Retirement Saving with Contribution Payments and Labor Income as a Benchmark for Investments[*]

Arjan Berkelaar[†] and Roy Kouwenberg

Econometric Institute Report EI-9946/A

Econometric Institute	Department of Finance
Faculty of Economics	Faculty of Economics
Erasmus University Rotterdam	Erasmus University Rotterdam
P.O.Box 1738, 3000 DR Rotterdam	P.O.Box 1738, 3000 DR Rotterdam
The Netherlands	The Netherlands

Abstract

In this paper we study the retirement saving problem from the point of view of a plan sponsor, who makes contribution payments for the future retirement of an employee. The plan sponsor considers the employee's labor income as investment-benchmark in order to ensure the continuation of consumption habits after retirement. We demonstrate that the demand for risky assets increases at low wealth levels due to the contribution payments. We quantify the demand for hedging against changes in wage growth and find that it is relatively small. We show that downside-risk measures increase risk-taking at both low and high levels of wealth.

Keywords: retirement saving, optimal asset allocation, discrete-time finance, dynamic programming.

JEL **Classifications Codes**: G11, G23.

---

1 Introduction

In this paper we study the problem of saving for retirement with contribution payments and labor income as a benchmark for investments. We consider the retirement saving problem from the point of view of a plan sponsor. The plan sponsor makes contribution payments to an investment fund in order to save for the future retirement of an employee. The goal is to ensure that the employee can continue his consumption pattern after retirement. As the consumption habits of the employee are related to his wages, the plan sponsor considers the labor income of the employee as a benchmark for investments.

Clearly, the plan sponsor is not only concerned about the welfare of the employee and also wants to minimize his contribution payments. We formalize this in a multi-period retirement saving model, where the plan sponsor makes a trade-off between the utility of low contribution payments and the utility of high fund values at retirement, relative to the labor income of the employee. The solution of the model reveals the optimal dynamic investment strategy and the optimal funding policy of the plan sponsor.

The retirement saving model presented in this paper includes both a defined contributions pension plan and a 'final pay' defined benefits plan as special cases. The generality of the model allows us to circumvent the large difference in pension schemes throughout the world: we focus on the core of the retirement saving problem. Clearly, not every plan sponsor or pension fund considers labor income as a benchmark for investments. However, we believe that it is in the interest of the employees to do so, without any adverse consequences.

An important assumption throughout the paper is that the labor income of the employee can not be replicated with the available assets: consequently, the financial market is incomplete. Moreover, we assume that the wage growth rate is partly predictable. Given the basic model setup, we derive optimal decision rules by applying dynamic programming. The optimal decision rules specify the asset weights and the contribution payment as a function of the state variables (the wealth-to-income ratio and the wage growth rate) and provide direct insight into the underlying problem.

We are particularly interested in the following issues, which are relevant for retirement saving and have not been studied adequately in the literature yet:

1. What is the magnitude of the demand for the hedge portfolio against random changes of the wage growth rate?
2. What is the impact of contribution payments on the optimal asset allocation?
3. What is the optimal multi-period investment strategy for investment objectives based on downside risk measures, which are very popular in practice?

With our implementation of the dynamic programming algorithm we were able to address these three main questions about optimal investment and funding in the retirement saving model. The conclusions to be drawn from our analysis and computational experiments are as follows:

1. Regardless of his utility function, the plan sponsor invests in a hedge-portfolio against random fluctuations of the employee's labor income. The hedge portfolio depends on the covariance of the asset returns with the wage growth rate. Whether the plan sponsor dynamically changes his holdings of the hedge portfolio depends on the relative risk aversion of the investment objective. Furthermore, the dynamic feature of the hedging strategy is solely driven by wealth effects as substitution effects are absent in our economy. The numerical results demonstrate that the correlation between asset returns and wage growth has a substantial influence on portfolio composition. Dynamic adjustments of the hedging strategy due to changes of the wage growth rate are relatively small.

2. Contribution payments change the optimal investment strategy considerably, even for a plan sponsor with constant relative risk aversion over fund value. The portfolio weights are no longer constant and there is a strong tendency to gamble at low levels of wealth. Intertemporal measurement of utility over wealth can reduce these gambling effects. We also find that contribution payments lead to a strong investment horizon effect. Moreover, additional constraints on the funding policy can completely alter investment strategies.

3. Plan sponsors increase the weight of risky assets in the portfolio at low levels of wealth, if they maximize the expected fund value subject to a penalty on downside-risk. This gambling policy can be attributed to the increasing relative risk aversion property of the downside-risk measure.

The insights gained from the basic retirement saving model studied in this paper may help plan sponsors to formulate dynamic investment policies and choose reasonable objectives. Furthermore, the optimal decision-rules derived here can be implemented in simulation-based systems for ALM, where additional market imperfections such as transaction costs could be added.

In order to place this paper in the literature, we could interpret the retirement saving problem with labor income as an asset-liability management (ALM) problem. In the literature many single-period ALM models have been studied (Sharpe and Tint 1990 and Leibowitz, Kogelman, and Bader 1994). These models indicate that investors should take the correlation between assets and liabilities into account, while deciding about the investment strategy. However, saving for retirement typically involves a long-term investment goal, and one-period models are therefore inappropriate due to stochastic opportunity sets and non-myopic preferences.

There is a large stream of literature about the application of stochastic programming methods for multi-period ALM (see Mulvey and Ziemba 1998 for an overview). Stochastic programming models formulate an accurate answer to the question: *how to invest today, given optimal recourse in the future?* However, the optimal policies do not explicitly reveal the relation between the decisions and the state-variables. Our main objective is to gain insight in optimal decision rules for portfolio choice and funding in a simple micro-economic model without market imperfections.[1]

Due to the focus on dynamic programming and micro-economic analysis this paper is related to the individual consumption-investment literature, which we review in Section 2.3. Our retire-

---

[1]The knowledge gained from a simple micro-economic ALM model might be applied for the formulation and interpretation of normative stochastic programming models with additional market imperfections.

ment saving model can be classified as a model with a stochastic opportunity set (predictable wage growth), with both negative and positive consumption (the net contribution payment) and without a riskless asset (due to market incompleteness). Koo (1999) analyzes a simple discrete-time model with labor income for power utility (constant relative risk aversion) and has to apply numerical techniques to solve it completely. As our model is more elaborate and as we additionally want to study more general objectives than power utility, we clearly need a numerical solution method.

Recently a number of consumption-investment models have been solved numerically in order to study the impact of relevant, but analytically complicated, issues such as predictability, transaction costs and parameter uncertainty (Brennan, Schwartz and Lagnado 1997, Balduzzi and Lynch 1999a, Barberis 1999). Unfortunately the numerical techniques applied are specifically intended for power utility functions. An alternative approach is the approximation technique of Campbell (1993), which can be used to solve general consumption-investment models quasi-analytically. However, the derived investment policies do not converge to the optimal solutions and could entail serious approximation errors (Campbell e.a. 1998).

In this paper we do not rely on simplifying assumptions or approximation techniques for the computations. We introduce a implementation of the dynamic programming algorithm with several improvements to increase the efficiency of the method. The improvements include transformation and interpolation of the value function, which are crucial to solve any investment model with power or HARA-utility. As we apply dynamic programming, the optimal policies are derived in feedback form and we acquire direct insight into the structure of the strategies.

This paper is organized as follows. In Section 2 we introduce the retirement saving model and the main assumptions. Moreover, we review the literature and introduce our implementation of the dynamic programming algorithm. In Section 3 we study the optimal investment strategies in detail for different investment objectives. We show that the widely used class of downside-risk measures leads to peculiar investment policies in a multi-period setting. In Section 4 we study the additional effect of funding payments on the investment strategy. Section 5 concludes and summarizes the paper.

2   Retirement Saving Model

2.1   Model Definition and Assumptions

In this section we introduce the model for retirement saving that will be studied in this paper. We assume that a plan sponsor pays contributions in order to finance the retirement of an employee at time T. The plan sponsor has established an investment fund with initial wealth W_t at time $t = 0$ in order to achieve the retirement goal at the planning horizon T. At the beginning of each period $t = 0, 1, ..., T - 1$ the plan sponsor decides either to add a contribution payment of $C_t^+ \geq 0$ to the fund or extract cash from the fund $C_t^- \geq 0$. We represent the net cashflow into the fund by $C_t = C_t^+ - C_t^-$.

After the cash in- and outflows a total amount of $(W_t + C_t)$ is left for investment in the capital market. There are $I + 1$ assets available for investment in the economy at time t, each providing a discretely compounded random return of $R_{i,t+1}$ at time $t+1$ for $i = 0, 1, ..., I$. Each period the plan sponsor invests a fraction w_{it} of the investment fund's wealth into asset $i = 1, ..., I$, while the remaining fraction $(1 - \sum_{i=1}^{I} w_t^i)$ is invested in asset 0.

The investment goal of the plan sponsor is to enable the employee to continue his consumption habits after retirement. As the employee's consumption pattern is likely to be related to his labor income, the plan sponsor considers the employee's wages at retirement I_T as a benchmark for investments. We assume that the plan sponsor measures wealth relative to the labor income of the employee and maximizes a utility function $U(F_T, T)$ over the relative fund value $F_T = W_T/I_T$ at retirement. The function $U(F_T, T)$ is increasing and concave in the fund value F_T.

Additionally, the plan sponsor tries to minimize his contribution payments by maximizing the utility function $V(-C_t, I_t, t)$ over the intertemporal payments C_t for $t = 0, 1, ..., T - 1$. The function $V(., I_t, t)$ is increasing and concave in its first argument. *Higher* contribution payments therefore *decrease* the utility of the plan sponsor and *increase* the marginal utility of a unit reduction in payments. We formalize the trade-off between fund value at retirement and contribution payments in the following retirement saving model:

(RS) Retirement Saving Model

$$(1) \qquad \max_{w_{it}, C_t} \quad E_0 \left[\sum_{t=0}^{T-1} V(-C_t, I_t, t) \ + \ \lambda U(W_T/I_T, T) \right]$$

$$(2) \qquad s.t. \qquad W_{t+1} = (\ W_t + C_t\ ) \left(\sum_{i=1}^{I} w_{it}(1 + R_{i,t+1}) \ + \ (1 - \sum_{i=1}^{I} w_{it})(1 + R_{0,t+1}) \right),$$
$$for \ t = 0, 1, ..., T - 1.$$

where $\lambda > 0$ is a parameter influencing the relative impact of utility over fund value $F_T = W_T/I_T$ at retirement T.

The general model for retirement saving needs a more detailed specification for a meaningful analysis. As a first building block, we will assume that the wage growth rate π_t follows a mean-reverting process:

Assumption 1 *Labor income I_t at time t grows with the continuously compounded wage growth rate π_{t+1}, which follows a mean-reverting process*

$$(3) \qquad I_{t+1} \ = \ I_t e^{\pi_{t+1}} = I_t \, (1 + \Pi_{t+1})$$
$$(4) \qquad \Delta \pi_{t+1} \ = \ \beta(\mu_\pi - \pi_t) + \epsilon_{t+1}, \ \epsilon_{t+1} \sim IID(0, \sigma_\pi^2)$$

where $\Delta \pi_{t+1} = \pi_{t+1} - \pi_t$, $0 < \beta < 1$ is the mean reversion parameter, μ_π is the long run mean wage growth rate and ϵ_{t+1} is an identically intertemporal independently distributed (IID)

innovation with expectation zero and variance σ_π^2. For ease of exposition we also define the discrete wage growth rate $\Pi_{t+1} = e^{\pi_{t+1}} - 1$.

Secondly, we assume that the asset returns are not correlated intertemporally and imperfectly correlated with the wage growth innovation:

Assumption 2 *The joint distribution of the continuously compounded asset returns $r_{i,t+1}$ and the wage growth innovation ϵ_{t+1} is equal in each period $t = 0, 1, ..., T - 1$. The vector of mean asset returns is $\mu = \{\mu_i\}_{i=0}^I$ and the covariance matrix is $\Omega = \{\sigma_{ij}\}_{i,j=0}^I$. The correlation vector of the asset returns with the wage growth innovation ϵ_{t+1} is denoted by $\sigma_{I\pi} = \{\sigma_{i\pi}\}_{i=0}^I$. The asset returns are not correlated intertemporally: $Cov(r_{i,t+1}; r_{i,t}) = 0$ for $t = 0, 1, ..., T - 1$. The discretely compounded asset return is defined as $R_{i,t+1} = e^{r_{i,t+1}} - 1$.*

Assumption 2 implies that there is no direct compensation for wage growth in the nominal asset returns, although the returns can still be imperfectly correlated with the wage growth rate. Consequently, this assumption entails that there is no perfect hedge for wage inflation: *the financial market is incomplete.*

Thirdly, we assume that the plan sponsor measures the utility of a contribution payment C_t relative to employee's wage income I_t at time t:

Assumption 3 *The utility V over contribution payments is a function of time and the ratio C_t/I_t only: $V(-C_t, I_t, t) = H(-C_t/I_t, t)$. We define $c_t = C_t/I_t$ as the contribution rate relative to the employee's labor income.*

Assumption 3 allows us to reduce the state space of the retirement saving problem from three variables (wealth W_t, income I_t and wage growth π_t) to two variables (the wealth-income ratio F_t and wage growth π_t). As a result the effort needed for solving the model reduces significantly. In our opinion this reduction of complexity makes up for the slight loss of generality. Moreover, the assumption of utility over contribution rates is also made by Boender (1997) in a widely used decision support system for pension funds.

