Robust Long-Term Interest Rate Risk Hedging in Incomplete Bond Markets☆

Sally Shen[a,c,*], Antoon Pelsser[b,c], Peter Schotman[b,c]

[a] *Global Risk Institute, 55 University Avenue, Toronto, ON M5J 2H7 Canada*
[b] *Department of Finance, Maastricht University, PO BOX 616, 6200 MD Maastricht, the Netherlands*
[c] *Netspar*

Abstract

Due to the scarcity of long-term market instruments, valuation of an ultra long-term pension liability under the market-consistent valuation framework must be model based. We develop a robust self-financing hedging strategy which adopts a min-max expected shortfall hedging criterion to replicate the long dated liability for agents who fear model misspecification. We introduce a backward robust least squares Monte Carlo method to solve this dynamic robust optimization problem. We find that both naive and robust optimal portfolios depend on the hedging horizon and the current funding ratio. The robust policy suggests to take more risk when the current funding ratio is low. The artificial yield curve constructed by the robust dynamic hedging portfolio is lower than the naive one but is higher than the model based yield curve in a low rate environment.
JEL classification: G11; C61; E43
Keywords: Interest Rate Risk; Model Misspecification; Robust Optimization; Least Squares Monte Carlo; Incomplete Market

---

1. Introduction

In Europe, market consistent valuation has been widely adopted, particularly as a valuation framework of Solvency II for calculating defined-benefit pension liabilities. When pension liabilities can be specified as a stream of risk-free cash flows, the underlying goal of market consistent valuation is to replicate these non-traded cash flows as much as

☆This version: April 4, 2018.
*Corresponding author Tel +1 647 448 8216
Email address: `sshen@globalriskinstitute.org` (Sally Shen)

possible by using cash flows of deeply liquid financial instruments. However, this task is challenged by the fact the liabilities of pension funds frequently extend beyond the term of available market instruments. The liquidity of financial market fades away beyond the term of 30 years, while pension funds usually face ultra long-term (30 years and longer) commitments with maturities of more than 70 years.

Due to the scarcity of the long-term market instruments, the valuation of the ultra-long-term pension liabilities requires an extrapolation of the yield curve beyond the last liquid point of the financial market. Yield curve extrapolation is mostly model based. Conditional on a perfectly accurate interest rate model, it is possible to replicate ultra-long term (30 years and longer) liabilities by using a duration matching strategy if short positions in assets are allowed.

However, with incomplete bond markets, pension liability valuation is challenged by both parameter uncertainty and model uncertainty. Parameter uncertainty is associated with the ambiguity about the value of the exact parameters given that the underlying model represents the true data generating process. Even in the liquid bond market, it is frequently the case that the deeply liquid instruments that promise the same amount of cash flows in the future are traded at different prices. Variations in issue size, coupon rate and other bond-specific factors may easily influence the estimation results. Further, many smoothing techniques such as Nelson and Siegel (1987) are based on minimizing sums of squared deviations and can give rise to more than one local solution. A small parameter misspecification in the interpolation phase can have a huge impact on the extrapolated yield curve.

Model uncertainty involves fragile beliefs of the probability distribution of the underlying process. Conditional on a perfectly accurate interest rate model, one can derive a term structure of the bond for all maturities. However, there are a large number of models that fit the bond prices up to their last liquid point while indicating very different prices for the longer-term bonds.

Fear of model misspecification motivates an investor to find a robust decision rule that works well over a set of models. The aim of this paper is to develop such a replication portfolio for an agent who faces uncertainty about the bond pricing model used for the ultra-long term liability valuation. Adapting the philosophy of Anderson et al. (2003), we consider a class of models by imposing distortions to the bond risk premium which is parameterized as a linear function of the spot rate. This framework enables us to modify the general model uncertainty problem of Anderson et al. (2003) to the uncertainty of a subclass of the bond pricing model, namely the risk premium parameters. However, our framework differs from a pure parameter uncertainty problem, since we allow for stochastic drift distortions, which carry fragile beliefs over the law of motion of the bond premium.

We adopt the robust control technique formulated by a min-max two players game to deal with the fear of parameter misspecification. An agent who faces parameter uncertainty wants to replicate the ultra-long dated cash flow as much as possible using the liquid fixed income instruments. With accessibility to short positions, we construct a dynamic hedging portfolio to minimize the expected shortfall of a long maturity commitment. Parameter uncertainty is modeled by an artificial agent who makes a decision on the risk premium distortions over the investment horizon to maximize the expected shortfall. We impose a per-period uncertainty set to restrict the size of distortions. To calibrate the bound, we link the distortions to the econometric parameter estimation error. The bounded distortions imposed to the bond premium are related to the Good Deal Bound of Cochrane and Saa-Requejo (2000), which restricts the maximum Sharpe ratio in an incomplete market. Our robust optimization framework differs from Hansen and Sargent (2007), in which they introduce a relative entropy[1] term to penalize the drift distortions and use detection error probabilities to calibrate the entropy parameter. The Hansen and Sargent framework implicitly imposes an "aggregate budget" constraint towards the distortions rather than a time homogeneous constraint.

We introduce a robust least squares Monte Carlo method to solve this dynamic hedging problem. It is a regression-based method that combines, in essence, the methods of

[1] Relative entropy is a statistical method to calculate the difficulty of distinguishing between two models.

Brandt et al. (2005), Carroll (2006) and Koijen et al. (2010). As an extension to the existing literature, we impose the preference for robustness to the algorithm. We derive the optimal solution for the max-player analytically, while the hedging position for the min-player is solved numerically. The essential idea behind the regression based numerical method is to describe the realized value function proxy as a polynomial expansion function of the risk factors.

Fear of parameter uncertainty induces a strong demand for the medium-term bonds when the solvency ratio is low. This result contradicts many other robust asset allocation studies such as Maenhout (2004) and Brennan (1998) who suggest a more conservative portfolio for robust investors. This is because the robust investor under our framework believes that the artificial agent would choose a lower bond premium than the model estimation. As a result, a risker portfolio with larger weight on medium-term bond is required for investors that are liability driven. The optimal hedging portfolio depends on the hedging horizon. The longer the hedging horizon, the more risk exposure to the medium-term bond markets is required.

We also construct an artificial yield curve that is based on the optimal dynamic hedging strategy. The robust yield curve is always lower than the naive one that does not consider parameter uncertainty. These lower discount rates imply that more initial wealth is needed to meet the shortfall target.

One strong assumption in this paper is that only bond risk premium parameters are subject to misspecification. In other word, we only allow misspecification under the risk-neutral $\mathcal{Q}$ measure, but not under the physical $\mathcal{P}$ measure. Another strong assumption of our work is that agents do not engage in any learning process. Brennan (1998) incorporates Bayesian learning with parameter uncertainty and finds that low risk-averse investors put more wealth into risky assets after learning while the high risk-averse investors are more conservative with their portfolio. Peijnenburg (2017) conduct a numerical study to show that both fear of model uncertainty and learning about the equity premium contribute to the low stock market participation rate over the life cycle.

In terms of hedging criterion, this paper is closely related to Föllmer and Leukert

(2000), in which investors also face an option-style objective function and aim to minimize the shortfall risk in incomplete markets while the underlying model is assumed accurate. As an extension, we impose model uncertainty to the hedging problem and we provide a semi-explicit solution in incomplete markets.

This paper is also related to the studies on long-term discount rate for pension liability valuation. Regulators propose a subjective method that extrapolates the liquid market interest rates such that they converge in the long run to an unconditional ultimate forward rate (UFR). The UFR is determined by long-term expectation of some macroeconomics factors. According to the Solvency II regulation (see CEIOPS (2010)), the UFR is set at 4.2% (valid until the end of 2017) which consists of a short-term real rate of 2.2% and a long-term inflation of 2%.[2] Broeders et al. (2014) argue that the UFR method is exposed to both parameter and model uncertainty, and they provide sensitivity analysis and some economic insights on how the choice of the UFR affects the pension funds hedging policy. In terms of ultra long-term liability valuation, these studies focus on finding a reasonable long-term discount rate via term structure extrapolation method, while we retrial this pricing problem by adopting the dynamic hedging technique.

The rest of the paper proceeds as follows. Section 2 describes the one-factor affine term structure model employed in our economy. Section 2.1 uses the GMM method to estimate the structural parameters. The dynamic robust optimization problem is explained in Section 3. Section 4 elaborates the regression-based techniques used in solving our dynamic programming problem. Section 5 discusses the robust optimal solution and we provide policy evaluation of the robust policy. Section 6 concludes.

2. Term Structure Model

The term structure model we use is based on the framework of Duffee (2002) and Duffie and Kan (1996). We assume that the spot rate r follows the one-factor Vasicek

[2]On 4 April 2017, EIOPA published its final decision on the methodology for the annual calculation of the UFR for each currency. For most currencies, the UFR will reduce to 4.05% on 1 January 2018.

model

$$dr = \kappa\left(\theta - r\right)dt + \sigma dW, \tag{1}$$

where θ is the unconditional mean, κ is the mean reversion, σ is a volatility parameter and dW is a univariate Brownian motion. For the purpose of hedging the risk of an ultra long-term liability using shorter-term bonds, one-factor affine term structure model is particularly suited since very long rates are mostly affected by the most persistent factor. Let $P(t,T)$ be the time t price of a discount bond maturing at time T. The Vasicek model implies that bond prices follow the diffusion

$$dP = rPdt + B\sigma P\left(\lambda dt + dW\right) \tag{2}$$

where $B = \frac{1-e^{-\kappa\tau}}{\kappa}$ is the volatility of long-term bond returns relative to the spot rate volatility and where λ is the market price of bond risk. Departing from the original Vasicek (1977) model, in which the price of risk parameter is assumed constant, we adopt Duffee (2002)'s framework by imposing an essential affine extension of the price of risk such that λ depends on the spot rate,

$$\lambda = \lambda_0 + \lambda_1 r \tag{3}$$

we parameterize λ as a linear affine function of r with constant coefficients λ_0 and λ_1.

