Robust hedging in incomplete markets

Citation for published version (APA):

Shen, S., Pelsser, A., & Schotman, P. (2019). Robust hedging in incomplete markets. *Journal of Pension Economics & Finance*, 18(3), 473-493. <https://doi.org/10.1017/S1474747218000069>

Document status and date:

Published: 01/07/2019

DOI:

[10.1017/S1474747218000069](https://doi.org/10.1017/S1474747218000069)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Robust hedging in incomplete markets

SALLY SHEN

Global Risk Institute, 55 University Avenue, Toronto, ON M5J 2H7, Canada and Network for Studies on Pensions, Aging and Retirement
(e-mail: sshen@globalriskinstitute.org)

ANTOON PELSSER AND PETER SCHOTMAN

Network for Studies on Pensions, Aging and Retirement and Department of Finance, Maastricht University, PO BOX 616, 6200 MD Maastricht, The Netherlands

Abstract

We considered a pension fund that needs to hedge uncertain long-term liabilities. We modeled the pension fund as a robust investor facing an incomplete market and fearing model uncertainty for the evolution of its liabilities. The robust agent is assumed to minimize the shortfall between the assets and liabilities under an endogenous worst-case scenario by means of solving a min–max robust optimization problem. When the funding ratio is low, robustness reduces the demand for risky assets. However, cherishing the hope of covering the liabilities, a substantial risk exposure is still optimal. A longer investment horizon or a higher funding ratio weakens the investor's fear of model misspecification. If the expected equity return is overestimated, the initial capital requirement for hedging can be decreased by following the robust strategy.

JEL CODES: G11, G13

Keywords: Model uncertainty, robust optimization, incomplete market, dynamic hedging, expected shortfall.

1 Introduction

After the 2008 global financial crisis, the performance of US pension funds has remained depressed. The poor solvency situation has been driven by a declining discount rate and also a fall in equity prices. Since 2012, funding ratios (asset values divided by projected benefit obligations) of the top 100 largest US corporate defined-benefit pension plans have not rebounded. More importantly, projected future funding ratios show a wide range of uncertainty for the next 2 years.¹ This raises the question of how to price and hedge downside risks when confronted with fragile beliefs about the likelihood of different funding ratio scenarios.

¹ Based on the data obtained from Milliman Pension Fund Index <http://us.milliman.com/pfi/us.milliman.com> for the Milliman 100 funding index ratio from the beginning of 2012 to July 2017 and projections from August 2017 to 2018.

Pricing and hedging pension or insurance liabilities faces two problems. First, the market is incomplete. Liability risks are typically not – or not actively – traded in the financial market. According to EIOPA (2011)², the two largest components of liability risks are market risk and life risk, which account for 67.4% and 23.7% of the diversified Basic Solvency Capital Requirement (SCR), respectively. However, these risks are not fully traded. Interest rate risk, one of the dominant market risks, is only partially traded in the financial market. Pension funds and insurance companies are often confronted with ultra-long-term commitments with maturities of more than 50 years. However, the longest dated government bonds even in developed markets such as the US, UK, and Canada are up to 30 years. In developing markets (such as Asia, Eastern Europe, and South America), long-term government bonds with maturities more than 10 years barely exist.

Life risk faces more serious market incompleteness problem, because mortality-linked securities in general have very low liquidity. For instance, longevity risk, the risk that insurers might live longer than anticipated, is the most important component of life risk. Turner (2006) shows that, in 2005, 2460 billion liabilities are associated with longevity risk in the UK. However, longevity risk had never been securitized until early 2000s. In the past decade, a limited number of mortality-linked products such as the longevity bond (see Blake and Burrows (2001)) have been proposed, while only a very small amount ($<1\%$)³ of longevity risk can be hedged.

The second problem concerns model parameter uncertainty in hedging liability risks. On the liability side of the balance sheet, longevity has been improving unprecedentedly in the past few decades in an unpredictable way (see, e.g. Benjamin and Soliman, 1993; McDonald *et al.*, 2006). An inaccurate mortality estimation makes pricing and hedging liability risks much more difficult and less reliable. Early work by Lee and Carter (1992) has been considered as the nucleus of modeling the dynamics of the mortality rate. Several plausible extensions of the Lee–Carter model, such as an incorporation of heterogeneous cell level (Li *et al.*, 2009), age-dependent factors (Cairns *et al.*, 2006), and structural changes (Coelho and Nunes, 2011; Van Berkum *et al.*, 2016), introduce substantial uncertainty on the trend in longevity and hence in the growth in liabilities.

On the other side of the balance sheet, the expected asset return is notoriously difficult to estimate from historical data. Merton (1980) argues that it is difficult to estimate expected returns from time series of realized stock return data. The standard deviation of the historical average return is $\sigma/\sqrt{T}$ where σ is the standard deviation of annual returns and T is the number of years. For example, if $T=100$ and $\sigma=16\%$, then the standard error of the equity premium is 1.6%, which leads to an approximate 95% confidence interval span of $6.3\%(\pm 1.96 \times 1.6\%)$. Although the interval shrinks with the square root of the sample size for estimation, it is difficult to maintain the same data generating process throughout the entire period. The investor is therefore exposed to estimation error in the expected asset returns.

² EIOPA (2011) is short for European Insurance and Occupational Pension Authority (2011). See https://eiopa.europa.eu/Publications/Reports/QIS5_Report_Final.pdf Report on the fifth quantitative impact study (QIS5) for Solvency II.

³ See Blake and Burrows (2001).

Doubts about the accuracy of the model makes an agent treat it as an approximation of an unknown true model. She wants her decision rules to work well over a set of models in the neighborhood of the approximating model. Our aim is to develop a hedging strategy for an agent who faces uncertainty about the expected return on the assets as well as uncertainty about the expected growth in liabilities. We adopted the robust control theory to deal with the fear of model uncertainty. The agent who worries about model misspecification looks for a prudent policy that is resilient to fragile beliefs about the likelihood of the state variables. Such decision rules are called robust policies. We introduced a robust hedging strategy along the lines of Hansen and Sargent (2007) to hedge undiversifiable downside risks.