In practice pension schemes are often categorized as either a defined contributions plan or a defined benefits plan. We would like to point out that our retirement saving model can include both schemes as special cases. If we fix the contribution payments C_t at a constant non-negative level, then we replicate a defined contributions scheme. The retirement saving model also includes a 'final pay' defined benefits scheme, where the plan sponsor guarantees to pay M times the employee's labor income at retirement (i.e. MI_T).[2] In the next sections we will discuss solution methods for the retirement saving model.

---

[2]In order to guarantee MI_T at retirement, the plan sponsor either has to pay an additional amount of $C_T^+ = \max\{0, MI_T - W_T\}$ or he receives the redundant fund value $C_T^- = \max\{0, W_T - MI_T\}$. Suppose that the plan sponsor maximizes a utility function $H(c_T, T)$ over the final contribution rate $c_T = (MI_T - W_T)/I_T = M - F_T$. We can model this equivalently with a utility function $U(F_T, T)$ over final fund value F_T, as in the objective (1) of the retirement saving model.

2.2 Dynamic Programming and First Order Conditions

A well-known solution technique for sequential decision problems under uncertainty is dynamic programming (see Bellman 1957 and Bertsekas 1976). It is based on the observation that an optimal policy for the entire sequential decision problem should also be optimal starting from any given state at a later date. This property is known as the principle of optimality. Given Assumption 1, 2 and 3, we can describe the state of the retirement fund at time t completely with the relative fund value $F_t = W_t/I_t$ and the wage growth rate π_t. Let the value function $J(F_t, \pi_t, t)$ denote the maximum expected utility obtainable for the financial institution starting at time t in state (F_t, π_t):

$$(5) \quad J(F_t, \pi_t, t) \;=\; \max_{w_{it}, c_t} E_t \left[\sum_{j=t}^{T-1} H(c_j, j) \;+\; \lambda U(F_T, T) \right]$$

$$(6) \quad s.t. \; F_{t+1} \;=\; (F_t + c_t) \left(\sum_{i=1}^{I} w_{it} \frac{(1 + R_{i,t+1})}{(1 + \Pi_{t+1})} \;+\; (1 - \sum_{i=1}^{I} w_{it}) \frac{(1 + R_{0,t+1})}{(1 + \Pi_{t+1})} \right),$$
$$for \; t = 0, 1, ..., T - 1.$$

The financial institution's initial planning problem is defined by $J(F_0, \pi_0, 0)$. Using the principle of optimality, the dynamic programming algorithm derives this function recursively by solving the following sequence of one-period problems:

$$(7) \quad J(F_t, \pi_t, t) \;=\; \max_{w_{it}, c_t} \{ \; H(c_t, t) \;+\; E_t \left[J(F_{t+1}, \pi_{t+1}, t + 1) \right] \; \},$$

$$(8) \quad s.t. \; F_{t+1} \;=\; (F_t + c_t) \left(\sum_{i=1}^{I} w_{it} \frac{(1 + R_{i,t+1})}{(1 + \Pi_{t+1})} \;+\; (1 - \sum_{i=1}^{I} w_{it}) \frac{(1 + R_{0,t+1})}{(1 + \Pi_{t+1})} \right),$$
$$for \; t = 0, 1, ..., T - 1.$$

where $J(F_T, \pi_T, T) = \lambda U(F_T, T)$.

We now state the first order conditions of optimality for the investment and funding policies of the retirement saving model, following from the dynamic programming equations (7):

$$(9) \quad 0 \;=\; H_c(c_t^*, t) - E_t \left[J_F(F_{t+1}, \pi_{t+1}, t + 1) \frac{(1 + R_{t+1}^*)}{(1 + \Pi_{t+1})} \right], \; for \; t = 0, 1, ..., T - 1.$$

$$(10) \quad 0 \;=\; E_t \left[J_F(F_{t+1}, \pi_{t+1}, t + 1) \frac{(R_{i,t+1} - R_{0,t+1})}{(1 + \Pi_{t+1})} \right], \; for \; i = 1, ..., I \; and \; t = 0, 1, ..., T - 1.$$

where $R_{t+1}^* = \sum_{i=1}^{I} w_{it}^* (R_{i,t+1} - R_{0,t+1}) + R_{0,t+1}$ denotes the optimal portfolio return.

In general it is not easy to solve to first order conditions for the optimal policies and in most cases numerical techniques are required. Before we present our numerical method, we will first discuss related literature on optimal portfolio choice.

2.3   Related Literature

The retirement saving model can be counted as a multi-period asset-liability management model, if we interpret the employee's labor income at retirement as a liability. In the literature many single-period ALM models have been studied (Sharpe and Tint 1990 and Leibowitz, Kogelman, and Bader 1994). These models indicate that investors should take the correlation between assets and liabilities into account, while deciding about the investment strategy. However, pension funds typically have a long investment horizon, and one-period models are therefore inappropriate in a context with stochastic opportunity sets or non-myopic preferences.

There is a large stream of literature about the application of stochastic programming methods for multi-period ALM (see Mulvey and Ziemba 1998 for an overview). Stochastic programming models formulate an accurate answer to the question: *how to invest today, given optimal recourse in the future?* However, the optimal policies do not explicitly reveal the relation between the decisions and the state-variables. Our main objective is to gain insight in optimal decision rules for portfolio choice and funding in a simple micro-economic model without market imperfections.

Given our focus on a dynamic programming and micro-economic analysis, the discrete-time individual consumption-investment literature is more closely related to this paper. Early references are Hakansson (1970), Mossin (1968) and Samuelson (1969), which study a multi-period consumption-investment problem for an investor with constant relative risk aversion (CRRA), i.e. power utility. Under the additional assumption of intertemporally independent asset returns they prove that the portfolio weights are equal in each period, independent of the individual's age and wealth. The investment policies are called myopic as the investor behaves identical to a single-period investor.[3]

In a continuous time framework, where the time-step between consecutive decisions becomes infinitesimally small, Merton (1969,1971) confirms the findings of Samuelson (1969) and Hakansson (1969,1971).[4] Recent papers by Duffie et al. (1997) and Koo (1998) apply the theory of viscosity solutions to study a continuous time model with labor income and bankruptcy constraints. A complete solution for HARA-utility (including CRRA as a special case) requires numerical techniques, equivalently to the discrete-time case studied by Koo (1999). Another important development in the field is the martingale methodology of Pliska (1986), Karatzas, Lehockzky and Shreve (1987) and Cox and Huang (1989). As the martingale methodology heavily relies on a complete markets setting, it is not suited for our model with non-traded wage income.[5]

---

[3]Hakansson (1969,1971) generalizes the multi-period model further and adds stochastic wage income, serially correlated asset returns and uncertainty about the individual's lifetime to the basic setup. All these results are limited to CRRA or the class of hyperbolic absolute risk aversion (HARA) utility functions with a negative threshold parameter.

[4]The closed-form solutions of Merton (1971) are limited to HARA-utility functions (including CRRA as a special case), under the additional assumption that the asset prices can be described with geometric Brownian motions. Sethi and Taksar (1988) have shown that some of the solutions in Merton (1971) are incorrect since bankruptcy problems are ignored. Explicit constraints are necessary to prevent bankruptcy, and closed-form solutions are more difficult to derive.

[5]Some progress has been made in extending this methodology to incomplete markets (see He and Pearson 1991), however the non-uniqueness of an equivalent martingale measure causes serious complications.

In the modern consumption-investment framework our retirement saving model can be classified as a model with a stochastic opportunity set (predictable wage growth), with both negative and positive consumption (the net contribution payment) and without a riskless asset (due to market incompleteness). In the defined contributions case it reduces to a model without consumption and with fixed labor income (the constant contribution rate). Koo (1999) analyzes a simple discrete-time model with labor income for power-utility, but has to apply numerical techniques to solve it completely. As our model is more elaborate and as we additionally want to study preferences outside the HARA class, we clearly need a numerical solution method.

Recently several authors have solved consumption-investment problems numerically in order to investigate research questions that can not be addressed analytically. These numerical schemes are often based on approximations or simulations and consider CRRA in order to facilitate calculations. Brennan, Schwartz, and Lagnado (1997) study the impact of time-varying expected returns under *power utility*. Brandt (1999) studies the impact of stochastic opportunity sets under *power utility* with an approximation based on the method of moments. Barberis (1999) considers parameter uncertainty and uses simulations for *power utility*. Balduzzi and Lynch (1999a,1999b) consider the impact of transactions costs and predictability using numerical techniques for *power utility*.

Campbell (1993) introduces a simple approximation technique which allows for quasi-analytical solutions of consumption-investment problems with recursive utility and stochastic opportunity sets. The approximation is based on a log-linearization of the Euler equation and the budget equation. As a result the optimal consumption and portfolio decisions are quadratic and linear in the state-variable respectively. Campbell and Viceira (1999a,1999b) and Viceira (1999) have applied this technique to solve interesting consumption-investment problems *approximately*. However, a disadvantage of the approximation method is that the derived solutions do not converge to the optimal strategies.[6]

Apart from serious approximation errors there are other drawback associated with the method introduced by Campbell (1993). First, all the consumption-investment models considered have an infinite horizon, which greatly simplifies the problem since the solutions no longer depend on time. Second, it does not seem straightforward to take additional restrictions into account such as limits on borrowing and constraints to prevent bankruptcy. We now propose a numerical solution technique that does not suffer from the drawbacks indicated above and moreover converges to the optimal solution.

2.4   Implementation of the Discretized Dynamic Programming Algorithm

We apply a numerical solution technique for dynamic programming that is based on a discretization of the state space, i.e. a grid for the values of the state variables F_t and π_t. For each point on this grid we solve the dynamic programming equation (7) numerically for the optimal investment and funding decisions. Note that the required computations easily get out of hand, as

---

[6]Campbell et al. (1998) show that the approximation error can indeed be quite large compared to a more general approximation technique based on polynomial decision rules for consumption and investment.

the expectation of next period's value function has to be integrated numerically. An additional complicating factor is the curse of dimensionality: the size of the state grid grows exponentially with the number of state variables.

We propose an implementation of the discretized dynamic programming algorithm with four crucial additions that increase the efficiency of the method:

1. *Variance reduction techniques in order to improve Monte-Carlo simulations.*
2. *Interpolation of the value function in order to reduce the size of the grid.*
3. *Transformation of the value function in order to facilitate interpolation methods.*
4. *Fast numerical optimizations using restarts from neighboring grid points.*

We apply Monte-Carlo simulation to calculate expectations over the return distributions, with an additional variance reduction technique to reduce the approximation error of the simulation method. We employ the method of antithetic sampling to replicate the mean of the underlying asset return distributions. Additionally we rescale the entire sample of returns in order to fit the variance of the returns exactly. This particular variance reduction method seems appropriate, as the mean and the variance of the return distributions play an essential role in portfolio selection.

Note that each simulated drawing from the distribution of the state variables tends to lie *in between the points on the grid* at time $t + 1$. In order to calculate the expectation at time t of the value function at time $t + 1$ in equation (7) we require an estimate of the value function in each simulated point. In the backward recursion of the dynamic programming algorithm we already know the value function in all points *on the grid* at time $t + 1$. We can pick a point on the grid that lies close to the simulated state, in order to obtain an estimate. Subsequently, the expectation of the value function can be computed as the sample mean of all the simulated estimates. Note that this simple version of the discretized dynamic programming algorithm approximates the value function at time $t + 1$ with a piecewise linear function.

A more smooth and accurate approximation of the value function can be achieved if we apply interpolation methods. In order to facilitate interpolation we first transform the value function to a roughly linear shape with a power function (see Appendix B for details). Next we interpolate the *transformed* value function in between grid points and apply the inverse transformation to obtain the approximation. The transformation of the value function is crucial to ensure accurate interpolation of utility functions with a vertical asymptote (e.g. power and HARA utility functions), which would otherwise require extremely fine-grained grids.

We interpolate the transformed value function with a local rational polynomial, which is a reliable and accurate method for nearly any functional form (see Press et al. 1989). Unfortunately the rational interpolation method is rather slow. Therefore, we calculate the scaled second derivative of the transformed value function to determine whether the function is locally linear. Depending on the value of this scaled second derivative we either apply rational, polynomial or linear interpolation (see Appendix B). We have tested this interpolation-rule on a number of functions and found that it provides a nice combination of reliability and efficiency.

