Cash-flow driven investment beyond expectations

July 25, 2021

Abstract

This paper makes significant advances in cashflow-driven investment where the aim is to find buy-and-hold portfolios whose future payouts cover given liability payments as well as possible. While current industry solutions are largely based on expected future cash flows, we use a stochastic optimization model that seeks portfolios that give the best possible match across time as well as scenarios. The optimized hedging strategies are able to employ any statistical connections between the liabilities and publicly quoted assets. Reinvestment risk is described by a stochastic model of an illiquid money market. While cashflow matching across scenarios is controlled by the risk aversion, the timing is controlled by the illiquidity factors. Besides optimal hedging strategies, we find the least cost of hedging which provides a market-consistent valuation based on the current quotes and the liquidity factors as well as the views and risk preferences of the investor/regulator. The approach is illustrated by pricing and hedging of defined benefit pension liabilities which depend on uncertain longevity developments and the consumer price index. The hedging strategies are constructed from 128 publicly quoted instruments including index-linked bonds and equities. Increasing the risk aversion and the illiquidity parameters, we find portfolios that hedge the liabilities with significantly lower risk but only slightly higher cost.

Keywords: Defined benefit pensions, Asset-liability management, Cashflow-driven investment, Convex stochastic optimization

1 Introduction

Cashflow-driven investment (CDI) has become a recommended approach for defined benefit (DB) pension funds who face an outflow of cash over the coming decades; see [HFW, Wil19, Pen19, cdi19, Exl17, KMB19, Mci19, Ins] for a small sample of practitioner-oriented papers. A CDI portfolio aims to cover the pension payments with the contractual payments of the involved assets when held to maturity. If one could achieve perfect cashflow matching, such a portfolio would provide the security of delivering liability payments without exposure to the reinvestment risks associated with future trading. This would be ideal for large DB schemes who face significant liquidity risks that may be difficult to control with dynamic asset allocation. Many more traditional ALM-approaches aim at matching valuations of liabilities with valuations of assets so they are sensitive to the employed valuation formulas which often have little to do with the true costs of delivering the cashflows. CDI can be seen as an economically consistent ALM technique as it focuses on the actual delivery of the pension benefits.

Most publicly available descriptions of CDI, however, focus on matching expected cashflows. This is inline with common actuarial practice, but it ignores the significant financial risks associated with DB-liabilities that often extend several decades into the future. The present paper develops a convex stochastic optimization model that incorporates uncertainty directly into CDI. Inclusion of risk into the optimization model allows for a consistent treatment of risky "growth assets" as well as hedging instruments whose contractual cashflows are uncertain by definition. The model can employ any publicly traded assets and it calibrates explicitly to the available quotes. Following the CDI principles, it seeks buy-and-hold portfolios in the quoted assets, but instead of matching expected cashflows, it seeks cheapest available portfolios that cover the uncertain liabilities at an acceptable level of risk. The risk is measured by a user-specified risk measure on the net terminal wealth.

As opposed to deterministic CDI models, real pension payments cannot be exactly matched with payouts of buy-and-hold portfolios. This gives rise to reinvestment risk of financing cashflow mismatches in the face of uncertain lending/borrowing costs and liquidity risk. This is the financial motivation for the aim of matching cashflows across time. We describe the reinvestment risk by stochastic money market rates with a user specified illiquidity parameter. While the liquidity risk drives the cashflow matching over time, the risk aversion drives cashflow matching across different scenarios. This allows the model to exploit the hedging potential of e.g. inflation-linked bonds which, in a deterministic model, would seem overpriced with respect to government bonds when only comparing the expected payouts.

Our optimization model seeks the cheapest buy-and-hold portfolios whose cashflows (together with reinvestment in the illiquid money market) cover the liabilities in all scenarios at an acceptable level of risk. The cost of a portfolio is calculated from available market quotes so the optimum value gives a natural hedging-based market-consistent valuation of the given liabilities. This can be seen as a natural extension of the classical no-arbitrage pricing principle to the incomplete market setting with bid-ask spreads, illiquidity effects and pension liabilities; see [Pen14] for a general study of contingent valuation in incomplete markets and [HKP11] for applications to pensions. This is in sharp contrast with the common actuarial valuations which are based on discounting expected future liability cashflows with a point estimate of future investment returns.

We illustrate the model by pricing and hedging defined benefit pension liabilities that are subject to longevity and indexation risks. As hedging instruments, we use live quotes for 32 gilts, 22 inflation-linked bonds and 35 zero-coupon bonds as well as equities with predetermined liquidations strategies. All the relevant risk factors are described by the probabilistic model of [AMAP21] that captures dependencies across time as well as the different risk factors. Optimal hedging strategies are constructed numerically by first discretizing the underlying probability measure and then solving the resulting finite-dimensional convex optimization problems by an interior point solver. The only dynamically updated decision variable is the money market position which is completely determined by the buy-and-hold portfolio and the development of the risk factors. It follows that the position is automatically adapted to the underlying stochastic processes so we can avoid using scenario trees in the discretization thus avoiding the complications that come with them; see [Sod05] for a survey of scenario tree-based ALM models and [MSZ$^+$08, GZ08] for applications to pensions.

Even with the above off-the-shelf computational approach, we find approximately optimal solutions within minutes. The quality of a solution is verified in out-of-sample simulations in less than a second on a desktop computer. The optimized hedging strategies achieve lower risk at a lower cost than the strategies obtained by matching expected cashflows. The risk and liquidity factors allow for an effective way to control the risk and portfolio composition. Increasing the risk aversion and/or the illiquidity parameters, the optimal portfolios become more diversified and they cover the pension payments well across time and scenarios with a moderate increase in the hedging cost. When the risk aversion is increased, the allocation shifts from equities towards inflation-linked bonds whose cashflows are more closely connected to the inflation-indexed pension benefits. When the money market liquidity decreases, there is a move from gilts to zero-coupon bonds that give more control over the timing of the payments.

The presented models and computational techniques are not limited to pensions and the hedging instruments employed in the computational examples. A similar approach could be taken in pricing and hedging of corporate debt, index-liked bonds or any other liability whose contractual payments can be described by a stochastic model. The payouts of corporate debt could be modelled e.g. by incorporating default intensity factors in the underlying stochastic model of the risk factors. Corporate bonds will be an important addition to the space of hedging instruments when the models presented in this paper are implemented in practice.

2 CDI under uncertainty

Consider a closed DB-scheme with outstanding future pension payments c_t over a finite time $t = 1, \dots, T$. Our aim is to design investment strategies that cover the future pension payments as well as possible. In the context of CDI, the most typical investment classes are fixed-rate and inflation-linked bonds as well as riskier instruments such as stocks and other "growth assets" aimed at achieving better returns in the face of long-dated liabilities. The basic CDI aims to construct buy-and-hold portfolios whose cash-flows match those of the pension liabilities over time. While investment grade fixed-rate bonds provide fairly predictable cash-flows, the pension payments c are subject to significant uncertainties due to longevity and indexation risks as the time horizon T is typically several decades in the future. Also, when employing index-linked bonds and growth assets as parts of the investment strategy, the investment income becomes highly uncertain as well. It follows that the cash-flows can be matched only partially and in each period t there will be either surplus or deficit that needs to be reinvested or paid by money market operations. If the money market was perfectly liquid and predictable, this wouldn't present problems as one could then roll over all mismatches over time and settle accounts at the end. In reality, however, the money market operations are subject to illiquidity and uncertainties that create an incentive to match the cash-flows across time and different future scenarios. This section proposes a mathematical model for optimization of CDI-strategies in the face of illiquidities and uncertainties.

Let K be a finite collection of assets the fund can buy or sell at time $t = 0$. We denote the cost of buying $z_k \in \mathbb{R}$ units of contract $k \in K$ by $p_k(z_k)$. As usual, negative purchases are interpreted as sales. If infinite quantities were available at the best bid and ask prices, we would simply have

$$p_k(z_k) = \begin{cases} p_k^a z_k & \text{if } z_k \geq 0, \\ p_k^b z_k & \text{if } z_k \leq 0, \end{cases}$$

where p_k^b and p_k^a are the bid- and ask-prices, respectively, of contract $k \in K$. As usual, buying negative units means selling. The cash-flow provided by one unit of $k \in K$ at time $t = 1, \ldots, T$ will be denoted by $c_{k,t}$. For example, if k is a fixed rate government bond with maturity T_k, then c_{k,T_k} would be the principal payment, $c_{k,t}$ for $t = 1, \ldots, T_k - 1$ would be the annual coupon payments while $c_{k,t} = 0$ for $t > T_k$. For index-linked bonds and stocks, the cash-flows would be stochastic; see Section 5 below.

If we hold $z^k \in \mathbb{R}$ units of contract $k \in K$ then the net investment income from all assets at time t is given by $\sum_{k \in K} z_k c_{k,t}$. In an idealized CDI, the net cash-flows from investments match the liability cash-flows $c = (c_t)_{t=0}^T$ perfectly so that $c_t = \sum_{k \in K} z_k c_{k,t}$. In reality, this can only be achieved approximately so the surplus/deficit needs to be covered by dynamically trading in the financial markets. The amount x_t of cash invested in the markets at time t evolves according to the equation

$$x_t = x_{t-1} + R_{t-1}(x_{t-1}) + \sum_{k \in K} z_k c_{k,t} - c_t, \quad t = 1, \ldots, T, \tag{1}$$

where $R_{t-1}(x_{t-1})$ is the interest received at time t when investing x_{t-1} units of cash over the period $[t-1, t]$. In the numerical implementations below, we will assume that

$$R_t(x) = \begin{cases} r_t^l x & \text{if } x \geq 0, \\ r_t^b x & \text{if } x \leq 0, \end{cases} \tag{2}$$

where r_t^l and r_t^b denote the lending and borrowing rates, respectively, at time t. Both $r^l = (r_t^l)_{t=0}^T$ and $r^b = (r_t^b)_{t=0}^T$ will be modelled as stochastic processes; see Section 5 below. In reality, the rates satisfy $r_t^b > r_t^l > -1$. In other words, the borrowing rate is always higher than the lending rate and the lending rate is greater than -100%. Violation of the second inequality would mean that money would become nondisposable waste. It should be noted that we do allow for strictly negative rates which would mean that money is subject to storage cost. At the time of writing, the interbank rates as well as the central bank overnight rate in the euro-zone are around -0.5%.

Given a portfolio z in the statically held assets K, equation (1) determines the development of the cash position of the fund uniquely in each scenario. The terminal wealth x_T is thus a random variable determined by z and the realization of all the risk factors that affect the cash-flows and money market returns; see Section 5 below. We will study the problem of finding the cheapest allocation $z \in \mathbb{R}^K$ that covers the pension payments over time and in all scenarios with an acceptable level of risk. Mathematically, the problem can be written as

$$\begin{array}{lll} \text{minimize} & x_0 + \sum_{k\in K} p_k(z_k) \quad \text{over} \quad z \in \mathbb{R}^K, x \in \mathcal{N} & \\ \text{subject to} & x_t \leq x_{t-1} + R_{t-1}(x_{t-1}) + \sum_{k\in K} z_k c_{k,t} - c_t, \quad t = 1,\dots,T, \; P\text{-a.s.} & \\ & \mathcal{R}(x_T) \leq 0, & \end{array} \tag{CDI}$$

where $\mathcal{N}$ is the linear space of adapted stochastic processes, $\mathcal{R}$ is a given risk measure and P denotes the probability measure that describes the agent's views concerning the relevant risk factors. We have written the budget constraint as an inequality to clarify that the problem is that of convex optimization. The inequality allows for throwing away cash but as long as the risk measure $\mathcal{R}$ is strictly decreasing in the sense that $\mathcal{R}(x^2) < \mathcal{R}(x^1)$ whenever $x^1 < x^2$ almost surely, the constraint will hold as an equality at any optimal solution. Indeed, since we are assuming that the lending rate is always greater than -100%, it is always rational to save available money for the future. Allowing for throwing away money, however, makes (CDI) a *convex* optimization problem which greatly facilitates its solution; see Section 6 below. Optimization of investment strategies in practice are driven by an investor's risk preferences and views concerning the future development of the market and the liability cash-flows. In problem (CDI), the risk preferences are described by the risk measure $\mathcal{R}$ and the views by the probability measure P that governs the behaviour of the relevant risk factors. It is important to note that the risk measure may depend on the measure P as is the case e.g. with the *entropic* risk measure

$$\mathcal{R}(x) = \frac{1}{\rho} \ln E e^{-\rho x}, \tag{3}$$

where $\lambda > 0$ is a given risk aversion parameter; see e.g. [FS16, Example 4.13]. The same is true of the conditional value at risk (CV@R) given by

$$\mathcal{R}(x) = \inf_{s\in\mathbb{R}} E\left[s + \frac{1}{1-\lambda}(x-s)\right],$$

where $\lambda \in (0,1)$ is a given parameter; see [UR01]. The CV@R measure focuses on the left tail of the distribution as its values do not depend on the distribution of x above a given quantile. Such a risk measure may be appropriate for a scheme sponsor who is liable to deliver the future pensions but who does not get to keep the possible upside. The entropic risk measure, on the other hand, takes into account the whole distribution so it may be more relevant e.g. for an annuity provider or a reinsurer who owns any residual wealth at time T.

The optimum value of problem (CDI) provides a market consistent valuation of the liabilities c. It is the least cost of "acceptable" hedging of the pension payments c when trading the instruments K and the money market. What is acceptable, is determined by the risk measure $\mathcal{R}$. The cost functions p_k are read off the current market quotes so the valuation is naturally calibrated to the market prices of other instruments. The valuation calibrates also to the given views concerning the risk factors such as money market rates, inflation and longevity developments.

