A Cautious 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

December 2012

Abstract

In this note, 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 averts 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.

JEL classification: G22; G32; J11.

Keywords: Natural Hedging, Longevity Risk, Mortality Forecasting, Non-Parametric 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.

[†] 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" (Wetzel and Zwiesler (2008), Tsai et al. (2010), Wang et al. (2010), and 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 (cf. Lee and Carter (1992)) or the CBD model (cf. 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 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 forecasting model on the effectiveness of natural hedging, in this note, 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 (2012). 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. The meager performance of natural hedging can be viewed as further evidence for endorsing market-based solutions for managing longevity risk.

The remainder of the note is structured as follows: Section 2 briefly introduces the considered

[1]Similar 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.

mortality forecasting models. Section 3 discusses the calculation of the economic capital for a stylized life insurance company. Section 4 revisits the natural hedging approach within our framework. Finally, Section 5 concludes.

2 Mortality Forecasting Models

We commence by introducing the mortality forecasting models that will be used throughout this note. We consider two representative models within the forward-mortality framework developed in Zhu and Bauer (2012), namely a parametric single-factor model as well as a nonparametric model for the annualized mortality innovations. Of course, other low-dimensional mortality factor models such as the Lee-Carter model or the CBD model theoretically also fit in the discussion here. However, 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* (cf. 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.

Assume that at times $t \in \{t_j, j = 1, 2, \ldots, N\}$, we are given (forward) survival probabilities $\{{}_\tau p_x(t) | \, (\tau, x) \in \mathcal{C}\}$ for a certain population, where ${}_\tau p_x(t)$ denotes the probability for an x-year old individual to survive for τ periods until time $t+\tau$ based on the information up to time t, and $\mathcal{C}$ denotes a (large) collection of term/age combinations $(\tau, x) \in [0, \infty) \times [0, \infty)$.[2] Mathematically, this is equivalent to

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

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.

Let $\tilde{\mathcal{C}} \subset \mathcal{C}$, $|\tilde{\mathcal{C}}| = K$, be a sub-collection of $\mathcal{C}$ such that for $(\tau, x) \in \tilde{\mathcal{C}}$, $(\tau + 1, x)$, $(\tau + t_{j+1} - t_j, x - t_{j+1} + t_j)$, and $(\tau + 1 + t_{j+1} - t_j, x - t_{j+1} + t_j) \in \mathcal{C}$ for all $j \in \{1, 2, \ldots, N - 1\}$. For

[2] As discussed in Zhu and Bauer (2012), relying on empirical observations of best-estimate generational/cohort life tables is problematic since these data are compiled rather infrequently and inconsistently. Thus, we instead use readily-available period mortality data and a fixed forecasting methodology to generate a time series of generational life tables for our applications. In particular, in this note we invariably rely on the Lee-Carter model and US female mortality experience for generating (deterministic) mortality forecasts. We refer to Zhu and Bauer (2012) for a discussion of the rationale for this choice as well as potential alternatives.

each $(\tau, x) \in \tilde{\mathcal{C}}$, we define

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

Hence, $F\left(t_j, t_{j+1}, (\tau, x)\right)$ 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 $\bar{F}(t_j, t_{j+1}) = \left(F(t_j, t_{j+1}, (\tau, x))\right)_{(\tau,x)\in\tilde{\mathcal{C}}}$, $j = 1, 2, \ldots, N-1$.

Proposition 2.1 in Zhu and Bauer (2012) shows that for a time-homogeneous (stochastic) evolution of the survival probabilities and with equidistant evaluation dates, i.e. $t_{j+1} - t_j \equiv \Delta$, the $\bar{F}(t_j, t_{j+1})$, $j = 1, \ldots, N-1$, are independent and identically distributed (iid). Therefore, in this case a non-parametric mortality forecasting methodology is immediately given by treating the observations $\bar{F}(t_j, t_{j+1})$, $j = 1, \ldots, N-1$, as a *bootstrap sample* (see e.g. Efron (1979)). More precisely, with Equation (1), we can generate simulations for the generational mortality tables at time t_{N+1}, $\{{}_{\tau}p_x(t_{N+1})\}$, by sampling (with replacement) $\bar{F}(t_N, t_{N+1})$ from $\left\{\bar{F}(t_j, t_{j+1}),\ j = 1, \ldots, N-1\right\}$ in combination with the known generational mortality tables at time t_N, $\{{}_{\tau}p_x(t_N)\}$. This serves as the algorithm for generating our non-parametric mortality forecasts.

