

A Cautionary Note on Natural Hedging of Longevity Risk *

Nan Zhu[†]

Department of Mathematics, Illinois State University
100 N University Street; Normal, IL 61790; USA
Email: nzhu@ilstu.edu

Daniel Bauer

Department of Risk Management and Insurance, Georgia State University
35 Broad Street, 11th Floor; Atlanta, GA 30303; USA
Email: dbauer@gsu.edu

November 2013

Abstract

In this paper, we examine the so-called “natural hedging” approach for life insurers to internally manage their longevity risk exposure by adjusting their insurance portfolio. In particular, unlike the existing literature, we also consider a non-parametric mortality forecasting model that avoids the assumption that all mortality rates are driven by the same factor(s).

Our primary finding is that higher order variations in mortality rates may considerably affect the performance of natural hedging. More precisely, while results based on a parametric single factor model—in line with the existing literature—imply that almost all longevity risk can be hedged, results are far less encouraging for the non-parametric mortality model. Our finding is supported by robustness tests based on alternative mortality models.

JEL classification: G22; G32; J11.

Keywords: Natural Hedging, Longevity Risk, Mortality Forecasting, Non-Parametric Model, Lee-Carter Model, Economic Capital.

*A previous version entitled “A Natural Examination of the Natural Hedging Approach” was presented at the Eighth International Longevity Risk and Capital Markets Solutions Conference. We are grateful for valuable comments from the conference participants. Moreover, we thank two anonymous referees whose suggestions greatly improved the paper. All remaining errors are our own.

[†]Corresponding author. Phone: +1-(309)-438-7838. Fax: +1-(309)-438-5866.

1 Introduction

Longevity risk, i.e. the risk that policyholders will live longer than expected, has recently attracted increasing attention from both academia and insurance practitioners. Different ways have been suggested on how to manage this risk, e.g. by transferring it to the financial market via mortality-linked securities (see e.g. Blake et al. (2006)). One approach that is particularly appealing at first glance since it can be arranged from within the insurer is “natural hedging”, i.e. adjusting the insurance portfolio to minimize the overall exposure to systematic mortality risk (longevity risk).

Cox and Lin (2007) first formally introduce this concept of mortality risk management for life insurers. They find that empirically, companies selling both life and annuity policies generally charge cheaper prices for annuities than companies with only single business line. Since then, a number of studies have occurred in the insurance literature showing “that natural hedging can significantly lower the sensitivity of an insurance portfolio with respect to mortality risk” (Bayraktar and Young, 2007; Wetzel and Zwiesler, 2008; Tsai et al., 2010; Wang et al., 2010; Gatzert and Wesker, 2012).

However, these contributions arrive at their positive appraisal of the natural hedging approach within model-based frameworks. That is, their conclusions rely on conventional mortality models such as the Lee-Carter model (Lee and Carter, 1992) or the CBD model (Cairns et al., 2006b). While these popular models allow for a high degree of numerical tractability and serve well for many purposes, they come with the assumption that all mortality rates are driven by the same low-dimensional stochastic factors. Therefore, these models cannot fully capture disparate shifts in mortality rates at different ages, which could have a substantial impact on the actual effectiveness of natural hedging.

To analyze the impact of the mortality forecasting model on the effectiveness of natural hedging, in this paper, we compare results under several assumptions for the future evolution of mortality in the context of a stylized life insurer. In particular, aside from considering deterministic mortality rates and a simple factor model as in previous studies, we also use a non-parametric forecasting model that arises as a by-product of the mortality modeling approach presented in Zhu and Bauer (2013). The advantage of a non-parametric model is that we do not make functional assumptions on the mortality model, especially the potentially critical factor structure indicated above.¹ Our results reveal that the efficiency of natural hedging is considerably reduced when relying on the non-parametric model—which underscores the problem when relying on model-based analyses for risk management decisions more generally.

We perform various robustness tests for this finding. In particular, we consider a setting without

¹Similar arguments can be found in other insurance related studies, e.g. Li and Ng (2010) use a non-parametric framework to price mortality-linked securities.

financial risk, and we repeat the calculations for alternative mortality models. While these analyses reveal additional insights, the primary result is robust to these modifications: Natural hedging only marginally reduces the exposure of the company to systematic mortality risk. This meager performance can be viewed as further evidence endorsing market-based solutions for managing longevity risk.

The remainder of the paper is structured as follows: Section 2 briefly introduces the considered mortality forecasting models. Section 3 discusses the calculation of the economic capital for a stylized life insurance company, while Section 4 revisits the natural hedging approach within our economic capital framework. Section 5 conducts the robustness tests. Section 6 concludes.

2 Mortality Forecasting Models

We commence by introducing the mortality forecasting models that will be primarily used in this paper. In particular, we consider two representative models within the forward-mortality framework developed in Zhu and Bauer (2013), namely a parametric single-factor model as well as a non-parametric model for the annualized mortality innovations. Employing two models from the same framework facilitates the interpretation of similarities and differences within certain applications. Moreover, as is detailed in Zhu and Bauer (2011), the use of conventional *spot mortality models* (Cairns et al., 2006a) will typically require so-called *nested simulations* in the numerical realizations within our Economic Capital framework, which in turn will considerably increase the computational difficulty of the optimization procedures described below. Since we are primarily interested in how the assumption of a low-dimensional factor structure—rather than the choice of any specific mortality forecasting model—affects the performance of the natural hedging approach in model-based analyses, we believe that our model choice serves well as a representative example in order to draw more general conclusions. Nevertheless, in Section 5 we conduct robustness tests of our results based on alternative factor and non-parametric models that are used in existing literature.

Underlying the approach is a time series of generation life tables for some population for years t_1, t_2, \dots, t_N . More precisely, in each of these tables labeled by its year t_j , we are given forward-looking survival probabilities ${}_{\tau}p_x(t_j)$ for ages $x = 0, 1, 2, \dots, 100$ and terms $\tau = 0, 1, 2, \dots, 101 - x$, where ${}_{\tau}p_x(t)$ denotes the probability for an x -year old to survive for τ periods until time $t + \tau$.² Mathematically, this is equivalent to

$${}_{\tau}p_x(t) \mathbf{1}_{\{\Upsilon_{x-t} > t\}} = \mathbb{E}^{\mathbb{P}} \left[\mathbf{1}_{\{\Upsilon_{x-t} > t + \tau\}} \mid \mathcal{F}_t \vee \{\Upsilon_{x-t} > t\} \right]$$

²In particular, in this paper, we use a maximal age of 101 though generalizations are possible.

for an $(x - t) > 0$ year old at time zero, where Υ_{x_0} denotes the (random) time of death or future lifetime of an x_0 -year old at time zero. In particular, ${}_{\tau}p_x(t)$ will account for projected mortality improvements over the future period $[t, t + \tau)$.