2.1. Model Calibration

Parameters of the model are estimated using standard GMM. We use Euro area nominal government bonds with triple A issuing ratings, obtained from the ECB statistical data warehouse. We use daily annualized data on 3 constant maturity zero rates with maturities of 3 months, and 5 and 10 years for the period from September 6, 2004 to November 15, 2013.[3]

[3] The yield data is available at sdw.ecb.europa.eu. Sample size is 2360, and there are 9.2 (approximately) years in our sample, hence the average yearly number of trading days is $\frac{2{,}360}{9.2} \approx 255$.

Table 1: Summary Statistics
Means, standard deviation and autocorrelations of daily yield curve spot rate with three different maturities and their difference. The variable r_t denotes three-month spot rate. ADF denotes the Augmented Dickey-Fuller unit root statistics with a 5% critical value of -3.43.

Variable	Mean	Std	Std Δ	ADF
r_t	1.62%	1.45%	3.28bp	-0.61
$Y_t(5)$	2.61%	1.10%	4.27bp	-0.40
$Y_t(10)$	3.36%	0.80%	3.97bp	-0.94

Table 2: GMM Estimates.
The term structure model is estimated by the GMM method using daily European data. The bond maturities used are three-month, five-year and ten-year maturities. The parameters are expressed in annual terms. We choose 20 as the number of lags applied in the Newey West estimator.

Parameter	Estimate	Std error	Correlation Matrix				
κ	17.16%	17.43%	1.0000	-0.5893	0.8674	0.0380	-0.9987
θ	1.62%	1.24%	-0.5893	1.0000	-0.6587	0.7808	0.5937
σ	0.52%	0.14bp	0.8674	-0.6587	1.0000	-0.1627	-0.8629
$\sigma\lambda_0$	-0.16%	0.36%	0.0380	0.7808	-0.1627	1.0000	-0.0363
$\sigma\lambda_1$	-14.19%	17.49%	-0.9987	0.5937	-0.8629	-0.0363	1.0000

Summary statistics are displayed in Table 1. The average three-month rate is 1.62% with a standard deviation of 1.45%. The 5-year rate has a mean of 2.61% with standard deviation of 1.1%. The 10-year bond has highest average rate but lowest volatility among the three. The ADF test shows that we cannot reject the unit root hypothesis for any of the three series.

In Table 2, we report the parameter estimation results for the full sample. The parameters are expressed in annual terms. The first two columns of Table 2 report the estimated value and standard deviation of each structural parameter. The unconditional mean of spot rate, θ, is about 162 basis points. The mean reversion coefficient κ implies half-life innovation of 4 years.

The rest of the table reports the correlation matrix of parameters. We find an extremely high correlation between κ and λ_1, and between θ and λ_0. The low volatility of σ implies an accurate estimation performance on spot rate volatility. Hence we ignore the parameter uncertainty on σ. Checking the eigenvalues of the covariance matrix of the estimates, we find two positive eigenvalues 0.0609, 0.0001 whereas the remaining three are very close to zero. From this we conclude that two sources of parameters estimation error exist.

2.2. Model Uncertainty

To introduce the preference for robustness to the bond pricing model (2), we borrow from the ideas of Anderson et al. (2003) and Hansen and Sargent (2007). Due to the preference for robustness, agents treat (2) as a reference model for the unknown true model and look for a robust decision rule that works well for a set of models. We consider a class of models indexed by a vector process c_t. The vector c_t represents model misspecification around the reference model with $c_t \equiv 0$ under the reference model. In a pure diffusion setting, Anderson et al. (2003) view c_t as an endogenous drift to the law of motions of the state variables and is formally measured by relative entropy, a statistical method to calculate the distance between two models.

In this paper, we assume that parameters σ, κ and θ can be estimated precisely, whereas the expected excess return parameters λ_0 and λ_1 are subject to estimation error. As such, the general model uncertainty problem of Anderson et al. (2003) and Hansen and Sargent (2007) is reduced to a subclass model uncertainty problem which leads to a problem of risk premium uncertainty. Hence, fear of model uncertainty only motivates an agent to find a robust decision rule that works well for a set of bond risk premia.

To formulate the subclass model misspecification, we impose the perturbation vector $c_t = \begin{pmatrix} c_0 \\ c_1 \end{pmatrix}$ to the vector of the estimated risk premium terms of (3) defined as δ, with

$$\delta = \begin{pmatrix} \lambda_0 \\ \lambda_1 \end{pmatrix} \tag{4}$$

and let δ_0 be the true value of risk premium parameters.

We assume that there is an artificial agent (" nature") who makes an instantaneous decision on c_t such that the true value of bond premium perturbs from the reference model. The distortions c_1 and c_2 shift the mean distribution of the bond diffusion process (2) by a unit of $B\sigma\,(c_0 + c_1 r)$. Hence they specify a set of alternative measures referring to different specifications of the stochastic process known as a Girsanov kernel.

The distortions are bounded by a time-homogeneous uncertainty set $\mathbf{S}$. A larger uncertainty set $\mathbf{S}$ indicates a greater fear of model uncertainty. To calibrate the bound on the specification error c_t, we relate the distortions to the econometric parameter estimation error. We start with mapping the distortion vector c_t to the estimation errors of δ,

$$\delta - \delta_0 = c_t. \tag{5}$$

The estimation errors are assumed to be asymptotically normal with mean zero and covariance matrix $\frac{\Omega}{\mathcal{N}}$, where Ω is the covariance matrix of δ and $\mathcal{N}$ is the length of the sample used for estimation. We obtain the uncertainty set based on the property that $c_t' \left(\frac{\Omega}{\mathcal{N}}\right)^{-1} c_t$ is a Chi-square distribution with two degrees of freedom, $\chi^2(2)$. This leads to the following uncertainty set $\mathbf{S}$

$$\mathbf{S} := \left\{ \begin{pmatrix} c_0 \\ c_1 \end{pmatrix} \mid \begin{pmatrix} c_0 \\ c_1 \end{pmatrix}' \Omega^{-1} \begin{pmatrix} c_0 \\ c_1 \end{pmatrix} \leq \gamma^2 \right\} \tag{6}$$

where $\gamma^2 = \frac{CV_\alpha}{\mathcal{N}}$, with CV_α denoting the critical value of $\chi^2(2)$ at α significance level. Our uncertainty set differs from the Hansen and Sargent framework in two ways. First, under their framework the perturbations are formally measured by relative entropy. The entropy term penalizes drift distortions that move away from the reference model. The relative entropy parameter plays the same role as a Lagrange multiplier of the Hamilton-Jacobian-Bellman equation for optimality. However, this Lagrange multiplier implicitly imposes a constraint on the following bound for the specification error

$$\mathbb{E}\left[\sum_0^T c_t \cdot c_t\right] \leq \eta_0, \tag{7}$$

where η_0 sets the average size of the potential model misspecification. The uncertainty set (7) is an aggregate-budget style present-value bound that restricts the distortions cumulatively over the decision horizon T. This present-value constraint (7) is time-inconsistent. However, our uncertainty set (6) is a time-homogenous constraint that we can impose cleanly on nature's future actions without keeping track of what has happened in the past.

The second difference is related to the uncertainty set calibration. Under the Hansen and Sargent framework, the boundary parameter η_0 is linked to relative entropy which in essence is a Lagrange multiplier over the misspecification bound (7). Hansen and Sargent (2007) use a Bayesian model selection problem (also called the detection error probability method) to calibrate this relative entropy parameter. The detection error probability method performs likelihood ratio tests under the two models based on available data. A higher degree of robustness leads to a lower detection error probability. We define the uncertainty set from the econometric estimation error.

We choose a significance level of $\alpha = 5\%$ at which the corresponding critical value of $\chi^2(2)$ is 5.99. We also assume that the minimum sample horizon for estimation is $\mathcal{N} = 200$ time units. Hence, as a benchmark, the misspecification bound is $\gamma = 0.17$ or $\gamma^2 = 0.03$. Either a longer data set or a higher significance level would reduce the bound parameter γ and results in a smaller uncertainty set for c_t, which indicates a gaining confidence over the reference model.

The impact of the spot rate r on the bond premium λ is governed by λ_1. If λ_1 is estimated with error, then the impact from the spot rate is also ambiguous. Panel (a) of Figure 1 presents the 20-year nominal bond risk premia for a realistic range of spot rates. The risk premium is increasing with the spot rate. Misspecification of market price of risk results in a perturbation of the bond premium of around 200 basis points either downwards or upwards.

The uncertainty set $\mathbf{S}$ has a circular shape in c_t space centered around zero, but turns to an ellipsoid shape when mapped to the δ space. Panel (b) of Figure 1 maps the uncertainty set of c_t to the (λ, λ_1) space centered by the point estimate $\left(\hat{\lambda}, \hat{\lambda}_1\right)$ with $r = 2\%$, where $\hat{\lambda} = \hat{\lambda}_0 + \hat{\lambda}_1 r$. The very narrowly shaped ellipsoid spaced is mainly driven by the large variation between the point estimate of λ_0 and λ_1.