3 Defined benefit pension liabilities

We will study cash-flow driven investment with defined benefit (DB) pension liabilities, where the yearly payments $(c_t)_{t=1}^T$ depend on the number of pensioners as well as the consumer price index (CPI). More precisely, the liability payment in year t is given by

$$c_t = \sum_{b \in B_t} F_t c_0^b,$$

where B_t is a set representing a population of pensioners, c_0^b are their nominal pension entitlements at time $t = 0$ and F_t is the accumulated pension adjustment, given by

$$F_t = \prod_{k=0}^{t} \left[1 + f_{adj} \left(\frac{I_k - I_{k-1}}{I_{k-1}} \right) \right], \tag{4}$$

where f_{adj} is a given function that determines how annual rate of price inflation affects the accrued pension entitlements. In the study, we adopt the adjustment policy of the Universities Superannuation Scheme (USS) [USS18], in which benefits are adjusted by the inflation rate, up to 5% inflation. Above this threshold, the scheme provides a top-up of half the inflation rate, up to a total adjustment of 10%.

Both the population sizes and the inflation are subject to significant uncertainties over the lifetime of the liabilities, i.e., the time it takes for the population size to converge to zero. Section 5 gives a brief description of the stochastic model used to describe these as well as other relevant risk factors in the model.

In the numerical examples below, we will study a hypothetical fund who is liable to pay the pensions for all members until they turn 100 years. The above liabilities should be taken just as an illustration as it would be straightforward to treat more complicated cash-flows with the techniques presented below.

4 Hedging instruments

Our CDI portfolios draw from a set K of hedging instruments that include fixed-rate bonds, inflation-linked bonds (ILB) and equities. As the present study focuses on UK pension funds, we will use UK government gilts, gilt strips (zero-coupon bonds), inflation-linked gilts and FTSE 100 exchange traded funds. We look for optimal buy-and-hold portfolios where all bonds are held to maturity and the investor collects the contractual coupon and principal payments. In the case of equities, we optimize over deterministic liquidation strategies where the number of shares liquidated over time $t = 1, \ldots, T$ does not depend on the scenarios. Such strategies can be expressed as linear combinations of T simple strategies in which all stock investment made at time $t = 0$ are liquidated at time t.

When an element $k \in K$ corresponds to an individual equity strategy, its cash-flows $c_{k,t}$ for the times $t = 1, \ldots, T$, will be given by

$$c_{k,t} = \begin{cases} S_t & \text{if } t = T_k, \\ 0 & \text{otherwise,} \end{cases}$$

where T_k corresponds to its predetermined liquidation date, and S_t is the value of the FTSE 100 index at time t. In the case of coupon paying bonds,

$$c_{k,t} = \begin{cases} C_k & \text{if } t < T_k, \\ 1 & \text{if } t = T_k, \\ 0 & \text{otherwise,} \end{cases}$$

with C_k as the annual coupon payments of k and T_k as its maturity date. Zero-coupon bonds are a special case in which $C_k = 0$. The cash-flows of inflation-linked bonds are given by

$$c_{k,t} = \begin{cases} C_k \frac{I_t}{I_k} & \text{if } t < T_k, \\ 1 \frac{I_t}{I_k} & \text{if } t = T_k, \\ 0 & \text{otherwise,} \end{cases}$$

where the term I_t/I_k is the inflation-adjustment. In the UK, ILBs are often indexed by the retail price index (RPI) instead of the consumer price index I; see e.g. [Uni05, Uni12]. In the numerical illustrations below, we will use the CPI as a proxy for the RPI since the former is readily incorporated in the stochastic model employed in this study.

The cost function p_k associated with the hedging instrument $k \in K$ is given by

$$p_k(z) = \begin{cases} p_k^a z & z \geq 0, \\ p_k^b z & z \leq 0, \end{cases}$$

where p_k^a and p_k^b are the best bid- and ask-prices, respectively, for k. In the numerical illustrations below, we use bid- and ask-prices provided by a Bloomberg terminal at a given point in time simultaneously for all the considered instruments.

Special care must be taken when working with bond quotes provided e.g. by Bloomberg. In most cases, the market quotes are not the actual cost for the settlement of a trade, which also include inflation adjustments and the "accrued interest"; see e.g. [Uni12, Uni05, BSB05]. Figure 1 gives a screenshot from a Bloomberg terminal displaying both the market quotes and the actual costs of trading an inflation-linked bond. Since our model uses yearly time increments, we have also accumulated the biannual coupon payments by rounding the payment dates to the nearest yearly increment. In the numerical study below, we used all the available bond quotes available on 08/04/2021. This included 31 gilts, 22 inflation-linked gilts and 35 zero-coupon bonds. In addition to this, we will use 35 basic equity strategies where 35 is the length T of the planning horizon after which the pension liabilities amortize.

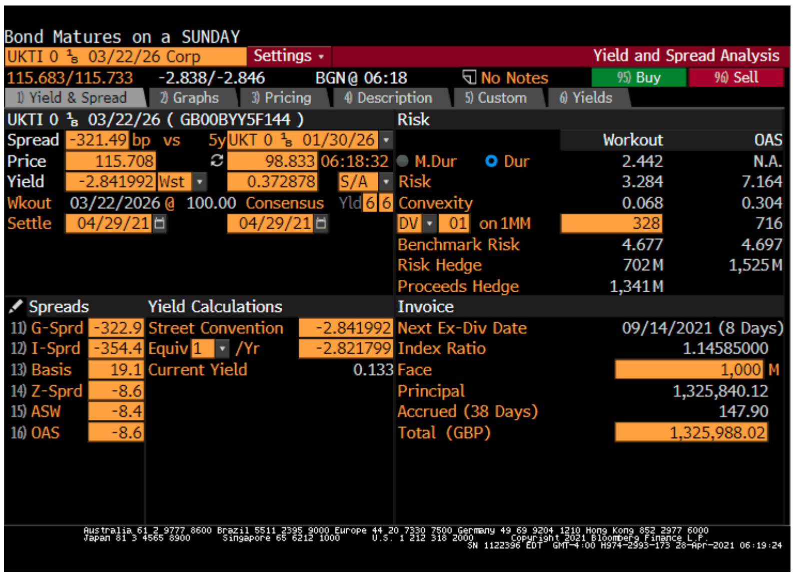

Figure 1: "Yield and Spread Analysis" screen from a Bloomberg terminal, illustrating the difference in quotes and settlement prices. The quoted "Price" (115.708) corresponds to a face value of 100 pounds, while the total cost (1,325,988.02) corresponds to a face value of one million pounds. The total cost includes the inflation adjustment ("Index Ratio 1.14585") and the accrued interest ("Accrued (38 Days): 147.90 pounds").

5 Stochastic modeling of the risk factors

A full description of the optimization model (CDI) requires a specification of the probability measure P governing the future values of the relevant risk factors. It is essential to describe statistical connections between the assets and liabilities, as that allows for the construction of investment portfolios with payouts that accommodate the liability payments not only across time but also across different scenarios. Inflation has a direct influence on both pension liabilities and inflation-linked bonds. Furthermore, there are statistical connections between the longevity and macroeconomic risk factors such as GDP and average earnings; see [AP14, HG18]. For the liabilities and the asset classes described in Sections 3 and 4, the relevant risk factors are longevity, price inflation, money market rates and the stock index.

In the computational examples below, we employ the multivariate stochastic model presented in [AMAP21], which accounts for the statistical connections mentioned above and describes the main macroeconomic, financial and longevity risk factors affecting longevity-sensitive financial products. In particular, the model describes yearly returns on equities and bonds (money market, government, inflation-linked and corporate). The model also describes stochastic survival probabilities for cohorts of both genders with ages between 18 and 105. This high-dimensional space of random vectors is modeled using only six longevity risk factors. Each realization of these six stochastic processes can be used to construct yearly survival probability curves. Figure 2 plots historical survival rates from 1922 to 2016. A more detailed description of the longevity side of the model can be found in [AP11].

The employed stochastic model is also able to incorporate short-term forecasts and long-term views of an investor; see [AMAP21, Section 4.2]. In this study, we calibrate the model to the long-term median values given Table 1. The values are based on the "Long-term economic determinants" published by the Office for Budget Responsibility (OBR) in [fBR21] as part of a report on its outlook for the UK economy in the coming years. In addition to the long-term views in Table 1, we specify the future median values of the money market rates according to the forward rates extracted from the ask-prices of the zero-coupon bonds. Denoting the ask-price of the zero coupon bond with maturity t by P_t, the forward rate F_t over $[t, t+1]$ is defined by

$$\frac{P_{t+1}}{P_t} = e^{F_t} \iff F_t = \Delta \ln P_{t+1}.$$

Market quotes for zero-coupon bonds and corresponding forward rates are illustrated in Figure 3. Of course, if one has different views concerning the future money market rates, one can use them instead of the forward rates. The forward rates can be thought of as market neutral views that seem appropriate when no extra information is available.

The money market is described by a single rate r so to describe the illiquidity of lending and borrowing, we add a margin on both sides so that the lending rate is $r_t^l = r_t - \delta$ and the borrowing rate $r_t^l = r_t + \delta$. Varying the illiquidity parameter δ will allow us to control the activity of the money market trading and control the time-matching of the liability cash-flows; see Section 7 and 8 below.

Simulated scenarios of the risk factors relevant to the (CDI) problem are illustrated in Figure 4. The same risk factors are used in the simulation of the pension liabilities illustrated in Figure 5. The numbers corresponds to a cohort of 1000 females with 65 years of age and an initial pension benefit of 1 GBP.

Long-term views	
Year	2070-71
Nominal GDP growth	3.9%
CPI inflation	2.0%
Average earnings growth	3.8%
Money market rates	4.1%
Gilt rates	4.1%
Stock index growth	6.0%

Table 1: Long-term views used in the stochastic modeling of the risk factors, based on the "Long-term economic determinants" published by the Office for Budget Responsibility (OBR) in the report "Economic and fiscal outlook - March 2021". The median value for growth of the stock index was chosen by the authors.

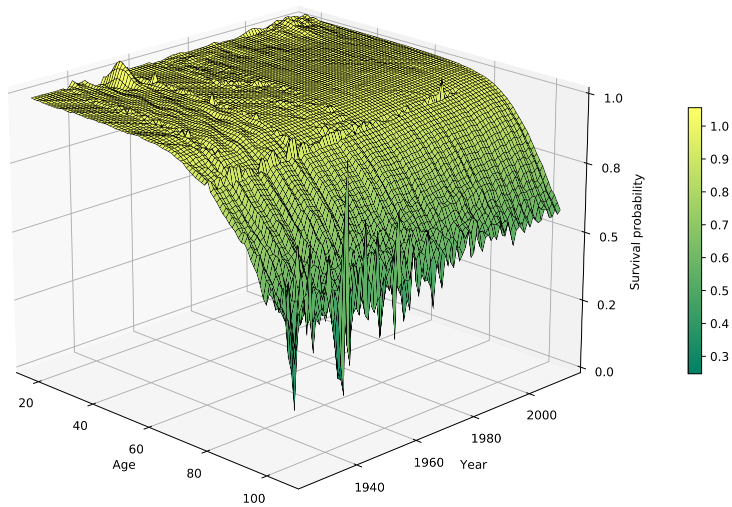

Figure 2: Yearly survival rates for females in the UK. Based on data from the Human Mortality Database [Uni].

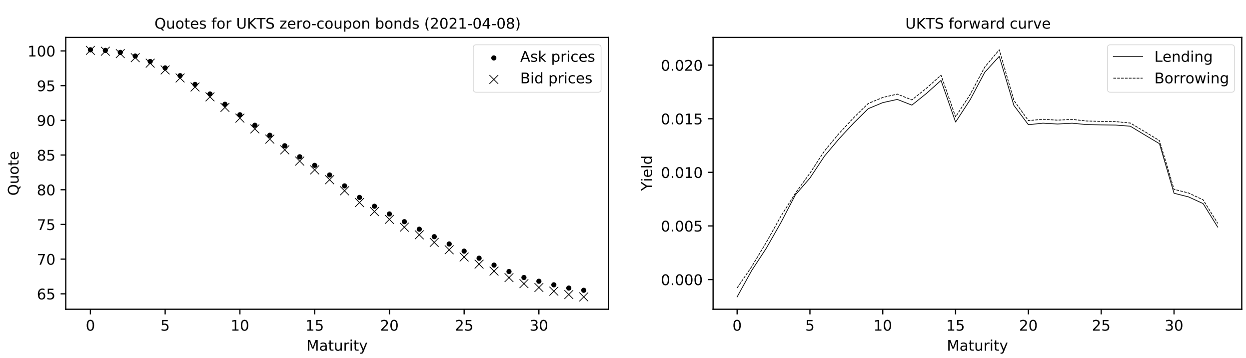

Figure 3: The plot on the left shows the bid- and ask-prices observed for a set of gilt coupon strips on 08/Apr/2021. The corresponding forward rates, used to calibrate the median values of the money market rates are shown on the right.