Now for each year t_j , $1 \leq j < N$ and for each term/age combination (τ, x) with $1 \leq x \leq 100$ and $0 \leq \tau \leq 100 - x$, we define:

$$F(t_j, t_{j+1}, (\tau, x)) = -\log \left\{ \frac{{}_{\tau+1}p_x(t_{j+1})}{{}_{\tau}p_x(t_{j+1})} \bigg/ \frac{{}_{\tau+1+t_{j+1}-t_j}p_{x-t_{j+1}+t_j}(t_j)}{{}_{\tau+t_{j+1}-t_j}p_{x-t_{j+1}+t_j}(t_j)} \right\}, \quad 1 \leq j < N. \quad (1)$$

Hence, $F(t_j, t_{j+1}, (\tau, x))$ measures the log-change of the one-year marginal survival probability for an individual aged x at time t_{j+1} over the period $[t_{j+1} + \tau, t_{j+1} + \tau + 1)$ from projection at time t_{j+1} relative to time t_j . Further, we define the vector $\bar{F}(t_j, t_{j+1}) = \text{vec}(F(t_j, t_{j+1}, (\tau, x)), 1 \leq x \leq 100, 0 \leq \tau \leq 100 - x)$, with $\dim(\bar{F}(t_j, t_{j+1})) = \frac{100 \times 101}{2} = 5,050$, $j = 1, 2, \dots, N - 1$.

Proposition 2.1 in Zhu and Bauer (2013) shows that under the assumption that the mortality age/term-structure is driven by a time-homogeneous diffusion and with equidistant evaluation dates, i.e. $t_{j+1} - t_j \equiv \Delta$, the $\bar{F}(t_j, t_{j+1})$, $j = 1, \dots, N - 1$, are independent and identically distributed (iid). Therefore, in this case a non-parametric mortality forecasting methodology is immediately given by bootstrapping the observations $\bar{F}(t_j, t_{j+1})$, $j = 1, \dots, N - 1$ (Efron, 1979). More precisely, with Equation (1), we can generate simulations for the generation life tables at time t_{N+1} , $\{{}_{\tau}p_x(t_{N+1})\}$, by sampling (with replacement) $\bar{F}(t_N, t_{N+1})$ from $\{\bar{F}(t_j, t_{j+1}), j = 1, \dots, N - 1\}$ in combination with the known generation life tables at time t_N , $\{{}_{\tau}p_x(t_N)\}$. This serves as the algorithm for generating our non-parametric mortality forecasts. A related approach that we consider in the robustness tests (Section 5) relies on the additional assumption that $\bar{F}(t_j, t_{j+1})$ are iid Gaussian random vectors. In this case, we can directly sample from a Normal distribution with the mean and the covariance matrix estimated from the sample.

To introduce corresponding factor models, it is possible to simply perform a factor analysis of the iid sample $\{\bar{F}(t_j, t_{j+1}), j = 1, \dots, N - 1\}$, which shows that for population mortality data, the first factor typically captures the vast part of the systematic variation in mortality forecasts. However, as is detailed in Zhu and Bauer (2013), factor models developed this way are not necessarily *self-consistent*, i.e. expected values derived from simulations of future survival probabilities do not necessarily align with the forecasts engrained in the current generation life table at time t_N .

To obtain self-consistent models, it is convenient to introduce the so-called *forward force of mortality* (Cairns et al., 2006a),

$$\mu_t(\tau, x) = -\frac{\partial}{\partial \tau} \log \{{}_{\tau}p_x(t)\},$$

so that we have

$${}_{\tau}p_x(t) = \exp \left\{ - \int_0^{\tau} \mu_t(s, x) ds \right\}. \quad (2)$$

Time-homogeneous (forward) mortality models can then be represented by an infinite-dimensional stochastic differential equation of the form:

$$d\mu_t = (A\mu_t + \alpha) dt + \sigma dW_t, \quad \mu_0(\cdot, \cdot) > 0, \quad (3)$$

where α and σ are sufficiently regular, function-valued stochastic processes, $A = \frac{\partial}{\partial \tau} - \frac{\partial}{\partial x}$, and (W_t) is a d -dimensional Brownian motion. Bauer et al. (2012a) show that for self-consistent models, we have the drift condition

$$\alpha(\tau, x) = \sigma(\tau, x) \times \int_0^{\tau} \sigma'(s, x) ds, \quad (4)$$

and for time-homogeneous, Gaussian models (where α and σ are deterministic) to allow for a factor structure, a necessary and sufficient condition is

$$\sigma(\tau, x) = C(x + \tau) \times \exp\{M\tau\} \times N,$$

for some matrices M , N , and a vector-valued function $C(\cdot)$. By aligning this semi-parametric form with the first factor derived in a factor analysis described above, Zhu and Bauer (2013) propose the following specification for the volatility structure in a single-factor model:

$$\sigma(\tau, x) = (k + c e^{d(x+\tau)}) (a + \tau) e^{-b\tau}. \quad (5)$$

Together with Equations (2), (3), and (4), Equation (5) presents the parametric factor mortality forecasting model employed in what follows. We refer to Zhu and Bauer (2013) for further details, particularly on how to obtain Maximum-Likelihood estimates for the parameters k , c , d , a , and b .

3 Economic Capital for a Stylized Insurer

In this section, we employ the mortality forecasting approaches outlined in the previous section to calculate the Economic Capital (EC) of a stylized life insurance company. We start by introducing the framework for the EC calculations akin to Zhu and Bauer (2011). Subsequently, we describe the data used in the estimation of the underlying models and resulting parameters. In addition to calculating the EC for a base case company with fixed investments, we derive an optimal static hedge for the financial risk by adjusting the asset weights.