3. Robust Optimal Portfolio Choice

3.1. Dynamic Replication Strategy

The agent with wealth X_t at time $t = 0$ aims to replicate an ultra long dated liability with payoff equal to one at a long end maturity time T by investing in two liquid fixed income instruments: a medium-term zero coupon bond with maturity $\tau_2 < T$; and a short term bond. The hedging horizon T is assumed beyond the term of the last liquid point of the financial market. The replication strategy is constructed based on dynamic hedging and the agent is only interested in eliminating the downside risk. Therefore the hedging criterion is defined over the expected shortfall $[1 - X_T]^+$ at time T. This hedging criterion allows the agent to replicate the ultra-long term cash flow as much as possible by active trading using the minimum amount of wealth. If the agent is not aware of parameter uncertainty, the hedging portfolio relies fully on the point estimator $\hat{\delta} = \begin{pmatrix} \hat{\lambda}_0 \\ \hat{\lambda}_1 \end{pmatrix}$ and we define such specification-error free trading strategy as a naive policy. The naive dynamic optimization problem is defined as

$$\min_{w_{t:0 \leq t < T}} \mathbb{E}\left[(1 - X_T)^+ \mid \mathcal{F}_t\right] \tag{8}$$

The fraction of wealth allocated to long-term bonds with maturity τ_2 at time t is indicated by w_t and short positions are allowed this setting. The law of motion of wealth is given by

$$dX = (r - wB(\tau_2)\sigma\lambda)\, Xdt - wB(\tau_2)\sigma X dW \tag{9}$$

where we omit the subscripts t.

Our dynamic trading strategy (8) is not affected by rollovers which occur when reinvesting wealth from matured bonds into newly issued bonds of the same feature. We assume that the invested fixed income securities are continuously issued and are liquidly traded during the entire decision horizon. Even when the remaining hedging horizon is shorter than then the marketable term, the investor still keeps on active trading. If we stop the active trading as soon as the market becomes complete when the remaining maturity of the liability equals to maturity of medium-term bond, then the expected shortfall hedging criterion would not be minimized and the required initial capital for hedging long-term liabilities would be higher.

3.2. Robust Hedging

The robust control method is formulated by a min-max expected utility framework which according to Hansen and Sargent (2007) resembles a two-player zero-sum game. The investor who fears model misspecification takes the role of min-player and seeks a robust hedging policy that minimizes the expected shortfall at maturity T under the worst-case scenario decided by the artificial agent. The artificial malevolent agent takes the role of max-player and makes a decision on distortions c_t given the min-player's choice. The equilibrium of the game gives us an instantaneous robust hedge. The robust optimization problem is given by

$$\min_{w_{t:0\leq t<T}} \max_{c_t \in \mathbf{S}} \mathbb{E}\left[(1-X_T)^+ \mid \mathcal{F}_t\right] \tag{10}$$

the artificial agent makes an instantaneous decision on c_0 and c_1 bounded by the time-homogeneous constraint $\mathbf{S}$, as thus controls the risk premium parameters λ_0 and λ_1.

4. Numerical Technique

Robust hedging with an expected shortfall objective function does not allow for an analytical solution, so we use a numerical approach instead. Our method in essence combines the methods proposed by Brandt et al. (2005) and Koijen et al. (2007) to approximate the conditional expectation that we encounter in solving the dynamic program by polynomial expansions in the state variables. We also follow Diris (2011)'s approach to parameterize the coefficients of the approximation function by a quadratic function of portfolio weights, such that we can calculate the optimal portfolio under each path analytically. Furthermore, this method allows us to achieve an accurate result using a very small grid of testing portfolios. We integrate robustness with the standard simulation based algorithm.

We start by simulating a large number of N sample paths with length T years of bond returns using Euler discretisation. We also choose an $M-$ dimensional grid for financial wealth values. The wealth grid points are indicated by X_j, $j = 1, \cdots, M$. In total, we have $(M \times N)$ grid points at each point in time.

The algorithm for the robust policy consists of two parts. We first solve nature's decision analytically. Nature's decision only influences the long-term bond premia. Therefore, under the assumption that $w_t \geq 0$, maximizing the expected shortfall boils down to minimizing bond premia. Therefore, we can analytically solve for nature's decision.

The optimal c_0, c_1 are the solution of following quadratic programming problem.

$$\begin{aligned} \min_{c_t} \quad & -B(\tau_2)\sigma\left(\hat{\delta} + c_t\right)' \alpha, \\ s.t \quad & c_t'\Omega c_t = \gamma^2 \end{aligned}$$

where $\hat{\delta} = \begin{pmatrix} \hat{\lambda}_0 & \hat{\lambda}_1 \end{pmatrix}'$ and $\alpha_t = \begin{pmatrix} 1 & r_t \end{pmatrix}$. We can easily find the optimal solution for c_t:

$$c_t^* = \gamma \frac{\Omega^{-1}\alpha_t}{\sqrt{\alpha_t'\Omega^{-1}\alpha_t}}$$

Therefore, the decision of the artificial agent at each rebalancing time step t depends exclusively on r_t. The min-max problem is now simplified to the minimization problem alone.

Figure 2 displays the robust optimal bond premium as a function of bond maturity against the point estimate. Nature intends to minimize the bond premium so as to maximize the expected shortfall at time T, therefore, the robust bond premium at any maturity is always lower than the point estimate.

In the second part of the algorithm, we use the *backward least squares Monte Carlo* (LSMC) method to solve for optimal portfolio. We summarize the LSMC algorithm in three steps.

Step 1. At time step t, i.e. starting from period $T-\Delta t$ and iterating to period 0, construct realized loss $V_T = (1 - X_T)^+$ for all simulated paths using the following information

- Robust optimal portfolio from previous time steps $w_s^*(s = t + \Delta t, \cdots, T - \Delta t)$
- Optimal natures decision from previous steps λ_s^* (this is analytically solved)
- A small grid of testing portfolio w_h, $h = 1, \cdots, H$
- Spot rate $r_{i,s}$ with $i = 1, \cdots, N$
- Current wealth $X_{j,t}$, $j = 1, \cdots, M$

so we have a cross-section size of $N \times M \times H$.

Step 2a. Run a cross-sectional regression by approximating the realized objective function $V_{T,ijh}$ calculated from Step 1 on trajectories of state variables as well as the testing portfolio w_h on a second-order polynomial expansion (including cross term) at a certain point in time.

$$V_{ijh,T}(\mathcal{F}_t) = \beta' f\left(X_{j,t}, r_{i,t}, w_h\right) + \epsilon_{ijh,t} \tag{11}$$

Step 2b. Now we parameterize the approximate conditional value function as a quadratic function of portfolio weights such that we can analytically calculate the optimal portfolio.

For a naive investor, the first part of the algorithm can be ignored since nature does not play a role in the naive hedging framework. Appendix A elaborates on the LSMC algorithm in more detail.

5. Long-term Investors and Bond Premia Uncertainty

In this section, we investigate the impact of model uncertainty on bond portfolio allocations. Section 5.1 analyzes the robust optimal portfolio. Section 5.2 discusses the property of the robust yield curve.

5.1. *Optimal Portfolio Choice*

Figure 3 plots the optimal portfolio of 20-year ($\tau_2 = 20$) nominal bonds as a function of the investment horizon based on three different hedging policies, namely naive policy, robust policy and delta hedging policy.

5.1.1. *Delta Hedging*

The delta of a long dated liability is defined as the rate of change of the price of liability with respect to the price of the underlying τ_2- year bond. The delta of long-term liability is

$$\Delta = \frac{\partial P(0,T)}{\partial P(0,\tau_2)} = \frac{B(T)}{B(\tau_2)},$$

assuming the Vasicek model is valid for all maturities. We use Delta hedge as a benchmark to investigate the optimal portfolio choice as a function of the solvency condition. We use the funding ratio, a fraction of the current wealth level X_0 and the hypothetical price of liability $P(0,T)$, to measure the solvency position of a fund. If the funding ratio is lower than one, then the fund has a higher expected value of liability than of its assets, which implies under-funding.

Suppose we have a funding ratio greater than one, and the underlying model is correctly specified, then the Delta position is indeed the optimal portfolio to hedge the long-dated liability. This is because when the current funding ratio is above one, downside risks fade away and the hedging problem boils down to a complete-market setting where the long-dated payoff can be fully replicated by traded bond portfolios. It is verified by Figure 3b that if the current funding ratio is above one, a delta neutral position is the optimal policy.

Figure 3 shows that the bond allocations are highly dependent on the hedging horizon and the present funding ratio for both the naive and the robust policies. Both hedging methods are identical when the horizon $T \leq 20$, because market completeness eliminates the model uncertainty. When the funding ratio is low (see Figure 3a), both policies suggest a risker position than the Delta hedge. In this case, following the Delta strategy may hedge the long-term interest rate risk but does not help to meet the long-dated commitment, due to insufficient current wealth.

Also, when the present funding ratio is low, a robust investor takes a more risky position than a naive investor. This is because a robust investor worries that the artificial agent may choose even lower bond premia than the model estimated, and therefore the robust method is even more aggressive. The result for the robust policy is the opposite of many results in the literature eg. Maenhout (2004). Usually, the robust policy is more conservative compared to the naive one. In Maenhout (2004)'s model, an investor aims to maximize her terminal wealth utility function. The preference for robustness suppresses the beliefs in risk premium. Without the liability constraint, a robust investor will therefore take less risk.

We now demonstrate how the optimal policy respond to changes in the funding ratio and the spot rate. Figure 4 displays the optimal asset allocation to 20-year bonds under the naive policy (Figure 4a) and the robust policy (Figure 4b) for a reasonable range of the funding ratio (FR). These figures are constructed by regressing the optimal policy along all trajectories of state variables used in the simulations at a certain point in time.

Figure 5 summarizes the insight of Figure 4. The lower the funding ratio, the riskier the position for hedging against the shortfall risk. This property holds true for both policies. However, the robust policy is risker due to the fear of underestimated bond premia. Figure 3 leads to the same conclusion.