The robust optimal hedging strategy that we propose takes both downside risks as well as market incompleteness into account for an agent who fears parameter uncertainty. The robust agent is assumed to minimize the shortfall between assets and liabilities under a statistically plausible worst-case scenario by means of solving a min-max robust optimization problem. The robust model includes three crucial elements. The first is downside risk, which we define as expected shortfall. In a static model, the expected shortfall between the assets and the liabilities can be valued as the payoff of an exchange option, which swaps the optimal value of the asset for the price of the liabilities. The second element is incomplete markets. We introduced two uncorrelated risk drivers in our model, one hedgeable and the other not hedgeable. The unhedgeable risk captures the incompleteness of the market. The asset market is exposed to hedgeable risk only, but the liability side is exposed to both types of risk. The third element is parameter misspecification. Following Anderson *et al.* (2003) and Maenhout (2004), we introduced drift distortions on the Brownian motions to represent parameter misspecification. These drift distortions perturb the true data generation process of approximate models. Economically, an additional drift on the Brownian motion can be understood as the unobservable market price of risk, which relates to Cochrane and Saa-Requejo (2000)'s concept of Good Deal Bounds (GDB). Technically, drift distortions measure the discrepancies between alternative probability distributions. A closely related idea appears in Cvitanić and Karatzas (1999), but in their model, liabilities only depend on the value of market instruments.

We solved the robust hedging problem in both a static and a dynamic environment. In both cases, the robust policy is more conservative than the naive policy. This result is in line with Brennan (1998). When the funding ratio is low, agents will increase the risk exposure to the stock market so as to gamble their way out of trouble (see also Ang *et al.*, 2013). The more the investor invests in the risky asset, the more she becomes exposed to estimation uncertainty. The robust agent is particularly afraid of a downside shock with the risky assets and hence she will put less wealth in the stock market compared with the agent who disregards the estimation uncertainty. We also found that for both the robust as well as the non-robust policy, the risky portion of the portfolio decreases with the hedging horizon when the funding ratio is low, and *vice versa* when the funding ratio is high. The impact of the preference for robustness depends on the hedging horizon as well as the funding ratio.

More importantly, we evaluated the robust policy by means of comparing its expected loss with the non-robust policy. The loss function is defined as the difference between the cost of hedging conditional on the estimated expected return and the true minimum cost. The benefits of a robust policy are twofold. First, the robust policy is less sensitive to the estimated parameters. Second, the robust policy has a lower hedging cost than a naive policy, under a range of alternative parameter values.

One strong assumption in this paper is that the investors only fear a subclass of model misspecification, namely the drift parameters of the state variables instead of the general model uncertainty problem. Similar to Maenhout (2004), we reduced the general model uncertainty problem to a first-moment parameter uncertainty problem.

Another strong assumption in our work is that the investors do not engage in any learning. Brennan (1998) incorporates learning with parameter uncertainty and finds that after learning, high-risk-averse investors are more conservative with their investment, but low-risk-averse investors allocate more wealth on risky assets. Wang (2009) incorporates income growth rate uncertainty with Bayesian learning for a consumption-saving and optimal portfolio choice problem and finds that learning induces additional precautionary saving.

Numerous studies deal with asset allocation problem for pension plans. Seminal work by Sharpe and Tint (1990) develops a surplus management approach in which funds care about assets minus liabilities. Detemple and Rindisbacher (2008) extend this framework to a dynamic setting. Ang *et al.* (2013) add an additional penalty function at the mean-variance framework of Sharpe and Tint (1990). The penalty function, which is the shortfall between the asset and liabilities, is the same as our objective function, to which we add model uncertainty.

A related study on model uncertainty by Garlappi *et al.* (2006) considers a mean-variance portfolio choice of a robust investor who has imperfect information on the expected return. They use multi-prior approach advocated by Gilboa and Schmeidler (1989), and they also find that allowing for parameter uncertainty reduces the portfolio weights on risky assets over time. Luo (2016) considers both model uncertainty and state uncertainty under decision-making. State uncertainty refers to incomplete information about the true value of the state due to sluggishness of the market. In this paper, we do not deal with state uncertainty, but the likelihood over the state variables.

2 Model

We considered a continuous-time incomplete market with a finite trading horizon $[0, T]$. The risk is modeled by a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$, on which are defined two uncorrelated risk factors, a hedgeable risk W_{1t} and an unhedgeable risk W_{2t} . Both W_{1t} and W_{2t} are univariate standard Brownian motions and we considered $\{\mathcal{F}_t : t \in [0, T]\}$ as the completion of the filtration generated by W_{1t} and W_{2t} . A hedgeable risk means we can replicate the payoff of this kind of risk perfectly. The payoff for an unhedgeable risk is not replicable because it is not traded.

2.1 Asset and liability model

On the asset side, we have a risk-free money-market account B_t , which earns a deterministic risk-free rate of interest r , so $dB_t = rB_t dt$. We also have a stock market. The stock price follows a geometric Brownian motion process $dS_t = \mu S_t dt + \sigma S_t dW_{1t}$. The agent can only invest in the money-market account and the stock market. Denote the value of the assets at time t by A_t . The investor puts an amount $w_t A_t$ in the stock market at time t . The remaining part of the assets $(1 - w_t)A_t$ is put into the money-market account. The asset diffusion process follows as

$$dA_t = (r + w_t(\mu - r))A_t dt + w_t \sigma A_t dW_{1t}, \quad (1)$$

where w_t is the possibly time-varying hedging strategy. We do not set a constraint on w_t , therefore short positions are allowed.

The liability is exposed to both hedgeable risk W_{1t} and unhedgeable risk W_{2t} . We assumed that the diffusion process of the liability L_t follows an exogenously given geometric Brownian motion with constant drift term and constant volatility,

$$dL_t = aL_t dt + bL_t \left(\rho dW_{1t} + \sqrt{1 - \rho^2} dW_{2t} \right), \quad (2)$$

where a is the drift of the liability and b is its volatility. The non-traded risk driver, dW_{2t} , represents the incomplete part of the market. We introduced a correlation parameter $\rho \in [-1, 1]$ between asset risk and liability risk. It controls the risk exposure to W_{2t} of the liability. If $\rho = \pm 1$, then the non-traded risk W_{2t} disappears from the liability side. The liability in this case can be perfectly hedged by a replicating portfolio. We are interested in the case when ρ is strictly between -1 and 1 .

2.2 Robust asset and liability model

We used the Hansen and Sargent's (2007) framework to integrate the preference for robustness to the asset-liability models (1) and (2). With a preference for robustness, the agent treats (1) and (2) as an approximate model for the unknown true state evolution of A_t and L_t . We limited the parameter uncertainty to the drift terms μ and a only, and assumed that the volatilities σ and b are known. The approximate model only provides an estimated value of the drift terms, but the growth rate of liabilities and the expected return are imprecisely estimated and subject to estimation error. However, the constant volatility parameter σ can potentially be estimated using high-frequency observations, and is therefore not subject to parameter estimation error.