3.1 EC Framework

Consider a (stylized) newly founded life insurance company selling traditional life insurance products only to a fixed population. More specifically, assume that the insurer's portfolio of policies consists of $n_{x,i}^{\text{term}}$ i -year term-life policies with face value B_{term} for x -year old individuals, $n_{x,i}^{\text{end}}$ i -year endowment policies with face value B_{end} for x -year old individuals, and n_x^{ann} single-premium life annuities with an annual benefit of B_{ann} paid in arrears for x -year old individuals, $x \in \mathcal{X}, i \in \mathcal{I}$. Furthermore, assume that the benefits/premiums are calculated by the Equivalence Principle based on the concurrent generation table and the concurrent yield curve without the considerations of expenses or profits. In particular, we assume that the insurer is risk-neutral with respect to mortality risk, i.e. the valuation measure \mathbb{Q} for insurance liabilities is the product measure of the risk-neutral measure for financial and the physical measure for demographic events.

Under these assumptions, the insurer's *Available Capital* at time zero, AC_0 , defined as the difference of the market value of assets and liabilities, simply amounts to its initial equity capital E . The available capital at time one, AC_1 , on the other hand, equals to the difference in the value of the insurer's assets and liabilities at time one, denoted by A_1 and V_1 , respectively. More specifically, we have

$$\begin{aligned} A_1 &= \left(E + B_{\text{ann}} \sum_{x \in \mathcal{X}} a_x(0) n_x^{\text{ann}} + B_{\text{term}} \sum_{x \in \mathcal{X}, i \in \mathcal{I}} \frac{A_{x:\overline{n}|}(0)}{\ddot{a}_{x:\overline{n}|}(0)} n_{x,i}^{\text{term}} + B_{\text{end}} \sum_{x \in \mathcal{X}, i \in \mathcal{I}} \frac{A_{x:\overline{n}|}(0)}{\ddot{a}_{x:\overline{n}|}(0)} n_{x,i}^{\text{end}} \right) \times R_1, \\ V_1 &= B_{\text{ann}} \sum_{x \in \mathcal{X}} \ddot{a}_{x+1}(1) (n_x^{\text{ann}} - \mathfrak{D}_x^{\text{ann}}(0, 1)) + B_{\text{term}} \sum_{x \in \mathcal{X}, i \in \mathcal{I}} \mathfrak{D}_{x,i}^{\text{term}}(0, 1) + B_{\text{end}} \sum_{x \in \mathcal{X}, i \in \mathcal{I}} \mathfrak{D}_{x,i}^{\text{end}}(0, 1) \\ &\quad + B_{\text{term}} \sum_{x \in \mathcal{X}, i \in \mathcal{I}} \left[A_{x+1:\overline{i-1}|}(1) - \frac{A_{x:\overline{i}|}(0)}{\ddot{a}_{x:\overline{i}|}(0)} \ddot{a}_{x+1:\overline{i-1}|}(1) \right] \times (n_{x,i}^{\text{term}} - \mathfrak{D}_{x,i}^{\text{term}}(0, 1)) \\ &\quad + B_{\text{end}} \sum_{x \in \mathcal{X}, i \in \mathcal{I}} \left[A_{x+1:\overline{i-1}|}(1) - \frac{A_{x:\overline{i}|}(0)}{\ddot{a}_{x:\overline{i}|}(0)} \ddot{a}_{x+1:\overline{i-1}|}(1) \right] \times (n_{x,i}^{\text{end}} - \mathfrak{D}_{x,i}^{\text{end}}(0, 1)). \end{aligned}$$

Here, R_1 is the total return on the insurer's asset portfolio. $\mathfrak{D}_{x,i}^{\text{con}}(0, 1)$ is the number of deaths between time zero and time one in the cohort of x -year old policyholders with policies of term i and of type $\text{con} \in \{\text{ann}, \text{term}, \text{end}\}$. And $\ddot{a}_x(t)$, $A_{x:\overline{n}|}(t)$, etc. denote the present values of the contracts corresponding to the actuarial symbols at time t —which are calculated based on the yield curve and the generation table at time t . For instance,

$$\ddot{a}_x(t) = \sum_{\tau=0}^{\infty} \tau p_x(t) p(t, \tau),$$

where $\tau p_x(t)$ is the time- t (forward) survival probability as defined in Section 2, and $p(t, \tau)$ denotes the time t price of a zero coupon bond that matures in τ periods (at time $t + \tau$).

The EC calculated within a one-year mark-to-market approach of the insurer can then be derived as (Bauer et al., 2012b)

$$EC = \rho \left(\underbrace{AC_0 - p(0, 1) AC_1}_L \right),$$

where L denotes the *one-period loss* and $\rho(\cdot)$ is a *monetary risk measure*. For example, if the EC is defined based on Value-at-Risk (VaR) such as the *Solvency Capital Requirement (SCR)* within the Solvency II framework, we have

$$EC = SCR = \text{VaR}_\alpha(L) = \arg \min_x \{ \mathbb{P}(L > x) \leq 1 - \alpha \}, \quad (6)$$

where α is a given threshold (99.5% in Solvency II). If the EC is defined based on the *Conditional Tail Expectation (CTE)*, on the other hand, we obtain

$$EC = \text{CTE}_\alpha = \mathbb{E}[L | L \geq \text{VaR}_\alpha(L)]. \quad (7)$$

In this note, we define the economic capital based on VaR (Equation (6)), and choose $\alpha = 95\%$.

3.2 Data and Implementation

For estimating the mortality models in this paper, we rely on female US population mortality data for the years 1933-2007 as available from the Human Mortality Database.³ More precisely, we use ages ranging between 0 and 100 years to compile 46 consecutive generation life tables (years $t_1 = 1963, t_2 = 1964, \dots, t_{46} = 2008$) based on Lee-Carter mean projections, with each generated independently from the mortality experience of the previous 30 years.⁴ That is, the first table (year $t_1 = 1963$) uses mortality data from years 1933-1962, the second table (t_2) uses years 1934-1963, and so forth.