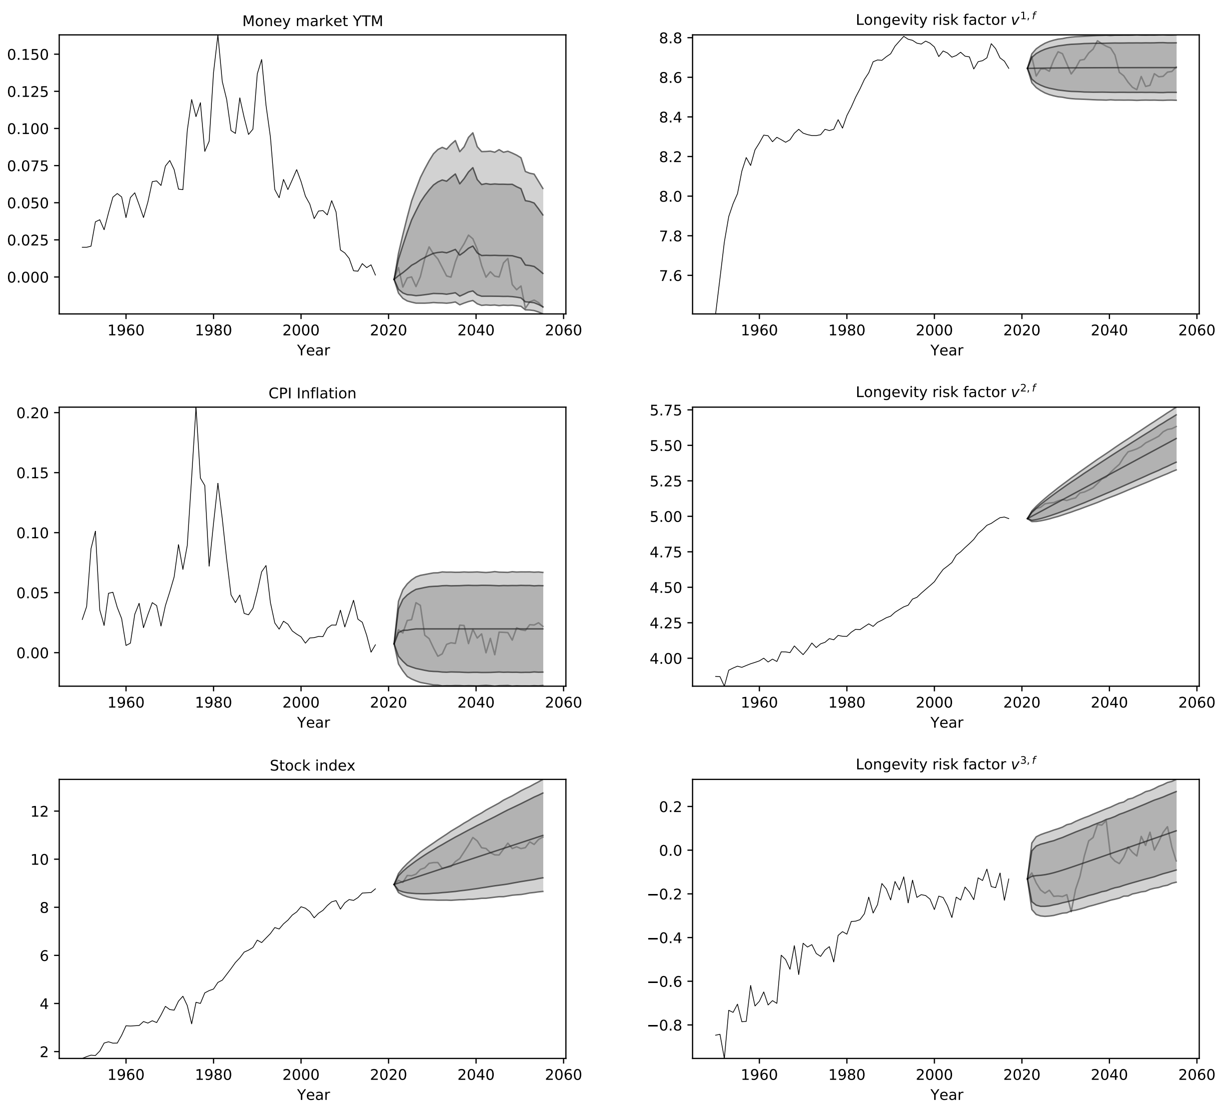

Figure 4: Historical and simulated values for the risk factors used in the computation of CDI portfolios. The 95% and 99% confidence bands obtained from a set of 100k simulated scenarios are illustrated in each plot along with a single simulated scenario. The rightmost risk factors correspond to longevity scenarios for females in the UK. Each realization of the three longevity risk factors reconstructs a survival probability surface such as the one illustrated in Figure 2.

Figure 5: Pension payments to a cohort of 1000 65 year-old females that receive an initial benefit of 1 GBP. The 95% and 99% confidence bands are plotted along with a single scenario. Benefits are adjusted yearly, following the rules of the University Superannuation Scheme.

6 Numerical solution

We will study problem (CDI) numerically in the case of sublinear money market rates (2) and the risk measure $\mathcal{R}$ defined as the entropic risk of the real terminal wealth, i.e.

$$\mathcal{R}(x) = \frac{1}{\rho} \ln E e^{-\rho \frac{I_0}{I_T} x},$$

where I is the consumer price index. As a first step, we approximate problem (CDI) by the finite-dimensional problem obtained when the probability measure P, described in Section 5, is approximated by a finitely supported measure of the form

$$P^N := \sum_{i=1}^{N} p^i \delta_{\xi^i}, \tag{5}$$

where $(\xi^i, p^i)_{i=1}^N$ is a finite collection of scenarios and associated probabilities. The problem then becomes:

$$\begin{array}{lll} \text{minimize} & x_0 + \sum_{k \in K} p_k(z_k) \quad \text{over} \quad z \in \mathbb{R}^K, x \in \mathcal{N}_N & \\ \text{subject to} & x_t^i \le x_{t-1}^i + R_{t-1}^i(x_{t-1}^i) + \sum_{k \in K} z_k c_{k,t}^i - c_t^i, \quad t = 1, \dots, T, \ i = 1, \dots, N & \text{(CDId)} \\ & \frac{1}{\rho} \ln \sum_{i=1}^{N} p^i e^{-\rho \frac{I_0}{I_T} x_T^i} \le 0, & \end{array}$$

where $\mathcal{N}_N$ denotes the finite dimensional space of N paths of the money market investments.

There are many different techniques for approximating the meaure P by a finitely supported measure of the form (5). The simplest option is Monte Carlo where one takes a random sample of N paths of the stochastic process and sets $p_i = 1/N$ for all $i = 1, \ldots, N$. Other options include various quasi-Monte Carlo methods such as Sobol sequences or lattice methods combined with the method of inversion. We will use "antithetic sampling" which constructs $N/2$ scenarios by randomly sampling from P and then obtains another set of $N/2$ scenarios by reflecting the scenarios. Antithetic sampling tends to reduce the variance of the expectation estimate when the integrand is monotonic with respect to the underlying random variables; see e.g. [Gla13].

Problem (CDId) is finite dimensional but it is given in terms of the nondifferentiable functions

$$p_k(z_k) = \begin{cases} p_k^a z_k & \text{if } z_k \geq 0, \\ p_k^b z_k & \text{if } z_k \leq 0, \end{cases}$$

and

$$R_t^i(x) = \begin{cases} r_t^{l,i} x & \text{if } x \geq 0, \\ r_t^{b,i} x & \text{if } x \leq 0, \end{cases}$$

where $r_t^{l,i}$ and $r_t^{b,i}$ are the values of the lending and borrowing rates, respectively, at time t in scenario i. We write the problem in standard form by making the substitution

$$z_k = z_k^+ - z_k^-,$$

and

$$x_t^i = x_t^{+,i} - x_t^{-,i}$$

where z_k^+, z_k^+, $x_t^{+,i}$ and $x_t^{-,i}$ are constrained to be nonnegative. This doubles the number of decision variables but the resulting problem is a convex optimization problem with a smooth objective and constraints. In fact, all the involved functions are either linear or a logarithm of the sum of exponentials. Such problems can be solved quite efficiently by modern interior point solvers. The numerical results presented below were obtained with the conic interior point solver of Mosek [ApS21]. The problem was formulated and communicated to Mosek using Python [VRD09] and CVXPY [DB16, AVDB18]. As an alternative to the splitting of variables and using of interior point methods, one could explore the use of the techniques developed in [BH05] in the present setting.

When using a small sample size N, one expects the optimized buy-and-hold portfolio in (CDId) to be somewhat infeasible with respect to the risk measure constraint in the original problem (CDI). We perform a simple numerical experiment to analyze the approximation error. For a given sample size N, we solve the optimization problem 50 times using independent random samples. For each of the 50 instances, we record the optimum value of (CDId) and evaluate the entropic risk measure of the corresponding buy-and-hold portfolio with an independent sample of 256k scenarios. The out-of-sample evaluation is computed by rolling the money market position over time along every scenario using the budget constraint in (CDId) for the given buy-and-hold portfolio.

We repeat the experiment with increasing sample sizes N (powers of 2) in order to study how the sample size affects the accuracy. The results are summarized in Table 2. As expected, optimized portfolios are infeasible in the sense that the out-of-sample evaluation of the risk measure is slightly positive. The infeasibility is reduced as the sample size increases. Similarly, the in-sample optimum values increase with N as more scenarios are added in the risk measure constraint of (CDId). With 32,768 scenarios, the numbers seem to have stabilized and the value of the risk measure seems acceptable given that it has units in cash and its value is about 10,000th of the initial wealth.

Table 3 gives the computation times for the above experiment on an AMD Ryzen Threadripper 1950x desktop with 128 GB of RAM. The out-of-sample simulations were implemented using PyCUDA [KPL$^+$12] and CUDA [NVI21], which reduced the computation times significantly. The numerical results of Sections 8 and 9 were obtained with 32,768 scenarios in the optimization of the buy-and-hold portfolios in CDId and 256k out-of-sample scenarios in the analysis of the optimized portfolios.

	Valuation (in-sample)				Risk measure (out-of-sample)			
Scenarios	Min	Max	Mean	SD	Min	Max	Mean	SD
1024	27.6438	28.3834	28.0708	0.1824	2.06e-04	3.09e-03	8.00e-04	5.41e-04
2048	27.9261	28.4241	28.1485	0.1016	8.52e-05	1.09e-03	4.99e-04	2.12e-04
4096	28.0152	28.3942	28.2406	0.0774	2.05e-04	1.70e-03	4.15e-04	2.47e-04
8192	28.1452	28.4361	28.2429	0.0494	1.59e-04	1.27e-03	3.83e-04	2.18e-04
16384	28.1879	28.3604	28.2767	0.0338	1.54e-04	1.95e-03	3.43e-04	2.69e-04
32768	28.2084	28.3436	28.2797	0.0352	1.75e-04	8.38e-04	2.96e-04	1.02e-04

Table 2: Monte Carlo approximation errors

	Optimization (in-sample)				Simulation (out-of-sample)			
Scenarios	Min	Max	Mean	SD	Min	Max	Mean	SD
1024	7.81	10.13	8.84	0.42	0.49	0.66	0.55	0.03
2048	18.20	28.86	20.27	1.50	0.53	0.87	0.60	0.07
4096	48.30	79.39	58.99	5.21	0.54	0.83	0.61	0.06
8192	107.69	203.25	180.38	21.18	0.51	0.68	0.60	0.04
16384	251.75	547.70	463.09	81.98	0.52	0.78	0.60	0.05
32768	540.16	1211.67	819.64	220.12	0.55	0.71	0.62	0.04

Table 3: Analysis of computational time

7 CDI based on expectations

We start numerical analysis of problem (CDI) by considering a completely deterministic model where all cash-flows and investment returns are assumed known at time $t = 0$. We will assume that all risk factors follow their median values. This corresponds to current industry practices where CDI analysis is often done against a single scenario. Deterministic models are commonly used also in actuarial liability valuations which are based on discounting a single forecast of the future payments using deterministic discount factors.

In the first instance, we only use zero-coupon bonds as the hedging instruments. In the deterministic setting, the liability cash-flows c can then be perfectly matched by buying c_t units of the zero-coupon bond with maturity t. We confirm this numerically by solving problem (CDI) with a single scenario that follows the median forecast. Figure 6 plots the yearly pension payments together with the cash-flows of the hedging portfolio. Pension payments are illustrated with the solid line and the portfolio payouts by the vertical bars positioned at payment dates. Since the portfolio is composed of zero-coupon bonds, the height of each bar corresponds to the amount of capital invested in each bond. The money market position over time (lending/borrowing) is represented by the dotted line. It should be noted that the optimal strategy involves zero trading in the money market. This is because the money market yields were calibrated to the quotes of the zero-coupon bonds as described in Section 5 so there is no incentive to invest in the money market. Moreover, the strictly positive spread between the lending and borrowing rates makes made zero-coupon bonds strictly better investment than the money market. In this example the margin between the lending and borrowing rates is $\delta = 0.01$ basis points (BPS).

Next, maintaining the margin of $\delta = 0.01$ BPS, we add coupon paying bonds to the problem. In this setting, the optimal strategy invests the whole initial wealth into a single gilt and finances most of the liability payments by trading in the money market, as shown in the top-left plot of Figure 7. The poor diversification is explained the lack of the risk-return trade-off in the deterministic model and the lower cost of gilts compared to gilt strips, used here as zero-coupon bonds. While the risk-return trade-off can only be adequately studied in the stochastic case, topic of Section 8, we can study different market conditions by changing the margin δ. When increased to $\delta = 50$ BPS, for example, the trading activity in the money market decreases, and the diversification in the optimal portfolio improves, as it now contains a variety of gilts and zero-coupon bonds, as represented in the middle-left plot of Figure 7. Perfect matching of the cash-flows is obtained when we increase the margin to $\delta = 200$ BPS. In this case, it is optimal to cover all payments by the statically held bonds and avoid borrowing from the money market.

We continue the experiment by adding stocks and inflation-linked bonds to the set of hedging instruments. The results should be interpreted with caution, however, as stocks and inflation-linked bonds are characterized by the uncertainty of their cash-flows which cannot be described by a deterministic model. In the deterministic case, one cannot see the risks of stock investments nor the hedging properties of inflation-linked bonds. Essentially, the lack of risk in the deterministic model makes stocks seem under priced while inflation-linked bonds seem overpriced. Nevertheless, we proceed with the experiment as this seems to be a common practice (more or less quantitatively) among investment advisors.

We start with a small margin δ that we later increase. The optimal hedging strategies are illustrated in the right column of Figure 7. As expected, the stock investments dominate in the deterministic model when the margin is small. The deterministic model cannot account for the downside risk, so the high level of return makes stocks seem like the best investment. With the small margin $\delta = 0.01$ BPS (top-right plot), the optimal portfolio invests everything in stocks that are liquidated at the end of the last period. All intermediate pension payments are financed by borrowing from the money market. When we increase the borrowing margin to $\delta = 500$ BPS, the diversification is increased in order to reduce borrowing from the money market (middle-right plot). To perfectly match the cash-flows, we increase the margin to $\delta = 1000$ BPS. In this extreme case, borrowing costs become prohibitive and the annual pension payments are covered by yearly liquidation of the equity investments (bottom-right plot).