Finally, we exploit Powell's method (Press et al. 1989) for the numerical optimizations in the dynamic programming equation (7). The computation-times can be cut drastically by using the optimal policy of a neighboring grid point as initial guess for the optimization method. Note that Powell's method does not explicitly handle bounds on the decision variables. However, we can still impose lower and upper bounds on the variables by simple transformations with an exponential or arctangent function.

Our implementation of the discretized dynamic programming algorithm converges to the optimal solution, given a number of mild continuity assumptions about the utility functions. Convergence in this context means that the numerical approximation of the value function at time $t = 0$ will get arbitrarily close to the actual value function, as the distance between points on the grid becomes smaller. We refer to Bertsekas (1976) for a convergence proof of the discretized dynamic programming algorithm without interpolation schemes. With the additional use of interpolation it can easily be shown that this proof still holds, if we ensure that the interpolation estimates stay within the boundaries of surrounding function values on the grid.

We have tested our implementation of the discretized dynamic programming algorithm by comparing numerical results with known analytical solutions and found satisfactory results. In practice the implementation allows us to solve problems with two state variables accurately and efficiently on a desktop PC. In the next two sections we will apply our algorithm in order to investigate the three main research questions which are listed in the Introduction.

3   Investment Policies for Retirement Saving

3.1   Portfolio Separation and Liability Hedging Credits

We try to gain insight in the optimal policies of the retirement saving model by analysing different parts in isolation. In this section we focus on optimal investment policies in a model without contribution payments, while in Section 4 we will investigate the additional effect of funding decisions. Mostly we will apply our numerical method to solve the model completely, except for the special case of a CRRA investor. Throughout the following sections we will assume that the discrete asset returns and the wage growth are log-normally distributed:

Assumption 4 *The joint distribution of the continuously compounded asset returns $r_{i,t+1}$ and the wage growth innovation ε_{t+1} is multi-variate normal in each period $t = 0, 1, ..., T - 1$, with a constant mean vector and a constant covariance matrix (see Assumption 2). Consequently, the discrete returns $(1 + R_{i,t+1}) = e^{r_{i,t+1}}$ and the discrete wage growth $(1 + \Pi_{i,t+1}) = e^{\pi_{t+1}}$ follow a multi-variate log-normal distribution.*

Before we apply our numerical technique to solve the retirement saving model completely, we first pay some attention to the general structure of investment strategies. First of all, the correlation of the asset returns with the wage growth rate is very important for investments in the retirement saving model. We can observe this by studying the variance of the asset returns

in excess of the wage growth $(r_{i,t+1} - \pi_{t+1})$, which drive the development of the relative fund value F_t:

(11) $\quad Var(r_{i,t+1} - \pi_{t+1}) = \sigma_i^2 - 2\sigma_{i\pi} + \sigma_\pi^2.$

The variance of the excess return decreases as an asset provides higher covariance with wage growth. This effect is related to the concept of liability hedging credits, introduced by Sharpe and Tint (1990): the investor is willing to accept lower expected returns in exchange for more protection against random fluctuations of the wage growth rate.

In order to provide more insight into the optimal investment strategies, we prove in Appendix C that three fund separation (12) holds for any investment objective, given Assumption 4, a linear approximation of continuously compounded returns and the assumption that asset 0 is riskless:

(12) $\quad w_t^* = -\dfrac{E_t[J_F]}{F_t E_t[J_{FF}]} \Omega^{-1}(\mu - \iota r_0) + \left(1 - \dfrac{E_t[J_{F\pi}]}{F_t E[J_{FF}]}\right) \Omega^{-1}\sigma_{I\pi}$

where w_t^* is the vector of optimal portfolio weights.

The optimal asset weights of the plan sponsor can be separated into a market portfolio, a hedge portfolio for the wage growth rate and a riskless portfolio. The fund separation result also holds for investors who minimize the variance of the fund value return $Var(F_{t+1}/F_t)$ subject to a given expected return $E[F_{t+1}/F_t]$, if we ignore the part $\frac{E_t[J_{F\pi}]}{F_t E[J_{FF}]}\Omega^{-1}\sigma_{I\pi}$ in (12). This particular part of the investment strategy represents the intertemporal hedge against predictable changes of the wage growth through time, which is clearly missing in a one-period model. In the next section we will quantify the investment strategy numerically, without the simplifying assumptions needed for the fund separation result.

3.2 Constant and Decreasing Relative Risk Aversion

The portfolio separation result holds for any investment objective, but does not provide a complete solution as the investment rule depends on the unknown value function $J(F_t, \pi_t, t)$. In order to learn more about the optimal investment strategies we will now study the special case of constant relative risk aversion. We assume that the plan sponsor maximizes a power utility function over the relative fund value F_T at the retirement date T:

(13) $\quad U(F_T, T) \ = \ (1+\rho)^{-T}\dfrac{1}{\alpha}(F_T)^\alpha,$

(14) $\quad R^U(F_T, T) \ = \ \dfrac{-F_T U_{FF}(F_T, T)}{U_F(F_T, T)} = 1 - \alpha$

where $\alpha < 1$ is the relative risk aversion coefficient and $\rho > 0$ is the intertemporal discount factor.

The function $R^U(F_T, T)$ in (14) is the Arrow-Pratt measure of relative risk aversion, which is constant for power utility. In a simple investment model with one risky asset and one riskless asset the portfolio weight of the risky asset increases (decreases) as a function of wealth, if the relative risk aversion is decreasing (increasing) in wealth. An investor with *constant* relative risk aversion does not change the risky asset weight and chooses a constant portfolio. It is not obvious whether these results also hold for the retirement saving model, as it additionally includes a stochastic opportunity set and contribution payments.

In Appendix A we derive the following result for the power utility objective (13), under the additional assumption of zero contribution rates $c_t = 0$ for $t = 0, 1, ..., T-1$:

$$(15) \quad \frac{dw_{it}^*}{dF_t} = 0 \ and \ \frac{dw_{it}^*}{d\pi_t} = 0, \ for \ t = 0, 1, ..., T-1 \ and \ i = 0, 1, ..., I.$$

where w_{it}^* denotes the optimal asset weight for $i = 0, 1, ..., I$.

We conclude from (15) that the plan sponsor does not adjust the portfolio weights if the fund value or the wage growth changes, in the case with a CRRA objective and without intertemporal contribution payments. Note that wage growth does not affect the drift rate and the volatility of the asset returns in our economy and therefore substitution between assets as a result of a change of the wage growth is needless. Wealth effects are present, but have no consequence for portfolio choice as we consider an investor with constant relative risk aversion.

We will now study utility functions with decreasing relative risk aversion (DRRA), in order to investigate whether the optimal investment strategies are indeed different. We introduce the following class of utility functions with DRRA, which can be applied in our economy without causing feasibility problems:[7]

$$(16) \quad U(F_T, T) \ = \ (1+\rho)^{-T} \left(\frac{1}{\alpha - p} F_T^{-p} \ + \ \frac{\theta}{\alpha} \right) F_T^\alpha,$$

$$(17) \quad R^U(F, T) \ = \ -F_T \frac{U_{FF}(F_T, T)}{U_F(F_T, T)} \ = \ \frac{(1 - \alpha + p) \ + \ (1 - \alpha)\theta F_T^p}{1 \ + \ \theta F_T^p}.$$

where $\alpha < 1$, $p > 0$ and $\theta > 0$ are constant parameters.

It is easy to prove that $U_F(F_T, T) > 0$, $U_{FF}(F_T, T) < 0$ and $R_F^U(F_T, T) < 0$ hold: U is a utility function with DRRA. If F_T is close to zero then the relative risk aversion coefficient equals $(1 - \alpha - p)$, while it decreases monotonically to $(1 - \alpha)$ as the fund value increases to infinity. Figure 1 shows a graph of the relative risk aversion function (17) for the parameter values $\alpha = -0.5$, $p = 3$ and $\theta = 3$.

---

[7]HARA-utility functions with a positive sustenance level $\eta > 0$ also have the property of DRRA, however we can not apply them in our incomplete market as they would cause bankruptcy problems (the constraint $F_t > \eta$ of the HARA-utility function can not be satisfied under all circumstances).

We would like to investigate the impact of DRRA on optimal investment strategies, however analytical solutions are not readily available under this particular utility specification and therefore we rely on our numerical method. First we specify the coefficients that will be used in all numerical examples from now on:

Assumption 5 *For the numerical cases we consider a retirement fund 10 years before the retirement date, i.e.* $T = 10$. *The discount rate* ρ *is 2% per year. The initial wealth of the fund is* $W_0 = 1$ *and the initial labor income of the employee is* $I_0 = 1$. *At the beginning of each period* $t = 0, 1, ..., T - 1$ *the plan sponsor can adjust the asset portfolio of the retirement fund. Two assets are available for investing, a risky bond and a risky stock. There is no riskless asset and short selling is not allowed. Consequently, the asset weights of the bond and the stock are within 0% and 100%. The asset returns are jointly log-normally distributed with expected returns* $\mu_B = 6.8\%$ *and* $\mu_S = 8.6\%$ *for the bond and the stock respectively, standard deviations* $\sigma_B = 5.9\%$ *and* $\sigma_S = 15.7\%$ *and a correlation of* $\rho_{BS} = +0.38$. *The long run mean wage growth rate is* $\mu_\pi = 5.7\%$, *the mean reversion parameter equals* $\beta = 0.31$ *and the standard deviation is* $\sigma_\pi = 3\%$. *We will make specific assumptions about the correlation between asset returns and wage growth later on.*

These parameters for the distribution of the asset returns and the wage growth rate are based on a Dutch yearly dataset from the period 1956-1994 (see Kouwenberg 1998). We only consider two risky assets in order to facilitate the interpretation of the results. We now apply the discretized dynamic programming algorithm and solve the retirement saving model with DRRA numerically:

Case 3.1 *We solve the retirement saving model with DRRA utility function (16) for* $\alpha = -0.5$, $p = 3$ *and* $\theta = 3$. *As a benchmark, we also solve the model with CRRA objective (13), for* $\alpha = -0.5$ *and* $\alpha = -3.5$. *In order to distinguish the impact of the objective we ignore uncertainty about labor income in this example: the wage growth rate is fixed at* $\pi_t = 3\%$ *annually. We apply the discretized dynamic programming algorithm on a grid consisting of 36 equally spaced fund values* $F = 0$ *to* $F = 3.5$ *(step size 0.1). Expectations over the return distribution are approximated with 10000 simulation points.*

Figure 2 shows the optimal proportion invested in stocks at time $t = 0$, as a function of initial fund value. The investment strategy is clearly myopic for the two constant relative risk aversion cases $\alpha = -0.5$ and $\alpha = -3.5$, demonstrating the convergence of the numerical algorithm. With decreasing relative risk aversion the weight of stocks increases from 28% to 90%. It is clear that this optimal investment strategy can be explained quite well by the corresponding decreasing relative risk aversion function in Figure 1.

We will now additionally investigate the impact of uncertainty about the wage growth rate. We already know that a plan sponsor with CRRA utility follows a myopic investment strategy, as there is no substitution between assets in our economy if the wage growth changes. Note that a change of the wage growth rate *does influence* the expected fund value. Consequently, we hypothesize that a plan sponsor with DRRA hedges intertemporally against changes of wage growth, as his investment strategy is influenced by changes in (expected) wealth. We will now demonstrate this effect numerically:

Case 3.2 *The plan sponsor maximizes DRRA utility function (16) with parameter values $\alpha = -0.5$, $p = 3$ and $\theta = 3$. Additionally, the wage growth rate π_t is stochastic, with the parameters specified in Assumption 5. For the correlation between wage growth and asset returns we consider two cases: 1. bond returns are positively correlated with wage growth, while stock returns are uncorrelated ($\rho_{B\pi} = +1/3$ and $\rho_{S\pi} = 0$) and 2. the opposite case ($\rho_{B\pi} = 0$ and $\rho_{S\pi} = +1/3$). We apply the discretized dynamic programming algorithm on a two-dimensional grid for the state variables consisting of 36 equally spaced fund values $F = 0$ to $F = 3.5$ (step size 0.1) and 17 wage growth rates $\pi = -0.10$ to $\pi = +0.22$ (step size 0.02). Expectations over the return distribution are approximated with 10000 simulation points.*

Figure 3 shows the optimal stock proportion at time 0 as a function of fund value and the initial wage growth rate, in the case that bond returns are positively correlated with wage inflation. We observe two major effects. First, the fraction invested in stocks increases as the fund value becomes larger: this is the familiar result for DRRA. Second, an increase in the expected wage growth leads to a larger proportion of bonds in the portfolio in order to hedge against the expected reduction of future wealth, as we hypothesized earlier.