Having obtained these generation tables $\{\tau p_x(t_j)\}$, $j = 1, \dots, N = 46$, we derive the time series of $\bar{F}(t_j, t_{j+1})$, $j = 1, 2, \dots, 45$, which serve as the underlying sample for our non-parametric forecasting methodology and as the basis for the maximum likelihood parameter estimates of our mortality factor model.⁵ In particular, time $t_N = 2008$ corresponds to time zero whereas time

³Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org or www.humanmortality.de.

⁴More precisely, for the estimation of the Lee-Carter parameters, instead of the original approach we use the modified weighted-least-squares algorithm (Wilmoth, 1993) and further adjust κ_t by fitting a Poisson regression model to the annual number of deaths at each age (Booth et al., 2002).

⁵Of course, the underlying sample of 45 realizations is rather small for generating a large bootstrap sample, which limits the scope of the approach for certain applications (such as estimating VaR for high confidence levels which is of practical interest). We come back to this point in our robustness tests (Section 5).

$t_{N+1} = 2009$ corresponds to time one in our EC framework. Table 1 displays the parameter estimates of the parametric factor model (5).⁶

<i>Parameters</i>				
k	c	d	a	b
2.3413×10^{-6}	3.3722×10^{-8}	0.1041	3.1210	0.0169

Table 1: Estimated parameters of the factor mortality forecasting model (5)

For the asset side, we assume that the insurer only invests in 5, 10, and 20-year US government bonds as well as an equity index (S&P 500) $S = (S_t)_{t \geq 0}$. For the evolution of the assets, we assume a generalized Black-Scholes model with stochastic interest rates (Vasicek model), that is, under \mathbb{P}

$$\begin{aligned} dS_t &= S_t(\mu dt + \rho \sigma_A dB_t^{(1)} + \sqrt{1 - \rho^2} \sigma_A dB_t^{(2)}), \quad S_0 > 0, \\ dr_t &= \kappa(\gamma - r_t) dt + \sigma_r dB_t^{(1)}, \quad r_0 > 0, \end{aligned} \quad (8)$$

where $\mu, \sigma_A, \kappa, \gamma, \sigma_r > 0$, $\rho \in [-1, 1]$, and $(B_t^{(1)})$ and $(B_t^{(2)})$ are independent Brownian motions that are independent of (W_t) . Moreover, we assume that the market price of interest rate risk is constant and denote it by λ , i.e. we replace μ by r_t and γ by $\gamma - (\lambda \sigma_r)/\kappa$ for the dynamics under the risk-neutral measure \mathbb{Q} .

We estimate the parameters based on US data from June 1988 to June 2008 using a Kalman filter. In particular, we use monthly data of the S&P 500 index,⁷ treasury bills (3 months), and government bonds with maturities of 1 year, 3 years, 5 years, and 10 years.⁸ The parameter estimates are displayed in Table 2.

Based on time-one realizations of the asset process, S_1 , and the instantaneous risk-free rate, r_1 , we have

$$R_1 = \omega_1 \frac{S_1}{S_0} + \omega_2 \frac{p(1, 4)}{p(0, 5)} + \omega_3 \frac{p(1, 9)}{p(0, 10)} + \omega_4 \frac{p(1, 19)}{p(0, 20)},$$

where $\omega_i, i = 1, \dots, 4$, are the company's proportions of assets invested in each category. A

⁶A Principal Component Analysis indicates that 85% of the total variation in the $\bar{F}(t_j, t_{j+1}), j = 1, 2, \dots, 45$, is explained by the leading factor for our dataset. Generally, the percentage of total variation explained is slightly larger for female data in comparison to male data (Zhu and Bauer, 2013), suggesting that for female populations a single factor model is more appropriate.

⁷Downloaded on 08/26/2012 from Yahoo! Finance, <http://finance.yahoo.com>.

⁸Downloaded on 08/26/2012 from the Federal Reserve Economic Data (FRED), <http://research.stlouisfed.org/fred2/>.

<i>Parameters</i>							
μ	σ_A	ρ	κ	γ	σ_r	λ	r_0 (06/2008)
0.0866	0.1372	-0.0078	0.0913	0.0123	0.0088	-0.7910	0.0188

Table 2: Estimated parameters of the capital market model

	x	i	$n_{x,i}^{\text{term/end/ann}}$	$B_{\text{term/end/ann}}$
<i>Term Life</i>				
	30	20	2,500	\$100,000
	35	15	2,500	\$100,000
	40	10	2,500	\$100,000
	45	5	2,500	\$100,000
<i>Endowment</i>				
	40	20	5,000	\$50,000
	45	15	5,000	\$50,000
	50	10	5,000	\$50,000
<i>Annuities</i>				
	60	(35)	2,500	\$18,000
	70	(25)	2,500	\$18,000

Table 3: Portfolio of policies for the stylized life insurer

procedure to generate realizations of S_1 , r_1 , and $p(t, \tau)$ with the use of Monte Carlo simulations is outlined in Zaglauer and Bauer (2008).

3.3 Results

Table 3 displays the portfolio of policies for our stylized insurer. For simplicity and without loss of generality, we assume that the company holds an equal number of term/endowment/annuity contracts for different age/term combinations and that the face values coincide—of course, generalizations are possible. The initial capital level is set to $E = \$20,000,000$. The insurer's assets and liabilities at time zero, A_0 and V_0 , are calculated at \$1,124,603,545 and \$1,104,603,545, respectively.

We consider three different approaches to modeling mortality risk: (i) a deterministic evolu-

	Deterministic Mortality	Factor Model	Non-Parametric Model
95% VaR (no hedging)	\$60,797,835	\$61,585,667	\$62,802,167
95% VaR (financial hedging)	\$3,201,921	\$9,871,987	\$10,049,401

Table 4: Economic capital for different investment strategies

tion of mortality given by the life table at time zero (2008), $\{\tau p_x(0)\}$; (ii) the parametric factor model (5); and, (iii), the non-parametric mortality model also introduced in Section 2. Within each approach, we use 50,000 simulations of the assets and liabilities to generate realizations of the loss (L) at time 1, where in addition to financial and systematic mortality risk, we also consider unsystematic mortality risk by sampling the number of deaths within each cohort. Finally, we can calculate the EC via the resulting empirical distribution functions and the given risk measure ρ . In particular, for VaR we rely on the empirical quantile. Table 4 displays the results for two assumptions regarding the insurer's investments.