In the Hansen and Sargent's framework, the robust model contains an unknown drift term on the Brownian motion. In our case, the Brownian motions dW_{1t} and dW_{2t} in (1) and (2) are replaced by $dW_{1t} + \lambda_{1t} dt$ and $dW_{2t} + \lambda_{2t} dt$. The two drift terms λ_{1t} and λ_{2t} are defined as two perturbation time-series processes that quantify the misspecification of the underlying model. The values of λ_{1t} and λ_{2t} shift the mean distribution of the asset and the liability diffusion process by a unit of $w_t \sigma \lambda_{1t}$ and $b \rho \lambda_{1t} + b \sqrt{1 - \rho^2} \lambda_{2t}$, respectively. Hence, they specify a set of alternative measures referring to different specifications of the stochastic process known as a Girsanov kernel. The misspecified expected return also generates an error in the market price of

risk. The perturbed evolution of the state variables is given by:

$$dA_t = (r + w_t(\mu - r))A_t dt + w_t \sigma A_t (dW_{1t} + \lambda_{1t} dt), \quad (3a)$$

$$dL_t = aL_t dt + bL_t \left(\rho(dW_{1t} + \lambda_{1t} dt) + \sqrt{1 - \rho^2}(dW_{2t} + \lambda_{2t} dt) \right). \quad (3b)$$

The perturbation of the model is bounded by an uncertainty set $\mathbb{S}$. The larger the uncertainty set $\mathbb{S}$, the more pessimistic the agent is about the accuracy of the underlying model. To describe the uncertainty set, we introduced some additional notation. Let δ be the vector of the estimated drift terms,

$$\delta = \begin{pmatrix} \mu \\ a \end{pmatrix}$$

and let δ_0 be the true drift. Then $\delta - \delta_0$ is the estimation error,

$$\delta - \delta_0 = \begin{pmatrix} \sigma & 0 \\ b\rho & b\sqrt{1 - \rho^2} \end{pmatrix} \begin{pmatrix} \lambda_{1t} \\ \lambda_{2t} \end{pmatrix} = \Gamma \lambda_t.$$

The estimation error $\delta - \delta_0$ is asymptotically normal with mean zero and covariance matrix (Σ/N) , where N is the length of a (hypothetical) sample used for estimation and

$$\Sigma = \Gamma \Gamma' = \begin{pmatrix} \sigma^2 & b\rho\sigma \\ b\rho\sigma & b^2 \end{pmatrix}.$$

We obtained the uncertainty set based on the property that $(\delta - \delta_0)'(\Sigma/N)^{-1}(\delta - \delta_0)$ is a χ^2 distribution with two degrees of freedom, $\chi^2(2)$. Denoting the critical value at α significance level as CV_α , we then have a probability of $1 - \alpha$ that

$$(\delta - \delta_0)' \Sigma^{-1} (\delta - \delta_0) \leq \kappa^2, \quad (4)$$

where $\kappa^2 = (CV_\alpha/N)$. Equation (4) provides a natural boundary of the perturbation parameters. Simplifying (4) further, we get

$$(\Gamma \lambda_t)' (\Gamma \Gamma')^{-1} (\Gamma \lambda_t) = \lambda_t' \lambda_t \leq \kappa^2.$$

Hence our uncertainty set is as follows,

$$\mathbb{S} = \{\lambda_{1t}, \lambda_{2t} | \lambda_{1t}^2 + \lambda_{2t}^2 \leq \kappa^2\}. \quad (5)$$

Our uncertainty set has a circular shape in λ_t space centered by zero. Given the estimates δ a credibility region for the true value, δ_0 can be constructed as

$$\delta_0 \in \{\delta + \Gamma \lambda_t | \lambda_t \in \mathbb{S}\}. \quad (6)$$

The true drift term δ_0 is constrained by an ellipsoid uncertainty set centered by δ and it can be at any point within this set. The size of the uncertainty set depends on the significance level α and the hypothetical sample size N . If the agent has infinite observations, then the uncertainty set shrinks to the point estimate δ .

Our stylized uncertainty set is related to the GDB proposed by Cochrane and Saa-Requejo (2000). Equation (4) can be understood as the GDB constraint in which we put a limit on the unobservable part of the market price of risks, λ_{1t} and λ_{2t} . The uncertainty set we proposed differs from the GDB in the way in which our uncertainty set is derived from the econometric estimation error. The uncertainty

set parameter κ depends only on the statistical quantities, α and N . However, the GDB method is inspired by an economic belief that the total market price of risk in an incomplete market has to have limits.

2.3 Robust optimization problem

Utility is defined as a function of the terminal value for A_T and L_T at the terminal date T . The optimal hedging strategy maximizes the utility function $\mathbb{E}[U(A_T, L_T)|\mathcal{F}_t]$. As a benchmark, we defined the naive policy w_{na} as the hedging strategy that does not consider model misspecification.

The uncertainty averse agent looks for a robust hedging policy that works well over a set of models. The robust hedging policy is defined as

$$\max_{w_t} \min_{\lambda_{1t}, \lambda_{2t} \in \mathbb{S}} \mathbb{E}[U(A_T, L_T)|\mathcal{F}_t]. \quad (7)$$

The max–min optimization problem is a two-player zero-sum game, see Anderson *et al.* (2003). This is a sequential game between the decision-maker and a malevolent nature. Player 1, the robust agent moves first by choosing investment decisions to maximize the utility function at time t , and then player 2 (the imaginary nature) picks the worst state of nature for player 1 by making an instantaneous choice of λ_{1t} and λ_{2t} , given player 1's choice. In other words, the agent is maximizing while nature is minimizing.

3 Static robust optimization

In the following two sections, we will show how to solve the robust optimization problem and how the robust solution differs from the naive one, and also how we can benefit from the robust decision. We started with the relatively simple static case, where both agent and nature only make decisions now at $t=0$ without rebalancing until the terminal date T . The static case is technically easy to solve, but still provides us with some intuition about the robust policy. However, the static solution may not be optimal. With a dynamic solution both w_t and $\lambda_{1t}, \lambda_{2t}$ are time-series processes.

Given the information at time T , our hedging strategy is defined over the hedging error $L_T - A_T$ at a predetermined time T . Our utility function takes the form of the shortfall risk $U(A_T, L_T) = -[L_T - A_T]^+$, which specifies the downside risk on the liability shortfall. The lower the shortfall risk, the higher the agent's utility will be. The naive optimization problem is given by

$$\min_{w_t} \mathbb{E}[(L_T - A_T)^+] \quad (8)$$

and the robust optimization is

$$\min_{w_t} \max_{\lambda_{1t}, \lambda_{2t} \in \mathbb{S}} \mathbb{E}[(L_T - A_T)^+]. \quad (9)$$

In the static case, the order of the two players is interchangeable. According to the saddle-point existence Theorem mentioned in Delbaen (2002) and Rockafellar (1976), the optimal solution of (9) is a saddle point, since both control variables

are constrained by a convex set and the value function is bilinear. Hence (9) and its dual problem $\max_{\lambda_{1t}, \lambda_{2t}} \min_{w_t} \mathbb{E}[(L_T - A_T)^+]$ have the same optimal solution.