Figure 4 displays the results for the case where stock returns are positively correlated with the wage growth rate. We clearly observe that the fraction invested in stocks increases at any level of fund value and wage inflation, compared to the situation in Figure 3. This parallel shift of the investment strategy can be explained by the increased liability hedging credits of stocks. If we concentrate on the dynamic hedging strategy, the results in Figures 3 and 4 are similar: an increase of expected wage growth reduces the proportion of stocks in the portfolio, regardless of the underlying correlation structure.

3.3 The V-Shaped Consequences of Downside-Risk Measures

In practice a popular objective is to maximize the expected fund value, subject to a penalty if wealth drops below the investment benchmark, i.e. a penalty on downside-risk. Downside-risk measures have been introduced as a substitute for standard risk measures such as variance, as it seems unreasonable to penalize negative and positive deviations from the benchmark equally. Popular measures include: shortfall probability, expected shortfall and semivariance (see Bawa and Lindenberg 1977).

In this section we study the optimal investment strategy if the plan sponsor maximizes the expected fund value at retirement, subject to a quadratic penalty if he fails to outperform the investment benchmark (i.e., if $F_T < 1$):

$$(18) \quad \max_{w_{it}} E_0[U(F_T, T)] = \frac{1}{(1+\rho)^T} \left(E_0[F_T] - \nu E_0[\min\{0, 1 - F_T\}^2] \right)$$

where $U(F_T, T) = (1 + \rho)^{-T} \left(F_T - \nu \min\{0, 1 - F_T\}^2 \right)$ and $\nu > 0$ is a constant risk aversion parameter .

In the normative ALM literature for pension funds the downside-risk concept is highly successful and seems to have replaced traditional utility functions. Investors that try to meet their liabilities usually apply a threshold of 1 for the funding ratio, i.e. the ratio of assets to liabilities. Even regulating authorities nowadays advocate the use of downside-risk measures to improve the solvency of institutional investors. Not much attention has been paid however to the optimal investment strategies for downside-risk objectives in a multi-period model.

Downside-risk averse investors can be considered as expected utility maximizers with a risk-neutral (linear) utility function above the threshold $F_T = 1$ and a utility function with increasing relative risk aversion (IRRA) below the threshold. The effect of IRRA below the threshold might be gambling behavior: a decrease of wealth increases the portfolio weight of stocks. Furthermore risk-neutrality above the threshold should lead to a strong increase in the demand for stocks.

Consequently, in a two-asset economy we expect a V-shaped figure if we draw the optimal stock weight as a function of fund value F_t. We will now solve the retirement saving model numerically in order to investigate these effects.

Case 3.3 *We solve the retirement saving model with downside-risk objective (18) for risk aversion values $\nu = 16$ and $\nu = 64$. In order to distinguish the impact of the investment objective clearly we ignore uncertainty about labor income: the wage growth rate is fixed at $\pi_t = 3\%$. We apply the discretized dynamic programming algorithm on a grid consisting of 36 equally spaced fund values $F = 0$ to $F = 3.5$ (step size 0.1). Expectations over the return distributions are approximated with 10000 simulation points.*

Figure 5 shows the optimal initial weight of stocks for risk aversion levels $\nu = 16$ and $\nu = 64$. The optimal investment strategy has a clear V-shape: at low wealth levels the weight of stocks decreases, while it increases at higher wealth levels. We find equivalent results for other downside-risk measures such as expected shortfall and shortfall probability, which have a similar relative risk aversion function. Given the wide-spread acceptance of downside-risk measures it is questionable whether investors are fully aware of the accompanying gambling behavior at low wealth levels.

4 Consumption and Funding Decisions

4.1 The Effect of Funding Payments

In the previous section we have ignored contribution payments by the plan sponsor, while we focused on optimal investment strategies for the retirement saving model with labor income as a benchmark. In this section we study the additional effect of funding payments on the retirement saving problem. As a first example, we consider the impact of a fixed non-negative contribution payment.

Case 4.1 *We solve the retirement saving model with a fixed contribution rate $c_t = 0.05$ for $t = 0, 1, ..., T - 1$. The utility over fund value at retirement is CRRA (13) with $\alpha = -3.5, -2$*

and −0.5. *We ignore uncertainty about labor income in this example: the wage growth rate is fixed at* $\pi_t = 3\%$ *annually. The grid of the discretized dynamic programming algorithm consists of 36 equally spaced fund values* $F = 0$ *to* $F = 3.5$ *(step size* 0.1*) and the return distributions are approximated with* 10,000 *simulation points.*

Figure 6 displays the optimal investment strategy at time $t = 0$ for the three levels of risk aversion. Due to the contribution payments, the portfolio weight of stocks is decreasing as a function of fund value for $\alpha = -2$ and -3.5 and the weight is always at the upper bound for $\alpha = -0.5$. The plan sponsor appreciates the fixed payment more at low fund values than at high fund values (i.e. marginal utility of the contribution payment decreases). Poor returns on investments therefore have less impact on utility at low fund levels than at high fund values, which explains the decreasing optimal portfolio weight of stocks for $\alpha = -2$ and -3.5.

We will now study the case where the plan sponsor can change the contribution payments intertemporally. Clearly, the plan sponsor has to make a trade-off between low contribution payments and high fund values at retirement. Given Assumption 3, we formalize this as a trade-off between the utility $H(c_t, t)$ of contribution rates c_t and the utility $\lambda U(F_T, T)$ of fund value at retirement. A widely used and intuitively attractive way for modelling preferences over contribution payments is to minimize the expected net present value. In our framework this can be specified as $H(-c_t, t) = -c_t/(1 + \rho)^t$, subject to the fixed bounds $\underline{c} \leq c_t \leq \overline{c}$. Note however that interior solutions never occur: the optimal contribution rate c_t^* is always at the lower or upper bound: $c_t^* = \underline{c}$ or $c_t^* = \overline{c}$.

Instead of the net present value rule, the plan sponsor can specify an increasing and strictly concave utility function $H(-c_t, t)$ over contribution rates. We know that if $H(-c_t, t)$ and $U(F_T, T)$ are both power utility functions with equal risk aversion coefficient, then the funding policy is linear and the investment strategy is myopic.[8] However, power utility enforces the constraint $-c_t > 0$ and the plan sponsor will only extract cash from the fund. We will investigate HARA-utility (19) for the contribution rates, as it allows the plan sponsor to make payments $c_t > 0$:

$$(19) \quad H(-c_t, t) \; = \; \frac{1 - \gamma}{\gamma} \left(\frac{-c_t + \overline{c}}{(1 - \gamma)} \right)^\gamma , \quad -\infty \leq c_t < \overline{c}$$

where $\gamma < 1$ is the risk aversion coefficient and $\overline{c} > 0$ is the upper bound of the contribution rate.

The HARA-utility function $H(-c_t, t)$ is increasing and concave in its first argument and enforces the upper bound $\overline{c} > 0$ on contribution rates. Moreover it has the property of increasing relative risk aversion (IRRA). We investigate in Case 4.2 whether this property influences the investment strategy:

Case 4.2 *We solve the retirement saving model with HARA-utility (19) for contribution rates, with risk aversion parameters* $\gamma = -1.5, -2$ *and* −2.5 *and upper bound* $\overline{c} = 0.1$. *The utility*

---

[8]This result holds in a consumption-investment model with constant opportunity sets and a risk free asset, see Hakansson (1969), Mossin (1968) and Samuelson (1969).

function over fund value at retirement is CRRA with $\alpha = -2$. The trade-off parameter between utility over contribution rates and fund value is $\lambda = 250,000$. We choose this value for λ to make sure that contribution rate is positive for funding levels $F \leq 1$.

Before we discuss the results, we first explain why the value $250,000$ is reasonable for λ. The parameter λ influences the trade-off between utility over contribution rates and fund value in the objective function. In numerical experiments we have found that the impact of λ on the optimal solution is limited to the slope of the contribution rates as a function of fund value. As positive contribution rates take on values between 0 and 0.10 (given upper bound $\overline{c} = 0.1$), we require a high value for the trade-off parameter λ to get positive contribution rates at fund values $F \leq 1$.

Figure 8 displays the optimal *contribution rate* c_t at time $t = 0$ as a function of initial fund value F_t, for the cases $\gamma = -1.5$, -2 and -2.5. The optimal contribution rates are decreasing linearly as function of fund value, while higher risk aversion γ over contribution rates leads to lower rates. Figure 7 shows the investment strategy at time $t = 0$ and we find the same decreasing stock weight pattern as in the case of a fixed contribution rate (see Figure 6), although more pronounced.

4.2 Horizon Effects and Gambling

So far we have displayed the optimal investment and contribution strategy at time $t = 0$. We now study to what extent the optimal strategies change as the retirement date T approaches. In the case of CRRA over fund value and without contribution payments, the asset weights are constant through time. For DRRA, Figure 9 shows the optimal investment strategy at times $t = 9$, 4 and 0 (based on Case 3.1). We find that the weight of stocks increases with the investment horizon for fund levels between $F = 0.4$ and $F = 2$.

In the case of a fixed contribution rate horizon effects are even present for a CRRA investor. Figure 10 shows the optimal stock weight at times $t = 9$, 4 and 0, for a CRRA objective with a fixed contribution rate $c_t = 0.05$ for $t = 0, 1, ..., T - 1$ (based on Case 4.1). One period before retirement ($t = 9$) the investment strategy is still constant, while we observe a drastic increase of the stock weight at low fund values for longer investment horizons ($t = 4$ and $t = 0$).

With flexible contribution rates and HARA-utility (19), the horizon effect is even stronger. Figure 11 shows the optimal stock weight at times $t = 9$, 4 and 0, for HARA-utility over contribution rates with $\gamma = -2$ and $\overline{c} = 0.1$ and with a CRRA objective over fund value with $\alpha = -2$ (based on Case 4.2). The corresponding contribution rules are displayed in Figure 12. We find that the contribution rates tend to decrease as we move further from the retirement date, while the stock weight at low fund levels increases.

We have shown that the payment of intertemporal contributions to the retirement fund leads to a gambling strategy: the plan sponsor prefers a high proportion of stocks at low fund levels. Moreover, the gambling becomes worse for longer horizons. This investment behavior might be

considered imprudent, most notably in the case of a defined benefits pension scheme. We will now demonstrate that the propensity to gamble can be reduced considerably if the plan sponsor cares about intertemporal fund values F_t at times $t = 1, 2, ..., T - 1$ before the planning horizon T. We first introduce the following objective for the plan sponsor, with intertemporal utility over fund values:

$$(20) \quad \max_{w_{it}, c_t} E_0 \left[\sum_{t=0}^{T-1} H(c_t, t) \; + \; \lambda \sum_{t=1}^{T} U(F_t, t) \right]$$

One might argue that it is not rational for the plan sponsor to worry about fund value before the actual retirement date T. Note however that solvency requirements could force the plan sponsor to do so, specially in the case of a defined benefits pension plan. In the next case we calculate investment strategies for an objective with intertemporal utility over fund values:

Case 4.3 *We solve the retirement saving model with CRRA utility for fund values, measured intertemporally according to (20), with $\alpha = -2$. The utility function over contribution rates is HARA (19), with risk aversion parameters $\gamma = -2$ and upper bound $\overline{c} = 0.1$. The trade-off parameter between utility over contribution rates and fund value is $\lambda = 250,000$. We choose this value for λ to make sure that contribution rate is positive for funding levels $F < 1$.*

Figure 13 and 14 show the optimal investment and contribution policy. From Figure 13 it is clear that the optimal weight of stocks is still slightly higher at low fund levels, however the effect has reduced considerably and can hardly be considered a pure gambling strategy anymore. Figure 14 is also interesting, as it show that the horizon effect for the contribution rates has reversed: as the retirement date T gets nearer, the plan sponsor pays less contributions. We conclude that a focus on fund value before the retirement date can reduce the gambling behavior induced by contribution payments.

4.3 Constraints on the Funding Policies

In practice the contribution rate might be constrained by explicit restrictions (see Boender 1997). The HARA-function (19) naturally imposes an upper bound on the funding payment as marginal utility approaches $-\infty$ near the level $\overline{c}$. We now consider the additional effect of a lower bound on funding payments ($c_t > \underline{c}$), which we enforce in the implementation of the discretized dynamic programming algorithm.

Case 4.4 *We solve the problem of Case 4.2 again, with the addition of a lower bound $\underline{c} = 0$ on funding payments (i.e. the plan sponsor can not extract cash from the fund). The upper bound equals $\overline{c} = 0.1$ and the risk aversion coefficients of the utility functions $H(c_t, t)$ and $U(F_T, T)$ are $\gamma = -2$ and $\alpha = -2$.*

With an additional lower bound the optimal contribution rate is decreasing linearly as a function of the fund value until it hits the lower bound $\underline{c} = 0$ at $F_0 = 1.5$ and from thereon it remains

constant. In contrast to this predictable result, the optimal investment strategy displayed in Figure 15 is more surprising. The stock weight first decreases to a minimum around $F_0 = 1.9$ and from this point onwards it starts to rise, up to 100%.