For the results in the first row of Table 4, we assume that the company does not optimize its asset allocation, but invests a fixed 30% in the equity index (see e.g. ACLI (2011)) and the rest in government bonds to match the duration of its liabilities (at 10.2560). Without stochastic mortality, we find EC levels of around \$60,000,000, which suggests that the current capital position of \$20,000,000 is not sufficient—i.e. the firm is undercapitalized. Surprisingly, including systematic mortality risk appears to have little influence on the results in this case: the EC increases by only \$787,832 (1.30%) or \$2,004,332 (3.30%) when introducing mortality risk via the factor mortality model or non-parametric mortality model, respectively.

However, this changes dramatically when we allow the insurer to pursue a more refined allocation strategy to better manage the financial risk exposure. In the second row of Table 4, we display the results when the insurer optimally chooses (static) asset weights in order to minimize the EC. The corresponding portfolio weights are displayed in Table 5. We find that while the EC level decreases vastly under all three mortality assumptions so that the company is solvent according to the 95% VaR capital requirement ($EC \leq AC_0$), the relative impact of systematic mortality risk now is highly significant. More precisely, the (minimized) EC increases by 208.31% (to \$9,871,987) and 213.86% (to \$10,049,401) if systematic mortality risk is considered via the factor model and non-parametric model, respectively. This underscores an important point in the debate about the economic relevance of mortality and longevity risk: While financial risk indices may be more

	Deterministic Mortality	Factor Model	Non-Parametric Model
Stock	0.2%	1.5%	0.9%
5-year Bond	2.5%	0.1%	0.5%
10-year Bond	87.3%	88.0%	90.8%
20-year Bond	10.0%	10.4%	7.8%

Table 5: Financial hedging—optimal weights

volatile and thus may dominate systematic mortality risk, there exist conventional methods and (financial) instruments to hedge against financial risk.

Of course, naturally the question arises if we can use a similar approach to protect against systematic mortality risk, *either* by expanding the scope of securities considered on the asset side toward mortality-linked securities *or* by adjusting the composition of the insurance portfolio on the liability side. The former approach has been considered in a number of papers (see e.g. Cairns et al. (2013), Li and Luo (2012), and references therein), but a liquid market of corresponding instruments is only slowly emerging. The latter approach—which is commonly referred to as *natural hedging* (Cox and Lin, 2007) and which is in the focus of this paper—has also received attention in the insurance literature and is reported to perform well (Wetzel and Zwiesler, 2008; Tsai et al., 2010; Wang et al., 2010; Gatzert and Wesker, 2012).

Before we explore this approach in more detail in the next section, it is helpful to emphasize that the results for the two mortality models—the non-parametric model and the parametric factor model—are very similar across both cases. This may not be surprising as these models originate from the same framework. Essentially, one can interpret the factor model as a parsimonious approximation of the non-parametric model that nevertheless captures the majority of the “important” variation—with resulting statistical advantages, e.g. in view of its estimation. However, there are also pitfalls for its application in the context of analyzing the performance of hedges as we will see in the next section.

4 Natural Hedging of Longevity Risk

Akin to the previous section, we consider the possibility of reducing the risk exposure by adjusting the portfolio weights. However, while there we adjusted *asset* weights in order to minimize the exposure to financial risk, here we focus on adjusting the composition of the *liability* portfolio in order to protect against mortality/longevity risk. More specifically, we fix the number of en-

dowment and annuity contracts in the life insurer’s portfolio to the same values as in the previous section (see Table 3) and vary the number of term-life policies n^{term} , where—for simplicity—we assume $n^{\text{term}} \equiv n_{x,i}^{\text{term}}$ is constant across age/term combinations (x, i) . For each n^{term} , the EC is then calculated analogously to the previous section under the assumption that the insurer hedges against financial risk, i.e. we determine the “optimal” asset allocation separately for each n^{term} . Finally, we determine the optimal number of term-life policies, $n^{\text{term}*}$, that minimizes the EC for the life insurer.⁹

We start by considering the factor mortality model and compare it to the case without stochastic mortality risk. Figure 1 shows the EC as a function of the number of term-life policies in the insurer’s portfolio n^{term} (we will refer to this as an “EC curve” in what follows). We first find that in the case of no systematic mortality risk (deterministic mortality), EC increases in the number of term policies. The reason is twofold: On the one hand, an increase leads to higher premiums and, thus, assets, which increases asset risk. On the other hand, although there is no systematic mortality risk, the number of deaths in each cohort is a random variable due to non-systematic mortality risk—which clearly increases in the number of policies. In contrast, under the stochastic factor mortality model, the EC is a convex function of n^{term} that initially decreases and then increases sharply, i.e. it is U-shaped. The optimal number of policies, $n^{\text{term}*}$, is approximately 60,000 and the corresponding minimal EC is \$4,165,973, which is only slightly larger than the corresponding EC level under deterministic mortality (\$4,128,345). Therefore, in line with other papers on natural hedging, it appears that an appropriately composed insurance portfolio can serve well for hedging against systematic mortality/longevity risk.

However, when repeating the same exercise based on the non-parametric forecasting model, the situation changes considerably. As is also depicted in Figure 1, in this case we can only observe a very mild effect of natural hedging when n^{term} is relatively small, and it is far less pronounced compared with the factor model. In particular, at the optimal term-insurance exposure $n^{\text{term}*} = 60,000$, under the factor model, the capital level is at \$13,872,739 for the nonparametric model, which is far greater than the corresponding capital level under deterministic mortality (\$4,128,345).

The intuition for this result is as follows: As indicated at the end of Section 3.3, the two models behave quantitatively alike in “normal” circumstances and particularly yield similar capital levels for the initial portfolio. This is not surprising since the rationale behind the single factor model—akin to other single factor models such as the Lee-Carter model—is that the majority of the variation across ages and terms can be explained by the leading factor (85% for our dataset).

⁹Note that we implicitly assume that the insurer can place arbitrarily many term-life insurance policies in the market place at the same price—which may be unrealistic for large $n^{\text{term}*}$. Moreover, we assume that underwriting profits and losses can be transferred between different lines of business and that there are no other technical limitations when pursuing natural hedging. However, such limitations would only cast further doubt on the natural hedging approach, so we refrain from a detailed discussion of these aspects.