Figure 6 presents the optimal asset allocation under different spot rates when the funding ratio is 80%. For both policies, the optimal allocation increases with the spot rate, as the spot rate is positively related to bond premia. A higher bond premium leads to a higher risk exposure on long-term bonds. The robust policy Figure 6b differs from the naive policy Figure 6a in two ways. First, the robust policy is less sensitive to the spot rate. This is an important and a desired feature of the robust policy. Second, the difference between the two policies gets larger for a lower value of the spot rate.

Figure 7 presents the optimal expected shortfall value as a function of the funding ratio and hedging horizon. Figure 8 summarizes the insights of Figure 7. First, without considering model uncertainty, the expected shortfall value is dramatically under-estimated when the funding ratio is low. Secondly, the mispricing error increases with the investment horizon. The increasing amount of mispricing is caused by an accumulative fear of the underestimated bond premia.

5.2. Robust Term Structure

An investor is interested in the optimal strategy that guarantees the liability with the lowest required initial wealth X^*, which is called the super-hedging strategy. Once we have determined the minimum wealth X^*, we can define the implicit $T-$ period discount rate as

$$y(T) = -\frac{1}{T} \ln X^* \tag{12}$$

For most models of the term structure, the super-hedging strategy will be extremely costly. In our model, a super-hedging strategy will not exist, since there will always be a small probability of underfunding because our interest rate process is Gaussian. We define approximately risk-free rates through the minimum assets required as having an expected shortfall of less than S,

$$y(T) = -\frac{1}{T} \ln X_S^*, \qquad \text{with} \quad S = \mathbb{E}_0 \left[(1 - X_T)^+ \mid X_0 = X_S^* \right] \tag{13}$$

Expected shortfall does not reward any upside. The lower S is, the more initial wealth is required to hedge against the targeting shortfall risk, hence the lower $y(T)$ will be. The lowest level of $y(T)$ is simply the underlying term structure model that by definition guarantees zero expected shortfall.

Figure 9 presents the implied yield curves under different hedging policies at different shortfall levels and initial spot rates. As a benchmark, we plot the two hedging policy based yields together against the Vasicek yield curve. We highlight two crucial insights from Figure 9. First, the robust yield is always lower than the naive yields regardless of the current spot rate or shortfall target. Due to the concern about an underestimated bond premium, a robust investor requires more initial wealth to meet the shortfall target.

Second, both policy-based yield curves are higher than the Vasicek curve when the spot rate is low (Figure 9a, 9b). When the spot rate is low, the present value of liability will be higher, hence the required initial wealth would be higher as well, in order to meet the long-term commitment. The Vasicek model is based on a zero shortfall requirement, and hence has a lower yield curve compared with the other two curves which allow for positive expected shortfall. We also find that the Vasicek yield curve overlaps the naive curve when the spot rate is high if S is sufficiently small (Figure 9c).

6. Conclusion

We have constructed an optimal bond portfolio to minimize the replication error of an ultra long-dated cash flow in the presence of incomplete bond markets and model uncertainty. The robust policy is not always more cautious than a naive policy. With a fixed long-term commitment, we find that the robust portfolio suggests holding riskier positions.

The robust hedging strategy is a powerful method for pricing an ultra long-dated liability. Although the robust policy requires more initial wealth to guarantee a certain level of shortfall risk, it provides a more successful and resilient hedge, especially in a disordered environment where the bond premium is misspecified.

There are at least two directions in which future research along these lines would be interesting. The first is to consider mean-reversion parameter uncertainty under the physical $\mathcal{P}$ measure. We find from calibration results that the mean reversion parameter κ which is under the $\mathcal{P}$ measure, is exposed to very higher estimation error. Further, κ is also highly correlated with other parameters, such as the long-run mean parameter and the volatility term as well as the market-price-of-risk parameters. The layout of the optimization problem in this study does not allow for parameter uncertainty of other dimensions. Mean-reversion is especially relevant in interest rate models. It not only influences the expected value of the bond return, but also determines the volatility of the bond diffusion process for longer term bonds. In our framework, we implicitly consider mean reversion uncertainty but is only limited under the risk-neutral $\mathcal{Q}$ measure. Since both λ_0 and λ_1 are subject to parameter uncertainty, the mean reversion under $\mathcal{Q}$, which is equal to $\kappa - \lambda_1 \sigma$ is uncertain. This is the most important parameter, since it determines the volatility of bond prices, namely $B(T)$ and $B(\tau_2)$. It does not affect the spot rate dynamics, but it does affect the medium-term bond.

Second, our robust model is applicable to any term structure model, and the Vasicek model used in this paper is simply a start. We can play our hedging problem under various reference models to see whether the model builder's decision also matters.

Acknowledgements

We are grateful for comments from Frank de Jong, Peter Vlaar, Hans Schumacher, participants in EEA-ESEM Toulouse 2014 conference and Actuarial and Financial Mathematics Conference 2014.

Appendix A Robust LSMC algorithm

In this section, we elaborate the numerical method used in this paper. We show how the regression-based method works in our robust hedging problem and we also analyze the accuracy of our algorithm. This section proceeds as follows, we first introduce the naive LSMC algorithm without considering model uncertainty. Second, we show how nature's choice can be solved analytically. Third, we explain the robust LSMC algorithm. Last, we analyze the accuracy of our algorithm.

A.1 Grid Generation

The hedging period is from 0 to T. We partition $[0, T]$ into m subintervals of length $\Delta t = \frac{T}{m}$. Hence the bond portfolio is rebalanced every Δt unit of time. We start by simulating N trajectories of m time periods of spot rate under $\mathbb{P}$ measure using discrete Euler approximation.

$$r_{t+\Delta t} = r_t + \kappa \left(\theta - r_t\right) \Delta t + \sigma \sqrt{\Delta t} Z$$

with Z a standard normal random variable. We indicate the spot rate at time t in trajectory i by $r_{i,t}$, $i = 1, \cdots, N$, $t = 0, \Delta t, \cdots, T - \Delta t, T$. At step 0, r_0 is also random with mean θ and volatility $\frac{\sigma}{\sqrt{2\kappa}}$.

Next, we need to generate an M dimensional grid of funding ratio. Funding ratio at time t is defined as a fraction of instantaneous wealth against the hypothetical bond price maturing at T, $P(t, T)$. The funding ratio grids are indicated by FR_j, $j = 1, \cdots M$. The reason for choosing funding ratio grid instead of wealth grid will be explained later in LSMC algorithm.

We also generate a small grid of testing portfolios, denoting w_h, $h = 1, \cdots, H$. We assume the portfolio grid is bounded between 0 and twice of delta hedge $\Delta = \frac{B(T)}{B(\tau_2)}$, $w_h \in [0, 2\Delta]$.

As a benchmark, we choose $\Delta t = 0.25$, $N = 10,000$, $M = 40$ and $H = 5$.

A.2 LSMC

The problem is solved by means of simulation based dynamic programming. We outline the general recursion. We first explain the naive case. In addition to the generated grids of the state variables, we also generate N paths of gross bond returns based on point estimates of bond premia $\hat{\Lambda}_0$ and $\hat{\Lambda}_1$,

$$R_{t+\Delta t} = 1 + \left(r_t(1 - B(\tau_2)\sigma\hat{\Lambda}_1) - B(\tau_2)\sigma\hat{\Lambda}_0 \right) \Delta t - B(\tau_2)\sigma\sqrt{\Delta t} Z$$

The gross bond return at time t in trajectory i is denoted by $R_{i,t+\Delta t}$.

Time $T - \Delta t$. The problem at time step $T - \Delta t$ can be summarized by

$$\min_{w_{T-\Delta t}} \mathbb{E}\left[(1 - X_T)^+ \mid \mathcal{F}_{T-\Delta t} \right]$$

We first generate a time-dynamic wealth grid denoting $X_{j,T-\Delta} = \mathrm{FR}_j \times P(T - \Delta t, T)$, where $P(T - \Delta t, T)$ is the hypothetical bond price at time $T - \Delta$ maturating at T depending on $r_{T-\Delta t}$.

The reason we choose a fixed funding ratio grid instead of a fixed wealth grid is because our value function depends on terminal wealth X_T, keeping wealth grid fixed cannot guarantee us a reasonable range of terminal wealth, typically for long hedging horizon and when recursion approaches to time step 0.

Next we construct the realized terminal wealth under each simulated path (i, j),

$$X_{ijh,T}(\mathcal{F}_{T-\Delta t}) = X_{j,T-\Delta}\left((1 + r_{i,T-\Delta})(1 - w_h) + R_{i,T}w_h\right)$$

The realized objective function is $V_{ijh,T}(\mathcal{F}_{T-\Delta t}) = (1 - X_{T,ijh}(\mathcal{F}_{T-\Delta t}))^+$. Next, we regress the realized value function with a polynomial expansion of the state variables.

$$\begin{aligned} V_{ijh,T}\left(\mathcal{F}_{T-\Delta t}\right) = {} & \left(a_0 + a_1 r_{i,T-\Delta t} + a_2 \mathrm{FR}_j + a_3 \mathrm{IFR}_j + a_4 \mathrm{FR}_j \mathrm{IFR}_j\right) \\ & + \left(b_0 + b_1 r_{i,T-\Delta t} + b_2 \mathrm{FR}_j + b_3 \mathrm{IFR}_j + a_4 \mathrm{FR}_j \mathrm{IFR}_j\right) w_h \\ & + \left(c_0 + c_1 r_{i,T-\Delta t} + c_2 \mathrm{FR}_j + c_3 \mathrm{IFR}_j + a_4 \mathrm{FR}_j \mathrm{IFR}_j\right) w_h^2 + \epsilon_{T-\Delta t} \end{aligned} \tag{14}$$

where IFR is an indicator function, with $\mathrm{IFR}_j = (1 - \mathrm{FR}_j)^+$.