The declining stock weight pattern before the contribution rate reaches the lower bound can be explained by the IRRA property of the HARA-function (19) over contribution rates. The increasing stock weight trend in the right side of the graph can be interpreted as a gambling policy: the potential losses of stocks become less important as the funding payment is stuck at an involuntary 'high' level (providing a cushion for possible losses) while the potential gains of stocks remain the same.

Another interesting constraint on the funding policy is one induced by a maximum allowable level of the fund value. Suppose that the plan sponsor decides that the value of the retirement fund should never exceed $\overline{F}$ times the labor income of the employee, as this fund value is already sufficient to provide a good retirement income. Consequently, if the fund value F_t exceeds the 'saturation' level $\overline{F}$, then the plan sponsor extracts cash from the fund (i.e. $c_t < 0$) in order to reduce the fund value back to $\overline{F}$.

Case 4.5 *We consider Case 4.2, with the addition of a maximum fund level at $\overline{F} = 2$. The risk aversion coefficients are $\alpha = \gamma = -2$ and the upper bound for funding $\overline{c} = 0.1$.*

The optimal funding policy in this case is piecewise linear: at fund values above $\overline{F} = 2$ the plan sponsor extract large amounts of cash from the fund. The effect on the optimal investment strategy is depicted in Figure 16. For low fund values the optimal stock weight is decreasing up to $\overline{F} = 2$, remaining at the minimum of 12% from thereon. As a consequence of the maximum fund level, the plan sponsor drastically reduces stock exposure at wealthy fund values. The gains of investing in stocks at high fund values are reduced by refunds, while the losses of investing in stocks still loom as large as before.

5   Conclusions

In this paper we studied a retirement saving model with labor income as a benchmark for investments and periodic contribution payments by the plan sponsor. Due to market incompleteness it is very hard to derive closed-form solutions for the optimal investment and contribution policies of the plan sponsor. We introduced an efficient implementation of the dynamic programming algorithm in order to solve the problem numerically. The main conclusions to be drawn from our analysis and computational experiments are as follows:

1. Regardless of his utility function, the plan sponsor invests in a hedge-portfolio against random fluctuations of the employee's labor income. The hedge portfolio depends on the covariance of the asset returns with the wage growth rate. Whether the plan sponsor dynamically changes his holdings of the hedge-portfolio depends on the relative risk aversion of the investment objective. Furthermore, the dynamic feature of the hedging strategy is

solely driven by wealth effects as substitution effects are absent in our economy. The numerical results demonstrate that the correlation between asset returns and wage growth has a substantial influence on portfolio composition. Dynamic adjustments of the hedging strategy due to changes of the wage growth rate are relatively small.

2. Contribution payments change the optimal investment strategy considerably, even for a plan sponsor with constant relative risk aversion over fund value. The portfolio weights are no longer constant and there is a strong tendency to gamble at low levels of wealth. Intertemporal measurement of utility over wealth can reduce these gambling effects. We also find that contribution payments lead to a strong investment horizon effect. Moreover, additional constraints on the funding policy can have a large impact on the optimal investment strategy.

3. Plan sponsors increase the weight of risky assets in the portfolio at low levels of wealth, if they maximize the expected fund value subject to a penalty on downside-risk. The gambling policy can be attributed to the increasing relative risk aversion property of the downside-risk measure.

The conclusions of this paper have important implications. The insights gained through the numerical computations in this paper may aid fund-managers to formulate dynamic policies and choose reasonable investment objectives. Furthermore, the decision-rules derived here could be applied in simulation-based models for asset-liability management, where additional market imperfections such as transaction costs and position limits can be taken into account.

A Investment under Constant Relative Risk Aversion

In order to learn more about the optimal investment strategies, we study the case of constant relative risk aversion in this appendix. We assume that the plan sponsor maximizes a power utility function (13) over the funding ratio F_T at the planning horizon T. As can be inferred from (14), the power utility function displays *constant relative risk aversion*. We will now prove that this feature implies that the fraction invested in each asset is independent of the funding ratio and the wage growth rate.

If we differentiate the first-order condition (10) with respect to F_t, then we obtain the following expression:

$$(21) \quad \frac{dw_{it}^*}{dF_t} = -\frac{E_t\left[J_{FF}(F_{t+1}, \pi_{t+1}, t+1)\frac{(R_{i,t+1}-R_{0,t+1})(1+R_{t+1}^*)}{(1+\Pi_{t+1})^2}\right]}{F_t E_t\left[J_{FF}(F_{t+1}, \pi_{t+1}, t+1)\frac{(R_{i,t+1}-R_{0,t+1})^2}{(1+\Pi_{t+1})^2}\right]}$$

Using the relative risk aversion function we can formulate the change in demand for asset i as follows:

$$(22) \quad \frac{dw_{it}^*}{dF_t} = \frac{1}{F_t^2}\frac{E_t\left[J_F(F_{t+1}, \pi_{t+1}, t+1)R^J(F_{t+1}, \pi_{t+1}, t+1)\frac{(R_{i,t+1}-R_{0,t+1})}{(1+\Pi_{t+1})}\right]}{E_t\left[J_{FF}(F_{t+1}, \pi_{t+1}, t+1)\frac{(R_{i,t+1}-R_{0,t+1})^2}{(1+\Pi_{t+1})^2}\right]}.$$

Similarly we can study the general effect of a change in the current inflation rate on the optimal investment strategy by studying the derivative $dw_{it}^*/d\pi_t$. Differentiating the first-order condition (10) with respect to π_t we obtain:

$$(23) \quad \frac{dw_{it}^*}{d\pi_t} = -(\beta-1)\frac{E_t\left[J_F(F_{t+1}, \pi_{t+1}, t+1)R^J(F_{t+1}, \pi_{t+1}, t+1)\frac{(R_{i,t+1}-R_{0,t+1})}{(1+\Pi_{t+1})}\right]}{F_t E_t\left[J_{FF}(F_{t+1}, \pi_{t+1}, t+1)\frac{(R_{i,t+1}-R_{0,t+1})^2}{(1+\Pi_{t+1})^2}\right]}$$

$$-(\beta-1)\frac{E_t\left[J_{F\pi}(F_{t+1}, \pi_{t+1}, t+1)\frac{(R_{i,t+1}-R_{0,t+1})}{(1+\Pi_{t+1})}\right]}{F_t E_t\left[J_{FF}(F_{t+1}, \pi_{t+1}, t+1)\frac{(R_{i,t+1}-R_{0,t+1})^2}{(1+\Pi_{t+1})^2}\right]}$$

We know that the value function at time T : $J(F_T, \pi_T, T) = H(F_T, T)$. Hence, we can conclude from (22) and the first order condition (10) that:

$$(24) \quad \frac{dw_{i,T-1}^*}{dF_{T-1}} = \frac{(1-\alpha)}{F_{T-1}^2}\frac{E_{T-1}\left[H_F(F_T, T)\frac{(R_{iT}-R_{0T})}{(1+\Pi_T)}\right]}{E_{T-1}\left[H_{FF}(F_T, T)\frac{(R_{iT}-R_{0T})^2}{(1+\Pi_T)^2}\right]} = 0$$

Moreover at time T the value function is independent of the inflation rate and hence from (23) and the first order condition (10):

$$(25) \quad \frac{dw_{i,T-1}^*}{d\pi_{T-1}} = -(1-\beta)\frac{(1-\alpha)}{F_{T-1}^2} \frac{E_{T-1}\left[H_F(F_T,T)\frac{(R_{iT}-R_{0T})}{(1+\Pi_T)}\right]}{E_{T-1}\left[H_{FF}(F_T,T)\frac{(R_{iT}-R_{0T})^2}{(1+\Pi_T)^2}\right]} = 0$$

Let $a(\pi_{T-1}) = E_{T-1}\left[\left(\frac{1+R_T^*}{1+\Pi_T}\right)^\alpha\right]$ and $R_T^* = \sum_{i=1}^I w_{i,T-1}^*(R_{iT} - R_{0T}) + R_{0T}$, then the value-function at time $T-1$ is given by

$$(26) \quad J(F_{T-1}, \pi_{T-1}, T-1) = a(\pi_{T-1})\frac{1}{\alpha}F_{T-1}^\alpha$$

Next, it is straightforward to show that $\frac{da(\pi_{T-1})}{d\pi_{T-1}} = -\alpha(1-\beta)a(\pi_{T-1})$ and hence we have the following relation:

$$(27) \quad J_{F\pi}(F_{T-1}, \pi_{T-1}, T-1) = -\alpha(1-\beta)J_F(F_{T-1}, \pi_{T-1}, T-1)$$

Hence from (22) and (23) we may conclude that

$$(28) \quad \frac{dw_{i,T-2}}{dF_{T-2}} = 0 \ and \ \frac{dw_{i,T-2}}{d\pi_{T-2}} = 0.$$

Recursively working backward and repeating the same steps we therefore derive that

$$(29) \quad \frac{dw_{it}}{dF_t} = 0 \ and \ \frac{dw_{it}}{d\pi_t} = 0, \ \ for \ t = 0, \dots, T-1.$$

Hence, the investor with constant relative risk aversion ignores future asset returns and does not hedge dynamically against fluctuations in the wage growth rate. Note that this conclusion differs from the standard result in the literature for investors with CRRA facing a stochastic opportunity set. However, in our economy the wage growth does not affect the drift rate and the volatility of the asset returns and therefore substitution between assets as a result of a change of the wage growth is needless. Wealth effects are present, but have no consequence for the portfolio choice of an investor with CRRA.

B Implementation of the Dynamic Programming Algorithm

In this appendix we specify the decision rules for interpolation and transformation that were used to solve the numerical cases in the paper with the discretized dynamic programming algorithm.

First we consider the cases without stochastic wage growth. Let $n = 0, 1, ..., N$ represent the points on the grid for the fund value F (wealth-to-income ratio), with a fixed step size of Δ_F between the $N + 1$ points and minimum fund value $\underline{F}$. Let $J(n, t)$ denote the value function in point $n = 0, 1, ..., N$ on the grid at time $t = 0, 1, ..., T$.

Suppose that we are at time t in the dynamic programming algorithm and want to know the value function at F^* at time $t + 1$ (in order to approximate the expectation of the value function numerically). Let n^* denote the point on the grid with fund value closest to F^*: $n^* = \lfloor (F^* - \underline{F})/\Delta_F \rfloor$. We apply interpolation with points surrounding n^*, as F^* tends to lie in between the available fund values $\underline{F} + n\Delta F$ on the grid for $n = 0, 1, ..., N$, . Moreover, in some cases F^* is larger than the maximum fund value on the grid ($F^* > \underline{F} + N\Delta_F$) and then we apply extrapolation.

In order to facilitate both interpolation and extrapolation for utility functions with a vertical asymptote, we first transform the value-function: $y(n, t) = H(J(n, t))$, where $H(.)$ is an invertible transformation-function. The goal of the transformation is to get rid of the asymptote at low fund values near $\underline{F}$ (in order to improve the quality of *interpolation*) and to get a roughly linear function at high fund values (in order to improve the quality of *extrapolation*).

Note that the utility function $U(F_T, T)$ over fund value at retirement is crucial for the asymptote of the value function and not the utility function over contribution rates $H(c_t, t)$. As the plan sponsor controls the contribution rate c_t he will never let the utility function over contribution rates go to minus infinity at a faster rate than the value function itself (at low levels of wealth $F \to \underline{F}$). Below is a list of the transformations applied for the cases studied in the paper:

1. Power and DRRA utility over fund values, see the specification in (16):[9]

   $y(n, t) = J(n, t)^{1/(\alpha - p)}$ for $n\Delta_F \leq 0.5$ and
   $y(n, t) = J(n, t)^{1/\alpha}$ for $n\Delta_F > 0.5$.

2. Expected fund value subject to a downside-risk penalty, see the specification in (18):

   no transformation required (no vertical asymptote)

Before we interpolate or extrapolate the transformed value function we would like to know its shape near F^*. Therefore we estimate the scaled second derivative of the transformed value function. Let $\Delta y(n, t + 1)$ denote a numerical estimate for the first derivative of $y(n, t + 1)$ at time $t + 1$ with respect to fund value F, while $\Delta^2 y(n, t + 1)$ is a numerical estimate of the second derivative of the transformed value function.[10] The estimate of the scaled second derivative

---

[9]Power utility is a special case of DRRA with $p = 0$ and $\theta = 0$.

[10]We use a mid-point formula to calculate numerical estimates for the points $n = 2, 3, .., N - 2$. For the points on the edge of the grid, $n = 0, 1$ and $n = N - 1, N$, we use a one-sided approximation of the derivatives (see Press e.a. 1989).