Figure 1: Optimal longevity hedging

Essentially, the residuals for lower ages are small in absolute terms and mostly unsystematic, whereas the residuals for higher ages (beyond 50) are relatively large in absolute terms and mostly systematic. And the latter are responsible for the high proportions of the variation explained in absolute terms. However, under the natural hedging approach, the large exposure in the term-life lines leads to a considerable rescaling of the profile of the residuals across terms and ages, so that this similarity breaks down. In particular, the residuals for lower-age groups become increasingly important, which in turn are considerably influenced by higher-order factors including but also beyond the second factor—some of which do not carry a systematic shape at all. Thus, for the analysis of the effectiveness of natural hedging, the consideration of higher order/non-systematic variation indeed might be important.

Again, we would like to emphasize that this is not general criticism of these models. For many applications, such as forecasting mortality rates, abstracting from these small and unsystematic variations is expedient. We solely challenge the reliance on low-dimensional factor models for the analysis of hedging performance. And, indeed, our results indicate that natural hedging may not be as effective as asserted in the existing literature.

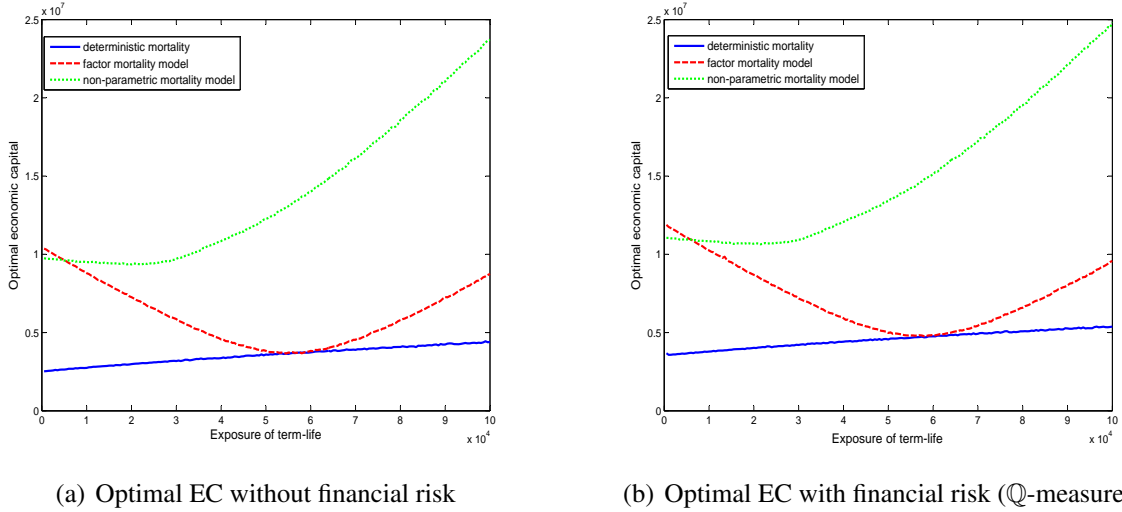
5 Robustness of the Results

Of course, the question may arise to what extent the results on the performance of natural hedging are driven by the details of our setup. Thus, in this section, we repeat the EC calculations under modified assumptions. In particular, we examine the impact of financial risk on the results, we consider modifications of our mortality models, and we derive EC curves for alternative mortality models.

5.1 The Impact of Financial Risk

To analyze the role played by financial risk in the results, we recalculate the EC levels for different term-life exposures under a deterministic evolution of the asset side—so that the results are solely driven by systematic and unsystematic mortality risk. More precisely, in our asset model (Equation (8)), we set both volatility terms σ_A and σ_r to zero, and we use the risk-neutral drift parameters throughout. In particular, the equity index S is now risk-less and returns the risk-free rate. Figure 2(a) displays the corresponding optimal EC curves.

We find that the EC levels under each of the mortality models are similar to the case with financial risk (Figure 1). For large values of n^{term} , the EC here even exceeds the corresponding values in the case with financial risk, which appears surprising at first glance. The reason for this observation is the change of probability measure in the financial setting and the associated risk premia



(a) Optimal EC without financial risk

(b) Optimal EC with financial risk (\mathbb{Q} -measure)

Figure 2: The impact of financial risk

paid over the first year. For comparison, we also plot the EC curves for all mortality assumptions when relying on the \mathbb{Q} -measure throughout in Figure 2(b). We find that the (hypothetical) EC, as expected, is always greater than without the consideration of financial risk, though the difference is not very pronounced. This indicates that the static hedging procedure eliminates most of the financial risk or, in other words, that financial risk does not contribute too much to the total EC.

5.2 Modifications of the Mortality Models

As indicated in Footnote 5 (Section 3.2), the relatively small size of the sample underlying our non-parametric forecasting approach may be problematic for certain applications such as estimating VaR for a high confidence level. To analyze the impact of the small sample size on our results, we follow the approach also described in Section 2 that relies on the additional assumption that $\bar{F}(t_N, t_{N+1})$ is Gaussian distributed. Then, rather than sampling $\bar{F}(t_N, t_{N+1})$ from the empirical realizations, we generate random vectors with the mean vector and the covariance matrix estimated from the underlying sample. Figure 3(a) shows the resulting EC curve in comparison to the deterministic mortality case. We find that the results are very similar to the non-parametric model underlying Figure 1. In particular, there is only a rather mild effect of natural hedging when n^{term} is small, and the economic capital levels considerably exceed those for the deterministic mortality case.