To introduce corresponding factor models, it is possible to simply perform a factor analysis of the iid sample $\left\{\bar{F}(t_j, t_{j+1}),\ j = 1, \ldots, N-1\right\}$, 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 (2012), 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 generational mortality table at time t_N.

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

$$\mu_t(\tau, x) = -\frac{\partial}{\partial \tau} \log\left\{{}_{\tau}p_x(t)\right\},\ (\tau, x) \in \mathcal{C}$$

so that we have

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

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

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

where α and σ are adequate, 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

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

and for a 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 (2012) propose the following specification for the volatility structure in a single-factor model:

$$\sigma(\tau, x) = k\, \frac{\exp\{c(x + \tau) + d\}}{1 + \exp\{c(x + \tau) + d\}}\, (a + \tau)\, e^{-b\tau}. \tag{5}$$

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 (2012) 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

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

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 con $\in \{\text{ann}, \text{term}, \text{end}\}$. And $\ddot{a}_x(t)$, $A_{x:\overline{i}|}(t)$, 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 (cf. Bauer et al. (2012b))

$$EC = \rho\, \underbrace{(AC_0 - p(0,1)\, AC_1)}_{L},$$

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

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

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

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

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 introduced in the previous section, 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 95 years to compile 46 consecutive generation life tables (years 1963-2008), with each generated independently from the mortality experience of the previous 30 years. That is, the first table (year 1963) uses mortality data from years 1933-1962, the second table uses years 1934-1963, and so forth.

Having obtained these generation tables $\{_\tau p_x(t_j)\}$, $j = 1, \ldots, N = 46$, we derive the time series of $\bar{F}(t_j, t_{j+1}), j = 1, 2, \ldots, 45$, which serve as the bootstrap sample for our non-parametric forecasting methodology and as the basis for the parameter estimates of our mortality factor model (cf. Zhu and Bauer (2012)). In particular, time t_N (2008) corresponds to time zero whereas time t_{N+1} (2009) corresponds to time one in our EC framework. Table 1 displays the parameter estimates of the factor mortality forecasting model.

Parameters				
k	c	d	a	b
0.1746	0.0637	-13.6507	30.2124	0.0030

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

For the asset side, we assume that the insurer only invests in 1, 3, 5, and 10-year US government bonds as well as an equity index (S&P 500) $S = (S_t)_{t\geq 0}$. For the evolution of these assets, we assume a generalized Black-Scholes model with stochastic interest rates (Vasicek model), that is,

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

under $\mathbb{P}$

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

where μ, σ_A, κ, γ, $\sigma_r > 0$, $\rho \in [-1, 1]$, and $(B_t^{(1)})$ and $(B_t^{(2)})$ are independent Brownian motions that are independent of (W_t). Moreover, 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,[4] treasury bills (3 months), and government bonds with maturities of 1 year, 3 years, 5 years, and 10 years.[5] The parameter estimates are displayed in Table 2.

Parameters							
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

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

$$R_1 = \omega_1 \frac{S_1}{S_0} + \omega_2 \frac{1}{p(0,1)} + \omega_3 \frac{p(1,2)}{p(0,3)} + \omega_4 \frac{p(1,4)}{p(0,5)} + \omega_5 \frac{p(1,9)}{p(0,10)},$$

where $\omega_i,\ i = 1, \ldots, 5$, are the company's proportions of assets invested in each asset. A 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

[4]Downloaded on 08/26/2012 from Yahoo! Finance, `http://finance.yahoo.com`

[5]Downloaded on 08/26/2012 from the Federal Reserve Economic Data (FRED), `http://research.stlouisfed.org/fred2/`.