(absolute risk aversion) is:

$$(30) \quad z(n, t+1) = \frac{\Delta^2 y(n, t+1)}{\Delta y(n, t+1)}, \ for \ n = 0, 1, ..., N.$$

Depending on the magnitude of the scaled second derivative we choose an interpolation method. Interpolation with a local rational polynomial is a good method for nearly any function, even with an asymptote (see Press e.a. 1989), however it is rather slow. If we know that the value function is locally linear then we rather apply a fast linear interpolation scheme with two points. Moreover, if the value function has an abrupt change in the second derivative due to a 'break' (as in the case of downside-risk measures near the threshold) it is safer to apply linear interpolation.

We use a decision rule for choosing the interpolation method, in order to increase the efficiency of the algorithm without sacrificing its reliability. Extensive numerical testing of efficiency and reliability resulted in the following specification:

1. Power and DRRA utility over fund values:

   If $F^* > \underline{F} + N\Delta_F$ then we apply linear *extrapolation* with 2 points.
   If $|z(n^*, t+1)| > 4.5$ then we apply rational interpolation with 10 points.
   If $0.5 < |z(n^*, t+1)| \leq 4.5$ then we apply polynomial interpolation with 4 points.
   If $|z(n^*, t+1)| \leq 0.5$ then we apply linear interpolation with 2 points.

2. Expected fund value subject to a downside-risk penalty:

   If $F^* > \underline{F} + N\Delta_F$ then we apply linear *extrapolation* with 2 points.
   If $F^* < 0.9$ and $|z(n^*, t+1)| > 0.5$ then we apply polynomial interpolation with 4 points.
   In other cases we apply linear interpolation.

After applying the interpolation or extrapolation above, we obtain y^* as an estimate of the transformed value function in the point F^*. As a last step, we use the inverse transformation $H^{-1}(y^*) = J^*$, to get the required estimate of the original value function in the point F^*.

Finally, if we additionally consider mean-reverting wage growth rates then the state space grid becomes 2-dimensional. In the wage growth dimension we use linear interpolation with two points, as there are no problems with asymptotes or abrupt changes of the value function. We apply the transformation and interpolation methods as described previously in the fund value dimension. We refer to Press e.a (1989) for an introduction to multi-dimensional interpolation.

C   Portfolio Separation in Discrete-Time

Given Assumption 4 and an additional linear approximation rule for discrete returns, we will prove that three-fund separation holds for any investment objective. Fund separation results are well known for log-normal return distributions in continuous time since Merton (1971), but to our knowledge this is the first proof for a discrete time model. In order to derive these results, we first apply the following linear approximation for the continuously compounded returns $r_{i,t+1}$:[11]

$$(31) \quad \frac{(1 + R_{i,t+1})}{(1 + \Pi_{t+1})} \; = \; e^{r_{i,t+1} - \pi_{t+1}} \approx 1 + r_{i,t+1} - \pi_{t+1}$$
$$\sim \; N\left(1 + \mu_i - (\pi_t + \beta(\mu_\pi - \pi_t)), \sigma_i^2 - 2\sigma_{i\pi} + \sigma_\pi^2\right), \; for \; i = 0, 1, ..., I$$

Moreover, we assume that the nominal return of asset 0 is riskless: $r_{0,t+1}$ is constant for $t = 0, 1, ..., T - 1$ and will be denoted by r_0. This assumption is purely for notational convenience and does not change our conclusions.

Consider the general first-order condition for investment (10) at time t in the dynamic programming recursion. We may formulate this equation as follows, using the definition of covariance:

$$(32) \quad 0 = Cov_t(J_F, r_{i,t+1} - r_0) \; + \; E_t[J_F]E_t[r_{i,t+1} - r_0], \; for \; i = 1, 2, ..., I$$

Given that the continuously compounded asset returns and wage growth rate are joint normally distributed we apply Stein's Lemma[12], yielding

$$(33) \quad 0 \; = \; F_t E_t[J_{FF}]Cov_t(r_{t+1}^* - \pi_{t+1}, r_{i,t+1} - r_0) \; + \; E_t[J_{F\pi}]Cov_t(\pi_{t+1}, r_{i,t+1} - r_0)$$
$$+ \; E_t[J_F]E_t[r_{i,t+1} - r_0] = 0, \; for \; i = 1, 2, ..., I$$

Solving this equation for the vector of optimal portfolio weights w_t^* we obtain:

$$(34) \quad w_t^* = -\frac{E_t[J_F]}{F_t E_t[J_{FF}]}\Omega^{-1}(\mu - \iota r_0) + \left(1 - \frac{E_t[J_{F\pi}]}{F_t E[J_{FF}]}\right)\Omega^{-1}\sigma_{I\pi}$$

---

[11]The proof presented here holds exactly, if the discretely compounded asset returns are generated by a multi-variate *normal* distribution. We have to apply an additional approximation because we assume a multi-variate *log-normal* distribution for the discrete asset returns.

[12]Stein's Lemma states that $Cov(f(X), Y) = E(f'(X))Cov(X, Y)$, for a differentiable function $f(\cdot)$ and X, Y bivariate normally distributed. A proof for the multi-variate case is available upon request from the authors.

References

Balduzzi, P. and A. Lynch, 1999a, Transaction Costs and Predictability: Some Utility Cost Calculations, Journal of Financial Economics 52, 47-78.

Balduzzi, P. and A. Lynch, 1999b, Predictability and Transaction Costs: The Impact on Rebalancing Rules and Behavior, forthcoming Journal of Finance.

Barberis, N., 1999, Investing for the Long Run when Returns are Predictable, forthcoming Journal of Finance.

Bawa V.S. and E.B. Lindenberg, 1977, Capital Market Equilibrium in a Mean-Lower Partial Moment Framework, Journal of Financial Economics 5, 189-200.

Bellman R., 1957, Dynamic Programming, Princeton University Press, Princeton.

Bertsekas D.P., 1976, Dynamic Programming and Stochastic Control, Academic Press, New York.

Boender G.C.E., 1997, A Hybrid Simulation/Optimisation Scenario Model for Asset/Liability Management, European Journal of Operational Research 99, 126-135.

Brandt, M.W., 1999, Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach, Journal of Finance 54, 1609-1646.

Brennan, M.J., E.S. Schwartz, and R. Lagnado, 1997, Strategic Asset Allocation, Journal of Economic Dynamics and Control 21, 1377-1403.

Campbell, J.Y., 1993, Intertemporal Asset Pricing Without Consumption Data, American Economic Review 83 (3), 487-512.

Campbell, J.Y. and L. Viceira, 1999, Consumption and Portfolio Decisions when Expected Returns are Time-Varying, forthcoming Quarterly Journal of Economics.

Campbell, J.Y. and L. Viceira, 1999, Who Should Buy Long Term Bonds, NBER Working Paper 6801.

Campbell, J.Y., J. Cocco, F. Gomes, P. Maenhout, and L. Viceira, 1998, Stock Market Mean Reversion and the Optimal Equity Allocation of a Long-Lived Investor, Manuscript, Harvard University.

Cox, J.C. and C.F. Huang, 1989, Optimum Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process, Journal of Economic Theory 49, 33-83.

Duffie, D., W. Fleming, M. Soner and T. Zariphopoulou, 1997, Hedging in Incomplete Markets with HARA Utility, Journal of Economic Dynamics and Control 21, 753-782.

Hakansson N.H., 1969, Optimal Investment and Consumption Strategies under Risk, An Uncertain Lifetime and Insurance, International Economic Review 10, 443-466.

Hakansson N.H., 1970, Optimal Investment and Consumption Strategies under Risk for a Class of Utility Functions, Econometrica 38, 587-607.

Hakansson N.H., 1971, On Optimal Myopic Portfolio Policies, With and Without Serial Correlation of Yields, Journal of Business 44, 324-334.

He, H. and N.D. Pearson, 1991, Consumption and Portfolio Policies with Incomplete Markets and Short Sale Constraints: The Infinite-Dimensional Case, Journal of Economic Theory 54, 259–304.

Karatzas, I., J. Lehockzky and S. Shreve, 1987, Optimal Portfolio and Consumption Decisions for a 'Small Investor' on a Finite Horizon, SIAM Journal on Control and Optimization 25, 1157-1186.

Koo, H. K., 1998, Consumption and Portfolio Selection with Labor Income: A Continuous-Time Approach, Mathematical Finance 8, 49-66.

Koo, H. K., 1999, Consumption and Portfolio Selection with Labor Income: A Discrete-Time Approach, Mathematical Methods of Operations Research 50, 219-243.

Kouwenberg R., 1998, Scenario Generation and Stochastic Programming Models for Asset-Liability Management, forthcoming European Journal of Operational Research.

Leibowitz M.L., S. Kogelman and L.N. Bader, 1994, Funding Ratio Return, Journal of Portfolio Management, Fall, 39-47.

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

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

Mossin J., 1968, Optimal Multiperiod Portfolio Policies, Journal of Business 41, 215-229.

Mulvey J.M. and W.T. Ziemba, 1998, Asset and Liability Management Systems for Long-Term Investors: Discussion of the Issues, in : J.M. Mulvey and W.T. Ziemba, eds., Worldwide Asset and Liability Modelling, (Cambridge University Press, Cambridge), 3-38.

Pliska. S.R., 1986, A Stochastic Calculus Model of Continuous Trading: Optimal Portfolios, Mathematics of Operations Research 11, 371-382.

Press W., B. Flannery, S. Teukolsky and W. Vetterling, 1989, Numerical Recipes: The Art of Scientific Computing, Cambridge University Press, Cambridge.

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

Sethi, S.P. and M. Taksar, 1988, A Note on Merton's 'Optimum Consumption and Portfolio Rules in a Continuous-Time Model', Journal of Economic Theory 46, 395-401.

Sharpe W.F. and L.G. Tint, 1990, Liabilities - A New Approach, Journal of Portfolio Management, Winter, 5-10.

Viceira, L., 1999, Optimal Portfolio Choice for Long Horizon Investors with Nontradable Labor Income, NBER Working Paper 7409.

Figure 1: Relative Risk Aversion Function for DRRA
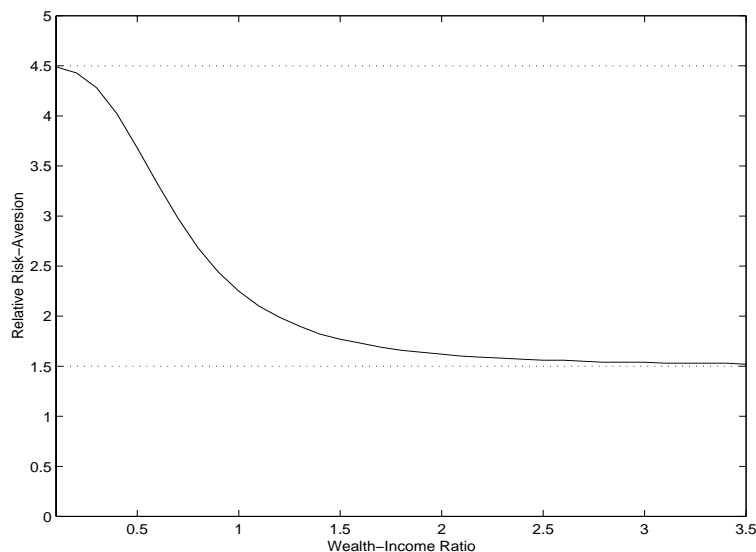
This figure shows the Arrow-Pratt measure of relative risk aversion (17) of utility function (16) as a function of the wealth-income ratio F, for parameter values $\alpha = -0.5$, $p = 3$ and $\theta = 3$.

Figure 2: Stock weight for DRRA
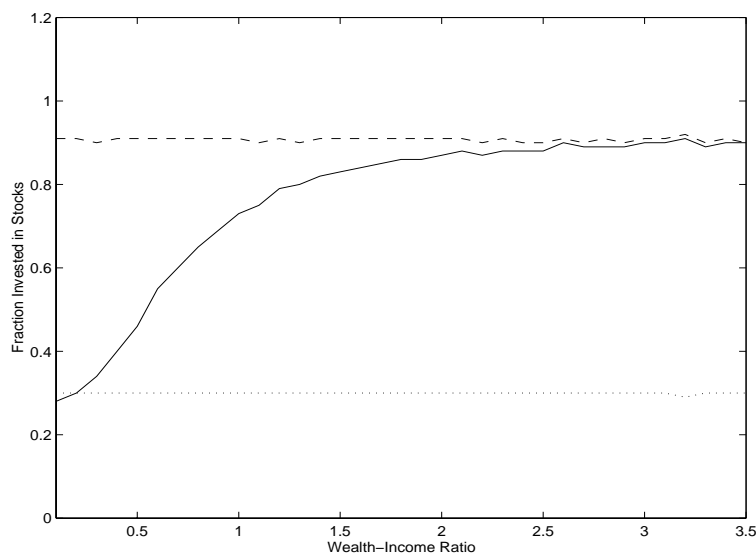
This figure shows the fraction invested in stocks for constant relative risk aversion levels $\alpha = -3.5$ (dotted line) and $\alpha = -0.5$ (dashed line) and decreasing relative risk aversion (solid curve). A more risk averse investor dedicates less to stocks (28% for risk aversion level $\alpha = -3.5$ compared to 90% for risk aversion level $\alpha = -0.5$). Moreover decreasing relative risk aversion implies that the investor desires less exposure to stocks when real wealth drops.