Despite being overly simplistic, the deterministic model allows us to illustrate some important features of problem (CDI). For convenience, results from the experiments are summarized in Table 4. When the borrowing margin increases, the diversification and cash-flow matching improve reducing the money market investments. Essentially, the illiquidity of the money market drives the matching of the cash-flows. With increased diversification, comes an increase in cost. When capital is allocated in a larger number of hedging instruments, instead of the most profitable one, the cost to cover liabilities increases. Costs decrease, however, with additional hedging instruments. From an optimization perspective, it is clear that the inclusion of new instruments can only improve optimal solutions.

We end this section by analyzing the performance of the deterministically optimized portfolios in the stochastic model. To this end, we use the model described in Section 5 to simulate 256k scenarios for the pension payments and asset returns. For a given portfolio, the terminal fund wealth is easily computed by following the budget constraints in (CDId). Figure 8 gives kernel density plots for the terminal wealth distributions of each deterministically optimized portfolio. It is clear that portfolios that provide perfect cashflow matching in a deterministic model can be dangerously risky in a stochastic (real) world.

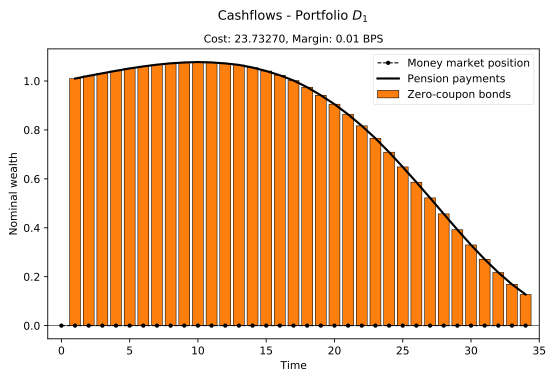

Figure 6: Deterministic pension payments along with cash-flows of an optimal hedging portfolio corresponding to a margin of $\delta = 0.01$ BPS. The perfect matching of cash-flows results in zero positions in the money market.

Portfolio	Asset classes	Margin	Valuation	Portfolio composition			
				ZCB	Gilts	ILB	Stocks
D_1	ZCB	0.01	23.73270	100%	–	–	–
D_2	ZCB, Gilts	0.01	23.34219	0%	100%	–	–
D_3	ZCB, Gilts	50	23.52629	7%	93%	–	–
D_4	ZCB, Gilts	200	23.55467	26%	74%	–	–
D_5	All	0.01	4.69419	0%	0%	0%	100%
D_6	All	500	13.16744	0%	0%	0%	100%
D_7	All	1000	13.24309	0%	0%	0%	100%

Table 4: Results obtained with the deterministic model in Section 7, showing that when the borrowing margin δ increases, the diversification in the hedging portfolio improves. With increased diversification, also comes an increase in cost (e.g. rows 2, 3 and 4; or 5, 6 and 7). In addition, costs decrease with additional hedging instruments (e.g. rows 1, 2 and 5).

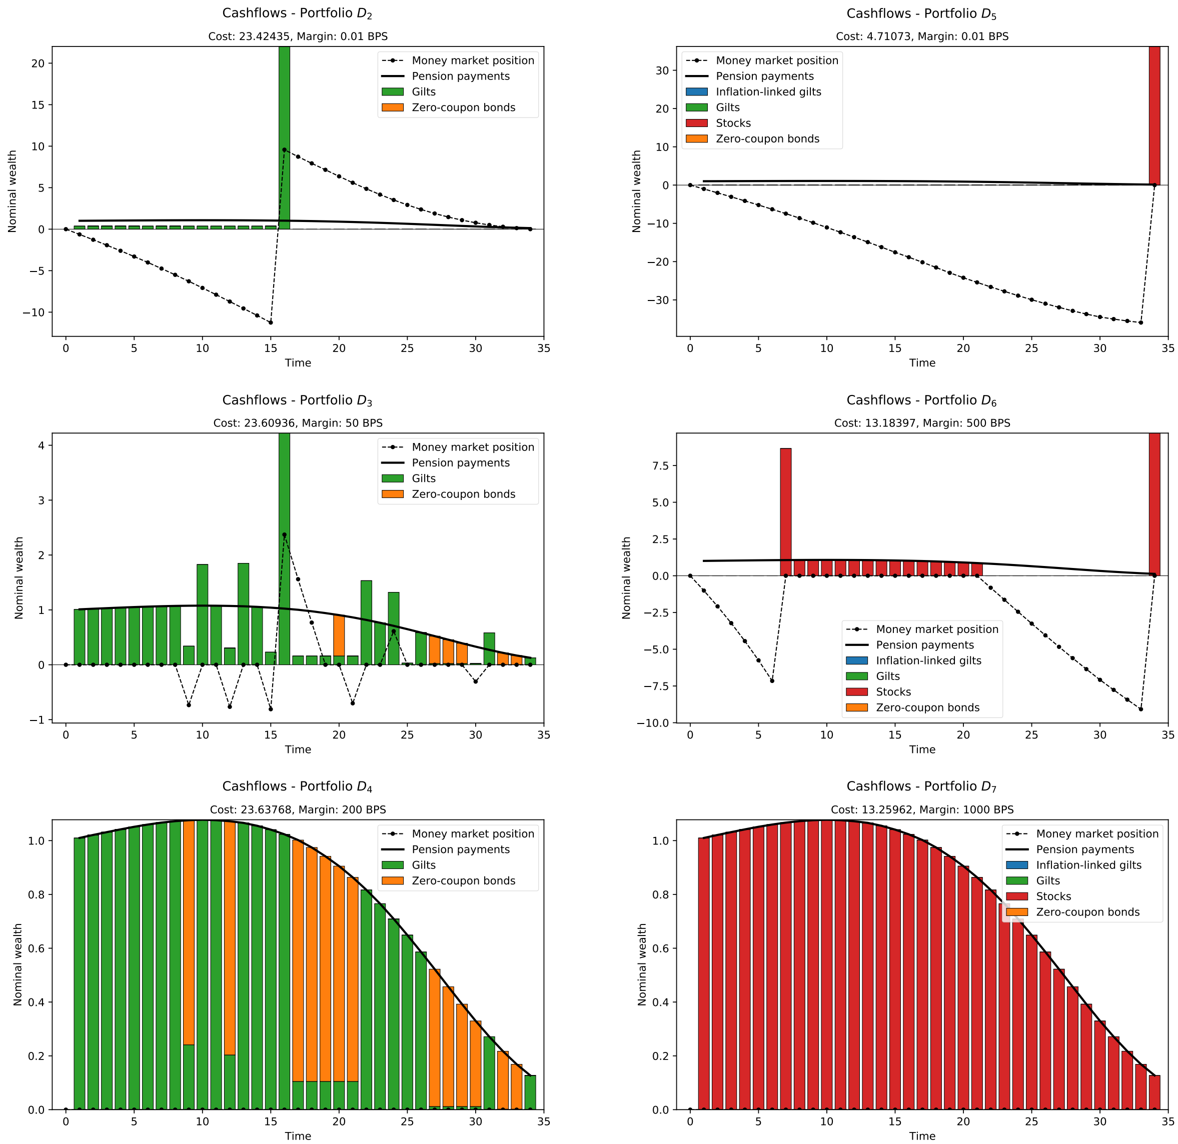

Figure 7: Illustration of the optimal portfolios in the deterministic model described in Section 7. Deterministic pension payments are illustrated along with aggregate yearly cash-flows of hedging portfolios. Optimal hedging portfolio in the top plots are concentrated on a single instrument; a gilt maturing at $t = 16$ (left) and stocks held until the last period (right). With the increase in the money market margin δ, we see more diversification in the portfolios, an increase in cost and a decrease in the money market positions.

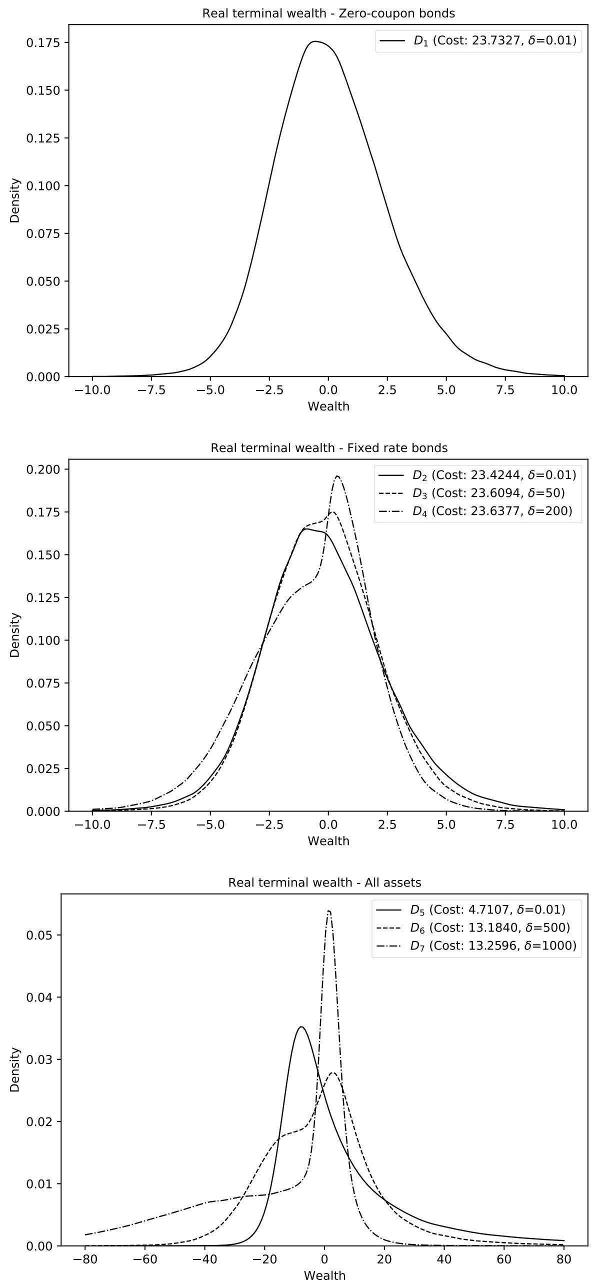

Figure 8: Distributions of the real terminal wealths of the deterministically optimized portfolios of Section 7 computed using 256k scenarios. The top plot is based on the portfolio of zero-coupon bonds (Figure 6). Middle plot on portfolios of fixed-rate instruments (left column in Figure 7), and bottom plot on portfolios with all asset classes (right column in Figure 7).

8 CDI under uncertainty

We now move to a fully stochastic version of (CDI), where the cashflows of both the liabilities and the hedging instruments are stochastic. This will allow us to quantify the risks and hedging potential of assets with cashflow uncertainties. Increasing the risk aversion will favour assets whose paoyouts hedge against the pension liabilities. A lower risk aversion, on the other hand, will favour assets with higher average payouts even if they are uncorrelated with the liability payments. In the numerical experiments below, we will see this in economically sensible allocations in equities and inflation-linked bonds that were not captured in the deterministically optimized portfolios.

Our first results, with margin $\delta = 100$ BPS and risk aversion $\rho = 5$, are illustrated in Figure 9. In the stochastic case, we cannot present the results using the simple cash-flow plots used in the deterministic case in Section 7. Instead, we create six different plots to illustrate the portfolios and their hedging properties in the stochastic case. In the top-left corner of the figure, we have a *wealth allocation plot* (not to be confused with the cash-flow plots of Section 7) where the vertical bars at time t represent the aggregated amount of capital invested in hedging instruments maturing at that time. In this particular example, we have an optimal portfolio S_1 that invests, approximately, 1k GBP in gilts maturing at times $t = 1, \ldots, 7, 13, 14$; over 4k GBP in inflation-linked gilts maturing at times $t = 10, 16, 20$ and other smaller positions after $t = 20$. The optimal portfolio also holds a position of approximately 0.5k GBP in equities that is only liquidated at the final period. The actual positions in the portfolio S_1 can be found in Table 7.

On the middle-left plot of Figure 9, we have a *portfolio composition plot*, in which stacked bands are used to illustrate the amount of capital invested in each asset class over time. In such plots, a drop occurs in one of the bands when one of the hedging instruments matures. Flat regions indicate positions being held fixed. Referring to the wealth allocation plot in the figure, for example, one can see a drop associated with the inflation linked gilt maturing at time $t = 10$. One can also see the flat region corresponding to the position in stocks. It spans the entire simulation period, as the equities are liquidated only at the end.

In the top-right corner of Figure 9, we have an *investment income plot*, which illustrates the aggregated cash-flows produced by the hedging strategy. The figures plot median values and 95% and 99% quantile bands of yearly cash-flows. Deterministic cash-flows, such as the ones from gilts and gilt strips, are represented as solid line segments. In this example, the optimal allocation has gilt positions maturing at times $t = 1, \ldots, 7$. The three spikes times at times $t = 10, 16$ and 20 correspond to the large positions in inflation-linked bonds. The cash-flows corresponding to the position in stocks are only found when they are liquidated at time $t = 34$. The width of the confidence bands show how uncertain stock returns can be after a few decades. Pension payments are illustrated in the plot as the dark shape in the background; see Figure 5 for a comparison.

The three remaining plots in Figure 9 illustrate the development of the net cash-flows and the money market position as well as the distribution of the terminal wealth. Again, we give the median values and the 95% and 99% quantile bands. The net cash-flows are obtained simply by taking the difference of the investment income and the pension outflow so a surplus increase the money market position and vice-versa. The net cash-flow is rolled over time through the money market account. Thus, the terminal wealth is simply the terminal position in the money market account.

