \boldsymbol{x}_t | \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\theta}}) &= \frac{1}{T}\sum_{t=1}^{T} \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\boldsymbol{x}_{t-1}, \boldsymbol{x}_t | \boldsymbol{\mu}, \boldsymbol{\theta}) \\ &+ \frac{1}{T}\sum_{t=1}^{T} \frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \log f(\boldsymbol{x}_{t-1}, \boldsymbol{x}_t | \boldsymbol{\mu}, \boldsymbol{\theta})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) + \frac{1}{T}\sum_{t=1}^{T} \frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\mu}'} \log f(\boldsymbol{x}_{t-1}, \boldsymbol{x}_t | \boldsymbol{\mu}, \boldsymbol{\theta})(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}). \end{aligned} \tag{102}$$

By applying again the Central Limit Theorem, we obtain

$$\sqrt{T}\frac{1}{T}\sum_{t=1}^{T}\frac{\partial}{\partial\boldsymbol{\theta}}\log f(\boldsymbol{x}_{t-1},\boldsymbol{x}_t|\boldsymbol{\mu},\boldsymbol{\theta}) \to N(0,H_1), \tag{103}$$

where H_1 is the so-called Fisher information defined by

$$H_1 \equiv \mathbb{E}\left[-\frac{1}{T}\sum_{t=1}^{T}\frac{\partial^2}{\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}'}\log f(\boldsymbol{x}_{t-1},\boldsymbol{x}_t|\boldsymbol{\mu},\boldsymbol{\theta})\right]. \tag{104}$$

Combining the first-order condition and the Taylor approximation leads to

$$\begin{aligned}&\sqrt{T}\frac{1}{T}\sum_{t=1}^{T}\frac{\partial^2}{\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}'}\log f(\boldsymbol{x}_{t-1},\boldsymbol{x}_t|\boldsymbol{\mu},\boldsymbol{\theta})(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}) = \\ &-\left[I \quad | \quad \tfrac{\sqrt{T}}{\sqrt{T+1}}\tfrac{1}{T}\textstyle\sum_{t=1}^{T}\tfrac{\partial^2}{\partial\boldsymbol{\theta}\partial\boldsymbol{\mu}'}\log f(\boldsymbol{x}_{t-1},\boldsymbol{x}_t|\boldsymbol{\mu},\boldsymbol{\theta})\right]\begin{bmatrix}\sqrt{T}\frac{1}{T}\sum_{t=1}^{T}\frac{\partial}{\partial\boldsymbol{\theta}}\log f(\boldsymbol{x}_{t-1},\boldsymbol{x}_t|\boldsymbol{\mu},\boldsymbol{\theta}) \\ \sqrt{T+1}(\hat{\boldsymbol{\mu}}-\boldsymbol{\mu})\end{bmatrix},\end{aligned} \tag{105}$$

where I is the identity matrix with size equal to the length of θ.

By applying the weak law of large numbers and Slutsky's theorem, we now have

$$\sqrt{T}(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}) \to N\left(0, H_1^{-1}\left[I \quad | \quad H_2\right]\Omega\left[I \quad | \quad H_2\right]'(H_1^{-1})'\right), \tag{106}$$

where H_2 is defined by

$$H_2 \equiv \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T}\frac{\partial^2}{\partial\boldsymbol{\theta}\partial\boldsymbol{\mu}'}\log f(\boldsymbol{x}_{t-1},\boldsymbol{x}_t|\boldsymbol{\mu},\boldsymbol{\theta})\right] \tag{107}$$

and Ω is defined by

$$\Omega \equiv \begin{bmatrix} H_1 & \Omega_{12} \\ \Omega_{12}' & S \end{bmatrix}, \tag{108}$$

with Ω_{12} the covariance between the moment conditions in step 1 (colums of Ω_{12}) and step 2 (rows of Ω_{12}).

Finally, by replacing the matrices S, H_1, H_2, and Ω_{12} with their sample equivalents, we obtain the standard errors for the parameter estimates $\hat{\Gamma}$ and $\hat{\Sigma}$ by taking the square root of the diagonal elements of the estimated asymptotic variances.

E No-arbitrage restrictions for price-of-risk parameters

Recall that for a non-dividend paying stock, absence of arbitrage implies

$$P_t = \mathbb{E}_t[M_{t+1}P_{t+1}].$$

Rewriting this yields

$$1 = \mathbb{E}_t[M_{t+1}(1 + R_{s,t+1})], \tag{109}$$

with $R_{s,t+1} \equiv \frac{P_{t+1}-P_t}{P_t}$. The log return is defined by $r_{s,t+1} \equiv \log(1 + R_{s,t+1})$.

Since we assumed that the return on the stock and the pricing kernel are lognormally distributed, we obtain

$$\begin{aligned} 0 &= \log(\mathbb{E}_t(M_{t+1}(1 + R_{s,t+1}))) = \mathbb{E}_t(\log(M_{t+1}) + r_{s,t+1}) + \frac{1}{2}\mathrm{Var}_t(\log(M_{t+1}) + r_{s,t+1}) \\ &= -y_t^{(1)} - \frac{1}{2}\boldsymbol{\lambda}_t'\boldsymbol{\lambda}_t + y_t^{(1)} + \boldsymbol{e}_s'\boldsymbol{\mu} + \boldsymbol{e}_s'\Gamma(\boldsymbol{x}_t - \boldsymbol{\mu}) + \frac{1}{2}\boldsymbol{\lambda}_t'\boldsymbol{\lambda}_t + \frac{1}{2}\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{e}_s - \boldsymbol{e}_s'\Sigma\boldsymbol{\lambda}_t \\ &= \boldsymbol{e}_s'\boldsymbol{\mu} + \boldsymbol{e}_s'\Gamma(\boldsymbol{x}_t - \boldsymbol{\mu}) + \frac{1}{2}\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{e}_s - \boldsymbol{e}_s'\Sigma(\boldsymbol{\Lambda}_0 + \Lambda_1\boldsymbol{x}_t), \end{aligned} \tag{110}$$