It is noticed that, along each path, the conditional variables are known hence the conditional expectation of the approximation regression is a quadratic function of the portfolio weight. Therefore, minimizing the conditional expectation boils down to minimizing a quadratic function of portfolio. We rewrite the conditional expectation of the value function as follows

$$\mathbb{E}\left[V_{ijh,T} \mid \mathcal{F}_{T-\Delta t}\right] = a_{ij} + b_{ij} w_{T-\Delta t} + c_{ij} w_{T-\Delta t}^2 \tag{15}$$

where

$$\begin{aligned} a_{ij}(\mathcal{F}_{T-\Delta t}) &= a_0 + a_1 r_{i,T-\Delta t} + a_2 \mathrm{FR}_j + a_3 \mathrm{IFR}_j + a_4 \mathrm{FR}_j \mathrm{IFR}_j \\ b_{ij}(\mathcal{F}_{T-\Delta t}) &= b_0 + b_1 r_{i,T-\Delta t} + b_2 \mathrm{FR}_j + b_3 \mathrm{IFR}_j + b_4 \mathrm{FR}_j \mathrm{IFR}_j \\ c_{ij}(\mathcal{F}_{T-\Delta t}) &= c_0 + c_1 r_{i,T-\Delta t} + c_2 \mathrm{FR}_j + c_3 \mathrm{IFR}_j + c_4 \mathrm{FR}_j \mathrm{IFR}_j \end{aligned}$$

The optimization problem boils down to solving for the root of Eq.(15), $w^*_{ij,T-\Delta t} = -\frac{b_{ij}}{2c_{ij}}$

if $c_{ij} > 0$.

Time $t = T - 2\Delta t, \cdots, 0$. We now discuss the general recursion of all any other point in time. Suppose we have optimized the hedging policy as of time $t + \Delta t$ onwards. The realized value function at time t is given by

$$V_{ijh,T}\left(\mathcal{F}_t\right) = \left[1 - X_{j,t}\left((1 + r_{i,t})(1 - w_h) + R_{i,t+\Delta t} w_h\right) \prod_{s=t+\Delta t}^{T-\Delta t} \left((1 + r_{i,s})(1 - w_s^*) + R_{i,s+\Delta t} w_s^*\right)\right]^+$$

where $X_{j,t} = \mathrm{FR}_j \times P(t, T)$. Next, we approximate conditional expectations $[V_{ijh,T} \mid \mathcal{F}_t]$ by functions of state variables and the testing portfolio

$$\mathbb{E}\left[V_{ijh,T} \mid \mathcal{F}_t\right] = \beta' f(X_{j,t}, r_{j,t}, w_h)$$

Last, we rewrite the conditional expectation of each path as a quadratic function of portfolio w_t and solving for the root.

$$\min_{w_t} \mathbb{E}\left[V_{ijh,T} \mid \mathcal{F}_t\right] = \min_{w_t} f(w_t)$$

We need to re-calculate the dynamic allocation $\frac{T}{\Delta t}$ times to retrieve the optimal decision now w_0 that we are interested in. Along this way, we have also obtained all other optimal portfolios for different hedging horizons $\tau < T$.

A.3 Closed-Form Nature's Choice

Next, we consider the robust case when bond premia are misspecified. A straight forward but more complicated method is to repeat the naive algorithm over a set of testing bond premia taking into uncertainty set constraint. Then we obtain a set of naive optimal policies under each path conditional on a testing bond premium. Then we could find the optimal nature's choice under each path either by grid search or using regression based method. However, neither methods are efficient nor accurate. If we use grid search method, we need to generate a fine grid of λ_0 and λ_1. This costs a huge computational memory. Unlike the portfolio weight, the quadratic approximation does not work in the nature's choice (λ_0, λ_1) due to low R^2.

A more efficient and accurate method we proposed is to approximate the expected shortfall by a function of bond return. Remark: this only works because we have simplified the problem so much that only Λ is uncertainty. We find that value function is monotonically decreasing with bond return. The approximation is sufficiently accurate since R^2 is nearly one. Therefore, we transform a maximization problem to a minimizing bond return problem, which can be analytically solved by linear programming.

To test the validity of the new methods, we first generate L pairs of testing market price of risk (MPR) denoting $(\lambda_{0,l}, \lambda_{1,l})$, with $l = 1, \cdots, L$ and $(\lambda_{0,l}, \lambda_{1,l}) \in \mathbf{S}$. Under each pair of testing MPR, we simulate N paths of gross bond returns over m periods. The gross return at time t in path i under testing MPR $(\lambda_{0,l}, \lambda_{1,l})$ is denoted by $R_{il,t+\Delta t}$

At step $T - \Delta t$, we construct realized value function $V_{ijl,T}$ based on mul-specification of bond returns $R_{il,T}$ using a random fixed portfolio weight, then we approximate the conditional value function as a function of $R_{il,T}$ and fund ratio

$$\mathbb{E}\left[V_{ijl,T}\left(w^*_{T-\Delta t}\right) \mid \mathcal{F}_{T-\Delta t}\right] = \beta_0 + \beta_1 R_{il,T} + \beta_2 R^2_{il,T} + \beta_3 FR_j + \beta_4 IFR_j$$

The goodness of fit is larger than 0.99 regardless of the portfolio weight we choose. We find that the conditional expectation of value function is a strict downward sloping convex function of $R_{il,t}$. This property holds for the entire hedging horizon if we recurse the algorithm backward till step 0.

Therefore, the global maximization problem boils down to minimizing bond return which is equivalent to minimizing bond premia since nature can only control over the drift term of bond diffusion process. Further, the ellipsoid uncertainty set $\mathbf{S}$ is convex, hence the minimum bond premium should locates on the ellipsoid.

The resulting nature's optimal decision depends only on the instantaneous spot rate r_t.

A.4 Robust LSMC

We can follow the naive LSMC algorithm to calculate the robust optimal portfolio expect that we first need to analytically solve for nature's choice at each backward step in time. Hence at time t, realized value function contains both optimal portfolios as well as nature's choice of bond premia from steps onwards

$$V_{ijh,T} = \left[1 - X_{j,t}\left((1+r_{i,t})(1-w_h) + R2^*_{i,t+\Delta t}w_h\right)\prod_{s=t+\Delta t}^{T-\Delta t}\left((1+r_{i,s})(1-w^*_s) + R^*_{i,s+\Delta t}w^*_s\right)\right]^+$$

where $R^*_{i,t+\Delta}$ indicates the optimal bond return at time t on path i with optimal MPR $\lambda^*_{0,t} + \lambda^*_{1,t}r_{i,t}$

A.5 Goodness of Fit

We investigate the accuracy of our algorithm by means of R^2. The R^2s of parametrization regression at step $T - \Delta t$ are higher than 0.995 for both naive and robust case. The goodness of fit by construction, has to decay backwards of time, because we are accumulating cross sectional information over each time period onward. The quality of the global quadratic approximation depends on R^2 at first approximation step $T - \Delta t$ and the speed of decaying. The longer the hedging horizon is, the lower R^2 will be at time 0. i.e. If $T \leq 20$ years, R^2 at time 0 is higher than 95%. If we set investment horizon extremely long, $T = 80$ years, R^2 drops to 0.83 at time step 0. This is still reasonably high, since the last step cross-sectional regression contains 320 time steps of cross-section information. Grid sizes or rebalancing frequency do not influence the goodness of fit.

References

Anderson, E., Hansen, L., Sargent, T., 2003. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. Journal of the European Economic Association 1 (1), 68–123.

Brandt, M., Goyal, A., Santa-Clara, P., Stroud, J., 2005. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. Review of Financial Studies 18 (3), 831.

Brennan, M. J., 1998. The role of learning in dynamic portfolio decisions. European Finance Review 1 (3), 295–306.

Broeders, D., Jong, F. d., Schotman, P., 2014. Interest rate models for pension and insurance regulation. Tech. rep., Working paper). Netspar Design Papers.

Carroll, C. D., 2006. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. Economics letters 91 (3), 312–320.

CEIOPS, D., 2010. Qis 5 technical specifications. European Commission, Annex to Call for Advice from CEIOPS on QIS5.

Cochrane, J. H., Saa-Requejo, J., 2000. Beyond arbitrage: Good-deal asset price bounds in incomplete markets. Journal of Political Economy 108 (1), 79–119.

Diris, B., 2011. Model uncertainty for long-term investors. Available at SSRN 1786587.

Duffee, G. R., 2002. Term premia and interest rate forecasts in affine models. The Journal of Finance 57 (1), 405–443.

Duffie, D., Kan, R., 1996. A yield-factor model of interest rates. Mathematical Finance 6 (4), 379–406.

Föllmer, H., Leukert, P., 2000. Efficient hedging: cost versus shortfall risk. Finance and Stochastics 4 (2), 117–146.

Hansen, L., Sargent, T., 2007. Robustness. Princeton University Press Princeton, NJ.

Koijen, R., Nijman, T., Werker, B., 2007. Appendix describing the numerical method used in'when can life-cycle investors benefit from time-varying bond risk premia?'. Available at SSRN 945720.

Koijen, R. S., Nijman, T. E., Werker, B. J., 2010. When can life cycle investors benefit from time-varying bond risk premia? Review of Financial Studies 23 (2), 741–780.

Maenhout, P., 2004. Robust portfolio rules and asset pricing. Review of Financial Studies 17 (4), 951.

Nelson, C., Siegel, A., 1987. Parsimonious modeling of yield curves. Journal of Business 60.

Peijnenburg, K., 2017. Life-cycle asset allocation with ambiguity aversion and learning. forthcoming Journal of Financial and Quantitative Analysis.

Vasicek, O., 1977. An equilibrium characterization of the term structure. Journal of Financial Economics 5 (2), 177–188.