As also indicated in Section 2, the competing model used in the calculations in Section 3 and 4, while originating from a factor analysis, presents a self-consistent, parametric approximation. In particular, the entire term structure is driven by only a handful parameters in this case, so that it is not immediately clear what aspects of the model are responsible for the results. Thus, as

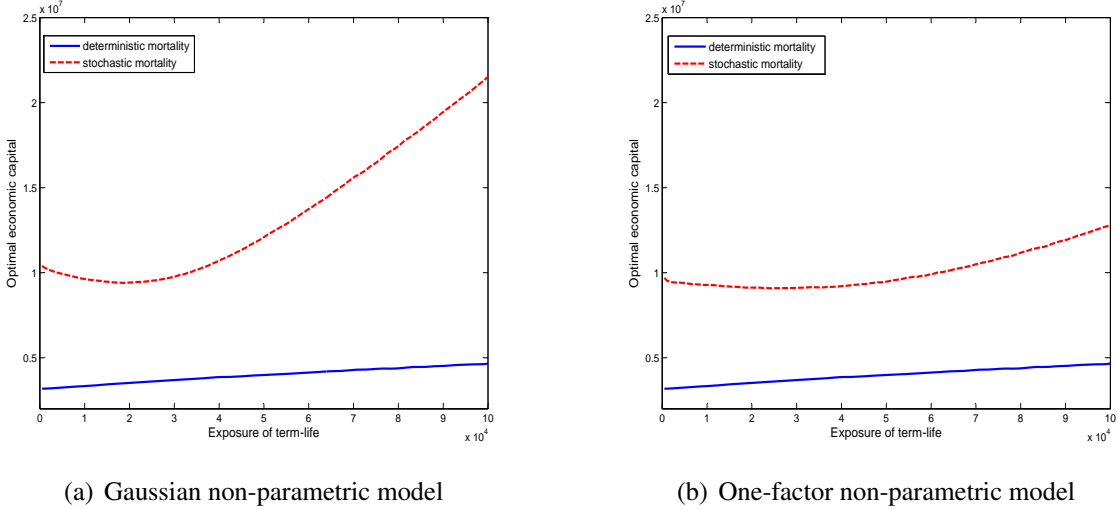


Figure 3: Modifications of the mortality models

an intermediary step, we also provide results based on a (high-dimensional) single-factor model. In particular, instead of directly relying on the leading principal component, we estimate a one-factor model following the approach from Bai and Li (2012), which allows for heteroscedasticity in the error term so that some variation will also be picked up for lower ages. More precisely, this approach posits a factor form

$$\bar{F}(t_j, t_{j+1}) = \alpha + \beta * \lambda_{t_j} + \delta_{t_j}, j = 1, \dots, N - 1,$$

where α and β are 5050×1 vectors, $\mathbb{E}[\delta_t] = 0$ and $E[\delta_t \times \delta_t'] = \Sigma_{error} = \text{diag}\{\sigma_1^2, \dots, \sigma_{5050}^2\}$, which is estimated via Maximum-Likelihood. Here, we employ the leading factor from the Principal Component Analysis as the starting value in the numerical optimization of the log-likelihood, and the resulting factors overall are very similar.

Figure 3(b) provides the EC curve based on this factor model. We find that the effect of natural hedging is far less pronounced than for the parametric factor model, and the Economic Capitals are considerably higher than for the deterministic mortality case throughout. However, we also observe that in contrast to the non-parametric approach in Figure 1, the EC curve is “flat” in the sense that the increased exposure to term life insurance only has little effect on the economic capital level. This indicates that the factor loadings for the lower age range are very small, or, in other words, that much of the variation is driven by higher order variations. For the parametric factor model, on the other hand, the parametric form is fit across all terms and ages, which appears to yield a more significant relationship between low and high ages in the first factor. Thus, the parametric nature of the model also seems to be an important driver for the positive appraisal of the natural hedging approach, at least in our setting.

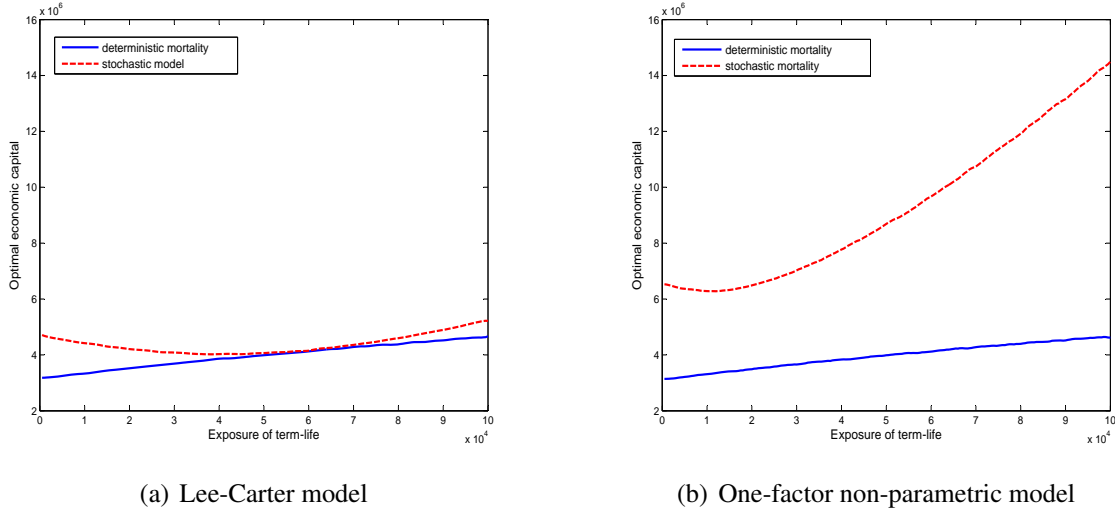


Figure 4: Alternative mortality models

5.3 Alternative Mortality Models

As a final robustness check, we repeat the calculations for alternative mortality models that do not fall within our framework. We start by providing results for the (stochastic) Lee-Carter model as another model where all variation is driven by a single factor. More precisely, we use the Lee-Carter parameters estimated at time t_N —which also serve for generating the corresponding generation life table that we use for the calculation of time zero premiums—and simulations of the $\kappa_{t_{N+1}}$ to generate life tables at time one with each based on the median projection starting from a simulated value of $\kappa_{t_{N+1}}$. Figure 4(a) presents the resulting EC curve.

We make two primary observations. On the one hand, the EC curve exhibits a U-shape similar to the parametric factor model in Figure 1, i.e. natural hedging again is found to be highly effective under this model. In particular, the optimal exposure to term life policies $n^{\text{term}*}$ again is around 60,000 with a corresponding minimum capital of \$4,150,010. This finding is not surprising since it was exactly this positive appraisal of the natural hedging approach in previous contributions that serves as the primary motivation of this paper. On the other hand, we observe that the magnitude of EC is considerably lower than in the models considered in Section 3 and 4. Again, this finding is not surprising since it is exactly the underestimation of the risk in long-term mortality trends that serves as the motivation for the underlying approach in Zhu and Bauer (2013).