3.1 Static solution

To facilitate calculation, let

$$\mu_S = \mu + \sigma\lambda_1, \quad (10a)$$

$$\mu_A = r + w(\mu_S - r), \quad (10b)$$

$$\mu_L = a + b\rho\lambda_1 + b\sqrt{1 - \rho^2}\lambda_2, \quad (10c)$$

represent the drift terms of the stock market, the asset and the liability, respectively. Note that μ_S depends on λ_1 ; μ_A depends on both w and λ_1 ; and μ_L depends on λ_1 and λ_2 .

In the static case, our criterion function $\mathbb{E}[(L_T - A_T)^+]$ is very similar to the value of an ‘exchange option’, which exchanges one asset for another at time T . This type of option has been valued in Margrabe (1978). The problem in our case is more complicated, because we are in an incomplete market, which means the equivalent martingale is not unique, or in other words, the so-called risk-neutral $\mathbb{Q}$ measure is not unique, but depending on λ_1 and λ_2 .

There are many ways to solve this static criterion function. We used the change of probability measure technique. The analytical solution of our objective function under the static case is given by

$$\mathbb{E}[(L_T - A_T)^+] = \bar{L}\Phi(-d_2) - \bar{A}\Phi(-d_1) = \bar{L}(\Phi(-d_2) - \bar{C}\Phi(-d_1)), \quad (11)$$

where

$$\bar{L} = L_0 \exp(\mu_L T),$$

$$\bar{A} = A_0 \exp(\mu_A T),$$

$$\bar{C} = C_0 \exp[(\mu_A - \mu_L)T],$$

and

$$d_1 = \frac{\ln \bar{C} + \frac{\sigma_C^2}{2} T}{\sigma_C \sqrt{T}},$$

$$d_2 = d_1 - \sigma_C \sqrt{T},$$

where $C_0 = (A_0/L_0)$ is the current funding ratio. The function Φ is the standard normal distribution function. If the funding ratio is less than one, the fund is facing a solvency risk. For given λ_1 and λ_2 , the optional hedge disregarding the preference for robustness is the solution of the first-order condition for maximizing $\mathbb{E}^L[(1 - C_T)^+]$ with respect to w ,

$$\frac{\partial[\Phi(-d_2) - \bar{C}\Phi(-d_1)]}{\partial w} = -\Phi(-d_1)\bar{C}(\mu - r)T + \bar{C}\phi(d_1)\sqrt{T} \frac{w\sigma^2 - b\rho\sigma}{\sigma_C} = 0, \quad (12)$$

where function ϕ denotes the standard normal density function. Note that $-\Phi(-d_1)$ is the δ of the Black–Scholes (BS) put-option that is always less than zero, and

$\bar{C}\phi(d_1)\sqrt{T}$ denotes the v of the BS option that is always positive. Therefore, we see from (12) that the optimal w strikes a balance between the ‘ δ effect’ that reduces the value of the option and the ‘ v effect’ that increases the value of the option. There is a special case when $\mu = r$ where the ‘ δ effect’ disappears and the optimal w is then given by the minimum variance solution $w = (b\rho/\sigma)$.

3.2 Static robust portfolio choice

Based on the analytical solution (11), we solved the static robust optimization problem numerically. As a benchmark scenario, we assumed $\mu = 0.04$, $\sigma = 0.16$, $r = 0$, $a = 0$, $b = 0.1$, $\rho = 0.5$. We assumed that the stock return μ is higher than the liability return a . As we discussed in Section 2.2, the uncertainty set parameter κ depends on the significance level α and the sample size N . Hence, it is fixed and state variable independent. For a significance level $\alpha = 0.05$, the corresponding χ^2 value with 2 degrees of freedom is 5.99. The choice of κ is also based on an implicit assumption that the risk premium $(\mu_S - r/\sigma)$ is always positive, which means $(\mu - r/\sigma) + \lambda_1 > 0$. Given the uncertainty set $\mathbb{S}$, the absolute value of λ_1 is bounded with $|\lambda_1| \in [-\kappa, \kappa]$, hence κ has to satisfy the condition that $\kappa \leq ((\mu - r)/\sigma) = 0.25$ in order to guarantee a positive risk premium. Therefore, we set $\kappa = 0.25$ for the benchmark scenario. Alternatively, using the significance level, the sample size N has to be larger than 96 years so as to satisfy this implicit assumption.

In Figure 1, we show the static optimal portfolio choice at time $t = 0$ as the function of the current funding ratio, C_0 . When there is underfunding, the robust and naive policies differ. Both take substantial risk betting on the chance to meet the liability, but the robust portfolio is more conservative than the naive one. For example, if the current funding ratio equals 80%, then the robust policy will reduce the risky asset exposure by approximately 6% relative to the naive policy. The robustness effect diminishes if C_0 goes up. The two curves converge to the minimum-variance hedging ratio $((b\rho)/\sigma) = 0.3125$ if C_0 is sufficiently large. The resulting volatility becomes $b(1 - \rho^2)$, which is the unhedgeable part of the liability risk. Also, this position neutralizes the λ_1 effect such that the misspecification of asset return does not influence the performance of hedges. Therefore, the robust hedges are not always more conservative than naive hedges. If the fund is already balanced – or even overfunded with $C_0 \geq 1$ – the two policies are almost identical.

The decision of nature is displayed in Figure 2. We showed λ_1 and λ_2 as a function of the present funding ratio C_0 . To facilitate the comparison, we put the two perturbations in one graph. We found that λ_1 is negative at any funding ratio level but is close to zero when C_0 is high; λ_2 is always positive and converges to κ . We also found that the optimal choice of λ_1 and λ_2 is always on the circle $\lambda_1^2 + \lambda_2^2 = \kappa^2$, which means the worst-case scenario is always at the boundary of the uncertainty set.

Figure 2 shows that a negative λ_1 and a positive λ_2 lead to the worst-case scenario. This is because the agent is afraid that the true expected asset return is lower than the estimated value, and the true liability return is higher than the estimated result. The resulting negative λ_1 represents the fear of an overestimated asset return. Hence, the absolute value of λ_1 is increasing with the exposure to the stock market, w . We

Figure 1. Static optimal portfolio choice. This figure compares the robust and naive static optimal hedging policies. The investor makes an investment decision at time $t=0$ with given current funding ratio C_0 so as to minimize the expected shortfall at time period T . The naive policy relies completely on the estimation parameters. The robust policy takes the parameter uncertainty into consideration and insures against the worst-case scenario. The horizontal axis depicts the present funding ratio. The results are based on the benchmark estimation parameters $\mu = 0.04$, $\sigma = 0.16$, $r = 0$, $a = 0$, $b = 0.1$, $\rho = 0.5$, $\kappa = 0.25$, and $T = 5$.

know from Figure 1 that risk exposure and the funding ratio are negatively related. The lower the funding ratio, the higher the risk exposure will be and therefore the more negative the value of λ_1 will be. In contrast, if the funding ratio is sufficiently high, both the λ_1 penalty as well as the weight in the risky asset are smaller. The penalty term λ_1 also plays a role in the liability return. A negative λ_1 can benefit the agent by reducing the expected liability return. To capitalize on the fear of an increase in the liability return, nature chooses a positive λ_2 so as to compensate for the negative effect from λ_1 and to increase the liability growth, making the liability more costly.