To understand a hedging strategy, we refer to the wealth allocation, investment income and money market position plots. The strategy illustrated in Figure 9, for example, relies primarily on gilts to cover pension payments in the first years. More specifically, until time $t = 7$. In this period, a surplus of capital is built, as indicated by the mostly positive values of the money market position. The strategy then resorts to borrowing from the money market, as the income from other investments is low in the next two years. Since no positions are maturing in that period, the income is from to coupon-payments of bonds maturing later. The two-year deficit in the money market position is then covered by the inflation-linked bond maturing at time $t = 10$. This creates a surplus of cash which is then increased by gilts maturing at times $t = 13$ and 14, after two years of low investment income. Next, at time $t = 15$, we find quantile bands that span both negative and positive positions in the money market position, indicating a good chance of covering payments without resorting to borrowing from the money market. Possible deficits are then be covered in the following year by the inflation-linked bond position maturing at time $t = 16$. The next inflows of capital then create enough surplus to cover pension liabilities in the majority of cases until the time $t = 30$. Notice that the median value of the money market position becomes negative after that time. Finally, the position in equities is liquidated at time $t = 34$.

Inspecting the terminal wealth plot for S_1, in Figure 9, we notice a distribution of outcomes that is skewed toward positive values. We compare this distribution to those of the perfectly matching portfolios obtained with the deterministic model (D_1, D_4 and D_7) in Figure 10. Clearly, the portfolio S_1 provides a better hedging to the pension liabilities, as the other portfolios present distributions that are approximately symmetric. The risk measure values in Table 5 provide a more rigorous check. In the table, we notice that S_1 presents the smallest level of risk among the portfolios and the higher cost as well. The risk measure values also show that the perfectly matching portfolios do not correspond to optimal solutions to (CDI) in the stochastic case, as their risk measures are significantly larger than zero. Putting it simply, matching deterministic cash-flows does not work.

Portfolio	Valuation (cost)	Risk measure	Portfolio composition			
			Stocks	Gilts	ILB	ZCB
S_1	28.1592	1.99e-04	2%	42%	56%	–
D_1	23.7327	4.69e-01	–	–	–	100%
D_4	23.6377	6.35e-01	–	74%	–	26%
D_7	13.2596	2.13e+00	100%	–	–	–

Table 5: Comparison of terminal wealth distribution

8.1 Effect of the lending and borrowing margin

We now look into two variations of the experiment described above. Still with risk aversion $\rho = 5$, we first increase the margin to $\delta = 1000$ BPS. Then, we lower it to $\delta = 0$ BPS and reoptimize. In the first case, with the increased margin, we notice more diversification in the optimal portfolio S_2, displayed in Table 8 and illustrated in the wealth allocation plot of Figure 11. We also find an improved matching of cash-flows in the investment income plot. Inspecting the money market position we see that, with an increased margin, the optimal strategy tends to avoid borrowing from the money market. Significant borrowing from the money market only occurs at the times $t = 15$ and $t = 27$. After the latter, the optimal portfolio is concentrated in stock investments, which explains the wide quantile bands at the end of the money market position and of the investment income plots. It is interesting to notice that the optimal portfolio S_2 now includes zero-coupon bonds. With higher money market costs, zero-coupon bonds provide a cheaper alternative for extra liquidity. The distribution of the terminal wealth still has a positive skew but it is now less risky than the portfolio S_1, illustrated in Figure 9. With the improved matching of cash-flows and lower risk, we also observe an increase in cost for S_2.

The case with zero margin, is illustrated in Figure 12. As expected, for this optimal portfolio S_3, we see much more activity in the money market and a reduced cost, but we also see poorer matching of the cash-flows and a less diversified portfolio, displayed in Table 9. All the results obtained in this section are summarized in Table 6.

Looking at the portfolio composition with fixed risk aversion and varying margin, we find that reduction of liquidity leads to increased allocation zero-coupon and inflation linked bonds and reduction in gilts and equities. This is inline with the findings of [APW14] who investigated the effects of liquidity risk on optimal portfolio composition. Their liquidity risk model was quite different from ours but the they found a similar shift from riskier assets to safer ones when the liquidity risk is increased.

8.2 Effect of the risk aversion

We now fix the margin at $\delta = 100$ BPS and analyse changes the effects of the risk aversion ρ. First, we lower the risk aversion from $\rho = 5$ to $\rho = 1$ and then increase it to $\rho = 10$, reoptimizing after each change. With the lower risk aversion, illustrated in Figure 13, we notice that the optimal portfolio S_4 has a decreased proportion of capital allocated to inflation-linked bonds and increased in stocks and gilts. The latter, in this case, became the main asset class in the portfolio, accounting for 76% of the invested capital. The wealth allocation plot shows that the hedging strategy expects to capitalize on the long-term returns of the stocks, as those are held until the end. The strategy also builds a surplus of capital in the early years, until $t = 16$. By increasing the allocation in stocks and the trading in the money market, S_4 obtains the smallest cost among the portfolios in this section; see Table 5.

Increasing the risk aversion to $\rho = 10$, we find a portfolio S_5 which has shifted capital from equities and gilts to inflation-linked bonds. The hedging cost has increased by 20%, and we find

Portfolio	Margin	Risk aversion	Valuation (cost)	Risk measure	Portfolio composition			
					Stocks	Gilts	ILBs	ZCBs
S_1	100	5	28.1592	1.99e-04	2%	42%	56%	–
S_2	1000	5	29.7009	1.11e-04	2%	27%	60%	11%
S_3	0	5	27.8031	4.33e-04	3%	38%	59%	–
S_4	100	1	24.1384	2.06e-04	11%	76%	13%	–
S_5	100	10	29.0350	4.31e-04	1%	37%	61%	2%

Table 6: Summary of results for the stochastically optimized hedging strategies

a more more diversification and less trading in the money market. When we compare S_5 to the original portfolio S_1, however, changes are not so dramatic. We find a 3% increase in the cost of the hedging portfolio and an increase of 5% in the capital allocated to inflation-linked bonds. Portfolio allocations for portfolios S_4 and S_5 are displayed in Tables 10 and 11, respectively.

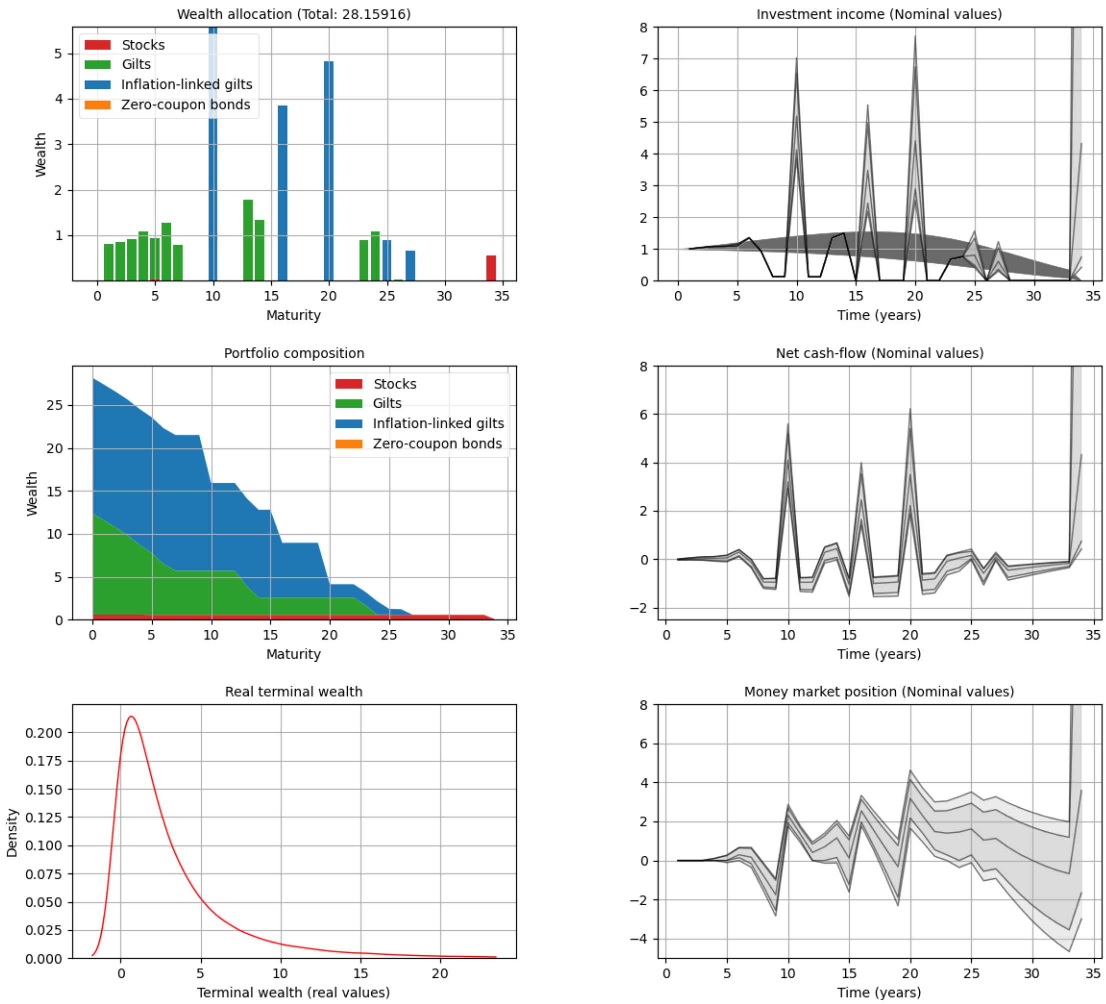

Figure 9: Optimal hedging strategy S_1 (risk aversion $\rho = 5$ and margin $\delta = 100$ BPS). Rightmost plots contain median values and 95% and 99% confidence bands. Leftmost plots illustrate the terminal wealth and the portfolio allocation. The detailed portfolio allocation can be found in Table 7.

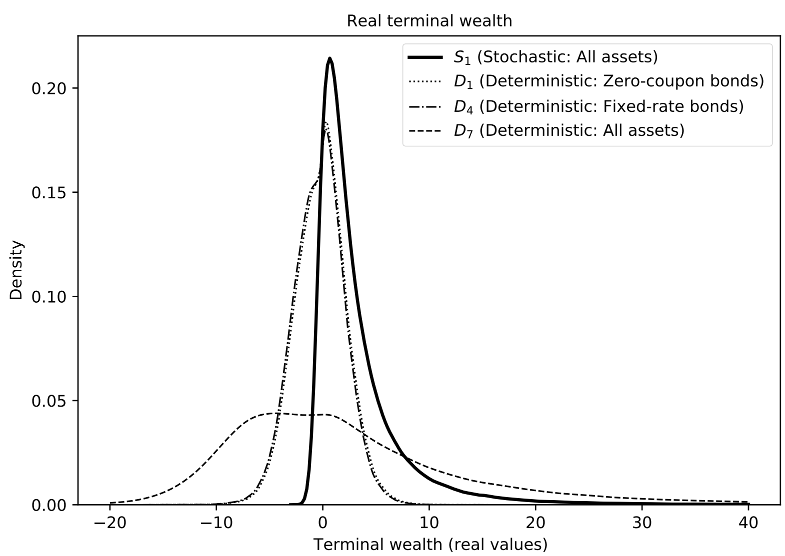

Figure 10: Comparison of terminal wealth distributions. See Table 5

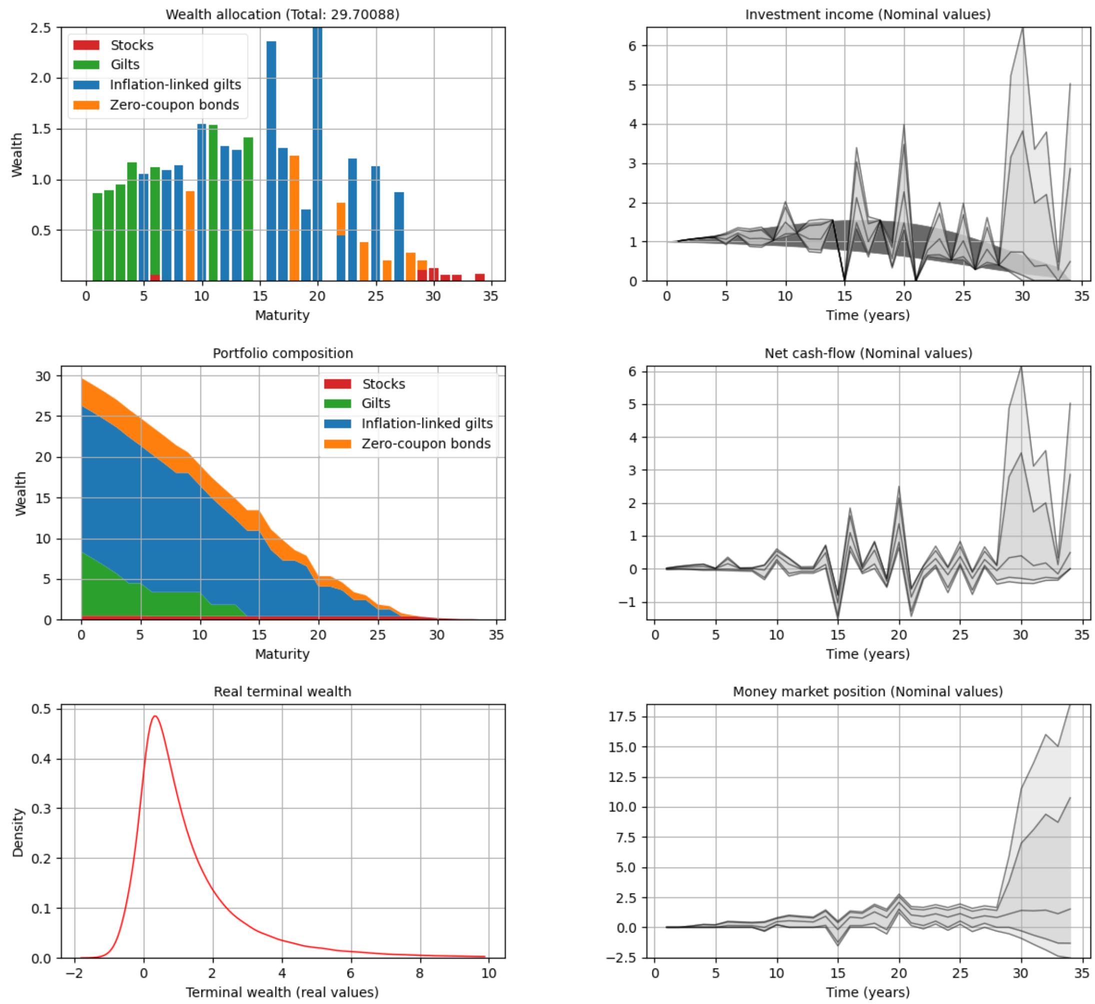

Figure 11: Optimal hedging strategy S_2 (risk aversion $\rho = 5$ and margin $\delta = 1000$ BPS). Rightmost plots contain median values and 95% and 99% confidence bands. Leftmost plots illustrate the terminal wealth and the portfolio allocation. The detailed portfolio allocation can be found in Table 8.