Figure 1: Bond Premia

The figure presents the 20-year bond risk premia for different spot rate under different choices of λ_0 and λ_1. The ambiguity of market price of risk constrained by set $\mathbf{S}$ is shown in (6) with $\gamma^2 = 0.03$. Panel 1a presents the minimum and maximum values of bond premia within the feasible region for different r. The point estimate premia $\hat{\lambda}$ are located in between the two extremes. Panel 1b plots the feasible region of bond premia for different choices of λ_1 under the unconditional expectation of spot rate.

(a)

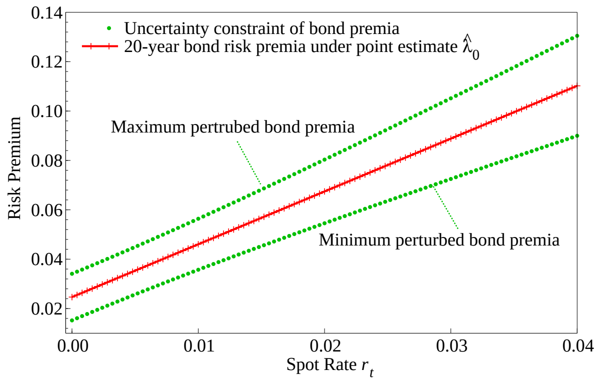

(b)

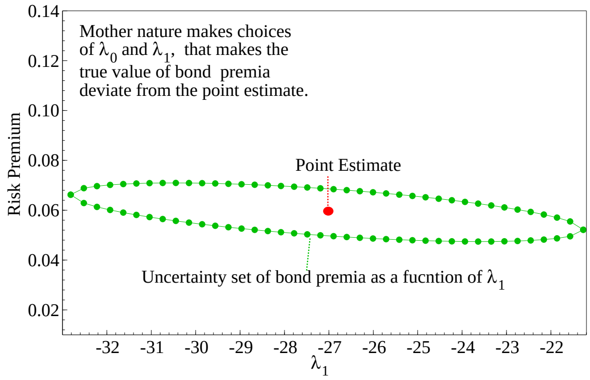

Figure 2: Nature's Decision on Bond Premium.

The figure presents nature's choice of bond risk premium against the naive estimate $-B(\tau)\sigma\hat{\lambda}$ under different bond maturities. The spot rate is equal to 2%

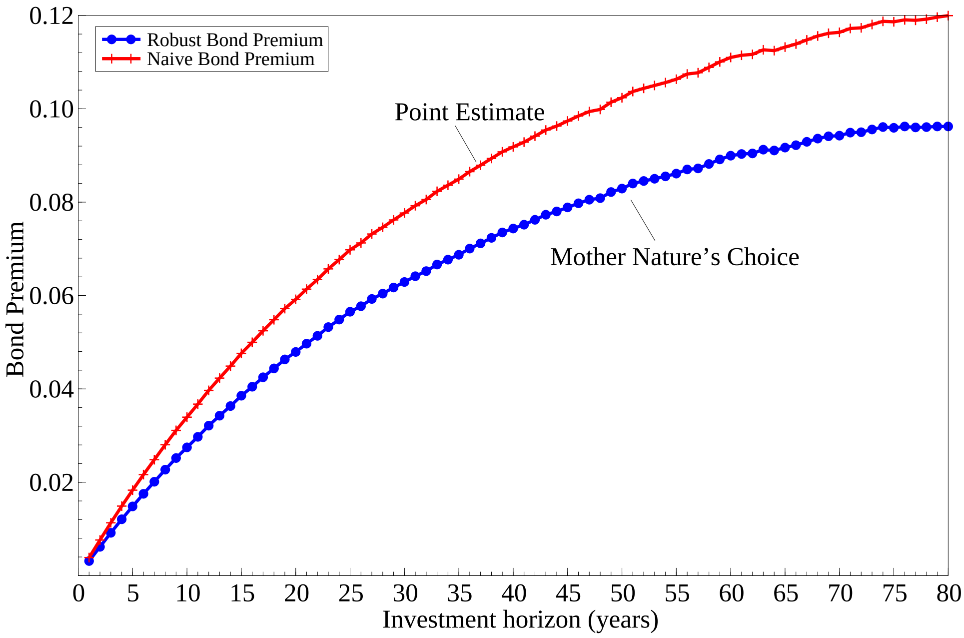

Figure 3: Optimal Bond Portfolio for Different Horizons
The figure plots 20-year bond weights against different investment horizons when current funding ratio is 80% (Figure 3a) and 110% (Figure 3b) for different hedging policies. The present spot rate is 2%. The weights are based on three different hedging policies: naive hedge, robust hedge and delta hedge. Results are based on 10,000 draws of predictive bond distribution.

(a) 80% Funding Ratio

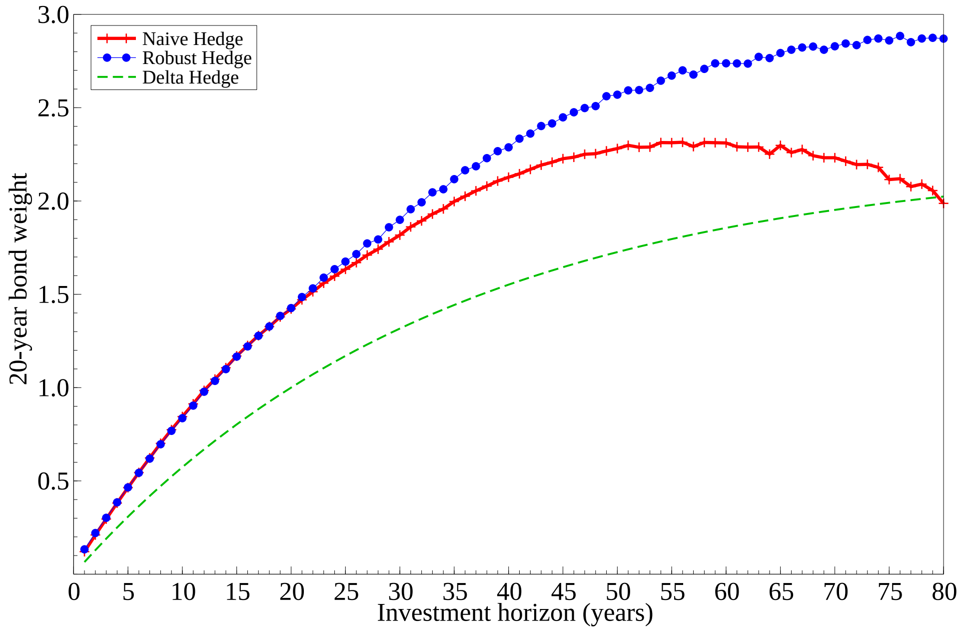

(b) 110% Funding Ratio

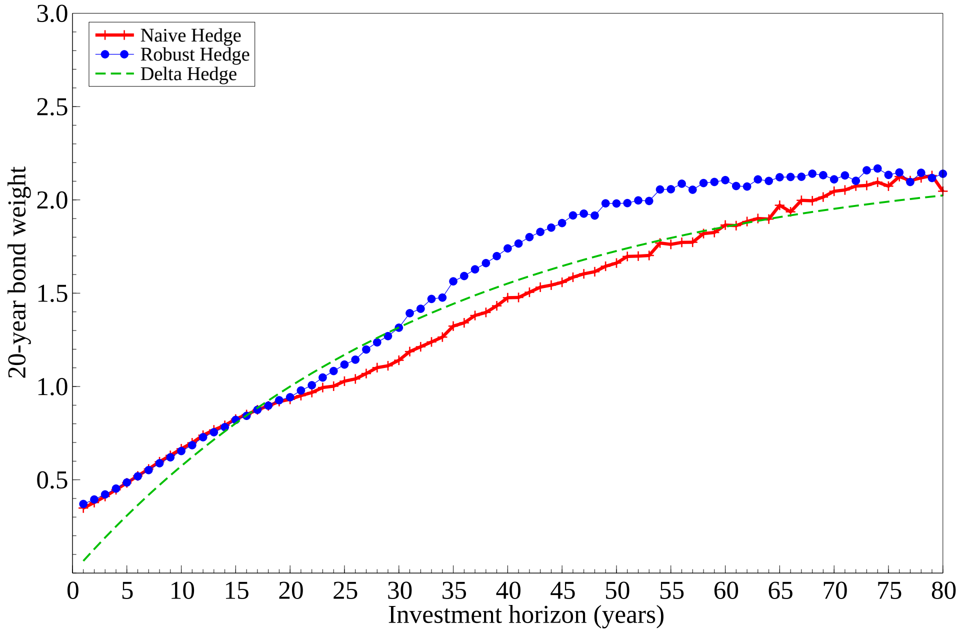

Figure 4: Optimal Portfolio under Different Funding Ratio (1)
The figure plots the optimal portfolio of 20-year bond as a function of investment horizon and current funding ratio with spot rate equal to 2%.