We further examined how the perturbation terms impact the expected returns. Figure 3 displays both the naive and robust mean rates of the stock return and the liability return as functions of C_0 . Without the preference for robustness, both drift terms are constant. However, if the investor is aware of the model misspecification, the perturbed expected stock return is dragged down by $|\sigma\lambda_1|$ due to the negative impact of λ_1 . Despite the mixed sign of λ_1 and λ_2 , the worst-case liability drift is pushed up by $|b\rho\lambda_1 + b\sqrt{1-\rho^2}\lambda_2|$ since the positive effect of λ_2 dominates the drift distortion. In general, the robust policy differs from the naive one in the sense that the robust agent requires an additional guarantee on top of the naive contract in order to neutralize the estimation error. In other words, the robust policy needs more capital to hedge downside risks.

Figure 2. Static optimal perturbations λ_1 and λ_2 . This figure depicts the optimal λ_1 and λ_2 as functions of the present funding ratio C_0 under the benchmark scenario with $\mu = 0.04$, $\sigma = 0.16$, $r = 0$, $a = 0$, $b = 0.1$, $\rho = 0.5$, $\kappa = 0.25$, and $T = 5$. Nature makes decisions of λ_1 and λ_2 at time 0 under the constraint $\lambda_1^2 + \lambda_2^2 \leq \kappa^2$ so as to maximize the expected shortfall at period T .

3.3 Policy evaluation

The robust policy is less sensitive to the parameter misspecification. In this section, we will show how and when the agent can benefit from the robust policy. Let $Q(w, \delta)$ be the cost of hedging following a particular policy w , where δ is the assumed value of the drift parameters. In our case, the cost of hedging is defined by

$$Q(w, \delta) = \mathbb{E}[(L_T - A_T)^+ | w, \delta]. \quad (13)$$

The optimal hedging policy has a cost $q(\delta) = \min_w Q(w, \delta)$ for given δ . Let δ_0 be the true value of δ , and denote $q(\delta_0)$ as the minimum hedging cost when the investor implements the associated optimal hedging policy w_0 under the true value δ_0 . Any other alternative hedging policies w_a ($w_a \neq w_0$) have higher expected shortfall.

Define the loss function $K(w_a | \delta_0)$ as the difference between the cost of hedging following a suboptimal policy w_a and the true minimum cost. The ‘cost of hedging’ here is defined as the initial wealth required to obtain a particular level of expected shortfall, denoted $Q(w_a, \delta_0)$, which gives the loss function

$$K(w_a | \delta_0) = Q(w_a, \delta_0) - q(\delta_0). \quad (14)$$

If $\delta \neq \delta_0$, the agent is facing estimation error, therefore $w_a \neq w_0$ and $K(w_a | \delta_0) > 0$.

The agent does not know the true value of the drift terms δ_0 . Given the estimated drift terms δ , she can choose between two alternative hedging policies, a robust policy w_{rob} and a naive policy w_{na} . At the benchmark scenario when the present funding ratio $C_0 = 80\%$, the solutions are $w_{rob} = 0.81$ and $w_{na} = 0.87$. When $C_0 = 90\%$, we found that $w_{rob} = 0.67$ and $w_{na} = 0.69$. The robust policy will perform better than the

Figure 3. Mean rate of stock and liability return with and without the preference for robustness. This figure displays the expected stock and liability returns before and after considering parameter uncertainty as functions of the present funding ratio. Panel 3a comparing the robust stock drift $\mu_S = \mu + \sigma\lambda_1$ with the naive drift term $\mu_S = \mu$. Panel 3b comparing the robust liability drift term $\mu_L = a + b\rho\lambda_1 + b\sqrt{1-\rho^2}\lambda_2$ with the naive one $\mu_L = a$ under the benchmark scenario.

naive policy, if

$$K(w_{rob}|\delta_0) < K(w_{na}|\delta_0). \quad (15)$$

We displayed the loss indifference curves in Figure 4 when the present funding ratio is 80% and 90%. The x -axis and y -axis represent the true value of liability return a_0 and asset return μ_0 , respectively. The point $[a = 0, \mu = 0.04]$ represents the estimated expected return δ . We also displayed the ellipsoid uncertainty set of the true drift term δ_0 in the figure.

When $K(w_{rob}|\delta_0) = K(w_{na}|\delta_0)$, the two policies require the same amount of wealth to hedge. In the region below the curve for both scenarios (when $C_0 = 80\%$ and 90%), the robust policy requires less initial wealth than the naive policy to hedge a certain amount of expected shortfall. We call this area the robust policy's 'beneficial region'. Hence, we can conclude that when the true drift term δ_0 is overestimated, the robust policy performs better.

This beneficial region is positively related to the present funding ratio C_0 . Since the additional cost of hedging by following a robust policy increases as C_0 decreases, a lower C_0 leads to a smaller beneficial region. When liabilities are covered, the difference between a robust and a naive policy is subtle and the robust investor's beneficial region should also be larger.

3.4 Sensitivity analysis

The correlation parameter ρ , representing the completeness of the market, plays an important role in the model. If $\rho = \pm 1$, and $\lambda_1 = \lambda_2 = 0$, then the market becomes complete and the unhedgeable risk driver W_2 does not play a role. In this section, we investigated how sensitive the optimal hedges are with respect to a change of ρ .

In Figure 5, we show an extreme case when $\rho = 1$. The non-traded risk driver W_2 disappears from the liability diffusion process, and the perturbation parameter λ_2 does not play a role either. Nature can only control λ_1 to maximize the expected

Figure 4. Loss function equivalent curves. The figure plots the indifference curve of the loss when $K(w_{rob}|\delta_0) = K(w_{na}|\delta_0)$. y-axis is the true value of the expected stock return μ_0 and x-axis is the true value of the liability drift a_0 . The estimated value is $\mu = 0.04$ and $a = 0$. The solid-dot indifference curve represents the case when $C_0 = 80\%$ and the open-dot curve is the when $C_0 = 90\%$. In the region below the curve, the robust policy outperforms the naive policy, and in the region above, it is the other way around.

shortfall at period T . The naive agent considers this as a complete market. However, the robust agent still faces another source of incompleteness, caused by model misspecification.