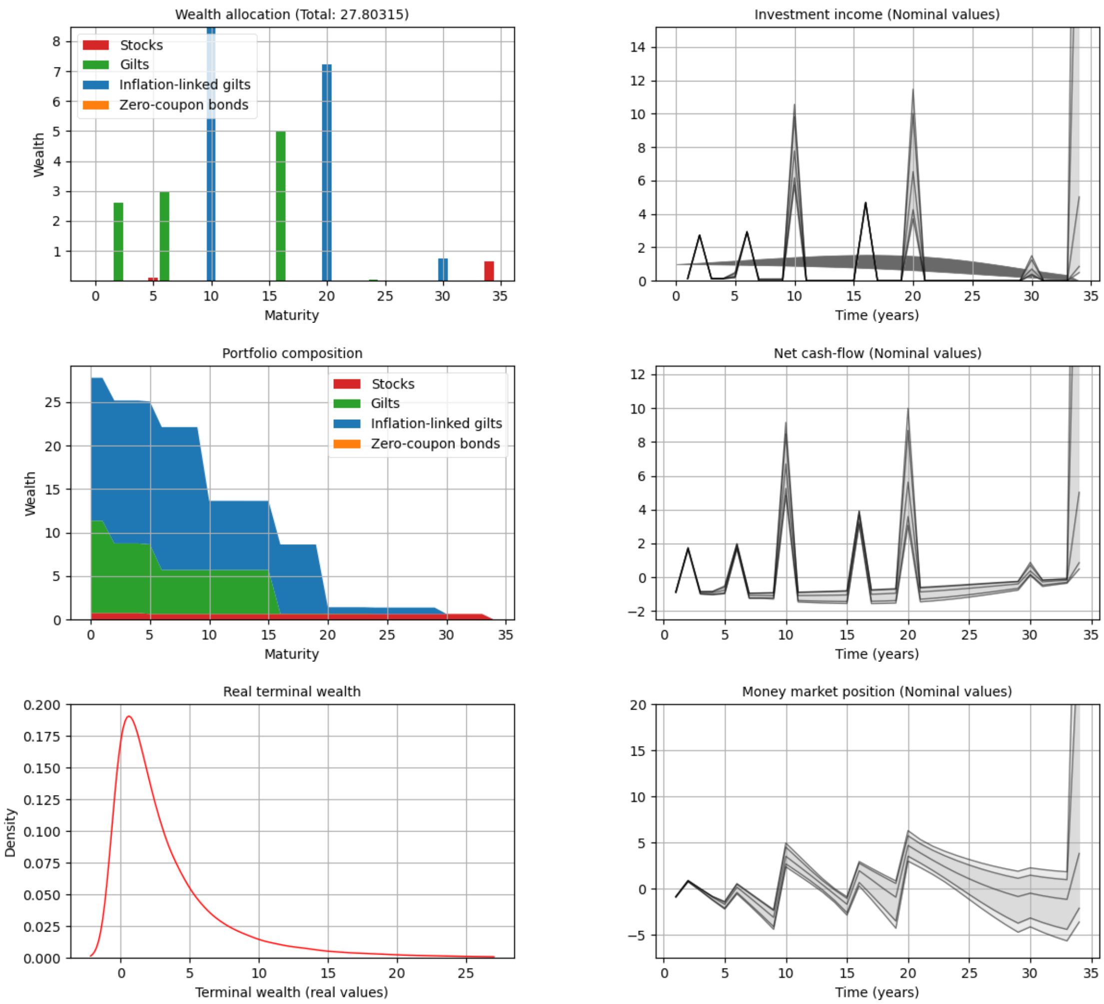

Figure 12: Optimal hedging strategy S_3 (risk aversion $\rho = 5$ and margin $\delta = 0$ BPS). Rightmost plots contain median values and 95% and 99% confidence bands. Leftmost plots illustrate the terminal wealth and the portfolio allocation. The detailed portfolio allocation can be found in Table 9.

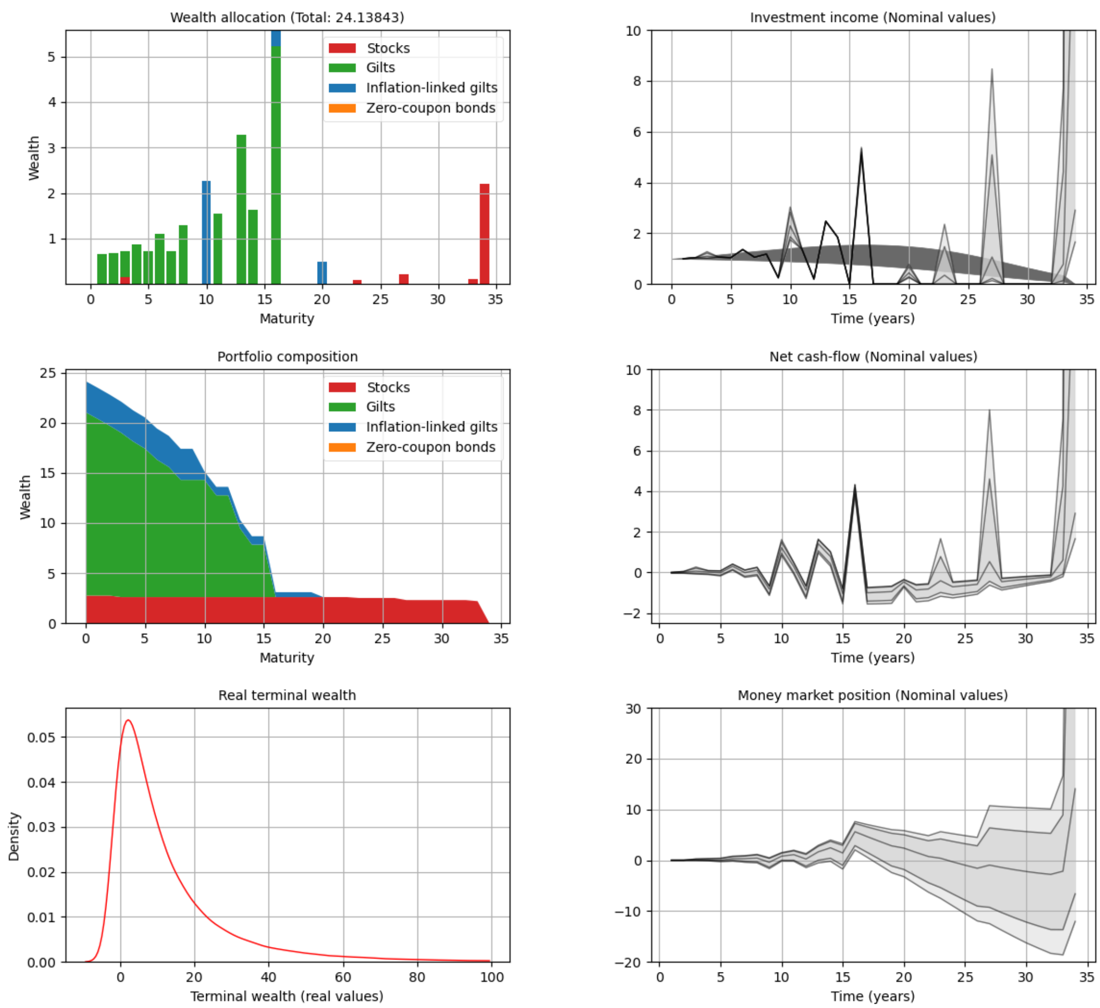

Figure 13: Optimal hedging strategy S_4 (risk aversion $\rho = 1$ and margin $\delta = 100$ BPS). Rightmost plots contain median values and 95% and 99% confidence bands. Leftmost plots illustrate the terminal wealth and the portfolio allocation. The detailed portfolio allocation can be found in Table 10.

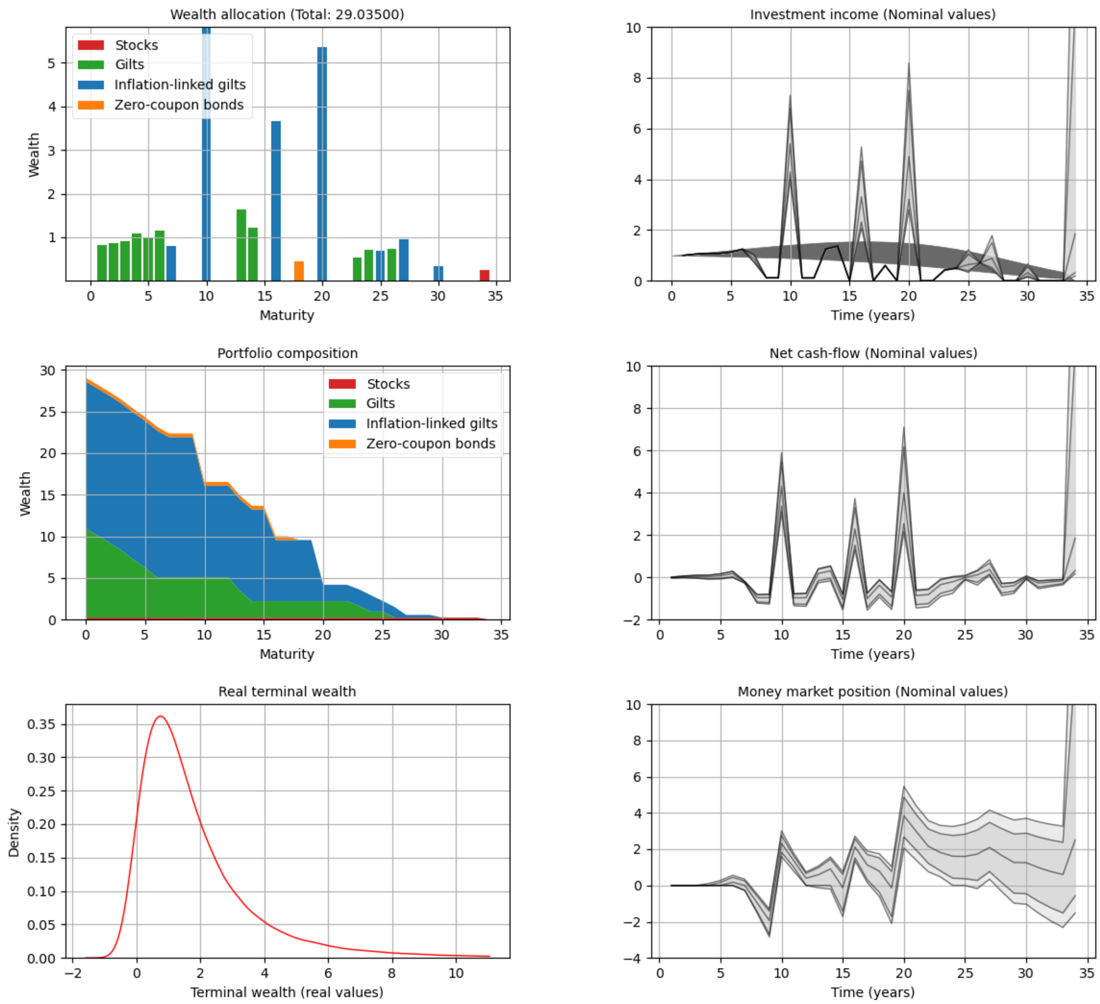

Figure 14: Optimal hedging strategy S_5 (risk aversion $\rho = 10$ and margin $\delta = 100$ BPS). Rightmost plots contain median values and 95% and 99% confidence bands. Leftmost plots illustrate the terminal wealth and the portfolio allocation. The detailed portfolio allocation can be found in Table 11.

9 Liability valuations

This section takes a closer look at the liability valuations as a function of the risk aversion ρ and the lending and borrowing margin δ. To that end, we find optimal strategies and valuations with the risk aversions $\rho = 0, 1, \ldots, 10$ and margins $\delta = 0, 100, \ldots, 1000$. The resulting valuations are illustrated in Figure 15 in the form of a surface plot. As expected, valuations increase when either the margin or the risk aversion increases. When the risk aversion is increased from zero, there is a sharp increase in the valuations but they seem to stabilize quickly. For $\rho \in [1, 10]$, the valuations lie between 23.28k GBP and 30.30k GBP.