As a second non-parametric modeling approach, we implement the model proposed by Li and Ng (2010) that relies on bootstrapping one-year mortality reduction factors $r_{x,t_j} = \frac{m_{x,t_{j+1}}}{m_{x,t_j}}$ for $j = 1, \dots, N - 1$ and $m_{x,t}$ being the central death rate for age x in year t . More precisely, as proposed in Li and Ng (2010), we use a block Bootstrap method with a block size of two to capture the serial dependency in the data. For consistency with the other mortality models in this

paper, we use 30 years of the historical data (1978-2007) and ages x ranging from 0 to 100. Figure 4(b) shows the resulting EC curves, where of course the deterministic curve is calculated based on a generation table compiled also using this model.

The capital levels are lower than for the approaches considered in Section 3 and 4 though larger than for the (stochastic) Lee-Carter model, which is due to differences in the model structures. However, the observations regarding natural hedging are in line with our results from Section 4. More precisely, the effect is not very pronounced under the non-parametric model—we only see a very mild U-shape—and the EC level increases considerably for higher values of n^{term} .

6 Conclusion

In this paper, we analyze the effectiveness of natural hedging in the context of a stylized life insurer. Our primary finding is that higher order variations in mortality rates may considerably affect the performance of natural hedging. More precisely, while results based on a parametric single factor model imply that almost all longevity risk can be hedged by appropriately adjusting the insurance portfolio (in line with the existing literature), the results are far less encouraging when including higher order variations via a non-parametric mortality forecasting model.

Of course, this is not a general endorsement of these more complicated models. Simple (or parsimonious) models may have many benefits in view of their tractability, their statistical properties, or their forecasting power. We solely show that relying on “simple” models for analyzing the performance of hedges may be misleading since they contain assumptions on the dependence across ages that are not necessarily supported by the data.

At a broader level, we believe our results call for more caution toward model-based results in the actuarial literature in general.

References

American Council of Life Insurers, 2011. Life Insurers Fact Book 2011.

<https://www.acli.com/Tools/IndustryFacts/LifeInsurersFactBook/Documents/2011FactBook.pdf>

Bai, J., Li, K., 2012. Statistical Analysis of Factor Models of High Dimension. *Annals of Statistics*, 40: 436-465.

Bauer, D., Benth, F.E., Kiesel, R., 2012a. Modeling the Forward Surface of Mortality. *SIAM Journal on Financial Mathematics*, 3: 639-666.

- Bauer, D., Reuss, A., Singer, D., 2012b. On the Calculation of the Solvency Capital Requirement Based on Nested Simulations. *ASTIN Bulletin*, 42: 453-499.
- Bayraktar, E., Young, V.R., 2007. Hedging life insurance with pure endowments. *Insurance: Mathematics and Economics*, 40: 435-444.
- Blake, D., Cairns, A.J.G., Dowd, K., 2006. Living with Mortality: Longevity Bonds and other Mortality-Linked Securities. *British Actuarial Journal*, 12: 153-197.
- Booth, H., Maindonald, J., Smith, L., 2002. Applying Lee-Carter under conditions of variable mortality decline. *Population Studies*, 56: 325-336.
- Cairns, A.J.G., Blake, D., Dowd, K., 2006a. Pricing Death: Frameworks for the Valuation and Securitization of Mortality Risk. *ASTIN Bulletin*, 36: 79-120.
- Cairns, A.J.G., Blake, D., Dowd, K., 2006b. A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration. *Journal of Risk and Insurance*, 73: 687-718.
- Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G..D., 2013. Longevity hedge effectiveness: A decomposition. *Quantitative Finance*, forthcoming.
- Cox, S.H., Lin, Y., 2007. Natural Hedging of Life and Annuity Mortality Risks. *North American Actuarial Journal*, 11: 1-15.
- Efron, B., 1979. Bootstrap Methods: Another Look at the Jackknife. *Annals of Statistics*, 7: 1-26.
- Gatzert, N., Wesker, H., 2012. The Impact of Natural Hedging on a Life Insurer's Risk Situation. *Journal of Risk Finance*, 13: 396-423.
- Lee, R.D., Carter, L.R., 1992. Modeling and Forecasting U. S. Mortality. *Journal of the American Statistical Association*, 87: 659-675.
- Li, J.S.-H., Ng, A.C.-Y., 2010. Canonical Valuation of Mortality-Linked Securities. *Journal of Risk and Insurance*, 78: 853-884.
- Li, J.S.-H., Luo, A., 2012. Key q-duration: A framework for hedging longevity risk. *ASTIN Bulletin*, 42: 413-452.
- Tsai, J.T., Wang, J.L., Tzeng, L.Y., 2010. On the optimal product mix in life insurance companies using conditional value at risk. *Insurance: Mathematics and Economics*, 46: 235-241.

- Wang, J.L., Huang, H.C., Yang, S.S., Tsai, J.T., 2010. An Optimal Product Mix for Hedging Longevity Risk in Life Insurance Companies: The Immunization Theory Approach. *The Journal of Risk and Insurance*, 77: 473-497.
- Wetzel, C., Zwiesler, H. J., 2008. Das Vorhersagerisiko der Sterblichkeitsentwicklung Kann es durch eine geeignete Portfoliozusammensetzung minimiert werden? *Blätter DGVM*, 29: 73-107.
- Wilmoth, J., 1993. Computational Methods for Fitting and Extrapolating the Lee-Carter Model of Mortality Change. Technical Report, University of California, Berkeley.
- Zaglauer, K., Bauer, D., 2008. Risk-neutral valuation of participating life insurance contracts in a stochastic interest rate environment. *Insurance: Mathematics and Economics*, 43: 29-40.
- Zhu, N., Bauer, D., 2011. Applications of Gaussian Forward Mortality Factor Models in Life Insurance Practice. *The Geneva Papers on Risk and Insurance – Issues and Practice*, 36: 567-594.
- Zhu, N., Bauer, D., 2013. Modeling and Forecasting Mortality Projections. Working paper, Georgia State University.