With a low funding ratio, the robust policy deviates from the naive one much more severely compared with the benchmark case. When the asset risk and the liability risk are perfectly correlated, nature will choose a more negative λ_1 so as to maximize the expected shortfall. Although a negative λ_1 reduces the expected liability return as well, the liability drift term is less sensitive to the change of λ_1 than the expected asset return, since $\sigma > b$. As a result, the robust investor's fear of an overestimated asset return is stronger than the benchmark level.

In the case of overfunding, the two policies are identical. The hedging error volatility becomes $\sigma_C^2 = (w\sigma - b\rho)^2$. The investor can fully replicate the liability by following a δ -neutral strategy $w = ((b\rho)/\sigma) = 62\%$ if she has sufficient assets. In that case, robustness does not play a role because the δ hedge neutralizes the λ_1 effect.

In Figure 6, we show the two hedging policies as a function of correlation parameter ρ . We displayed two scenarios, one when $C_0 = 80\%$ and the other when $C_0 = 90\%$. The relation between the optimal portfolios and ρ is not monotone but is hump shaped. This is because the volatility of the value function σ_C is a quadratic function of ρ .

The optimal portfolio initially increases with ρ for both policies because the liability is more exposed to the tradable risk driver W_1 . Therefore, the risky portfolio has to increase as well, in order to hedge the traded liability risk. The optimal portfolio reaches the peak where ρ maximizes the total volatility σ_C . After the peak, the risky

Figure 5. Sensitivity analysis with $\rho = 1$. The figure depicts the optimal portfolio choice when $\rho = 1$. The remaining parameters stay at the benchmark level. The solid-dot line represents the robust policy and the empty-dotted curve is the naive policy. The naive agent considers such an economy a complete market, since the non-tradable risk driver W_2 is gone. However, the robust agent still stays in the incomplete market, because the model misspecification ($\lambda_1 \neq 0$ $\lambda_2 \neq 0$) is also considered as another source of market incompleteness.

portfolio goes down with ρ , because after the peak, any higher level of correlation will reduce σ_C . From Figure 6, we can also see that the difference between the two policies under the lower funding ratio is wider than under the higher C_0 .

4 Dynamic robust optimization

In this section, we will extend the problem to a dynamic strategy. The robust investor still aims to minimize the final-period expected shortfall under the worst-case scenario, but instead of making a static portfolio choice, she is now considering a dynamic optimal portfolio. Nature also can rebalance her choice of $(\lambda_{1t}, \lambda_{2t})$ instantaneously given the intertemporal decision of w_t . We employed dynamic programming to solve this robust optimization problem.

4.1 Dynamic programming

Define the indirect utility function $V(A_t, L_t)$, which follows the min–max expected utility given by (9). Both the investor and nature have a planning horizon of T . We omitted the time subscript t for notation convenience. Using Feynman–Kac, we can derive the Hamilton–Jacobi–Bellman equation (henceforth HJB) or partial differential equation (pde) for the investor’s min–max problem:

$$0 = \min_w \max_{\lambda_1, \lambda_2} V_t + V_A A(r + w(\mu - r) + w\sigma\lambda_1) + V_L L(a + b\rho\lambda_1 + b\sqrt{1 - \rho^2}\lambda_2) + \frac{1}{2} V_{AA} w^2 \sigma^2 A^2 + \frac{1}{2} V_{LL} b^2 L^2 + V_{AL} b\rho w\sigma AL - \frac{1}{2} v(\lambda_1^2 + \lambda_2^2 - \kappa^2), \quad (16)$$

Figure 6. Sensitivity analysis with respect to ρ . The figure plots the optimal naive and robust hedging policies as a function of correlation parameter ρ . We show two pairs of comparison: one with the present funding ratio C_0 of 80%, and the other with $C_0 = 90\%$. The solid-dot curves represent the robust policy and the empty-dot curves are the naive policy.

where the partial derivative with respect to x is denoted as V_x . We formed a Lagrangian function with multiplier ν over the boundary condition: $\lambda_1^2 + \lambda_2^2 \leq \kappa^2$.⁴

By solving a linear system of equations based on the first-order condition of (16) with respect to the strategy variables w , λ_1 , and λ_2 , we have

$$w^* = -\frac{(\mu - r)V_A A \nu}{\sigma^2(V_{AA}A^2\nu + V_A^2A^2)} - \frac{V_{AL}ALb\rho\sigma\nu + V_LV_AALb\rho\sigma}{\sigma^2(V_{AA}A^2\nu + U_A^2A^2)}, \quad (17a)$$

$$\lambda_1^* = -\frac{(\mu - r)V_A^2A^2\sigma}{\sigma^2(V_{AA}A^2\nu + V_A^2A^2)} - \frac{b\rho(V_{AL}ALV_{AA} - V_LLV_{AA}A^2)}{(V_{AA}A^2\nu + V_A^2A^2)}, \quad (17b)$$

$$\lambda_2^* = \frac{V_LLb\sqrt{1 - \rho^2}}{\nu}. \quad (17c)$$

This is a partial solution. The Lagrange multiplier ν , as well as V_A , V_{AA} , still need to be solved numerically.

The sign of the optimal λ_2 must be positive since it increases the expected liability return but does not influence the pension asset. The sign of λ_1 is ambiguous. A positive λ_1 not only increases the liability but also the asset, but the net effect depends on the value of other input variables.

⁴ The Lagrangian multiplier ν relates to the time-consistent per-period constraint (6), which is different from the setup of Anderson *et al.* (2003). In Anderson *et al.* (2003), the last term of the HJB equation is replaced by a relative entropy function, $1/2\theta(\lambda_1^2 + \lambda_2^2)$ which penalizes drift distortions. Although the two HJB equations look similar, θ implicitly imposes an aggregate-budget style uncertainty constraint rather than a per-time step constraint. Hansen and Sargent (2007) employ the detection-error-probability methods to calibrate θ . The detection error probability method performs likelihood ratio tests under the two models based on available data. By linking the relative entropy parameter to a Bayesian model selection function, one can derive the value of the relative entropy.

The solution (17) has an interesting structure. The dynamic optimal investment strategy w^* is a trade-off between hedging and speculation. We can see this by considering the extreme case when $v \rightarrow 0$ and $v \rightarrow \infty$.

For $v \rightarrow 0$, the discrepancy parameters λ_1 and λ_2 have more freedom to choose an arbitrarily large aversion pair of drift for the Brownian motions, or in other words, the agent is extremely pessimistic about the approximation model. When $v \rightarrow 0$, we have

$$w_{v \rightarrow 0}^* = -\frac{V_L L b \rho}{V_A A \sigma}. \quad (18)$$

This is a pure hedging portfolio, where the agent invests an amount in risky assets such that the change in the value function due to L is (as much as possible) offset by a change in value due to a . It is not possible to completely eliminate the volatility of L . This is because the liabilities are exposed both to hedgeable risk W_1 and unhedgeable risk W_2 , but only the hedgeable part W_1 can be eliminated.