Figure 3: Stock weight for DRRA with stochastic wage growth ($\rho_{S\pi} = 0$ and $\rho_{B\pi} = 0.33$)
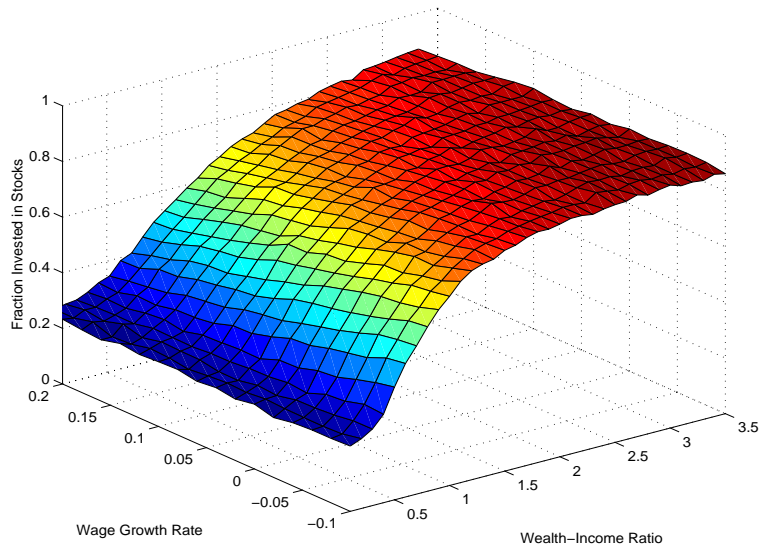
This figure shows the fraction invested in stocks when returns on bonds are positively correlated with wage growth. The percentage of stocks in the optimal portfolio slightly decreases for increasing wage growth levels.

Figure 4: Stock weight for DRRA with stochastic wage growth ($\rho_{S\pi} = 0.33$ and $\rho_{B\pi} = 0$)
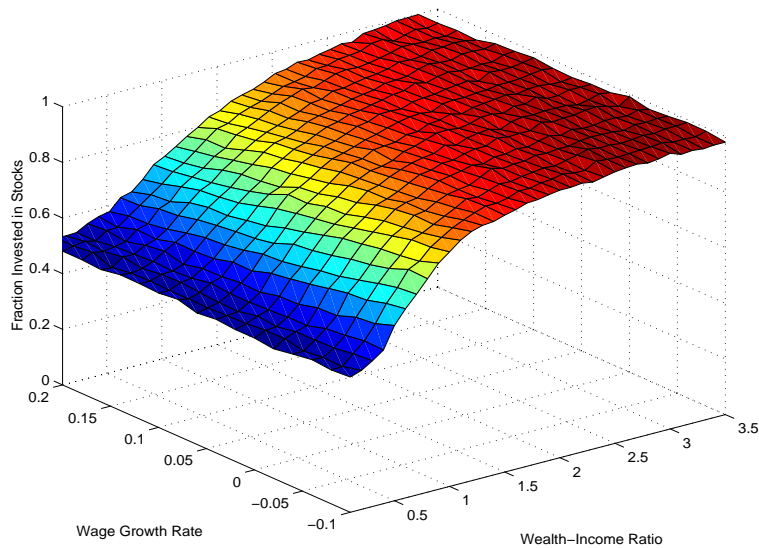
This figure shows the fraction invested in stocks when returns on stocks are positively correlated with wage growth. Compared to Figure 3 the graph is shifted only in a parallel way, i.e. the weight increases generally when stock returns are positively correlated with wage growth. Adjustments of the portfolio as a function of the wage growth rate are small and comparable to Figure 3.

Figure 5: Stock weight for Downside-Risk Objective
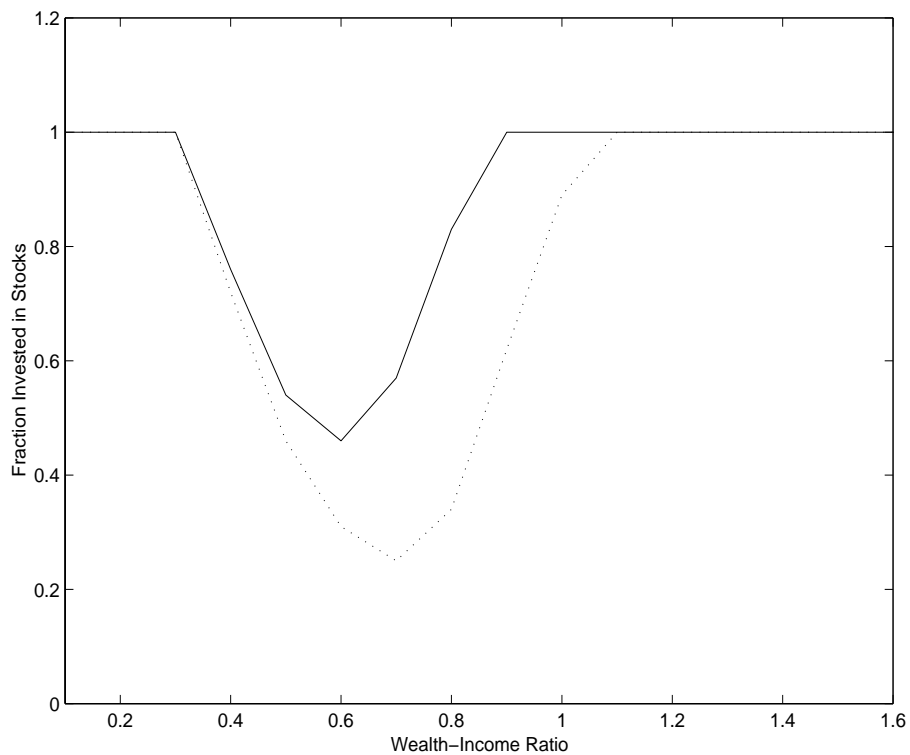
This figure shows the fraction invested in stocks at time $t = 0$ as a function of the wealth-income ratio for risk aversion parameters $\nu = 16$ (solid curve) and $\nu = 64$ (dotted line). Note that the investor cannot dedicate more than 100% to stocks due to a borrowing constraint.

Figure 6: Stock weight for CRRA with fixed contribution rates
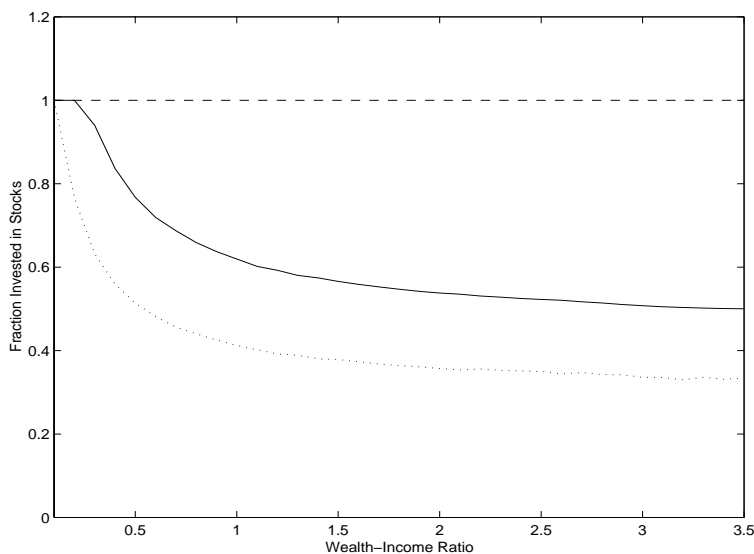
The figure shows the stock weight at time $t = 0$ for power-utility with risk aversion $\alpha = -0.5$ (dashed line), $\alpha = -2$ (solid line) and $\alpha = -3.5$ (dotted line), with a fixed contribution rate of $c_t = 0.05$ for $t = 0, 1, ..., T - 1$. Due to the fixed contribution rate the portfolio weight of stocks increases at low levels of fund value (wealth-to-income ratio).

Figure 7: Stock weight with HARA-utility over contribution rates
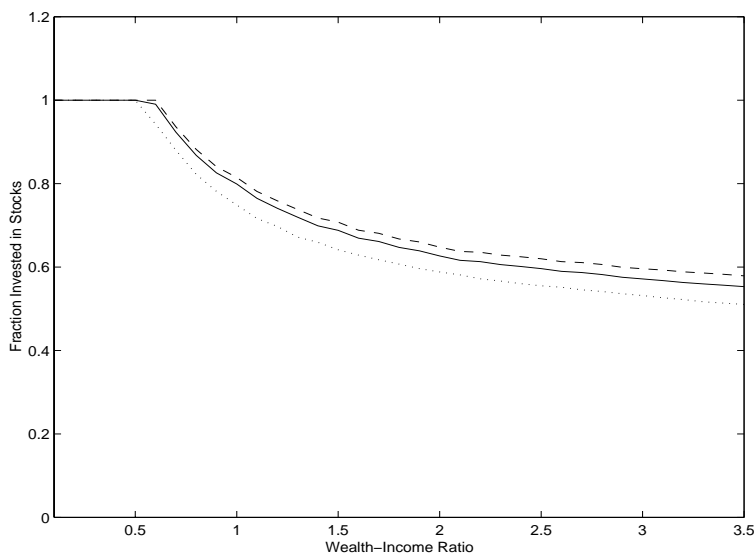
The figure shows the fraction invested in stocks at time $t = 0$ for power-utility with $\alpha = -2$ over fund value and HARA-utility over contribution rates with $\gamma = -2.5$ (dashed line), -2 (solid line) and -1.5 (dotted line). The weight of stock increases to the upper bound 100% at low levels of fund value, due to the increasing marginal value of the contribution payment.

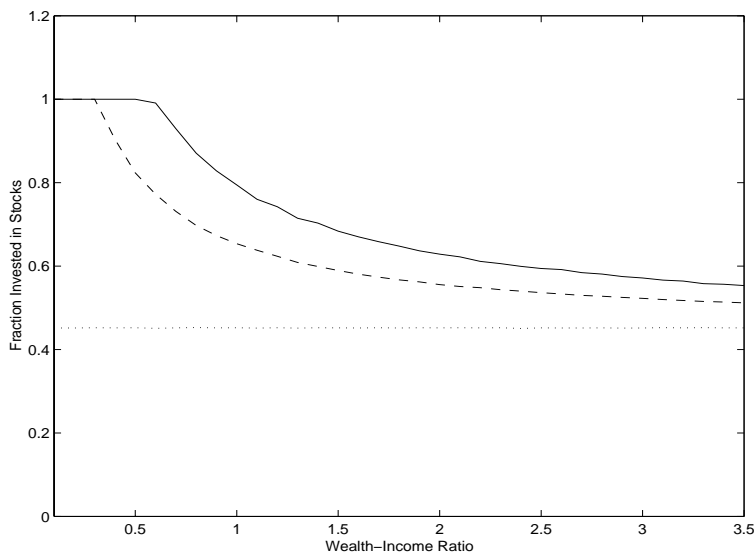Figure 8: Contribution rates with HARA-utility
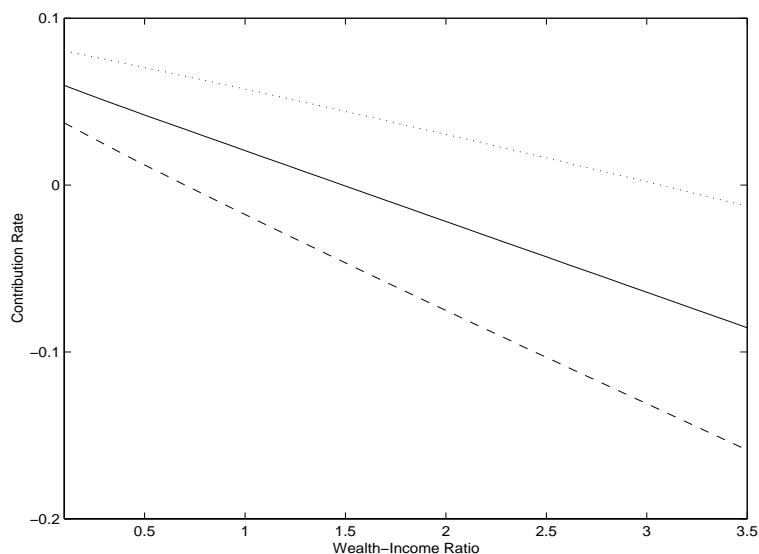
The figure shows the contribution rate at time $t = 0$ for power-utility with $\alpha = -2$ over fund value and HARA-utility over contribution rates with $\gamma = -2.5$ (dashed line), -2 (solid line) and -1.5 (dotted line). The contribution rates decrease approximately linear as a function of fund value.