Figure 16 illustrates the dependence of the optimal portfolios on the margin and the risk aversion. Each plot in the figure represent the optimal allocations as functions of the risk aversion for a given margin. For risk aversion close to zero, stocks dominate the allocations but, as soon as the risk aversion ρ is increased, the stock allocations decreasing sharply. When the risk aversion approaches 10, the allocations seem to stabilize with a small but nonzero investment in stocks. Gilts are added to the portfolios when the positions in stocks decrease. When the risk aversion is increased further, inflation-linked bonds begin to dominate. It seems that with low risk aversion, the hedging potential of the inflation-linked bonds does not seem worth the higher price. The situation is reversed when the risk aversion increases. Zero-coupon bonds are added to the hedging portfolios when the margin is high (bottom plots in Figure 16). Overall, the percentages associated with each asset class seem relatively stable for higher values of the risk aversion (e.g. $\rho \geq 5$).

The plots in Figure 17 show the distributions of real terminal wealth for fixed values of the margin and for increasing values of the risk aversion. As expected, the dispersion of the terminal wealth and the risk of negative outcomes decrease as the risk aversion increases. Interestingly, the same happens when the margin is increased. This phenomenon is more easily seen in Figure 18 which superimposes the distributions for fixed risk aversions. Hedging strategies optimized with a higher margin provide better cashflow matching thus reducing the reinvestment risk.

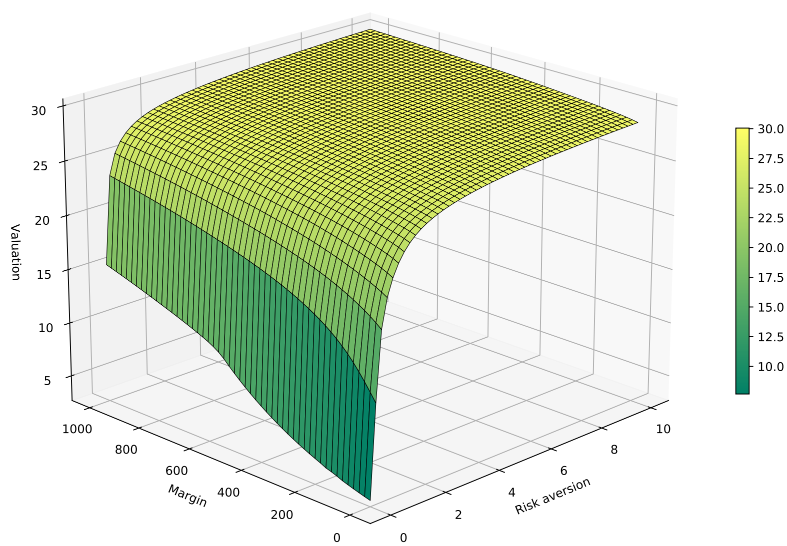

Figure 15: Liability valuations as a function of the risk aversion and the lending and borrowing margin

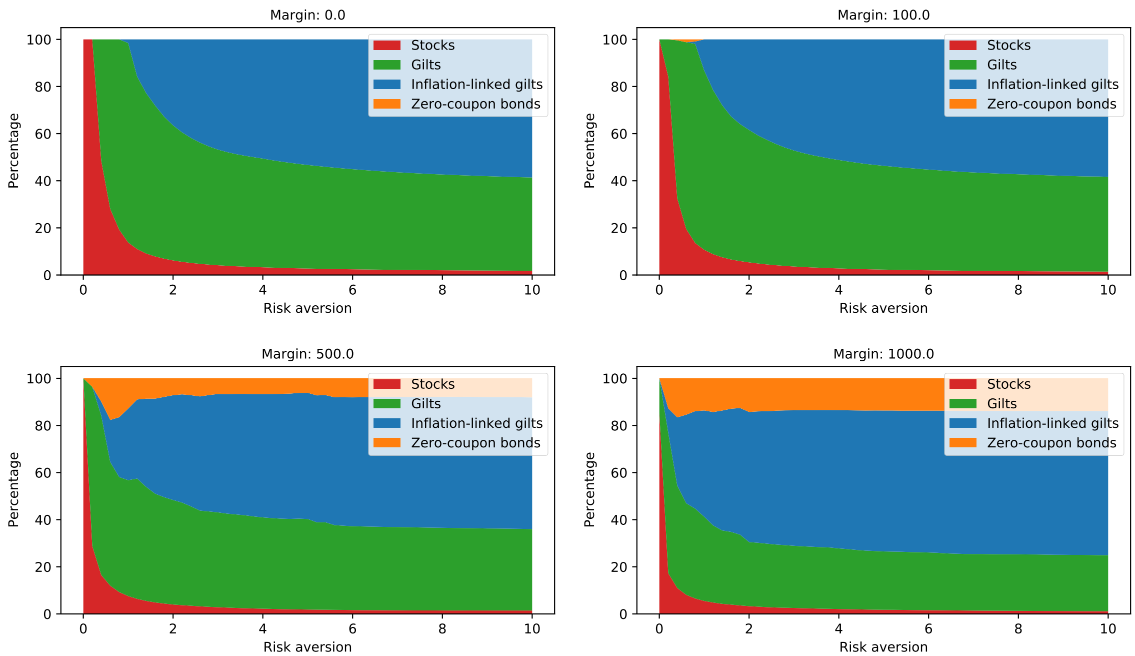

Figure 16: Portfolio allocations (in percentages) for different margins

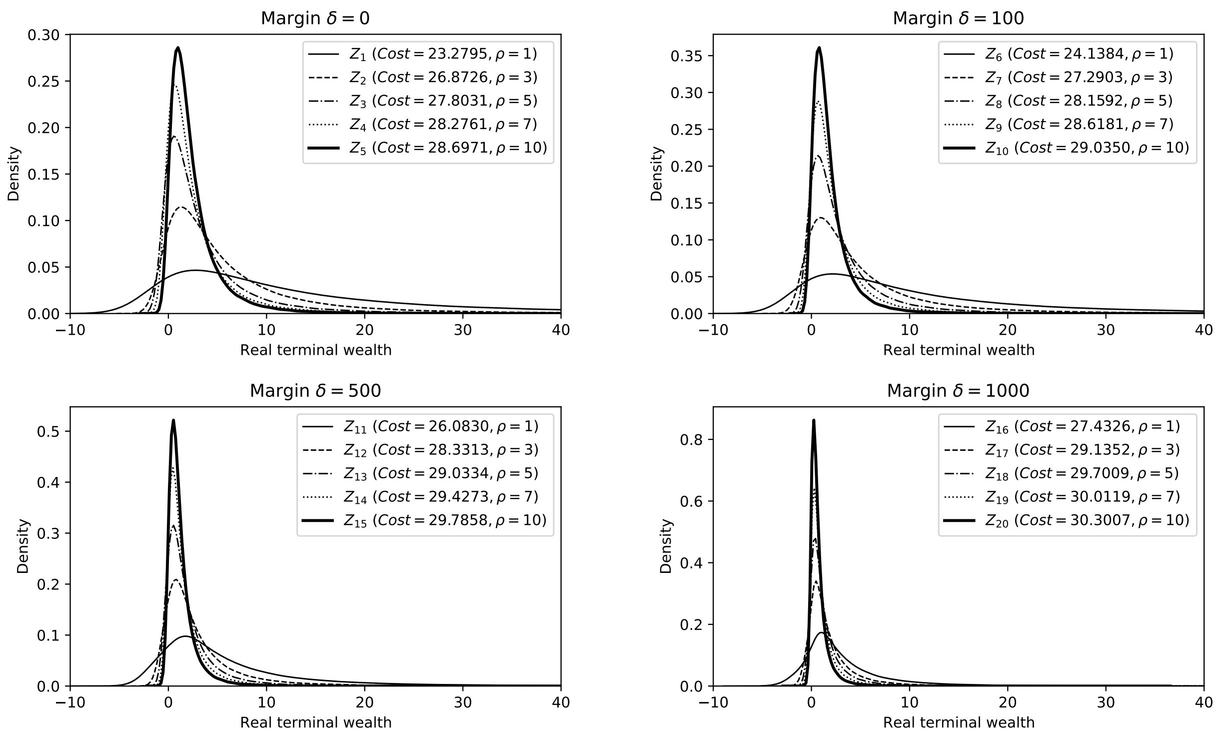

Figure 17: Terminal wealth distributions with varying margins and risk aversions

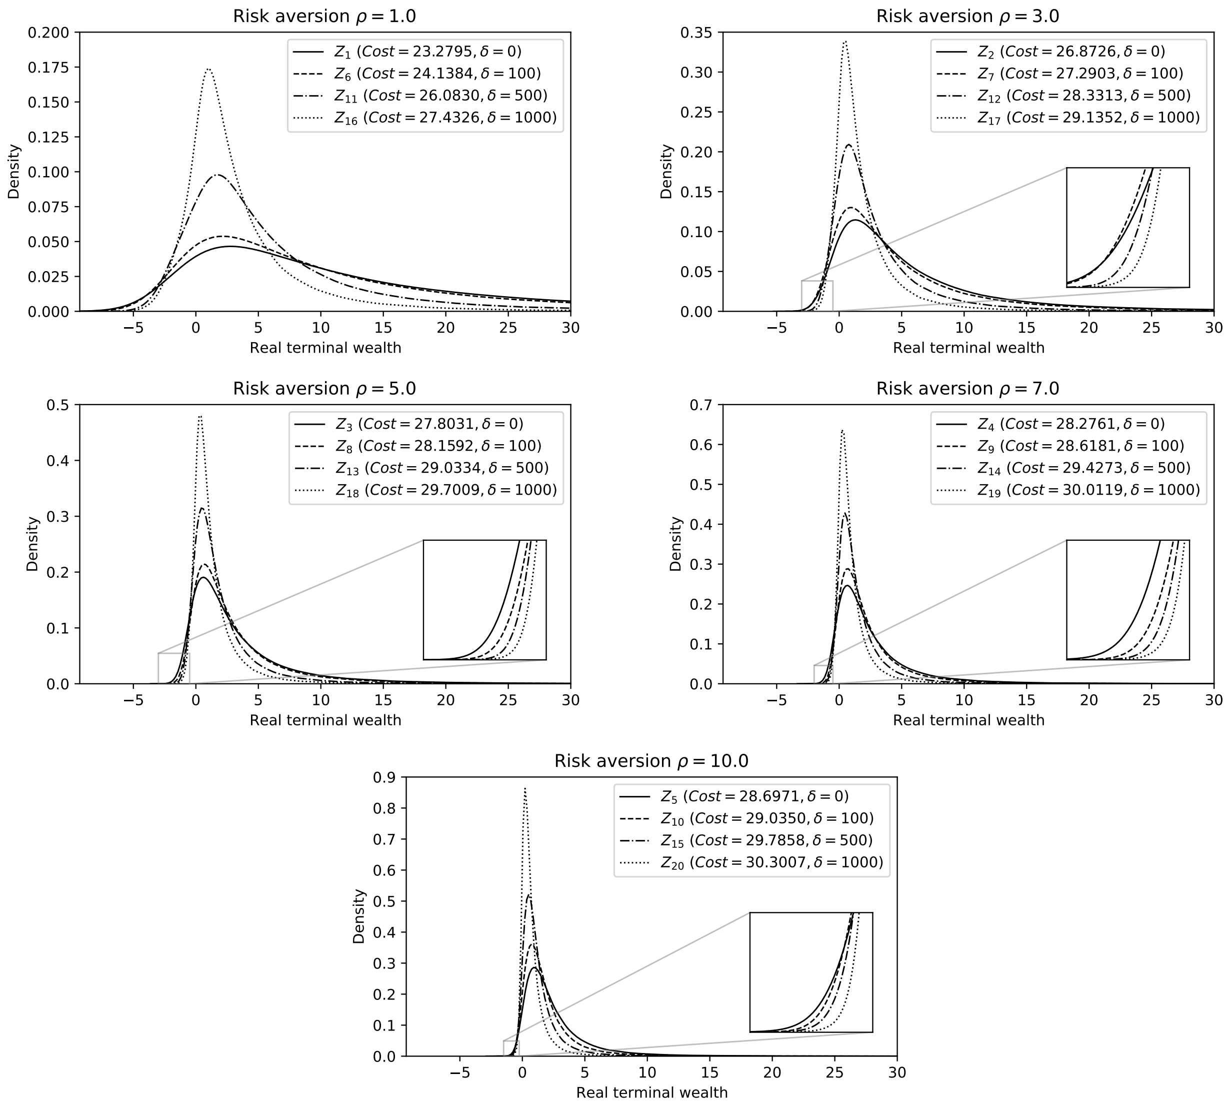

Figure 18: Terminal wealth distributions with varying margins and risk aversions

References

[AMAP21] Sergio Alvares Maffra, John Armstrong, and Teemu Pennanen. Stochastic modeling of assets and liabilities with mortality risk. *Scandinavian Actuarial Journal*, pages 1–31, 2021.

[AP11] H. Aro and T. Pennanen. A user-friendly approach to stochastic mortality modelling. *European Actuarial Journal*, 1(2):151–167, 2011.

[AP14] H. Aro and T. Pennanen. Stochastic modelling of mortality and financial markets. *Scandinavian Actuarial Journal*, 2014(6):483–509, 2014.

[ApS21] MOSEK ApS. *MOSEK Optimizer API for Python 9.2.46*, 2021.

[APW14] Andrew Ang, Dimitris Papanikolaou, and Mark M Westerfield. Portfolio choice with illiquid assets. *Management Science*, 60(11):2737–2761, 2014.

[AVDB18] Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.

[BH05] Michael J Best and Jaroslava Hlouskova. An algorithm for portfolio optimization with transaction costs. *Management Science*, 51(11):1676–1688, 2005.

[BSB05] G. V. Boyles, T. W. Secrest, and R. B. Burney. The pricing of bonds between coupon payments: From theory to market practice. *Journal of Economics and Finance Education*, 4(2):61, 2005.

[cdi19] A holistic study into cash flow driven investment. Available at `http://tiny.cc/jinvtz`, December 2019.

[DB16] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

[Exl17] J. Exley. Cashflow driven investment or clever distribution investing? Available at `http://tiny.cc/9mmvtz`, September 2017.