The optimal value for λ_1^* when $v \rightarrow 0$ is given by

$$\lambda_{1,v \rightarrow 0}^* = -\frac{\mu - r}{\sigma} - \frac{b \rho (V_{AL} V_A - V_L V_{AA}) L}{V_A^2}, \quad (19)$$

which contains two terms. The first term is the observable market price of risk, which we can see from the BS setup. The second term is more interesting. Since

$$-\frac{b \rho (V_{AL} V_A - V_L V_{AA}) L}{V_A^2} = \sigma \frac{\partial (w_{v \rightarrow 0}^* A)}{\partial A} = w_{v \rightarrow 0}^* \sigma + \sigma A \frac{\partial w_{v \rightarrow 0}^*}{\partial A},$$

this reflects to what extent the agent's best possible hedging strategy is influenced by the instantaneous wealth level A_t .

At the other extreme, when $v \rightarrow \infty$, both λ_1 and λ_2 shrink to zero, so $\kappa = 0$. This corresponds to the case when the agent faces no model misspecification. Hence, we recovered the 'classical' Merton's solution for the optimal portfolio choice:

$$w_{v \rightarrow \infty}^* = -\frac{\mu - r}{\sigma^2} \frac{V_A}{V_{AA} A} - \frac{V_{AL} L b \rho}{V_{AA} A \sigma}. \quad (20)$$

The first term is a speculative portfolio, where the agent invests in the stock market to obtain the optimal trade-off between the observable market price of risk $((\mu - r)/\sigma^2)$ and the local risk aversion $-((V_A/V_{AA} A))$. The second term is the intertemporal hedging component, but the optimal amount to hedge is now measured in terms of the 'CAPM- β '. That is, the optimal hedge is the local covariance term $b \rho \sigma$ divided by local variance term σ^2 , i.e. the stock market investment that minimizes locally the (unhedgeable) variance in the portfolio.

4.2 Numerical solution

As we cannot solve the PDE analytically, we will present numerical results for the dynamic optimization problem.

In Figure 7, we show the dynamic robust investment policy as a function of the instantaneous funding ratio C_t and hedging horizon T . The optimal weight on the

Figure 7. Dynamic robust optimal hedging strategy. This figure displays the robust optimal investment policy as a function of the instantaneous funding ratio C_t with benchmark input parameters under different hedging horizons. Panel 7a plots the robust portfolio choice as a function of the instantaneous funding ratio and the investment horizon T . Panel 7b depicts the solutions when investment horizon is $T = 1, 3, 5$. Due to technical limitations, our grid searching interval for the risky portfolio w has to be smaller than 1.95, otherwise we will confront a negative probability problem in some trinomial trees.

risky asset depends both on the solvency condition and the investment horizon. If the funding ratio is low, a longer term investor takes less risk than a shorter term investor. In other words, the risk exposure to the stock market decreases with the hedging horizon. When underfunded, an investor would take an aggressive risk position, betting on the chance of avoiding a shortfall, exactly as we have seen in the static case. A shorter planning horizon triggers a stronger intention to cover the liabilities, hence leads to a riskier position. However, when overfunded ($C_t > 1$), the longer the investment horizon is, the more risk can be taken. The optimal portfolio converges to the hedging ratio $\delta ((bp)/\sigma)$ when the hedging horizon T is close to zero.

Next, we investigated the difference between the robust and the naive dynamic policies. In Figure 8, we present the two investment policies as a function of the instantaneous funding ratio under two hedging horizons, $T = 5$ and $T = 3$. We highlight two findings from the figure. First, the robust policy is less risky than the naive one as long as the instantaneous funding ratio is lower than 1. Second, the difference between the two policies decreases with the investment horizon. As the risk exposure decreases with hedging horizon, so does the fear of uncertainty. Compared with static hedges (Figure 1), dynamic hedges (Figure 8) take riskier positions under both robust and naive policies.

Figure 9 shows the dynamic optimal λ_1, λ_2 as functions of the funding ratio at three different investment horizons. It is still the case that λ_1 is always negative and λ_2 is always positive (see also Figure 2). We now focus on the dynamic effect of the processes.

When underfunded ($C_t < 1$), the absolute value of λ_1 decreases when the hedging horizon increases, since the longer term investor is less exposed to the stock market (see Figure 7b) than the shorter term investor. Therefore, nature becomes less effective in distorting the asset model when the hedging horizon increases. The optimal value of λ_2 increases with the investment horizon so as to offset the diminished effect of λ_1 .

Figure 8. Dynamic robust and naive optimal hedging strategy as a function of instantaneous funding ratio at selected hedging horizon. In this figure, we display both robust and naive investment policies as functions of the instantaneous funding ratio. Panel 8a plots the solution when $T = 5$. Panel 8b shows the result when $T = 3$.

Figure 9. Dynamic optimal perturbation processes. In this figure, we show the optimal perturbation processes λ_1 and λ_2 as functions of the instantaneous funding ratio when hedging horizon equals to $T = 1, 3, 5$. The solid lines are the movement of λ_1 and λ_2 when $T = 5$, the dashed curves are for the case $T = 3$, and the dotted curves are for $T = 1$. The upper panel with positive perturbations gives the optimal results of λ_2 . The negative portion of the figure gives the optimal solutions of λ_1 .

We moved on to analyze the dynamic perturbed drift terms displayed in Figure 10. Panel 10a plots the perturbed expected stock return process μ_S as a function of C_t under three different investment horizons. Since $\mu_S = \mu + \sigma\lambda_1$ is a linear function of λ_1 , it shares common characteristics with λ_1 shown in Figure 9. After all, μ_S increases with hedging horizon when underfunded and *vice versa* if $C_t > 1$. Panel 10b shows the movement of μ_L . When C_t is low, the perturbed expected liability return μ_L increases with T to offset the diminishing distortions from the asset side.

Figure 10. Dynamic perturbation effect on drift terms. In this figure, we plot the dynamic movement of the perturbed drift terms as functions of the instantaneous funding ratio when hedging horizon is $T = 1, 3, 5$. Panel 10a depicts the movement of $\mu_S = \mu + \sigma\lambda_1$ and panel 10b shows $\mu_L = a + b\rho\lambda_1 + b\sqrt{1-\rho^2}\lambda_2$.