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

Table 3: Portfolio of policies for the stylized life insurer

contracts for different age/term combinations and that the face values coincide—of course, generalizations are possible. The initial capital level is set to $E = \$20,000,000$.

We consider three different approaches for modeling mortality risk: (i) a deterministic evolution of mortality given by the life table at time zero (2008), $\{{}_{\tau}p_x(0)\}$; (ii) the mortality factor model (5); and, (iii), the non-parametric mortality model introduced in the previous section. 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 simply invests in the equity market as well as each type of government bond at equal proportions, i.e. $\omega_i \equiv 20\%$. We find EC levels of around \$100,000,000, which implies that the current equity 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 $\$714,810$ (0.72%) or $\$2,284,554$ (2.29%) when introducing

	Deterministic Mortality	Factor Model	Non-Parametric Model
95% VaR (no hedging)	$\$99,870,175$	$\$100,584,985$	$\$102,154,729$
95% VaR (financial hedging)	$\$8,519,083$	$\$15,507,323$	$\$15,148,969$

Table 4: Economic capital for different investment strategies

mortality risk via the factor mortality model or non-parametric mortality model, respectively.[6]

However, this changes dramatically when we consider an insurance company that pursues an active allocation strategy by matching assets and liabilities. In the second row of Table 4, we display the results under the assumption that the insurer chooses (fixed) asset weights in order to minimize EC. The corresponding portfolio weights are illustrated in Table 5.[7] We find that while the EC level decreases vastly under all three mortality assumptions so that the company is solvent in all cases ($EC \leq AC_0$), the relative impact of systematic mortality risk now is highly significant. More precisely, the (minimized) EC increases by 82.03% (to $\$15,507,323$) and 77.82% (to $\$15,148,969$) if systematic mortality risk is considered via the factor model and non-parametric model, respectively. This underscores a very important point in the debate about the economic relevance of mortality and longevity risk: While financial risk indices may be more 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. (20012), 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

[6]Although the EC levels roughly accord with the corresponding calculations in Zhu and Bauer (2011), including mortality risk there did increase EC more considerably. This difference is due to the underlying mortality data in combination with the Lee-Carter method for generating deterministic forecasts: While Lee-Carter squares relatively well with US female mortality, for England/Wales male data we obtain a relatively high volatility due to inconsistencies of the Lee-Carter based forecasts and the observations. We refer to Zhu and Bauer (2012) for more details.

[7]The observation that—in this case—it is optimal to invest the vast majority of the assets in the longest-termed bond is a peculiarity of the financial market parameters in combination with the insurance portfolio. Sensitivity analyses show that generally, under the one-factor interest model considered here, optimal strategies involve large positions in (up to) two of the bonds and almost no funds invested in the equity funds.

	Deterministic Mortality	Factor Model	Non-Parametric Model
Stock	0.5%	0.9%	0.7%
1-year Bond	0.6%	0.0%	0.0%
3-year Bond	0.0%	0.0%	0.0%
5-year Bond	1.0%	0.0%	0.0%
10-year Bond	97.9%	99.1%	99.3%

Table 5: Financial hedging—optimal weights

natural hedging (cf. Cox and Lin (2007)) and which is in the focus of this note—has also received attention in the insurance literature and is reported to perform well (see e.g. Wetzel and Zwiesler (2008), Tsai et al. (2010), Wang et al. (2010), or Gatzert and Wesker (2012)). Before we explore this opportunity in more detail in the next section, it is helpful to note that the similarities between the results for the two considered forecasting models to some extent are attributable to the fact that 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.

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 endowment and annuity contracts in the life insurer's portfolio to the same values as in the previous section (cf. Table 3) and vary the number of term-life policies n^{term}, where—for simplicity—we assume $n^{\text{term}} \equiv n^{\text{term}}_{x,i}$ 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.[8]

[8] Note that we implicitly assume that there are no demand-side effects when increasing the supply of life insurance. 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.