(a) Naive Policy

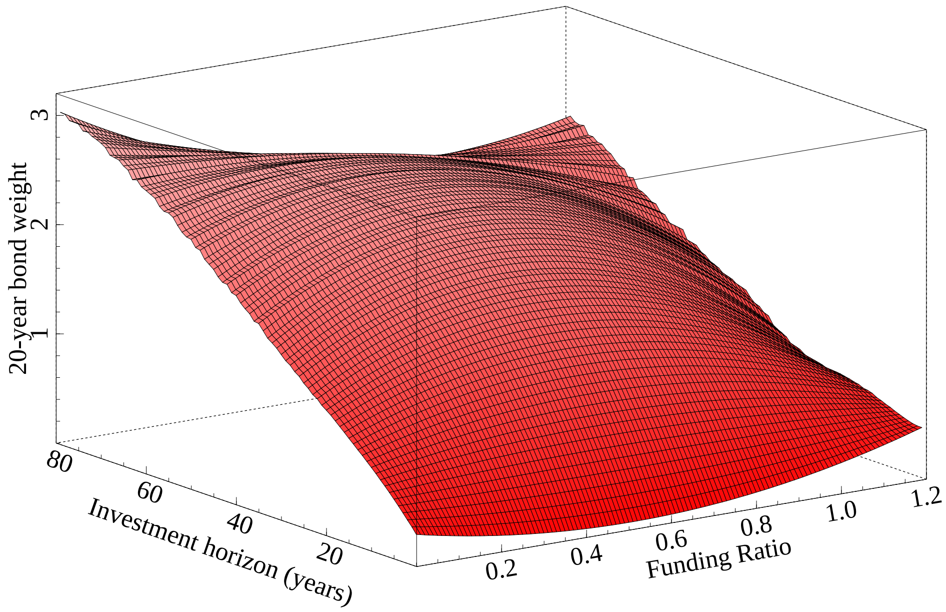

(b) Robust Policy

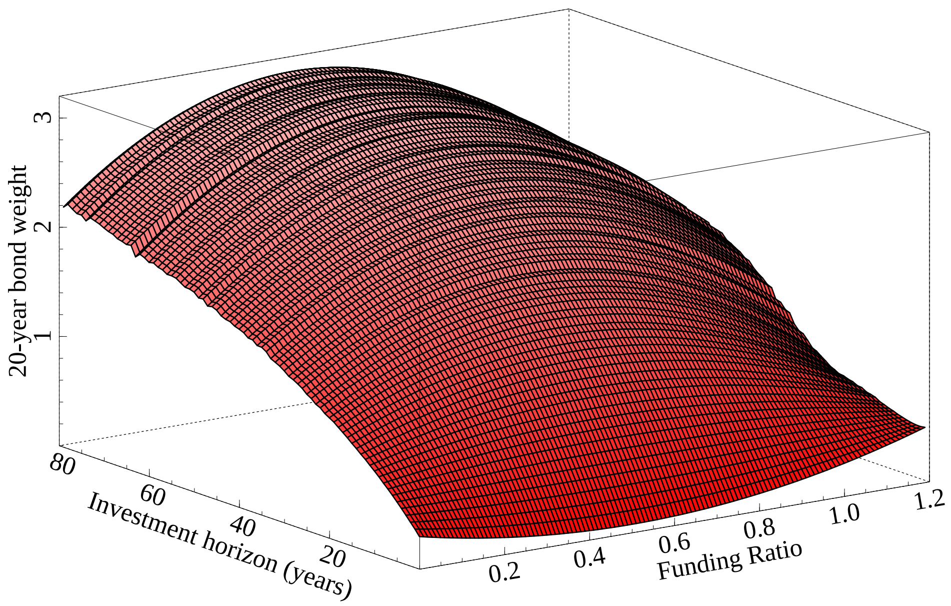

Figure 5: Optimal Portfolio under Different Funding Ratio (2).
The figure summarize the key insight of Figure 4 under three reasonable funding ratio levels.

(a) Naive Policy

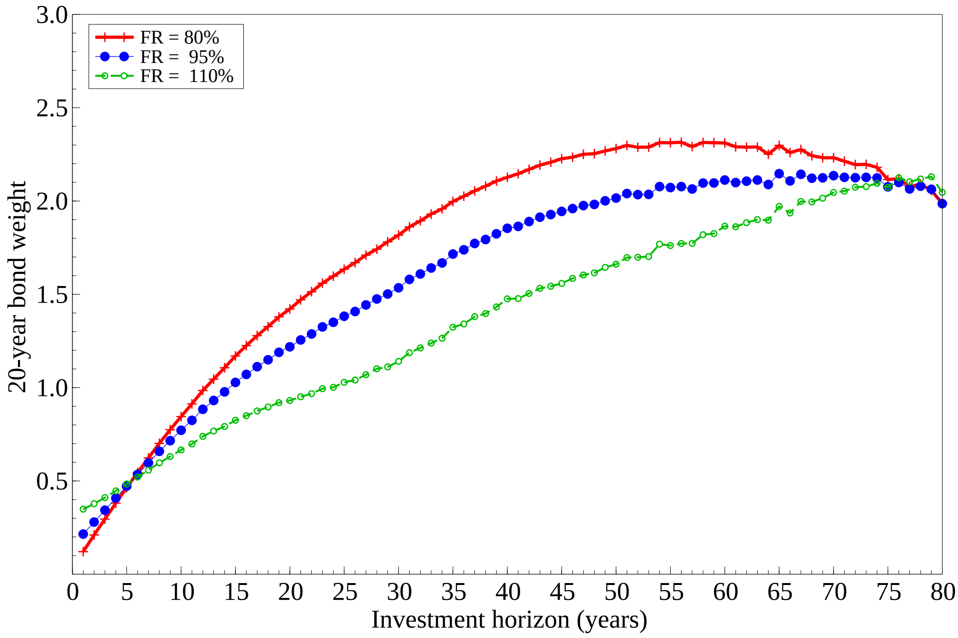

(b) Robust Policy

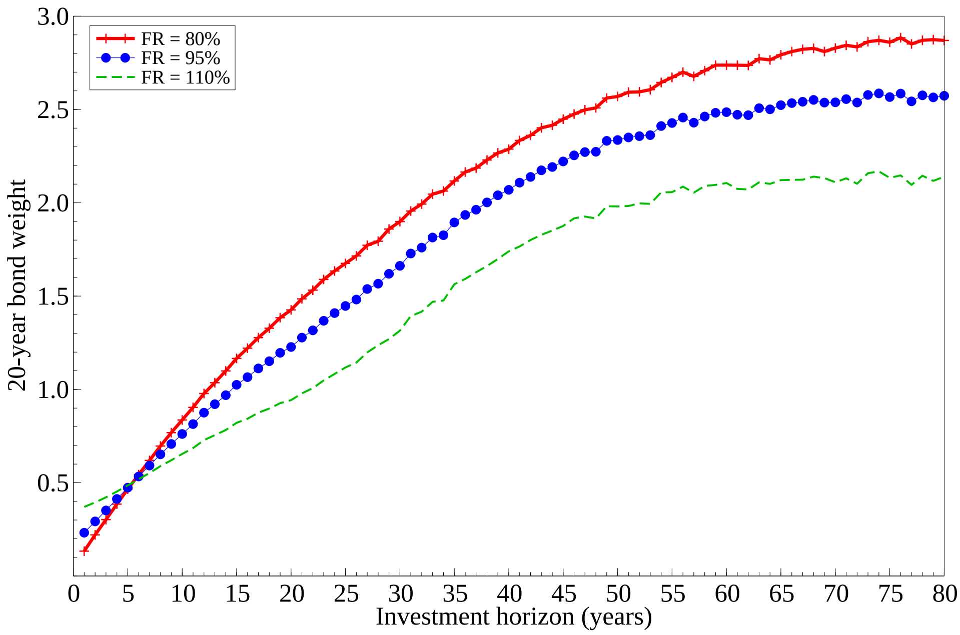

Figure 6: Robust Policy under Different Spot Rate.
The figure plots naive (Figure 6a) and robust (Figure 6b) optimal bond portfolio as a function of spot rate and hedging horizon when current funding ratio is 80%.