Figure 11. Dynamic loss function equivalent curve. The figure shows the policy indifference curve at a function of the true drift terms (μ_0, a_0) at three different horizons $T = 1, 3, 5$. The left panels plots the case when the instantaneous funding ratio equals to 80%, and the right panel is the case when $C_t = 90\%$. The robust policy are better off in the area beneath the indifference curves. The dynamic policies w_{rob} and w_{na} are determined based on the estimated drift terms with values $\mu = 0.04$ and $a = 0$.

4.3 Dynamic policy evaluation

From the static case, we know that the robust policy performs better when the drift terms are overestimated. In this section, we will investigate the effect of the hedging horizon.

Figure 11 displays the policy indifference curve under different hedging horizons T . The area beneath the indifference curves represents the scenarios that require less initial wealth to hedge a given amount of downside risks by following a robust policy. Different from the static case (Figure 4), the beneficial region of the robust policy in the dynamic setting is smaller than it is in the static case. This means naive dynamic hedging is less sensitive to parameter uncertainty. Additionally, we found that the beneficial region increases with the hedging horizon. Therefore, long-term investors should be more inclined to follow a robust investment strategy than short-term

investors. The horizon effect is weaker when the instantaneous funding ratio is relatively high. As a robustness check, we also conducted sensitivity analysis on ρ in the dynamic hedging environment when the instantaneous funding ratio is low. The hedging portfolios for both policies are positively related to the correlation factor ρ and are lower under longer term investment.

5 Conclusion

We analyzed a robust hedging strategy under the condition that the market is incomplete and the underlying model can be misspecified. We employed and simplified the general model uncertainty problem of Hansen and Sargent (2007) to uncertainty about the drift terms. The robust policy requires an extra cost of capital, or lower liability discount rate, to guarantee against model uncertainty. That is the price to pay for coping with the parameter uncertainty. If the model is truly misspecified, the hedging will be more successful.

From our analysis, we summarize two major characteristics of the robust policy. We first found that the robustness effect strongly depends on the instantaneous funding ratio. The preference for robustness only influences the hedging policy when the funding ratio is low; if the fund's assets are large enough to cover the liability payoff, then the robust and the naive policies are identical. Second, the robust policy also becomes more valuable for longer investment horizons.

The investor can benefit from the robust policy when the expected return is overestimated. That means, with a given expected-shortfall hedging target, the robust policy requires less initial wealth to obtain a successful hedge than the naive policy if the true expected stock return is lower than the estimated value.

Acknowledgements

The authors thank two anonymous referees and the editor Clemens Sialm for extremely valuable suggestions and comments. The authors are also grateful for comments from Jaap Bos, Alexey Rubtsov, Bertrand Melenberg, Marcel Rindisbacher, Hans Schumacher and Michel Vellekoop. This research is financially supported by Netspar. All errors are our own.

References

- Anderson, E., Hansen, L., and 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.
- Ang, A., Chen, B., and Sundaresan, S. (2013) Liability-driven investment with downside risk. *Journal of Portfolio Management*, **40**(1): 71.
- Benjamin, B. and Soliman, A. S. (1993) *Mortality on the Move: Methods of Mortality Projection*. Oxford: Actuarial Education Service.
- Blake, D. and Burrows, W. (2001) Survivor bonds: helping to hedge mortality risk. *Journal of Risk and Insurance*, **68**: 339–348.
- Brennan, M. J. (1998) The role of learning in dynamic portfolio decisions. *European Finance Review*, **1**(3): 295–306.

- Cairns, A. J., Blake, D., and Dowd, K. (2006) A two-factor model for stochastic mortality with parameter uncertainty: theory and calibration. *Journal of Risk and Insurance*, **73**(4): 687–718.
- Cochrane, J. H. and Saa-Requejo, J. (2000) Beyond arbitrage: good-deal asset price bounds in incomplete markets. *Journal of Political Economy*, **108**(1): 79–119.
- Coelho, E. and Nunes, L. C. (2011) Forecasting mortality in the event of a structural change. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **174**(3): 713–736.
- Cvitanić, J. and Karatzas, I. (1999) On dynamic measures of risk. *Finance and Stochastics*, **3**(4): 451–482.
- Delbaen, F. (2002) *Coherent Risk Measures on General Probability Spaces*. Advances in Finance and Stochastics. Berlin, Heidelberg: Springer, pp. 1–37.
- Detemple, J. and Rindisbacher, M. (2008) Dynamic asset liability management with tolerance for limited shortfalls. *Insurance: Mathematics and Economics*, **43**(3): 281–294.
- EIOPA (2011) EIOPA Report on the fifth Quantitative Impact Study (QIS5) for Solvency II. EIOPA-TFQIS5-11/001, 14 March 2011. European Insurance and Occupational Pensions Authority.
- Garlappi, L., Uppal, R., and Wang, T. (2006) Portfolio selection with parameter and model uncertainty: a multi-prior approach. *The Review of Financial Studies*, **20**(1): 41–81.
- Gilboa, I. and Schmeidler, D. (1989) Max–min expected utility with non-unique prior. *Journal of Mathematical Economics*, **18**(2): 141–153.
- Hansen, L. and Sargent, T. (2007) *Robustness*. Princeton, NJ: Princeton University Press.
- Lee, R. and Carter, L. (1992) Modeling and forecasting US mortality. *Journal of the American Statistical Association*, **87**(419): 659–671.
- Li, J. S.-H., Hardy, M. R., and Tan, K. S. (2009) Uncertainty in mortality forecasting: an extension to the classical Lee-Carter approach. *Astin Bulletin*, **39**(1): 137–164.
- Luo, Y. (2016) Robustly strategic consumption–portfolio rules with informational frictions. *Management Science*, **63**(12): 4158–4174.
- Maenhout, P. (2004) Robust portfolio rules and asset pricing. *Review of Financial Studies*, **17**(4): 951.
- Margrabe, W. (1978) The value of an option to exchange one asset for another. *The Journal of Finance*, **33**(1): 177–186.
- McDonald, R., Cassano, M., and Fahlenbrach, R. (2006) *Derivatives Markets*, Volume 2. Boston: Addison-Wesley.
- Merton, R. C. (1980) On estimating the expected return on the market: an exploratory investigation. *Journal of Financial Economics*, **8**: 323–361.
- Rockafellar, R. T. (1976) Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, **14**(5): 877–898.
- Sharpe, W. F. and Tint, L. G. (1990) Liabilities – a new approach. *The Journal of Portfolio Management*, **16**(2): 5–10.
- Turner, A. (2006) Pensions, risks, and capital markets. *Journal of Risk and Insurance*, **73**(4): 559–574.
- Van Berkum, F., Antonio, K., and Vellekoop, M. (2016) The impact of multiple structural changes on mortality predictions. *Scandinavian Actuarial Journal*, **2016**(7): 0 581–603.
- Wang, N. (2009) Optimal consumption and asset allocation with unknown income growth. *Journal of Monetary Economics*, **56**(4): 524–534.