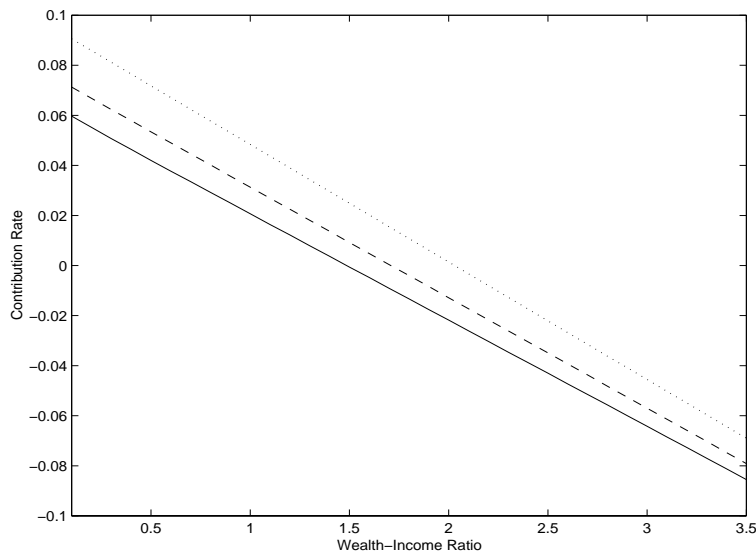Figure 9: Stock weight at time $t = 9,4$ and 0 for DRRA utility
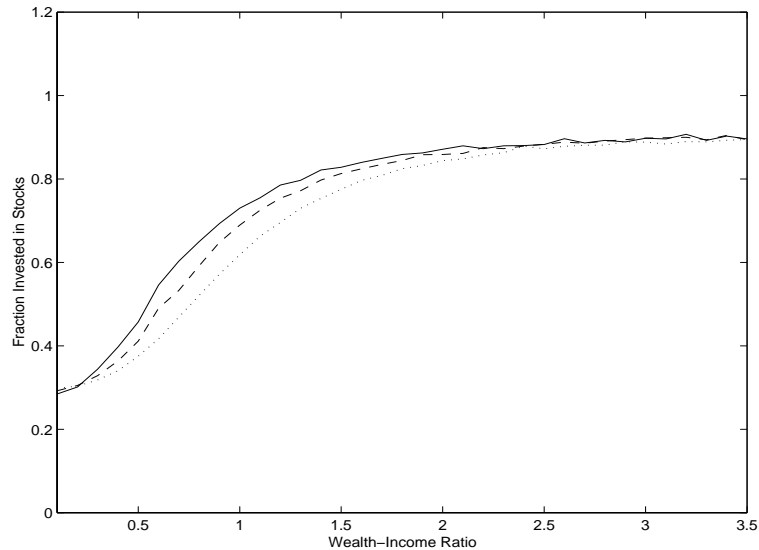
The figure shows the fraction invested in stocks for different investment horizons, given DRRA utility with parameter values $\alpha = -0.5$, $p = 3$ and $\theta = 3$. The dotted line represents the optimal fraction with only one year left ($t = 9$). The dashed line correspond to an intermediate period ($t = 4$) and the solid line represents the initial portfolio choice ($t = 0$).

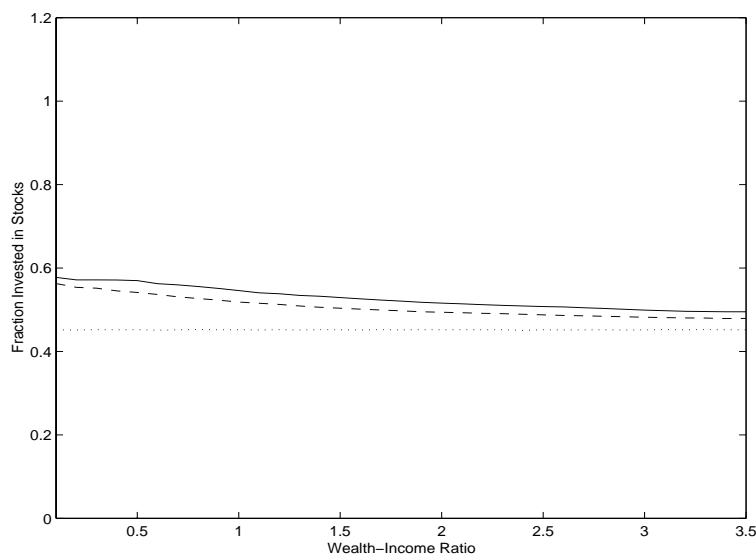Figure 10: Stock weight at time $t = 9,4$ and 0 for CRRA utility, with a fixed contribution rate
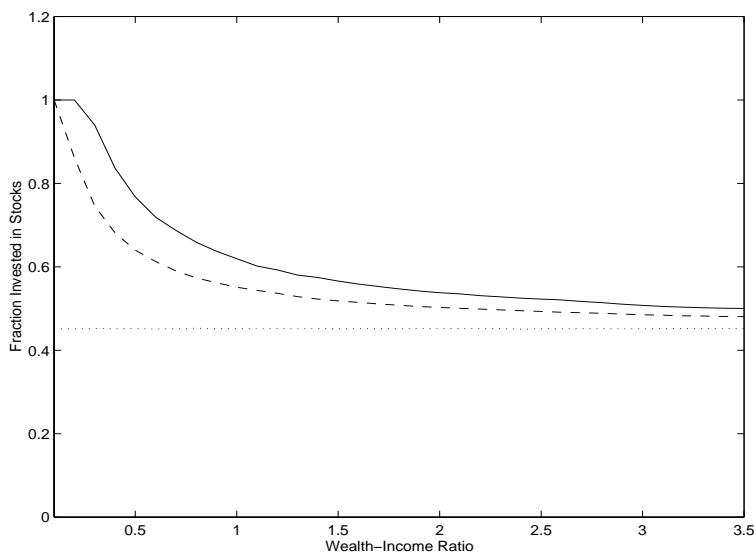
The figure shows the fraction invested in stocks for different investment horizons, given CRRA utility with risk aversion $\alpha = -2$ and a fixed contribution rate $c_t = 0.05$ for $t = 0, 1, ..., T - 1$. The dotted line represents the optimal fraction with only one year left ($t = 9$). The dashed line correspond to an intermediate period ($t = 4$) and the solid line represents the initial portfolio choice ($t = 0$).

Figure 11: Stock weight at time $t = 9,4$ and 0, with HARA utility over contribution rates

The figure shows the fraction invested in stocks for different investment horizons, given CRRA utility with risk aversion $\alpha = -2$ and HARA-utility over contribution rates with $\gamma = -2$. The dotted line represents the optimal fraction with only one year left ($t = 9$). The dashed line correspond to an intermediate period ($t = 4$) and the solid line represents the initial portfolio choice ($t = 0$).

Figure 12: Contribution rate at time $t = 9,4$ and 0, with HARA utility over contribution rates

The figure shows the contribution rate for different investment horizons, given CRRA utility with risk aversion $\alpha = -2$ and HARA-utility over contribution rates with $\gamma = -2$. The dotted line represents the optimal contribution rate with only one year left ($t = 9$). The dashed line correspond to an intermediate period ($t = 4$) and the solid line represents the initial portfolio choice ($t = 0$).

Figure 13: Stock weight at time $t = 9,4$ and 0, with intertemporal utility over fund value

The figure shows the fraction invested in stocks for different investment horizons, given *intertemporal measurement* of CRRA utility over fund values with risk aversion $\alpha = -2$ and HARA-utility over contribution rates with $\gamma = -2$. The dotted line represents the optimal fraction with only one year left ($t = 9$). The dashed line correspond to an intermediate period ($t = 4$) and the solid line represents the initial portfolio choice ($t = 0$).

Figure 14: Contribution rate at time $t = 9, 4$ and 0, with intertemporal utility over fund value
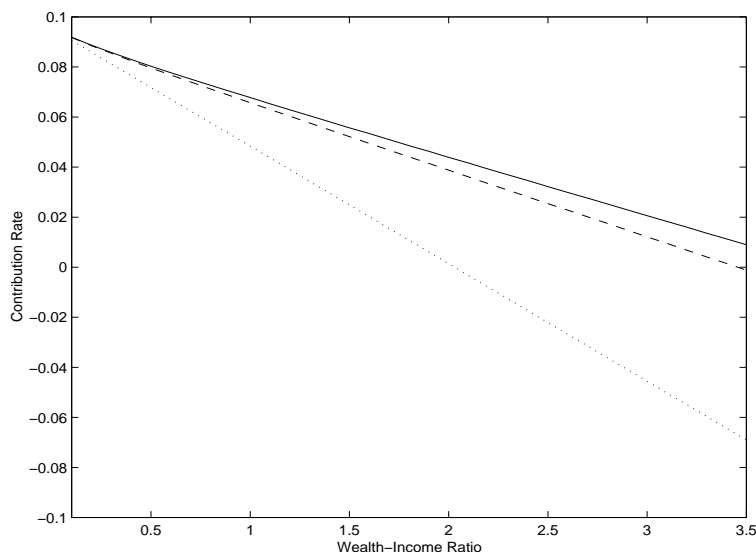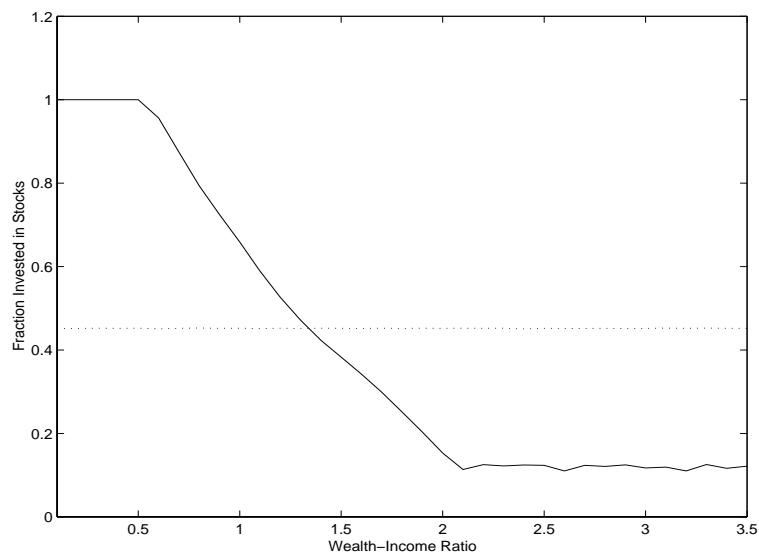

The figure shows the optimal contribution rate for different investment horizons, given *intertemporal measurement* of CRRA utility over fund values with risk aversion $\alpha = -2$ and HARA-utility over contribution rates with $\gamma = -2$. The dotted line represents the optimal contribution rate with only one year left ($t = 9$). The dashed line correspond to an intermediate period ($t = 4$) and the solid line represents the initial portfolio choice ($t = 0$). Due to the intertemporal measurement of utility over fund value, the contribution rates increase for longer investment horizons (contrary to the results in Figure 12).

Figure 15: Stock weight for CRRA with a lower bound on funding payments
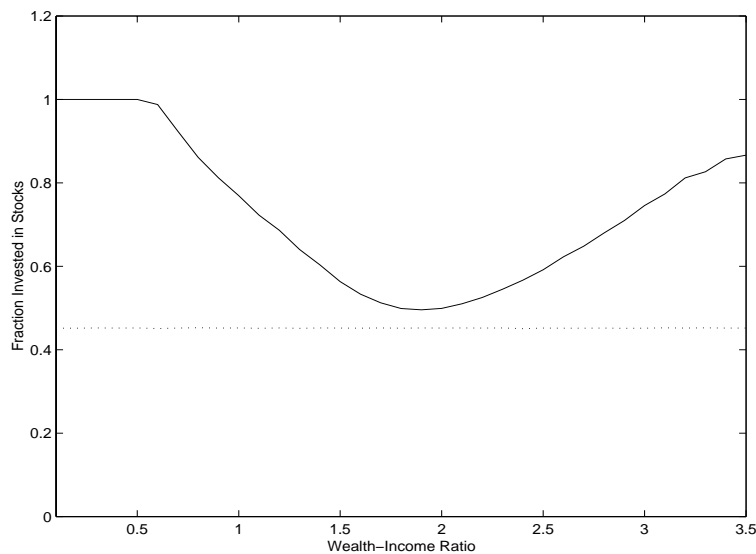
This figure shows the fraction invested in stocks with lower bound $\underline{c} = 0$ on contribution rates. The upward part of the smile is due to the lower bound on funding payments, which provide a cushion against losses.

Figure 16: Stock weight for CRRA with an upper bound on fund value

This figure shows the fraction invested in stocks with an upper bound $\overline{F} = 2$ on fund values, enforced with negative contribution rates (i.e. refunds). Beyond the maximum allowed fund value $\overline{F} = 2$ the stock weight becomes constant at a low level.