[fBR21] Office for Budget Responsibility. Economic and fiscal outlook – March 2021. Available at `http://tiny.cc/sdjytz`, March 2021.

[FS16] H. Föllmer and A. Schied. *Stochastic finance*. De Gruyter Graduate. De Gruyter, Berlin, 2016. An introduction in discrete time, Fourth revised and extended edition of [MR1925197].

[Gla13] Paul Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer Science & Business Media, 2013.

[GZ08] Alois Geyer and William T. Ziemba. The innovest Austrian pension fund financial planning model InnoALM. *Oper. Res.*, 56(4):797–810, 2008.

[HFW] K. Heaven, M. Fazal, and A. White. Destination endgame. Available at `http://tiny.cc/3lmvtz`.

[HG18] Matt Hamilton-Glover. Trends in the uk mortality experience. MPhil Dissertation, King's College London, February 2018.

[HKP11] P. Hilli, M. Koivu, and T. Pennanen. Cash-flow based valuation of pension liabilities. *European Actuarial Journal*, 1:329–343, 2011.

[Ins] Insight Investment. Cashflow driven investment (CDI). Available at `http://tiny.cc/co0wtz`.

[KMB19] S. Kazziha, R. Martel, and M. A. Burns. Cashflow driven investing: Does it make sense for uk db schemes? Available at `http://tiny.cc/oinvtz`, June 2019.

[KPL+12] Andreas Klöckner, Nicolas Pinto, Yunsup Lee, B. Catanzaro, Paul Ivanov, and Ahmed Fasih. PyCUDA and PyOpenCL: A Scripting-Based Approach to GPU Run-Time Code Generation. *Parallel Computing*, 38(3):157–174, 2012.

[Mci19] K. Mcinally. What is cashflow driven investment? Available at `http://tiny.cc/hmmvtz`, February 2019.

[MSZ+08] John M Mulvey, Koray D Simsek, Zhuojuan Zhang, Frank J Fabozzi, and William R Pauling. Or practice—assisting defined-benefit pension plans. *Operations research*, 56(5):1066–1078, 2008.

[NVI21] NVIDIA Corporation. *CUDA Toolkit Documentation*. NVIDIA Corporation, 2021.

[Pen14] T. Pennanen. Optimal investment and contingent claim valuation in illiquid markets. *Finance and Stochastics*, 18(4):733–754, 2014.

[Pen19] Pensions and Lifetime Savings Association. Cashflow driven investment made simple. Available at `http://tiny.cc/linvtz`, October 2019.

[Sod05] ManMohan S Sodhi. Lp modeling for asset-liability management: A survey of choices and simplifications. *Operations Research*, 53(2):181–196, 2005.

[Uni] University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Human mortality database. Available at `www.mortality.org`.

[Uni05] United Kingdom Debt Management Office. Formulae for calculating gilt prices from yields. Available at `http://tiny.cc/ldsvtz`, March 2005.

[Uni12] United Kingdom Debt Management Office. Uk government securities: a guide to 'Gilts'. Available at `http://tiny.cc/ojpvtz`, June 2012.

[UR01] Stanislav Uryasev and R. Tyrrell Rockafellar. Conditional value-at-risk: optimization approach. In *Stochastic optimization: algorithms and applications (Gainesville, FL, 2000)*, volume 54 of *Appl. Optim.*, pages 411–435. Kluwer Acad. Publ., Dordrecht, 2001.

[USS18] USS. Your guide to the universities superannuation scheme. Available at `http://tiny.cc/fazvty`, April 2018.

[VRD09] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.

[Wil19] Willis Towers Watson. Cashflow driven investment: Ultimate solution to the pensions problem, or just glorified liability-driven investment? Available at `http://tiny.cc/olmvtz`, April 2019.

10 Appendix

	Asset class	Bloomberg global ID	Coupon	Coupon frequency	Maturity	Capital
1	UKT	BBG0000TFSH9	4.000	2	2022-03-07	0.81874
2	UKT	BBG00SSK0M13	0.125	2	2023-01-31	0.84739
3	UKT	BBG00LG783Z7	1.000	2	2024-04-22	0.90645
4	UKT	BBG00004SL94	5.000	2	2025-03-07	1.09001
5	UKT	BBG00SSK93Y0	0.125	2	2026-01-30	0.91568
6	UKT	BBG00G5GVKZ6	1.250	2	2027-07-22	1.27297
7	UKT	BBG00V94C599	0.125	2	2028-01-31	0.78661
8	UKT	BBG0000D14P3	4.500	2	2034-09-07	1.77452
9	UKT	BBG00X2TXPF3	0.625	2	2035-07-31	1.34168
10	UKT	BBG003H46RK8	3.250	2	2044-01-22	0.89548
11	UKT	BBG006N6HZM7	3.500	2	2045-01-22	1.08927
12	UKT	BBG00DSPQPB9	1.500	2	2047-07-22	0.03635
13	UKTI	BBG00YZD3PJ9	0.125	2	2031-08-10	5.58410
14	UKTI	BBG00CCP0450	0.125	2	2036-11-22	3.84446
15	UKTI	BBG00L991NQ4	0.125	2	2041-08-10	4.81502
16	UKTI	BBG009CQDGT8	0.125	2	2046-03-22	0.88867
17	UKTI	BBG00J2D03W5	0.125	2	2048-08-10	0.66255
18	STOCK	--	0.000	0	2026-04-08	0.02763
19	STOCK	--	0.000	0	2055-04-08	0.56157

Table 7: Capital allocation for the portfolio S_1, also illustrated in Figure 9.

	Asset class	Bloomberg global ID	Coupon	Coupon frequency	Maturity	Capital
1	UKT	BBG0000TFSH9	4.000	2	2022-03-07	0.86210
2	UKT	BBG00SSK0M13	0.125	2	2023-01-31	0.89226
3	UKT	BBG00LG783Z7	1.000	2	2024-04-22	0.95049
4	UKT	BBG00004SL94	5.000	2	2025-03-07	1.16440
5	UKT	BBG00G5GVKZ6	1.250	2	2027-07-22	1.06139
6	UKT	BBG00001YK86	4.250	2	2032-06-07	1.53335
7	UKT	BBG00X2TXPF3	0.625	2	2035-07-31	1.41480
8	UKT	BBG0000547V3	4.250	2	2036-03-07	0.00001
9	UKTI	BBG009KCJHC0	0.125	2	2026-03-22	1.05687
10	UKTI	BBG00L4P50W2	0.125	2	2028-08-10	1.08573
11	UKTI	BBG002802G17	0.125	2	2029-03-22	1.13997
12	UKTI	BBG00YZD3PJ9	0.125	2	2031-08-10	1.54290
13	UKTI	BBG0000VN0M0	1.250	2	2032-11-22	1.32641
14	UKTI	BBG001PNKYF6	0.750	2	2034-03-22	1.29174
15	UKTI	BBG00CCP0450	0.125	2	2036-11-22	2.35762
16	UKTI	BBG0000GRYH1	1.125	2	2037-11-22	1.30824
17	UKTI	BBG0000GW3R2	0.625	2	2040-03-22	0.70541
18	UKTI	BBG00L991NQ4	0.125	2	2041-08-10	2.50022
19	UKTI	BBG0000G0PJ7	0.625	2	2042-11-22	0.44811
20	UKTI	BBG0036PB7W8	0.125	2	2044-03-22	1.20397
21	UKTI	BBG009CQDGT8	0.125	2	2046-03-22	1.12476
22	UKTI	BBG00J2D03W5	0.125	2	2048-08-10	0.86871
23	UKTS	BBG000021SM7	0.000	0	2029-03-07	0.00001
24	UKTS	BBG000021ZH7	0.000	0	2030-03-07	0.87823
25	UKTS	BBG0000D7P61	0.000	0	2039-03-07	1.23073
26	UKTS	BBG0000BSSB0	0.000	0	2043-06-07	0.32134
27	UKTS	BBG0000BSSL9	0.000	0	2045-06-07	0.38369
28	UKTS	BBG0000BSRS4	0.000	0	2047-06-07	0.20508
29	UKTS	BBG0000BSS84	0.000	0	2049-06-07	0.27496
30	UKTS	BBG0000BSSX6	0.000	0	2050-06-07	0.09403
31	STOCK	--	0.000	0	2027-04-08	0.06138
32	STOCK	--	0.000	0	2050-04-08	0.10465
33	STOCK	--	0.000	0	2051-04-08	0.12102
34	STOCK	--	0.000	0	2052-04-08	0.05730
35	STOCK	--	0.000	0	2053-04-08	0.05834
36	STOCK	--	0.000	0	2054-04-08	0.00667
37	STOCK	--	0.000	0	2055-04-08	0.06394

Table 8: Capital allocation for the portfolio S_2, also illustrated in Figure 11.

	Asset class	Bloomberg global ID	Coupon	Coupon frequency	Maturity	Capital
1	UKT	BBG0000TFSH9	4.000	2	2022-03-07	0.00001
2	UKT	BBG00SSK0M13	0.125	2	2023-01-31	2.61210
3	UKT	BBG00LG783Z7	1.000	2	2024-04-22	0.00002
4	UKT	BBG00G5GVKZ6	1.250	2	2027-07-22	2.94432
5	UKT	BBG00F5MW006	1.750	2	2037-09-07	5.01497
6	UKT	BBG003H46RK8	3.250	2	2044-01-22	0.00004
7	UKT	BBG006N6HZM7	3.500	2	2045-01-22	0.03947
8	UKTI	BBG00YZD3PJ9	0.125	2	2031-08-10	8.47865
9	UKTI	BBG00CCP0450	0.125	2	2036-11-22	0.00017
10	UKTI	BBG00L991NQ4	0.125	2	2041-08-10	7.21289
11	UKTI	BBG00Z0Y5WF5	0.125	2	2051-03-22	0.73502
12	STOCK	--	0.000	0	2026-04-08	0.11237
13	STOCK	--	0.000	0	2055-04-08	0.65310

Table 9: Capital allocation for the portfolio S_3, also illustrated in Figure 12.

	Asset class	Bloomberg global ID	Coupon	Coupon frequency	Maturity	Capital
1	UKT	BBG0000TFSH9	4.000	2	2022-03-07	0.65200
2	UKT	BBG00SSK0M13	0.125	2	2023-01-31	0.67582
3	UKT	BBG00LG783Z7	1.000	2	2024-04-22	0.57006
4	UKT	BBG00004SL94	5.000	2	2025-03-07	0.86371
5	UKT	BBG00SSK93Y0	0.125	2	2026-01-30	0.71308
6	UKT	BBG00G5GVKZ6	1.250	2	2027-07-22	1.09643
7	UKT	BBG00V94C599	0.125	2	2028-01-31	0.73179
8	UKT	BBG00005CXH3	6.000	2	2028-12-07	1.28664
9	UKT	BBG00Y3L4KR7	0.250	2	2031-07-31	0.00001
10	UKT	BBG00001YK86	4.250	2	2032-06-07	1.54044
11	UKT	BBG0000D14P3	4.500	2	2034-09-07	3.28801
12	UKT	BBG00X2TXPF3	0.625	2	2035-07-31	1.63449
13	UKT	BBG00F5MW006	1.750	2	2037-09-07	5.22611
14	UKTI	BBG00YZD3PJ9	0.125	2	2031-08-10	2.25568
15	UKTI	BBG00CCP0450	0.125	2	2036-11-22	0.35556
16	UKTI	BBG00L991NQ4	0.125	2	2041-08-10	0.48803
17	STOCK	--	0.000	0	2024-04-08	0.15255
18	STOCK	--	0.000	0	2044-04-08	0.08760
19	STOCK	--	0.000	0	2048-04-08	0.21066
20	STOCK	--	0.000	0	2049-04-08	0.00072
21	STOCK	--	0.000	0	2054-04-08	0.10798
22	STOCK	--	0.000	0	2055-04-08	2.20105

Table 10: Capital allocation for the portfolio S_4, also illustrated in Figure 13.

	Asset class	Bloomberg global ID	Coupon	Coupon frequency	Maturity	Capital
1	UKT	BBG000TFSH9	4.000	2	2022-03-07	0.83173
2	UKT	BBG00SSK0M13	0.125	2	2023-01-31	0.86874
3	UKT	BBG00LG783Z7	1.000	2	2024-04-22	0.91916
4	UKT	BBG00004SL94	5.000	2	2025-03-07	1.09615
5	UKT	BBG00SSK93Y0	0.125	2	2026-01-30	0.98927
6	UKT	BBG00G5GVKZ6	1.250	2	2027-07-22	1.14964
7	UKT	BBG0000D14P3	4.500	2	2034-09-07	1.64504
8	UKT	BBG00X2TXPF3	0.625	2	2035-07-31	1.22186
9	UKT	BBG003H46RK8	3.250	2	2044-01-22	0.54058
10	UKT	BBG006N6HZM7	3.500	2	2045-01-22	0.70275
11	UKT	BBG00DSPQPB9	1.500	2	2047-07-22	0.72867
12	UKTI	BBG00L4P50W2	0.125	2	2028-08-10	0.80999
13	UKTI	BBG00YZD3PJ9	0.125	2	2031-08-10	5.82667
14	UKTI	BBG00CCP0450	0.125	2	2036-11-22	3.66239
15	UKTI	BBG00L991NQ4	0.125	2	2041-08-10	5.37021
16	UKTI	BBG009CQDGT8	0.125	2	2046-03-22	0.68896
17	UKTI	BBG00J2D03W5	0.125	2	2048-08-10	0.96082
18	UKTI	BBG00Z0Y5WF5	0.125	2	2051-03-22	0.32551
19	UKTS	BBG0000D7P61	0.000	0	2039-03-07	0.44634
20	STOCK	--	0.000	0	2027-04-08	0.00971
21	STOCK	--	0.000	0	2055-04-08	0.24078

Table 11: Capital allocation for the portfolio S_5, also illustrated in Figure 14.