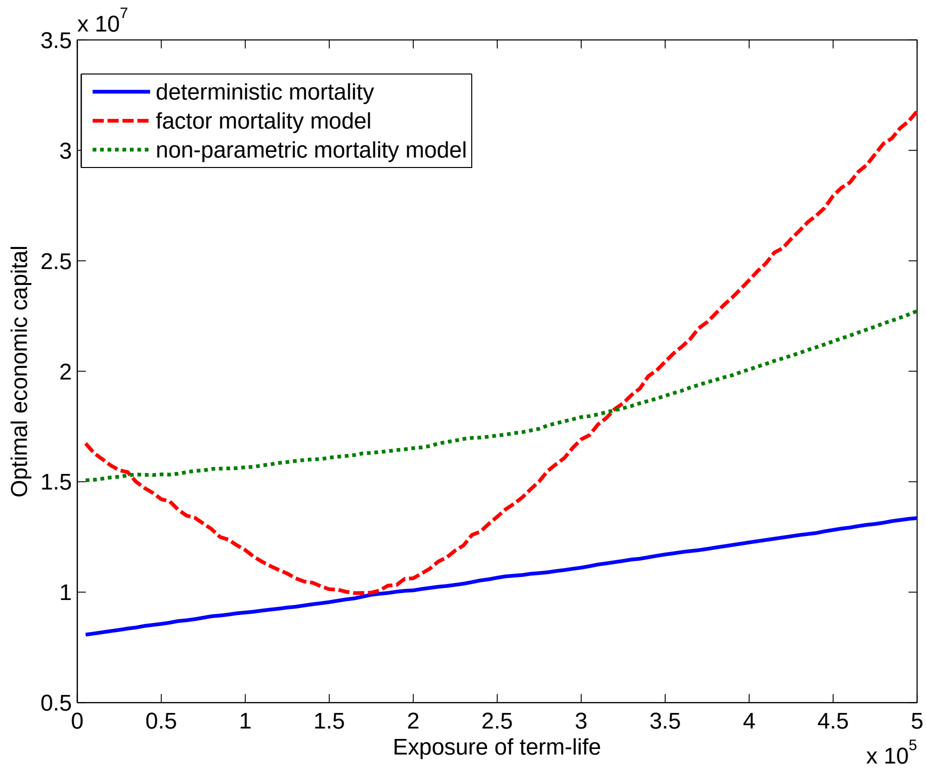

Figure 1: Optimal longevity hedging

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 find that even 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. The optimal number of policies, $n^{\text{term}*}$, approximately equals $16,500$ and the corresponding minimal EC is $\$9,954,353$, which is only slightly larger than the corresponding EC level under deterministic mortality ($\$9,909,874$). 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 EC again is strictly increasing in n^{term}. Although we observe a very mild effect of natural hedging as the distance between the capital level for the non-parametric model and deterministic mortality first decreases and then increases, it is far less pronounced than for the factor model. In particular, at the optimal term-insurance exposure $n^{\text{term}*} = 16,500$ under the factor model, the capital level is at $\$16,259,641$ for the nonparametric model, which is far greater than the corresponding capital level under deterministic mortality ($\$9,909,874$).

The rational for this observation is quite straightforward: Our single factor model (akin to other single factor models such as the Lee-Carter model) assumes that all the variation in mortality rates is driven by a single source of randomness such that mortality rates for all ages evolve in lockstep—and thus natural hedging is very effective. As indicated before, this assumption is warranted by a factor analysis that implies that the majority of the "important" variation when generating forecasts is driven by a single factor. However, for the analysis of the effectiveness of natural hedging, the consideration of higher order/non-systematic variation indeed is *important* since it may imply disparate movements of mortality rates at different ages—particularly because it makes up a significant part of all the observed variation in mortality rates for young ages. Hence, natural hedging may not be as effective as indicated in the existing literature.

5 Conclusion

In this note, 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. The key point is that "simple" mortality models—which may serve very well for certain purposes—strip away certain aspects in the data that are important for the analysis of natural hedging. At a broader level, we believe our results call for more caution toward model-based results in the actuarial literature in general.

References

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.

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.

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., 2012. 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, forthcoming.

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, forthcoming.

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 DGVFM, 29: 73-107.

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., 2012. Coherent Modeling of the Risk in Mortality Projections: A Semi-Parametric Approach. Working paper, Georgia State University.