(a) Naive Policy

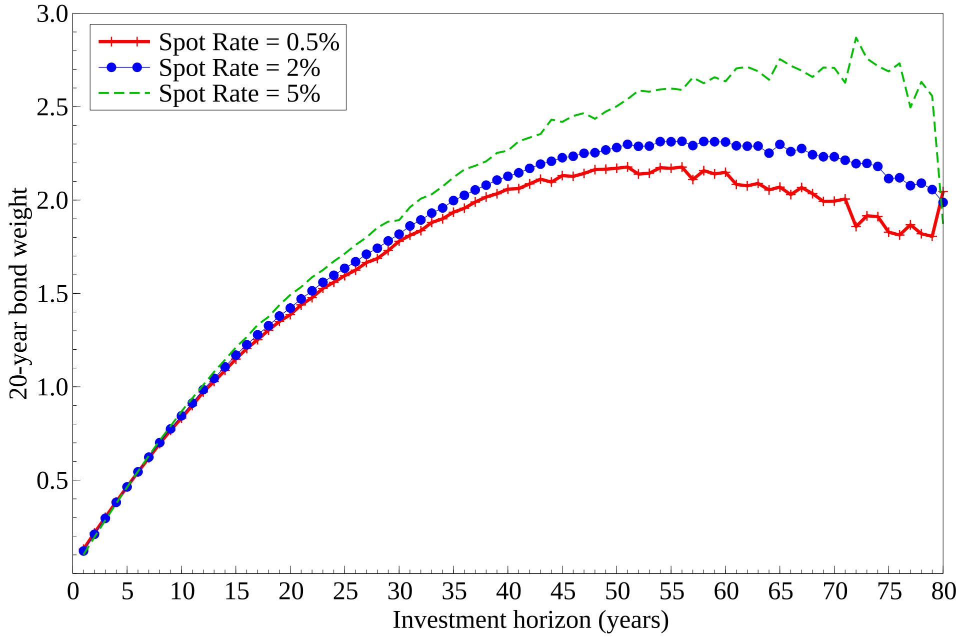

(b) Robust Policy

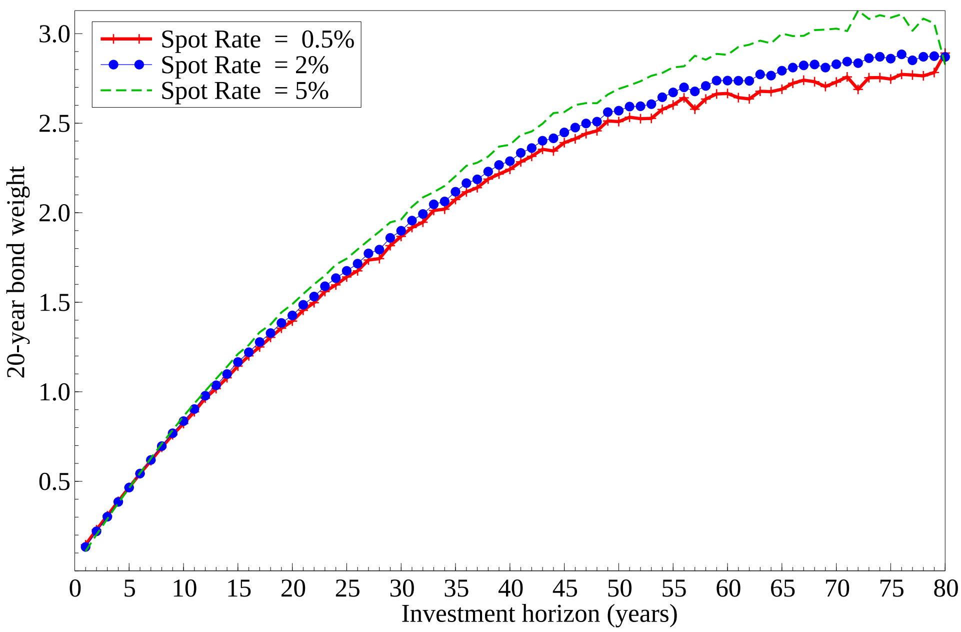

Figure 7: Optimal Expected Shortfall (1).
The figure plots the expected shortfall as a function of hedging horizon and funding ratio under naive (Figure 7a) and robust (Figure 7b) policies when spot rate is 2%.

(a) Naive Expected Shortfall

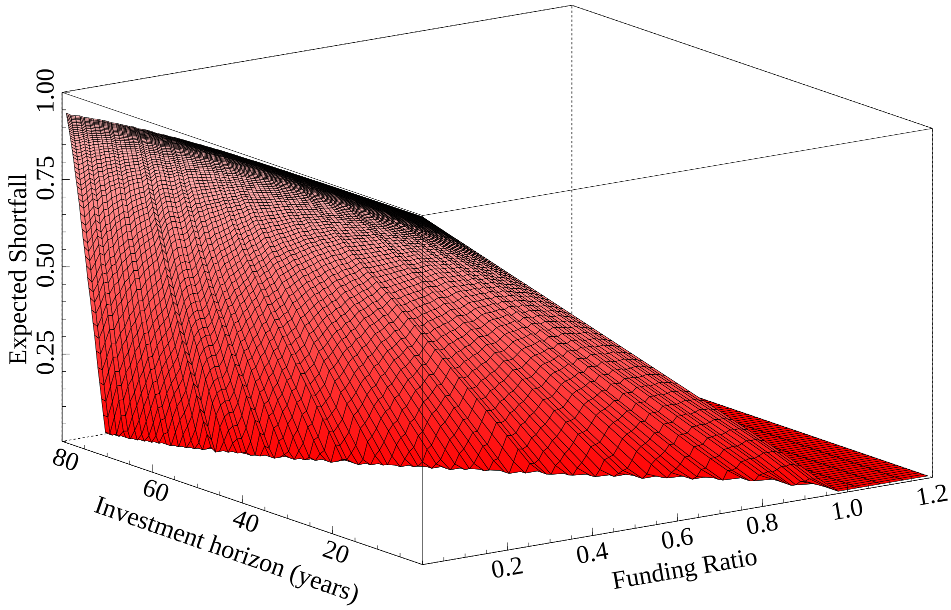

(b) Robust Expected Shortfall

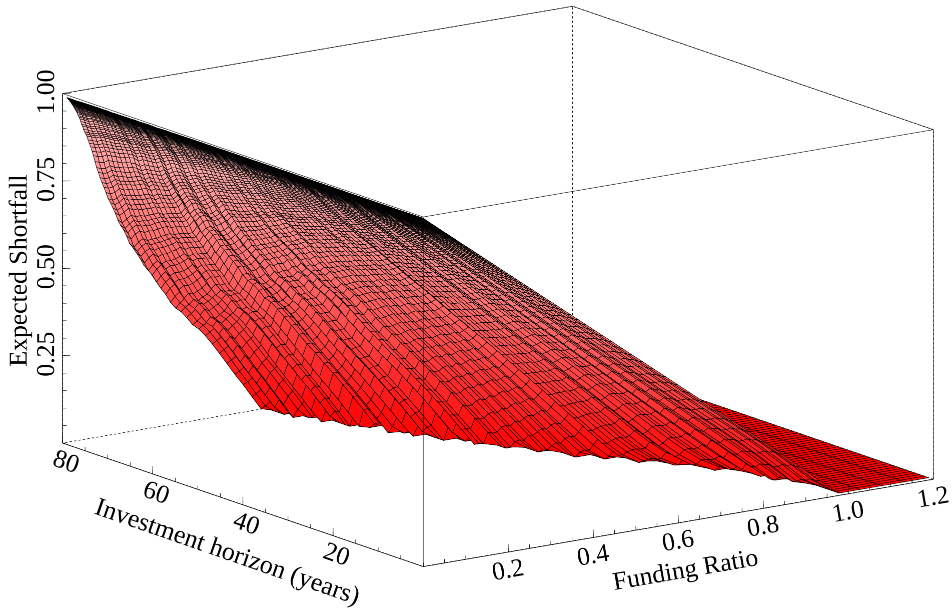

Figure 8: Optimal Expected Shortfall (2).
The figure plots the robust expected shortfall against the naive one as a function of funding ratio when hedging horizon is 40 years (Figure 8a) and when the horizon is 60 years (Figure 8b).

(a) T = 40

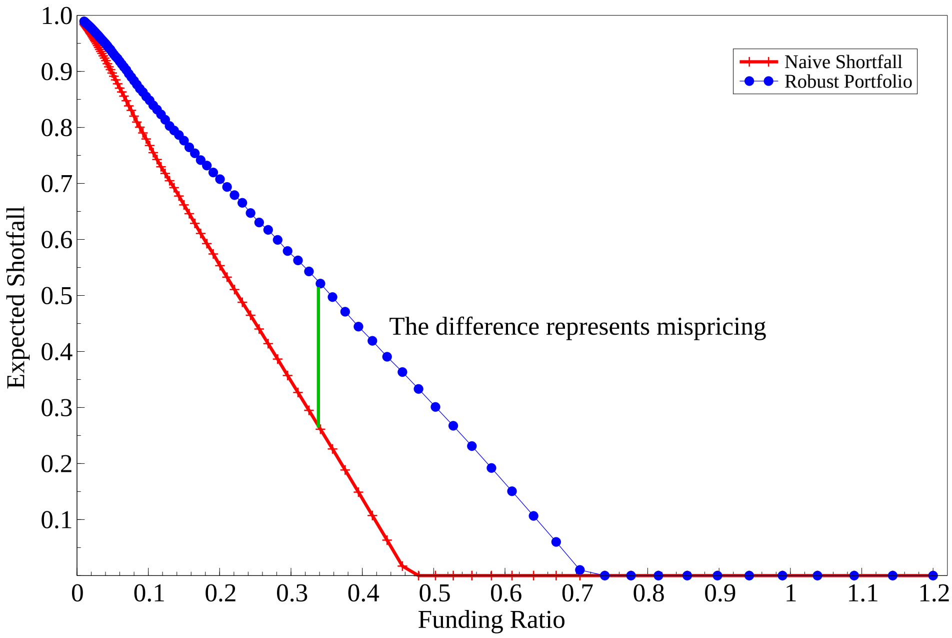

(b) T = 60

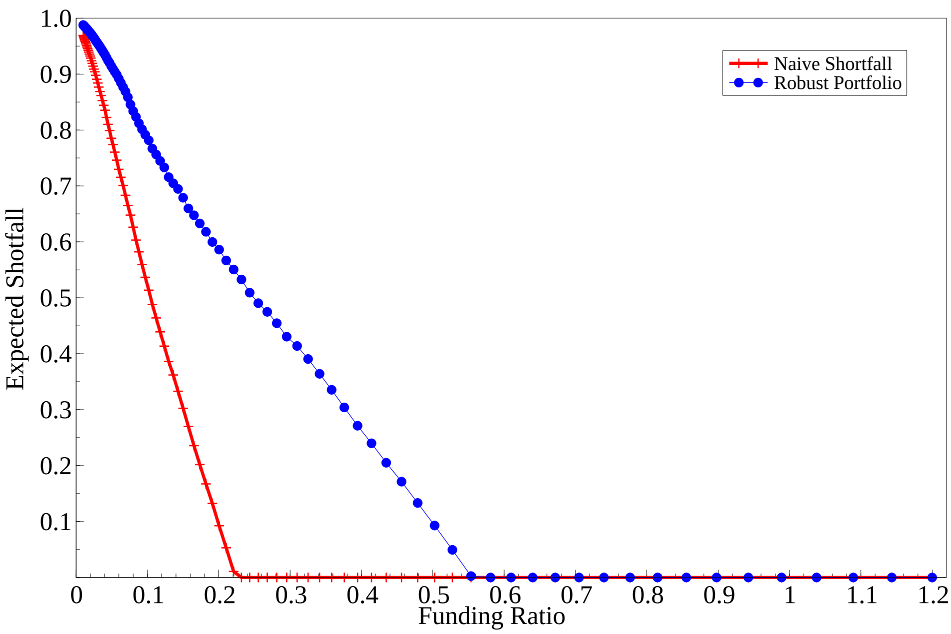

Figure 9: Robust Yield Curve
The figure plots policy-based yield curve again the benchmark Vasicek yield curve under different spot rate and shortfall target. The benchmark yield curve by definition has zero expected shortfall.