This yields the result (39):

$$\begin{aligned} \boldsymbol{e}_s'\Sigma\boldsymbol{\Lambda}_0 &= \boldsymbol{e}_s'(I - \Gamma)\boldsymbol{\mu} + \frac{1}{2}\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{e}_s \\ \boldsymbol{e}_s'\Sigma\Lambda_1 &= \boldsymbol{e}_s'\Gamma. \end{aligned}$$

F No-arbitrage valuation of human capital

Recall that at time t, the amount $\rho_t L_t$ is contributed to the pension scheme. Let $H_t^{(n)}$ be the value at time t of the contribution $\rho_{t+n} L_{t+n}$ at time $t+n$. We show that $H_t^{(n)}$ satisfies:

$$H_t^{(n)} = \rho_{t+n} \Pi_t \exp(G_{t+n} + \nu_t - A^L(n) - \boldsymbol{B}^L(n)' \boldsymbol{x}_t). \tag{111}$$

First of all, we determine the value of the contribution in one year from now:

$$\begin{aligned}
\log(H_t^{(1)}) &= \log(\mathbb{E}_t[M_{t+1}\rho_{t+1}L_{t+1}]) \\
&= \mathbb{E}_t(\log(M_{t+1}) + \log(\rho_{t+1}L_{t+1})) + \frac{1}{2}\text{Var}_t(\log(M_{t+1}) + \log(\rho_{t+1}L_{t+1})) \\
&= \log(\rho_{t+1}) + \log(\Pi_t) + G_{t+1} + \nu_t - y_t^{(1)} + \boldsymbol{\delta}_\pi'(\boldsymbol{\mu} + \Gamma(\boldsymbol{x_t} - \boldsymbol{\mu})) + \frac{1}{2}\boldsymbol{\delta}_\pi'\Sigma\Sigma'\boldsymbol{\delta}_\pi \\
&\quad + \frac{1}{2}\sigma_u^2 - \boldsymbol{\delta}_\pi'\Sigma\boldsymbol{\lambda}_t - k\boldsymbol{e}_s'\Sigma\boldsymbol{\lambda}_t \\
&= \log(\rho_{t+1}\Pi_t) + G_{t+1} + \nu_t - (A^R(1) + k\boldsymbol{e}_s'\Sigma\boldsymbol{\Lambda}_0 - \frac{1}{2}\sigma_u^2) - (\boldsymbol{B}^R(1) + k\Lambda_1'\Sigma'\boldsymbol{e}_s)'\boldsymbol{x}_t \\
&= \log(\rho_{t+1}\Pi_t) + G_{t+1} + \nu_t - A^L(1) - \boldsymbol{B}^L(1)'\boldsymbol{x}_t,
\end{aligned} \tag{112}$$

where

$$\begin{aligned}
A^L(1) &= A^R(1) + k\boldsymbol{e}_s'\Sigma\boldsymbol{\Lambda}_0 - \frac{1}{2}\sigma_u^2 \\
\boldsymbol{B}^L(1) &= \boldsymbol{B}^R(1) + k\Lambda_1'\Sigma'\boldsymbol{e}_s.
\end{aligned} \tag{113}$$

Now, we assume that the equation holds for all maturities up to $n-1$. We then find

$$\begin{aligned}
\log(H_t^{(n)}) &= \mathbb{E}_t\left(\log(M_{t+1}) + \log\left(H_{t+1}^{(n-1)}\right)\right) + \frac{1}{2}\text{Var}_t\left(\log(M_{t+1}) + \log\left(H_{t+1}^{(n-1)}\right)\right) \\
&= \log(\rho_{t+n}) + \log\Pi_t + G_{t+n} + \nu_t - y_t^{(1)} - A^L(n-1) \\
&\quad -(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi)'(\boldsymbol{\mu} + \Gamma(\boldsymbol{x}_t - \boldsymbol{\mu})) \\
&\quad +\frac{1}{2}(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\Sigma'(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi) + \frac{1}{2}\sigma_u^2 - k\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{B}^L(n-1) \\
&\quad +(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\boldsymbol{\lambda}_t - k\boldsymbol{e}_s'\Sigma\boldsymbol{\lambda}_t \\
&= \log(\rho_{t+n}) + \log\Pi_t + G_{t+n} + \nu_t - A^L(n) - \boldsymbol{B}^L(n)'\boldsymbol{x}_t,
\end{aligned} \tag{114}$$

where

$$\begin{aligned}
A^L(n) &= A^L(n-1) + (\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi)'(I - \Gamma)\boldsymbol{\mu} - (\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\boldsymbol{\Lambda}_0 + k\boldsymbol{e}_s'\Sigma\boldsymbol{\Lambda}_0 \\
&\quad -\frac{1}{2}(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi)'\Sigma\Sigma'(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi) - \frac{1}{2}\sigma_u^2 + k\boldsymbol{e}_s'\Sigma\Sigma'\boldsymbol{B}^L(n-1) \\
\boldsymbol{B}^L(n) &= (\Gamma - \Sigma\Lambda_1)'(\boldsymbol{B}^L(n-1) - \boldsymbol{\delta}_\pi) + k\Lambda_1'\Sigma'\boldsymbol{e}_s + \boldsymbol{\delta}_y.
\end{aligned} \tag{115}$$

By induction, it follows that the pricing relation holds. The initial conditions are given by (113).