SOCIETY OF ACTUARIES

LIVING to 100
SOCIETY OF ACTUARIES INTERNATIONAL SYMPOSIUM

Mortality and Longevity

Aging and Retirement

Quantification and Management of Longevity Risk in China

January 2020

Quantification and Management of Longevity Risk in China

AUTHOR

Johnny Li, FSA, ACIA, Ph.D.
University of Melbourne

Kenneth Q. Zhou, ASA, ACIA, Ph.D.
Arizona State University

Wai-Sum Chan, FSA, CERA, Ph.D.
The Chinese University of Hong Kong

SPONSOR

Society of Actuaries:
Aging and Retirement Research
Committee on Life Insurance Research
Financial Reporting Section
Mortality and Longevity Research
Product Development Section

Quantification and Management of Longevity Risk in China

Johnny S.-H. Li[1,2], Kenneth Q. Zhou[2,3], and Wai-Sum Chan[4]

[1] Department of Economics, The University of Melbourne, Australia

[2] Department of Statistics and Actuarial Science, University of Waterloo, Canada

[3] School of Mathematical and Statistical Science, Arizona State University, USA

[4] Department of Finance, The Chinese University of Hong Kong, Hong Kong

Abstract: As the annuity market in China develops, the Chinese insurance industry is increasingly exposed to longevity risk. The recently introduced China Risk Oriented Solvency System (C-ROSS) requires both domestic and foreign insurers operating in China have to hold solvency risk capital for the longevity risk they take. In this paper, we study how the longevity risk facing insurers operating in China can be quantified using stochastic mortality models that are developed to suit the demographic situation in China. We also study how longevity risk management solutions such as securitization may reduce an insurer's C-ROSS solvency capital.

Keywords: *C-ROSS; Delta hedging; Solvency risk capital; Stochastic mortality modeling*

1 Introduction

Since the launch of its reformation and opening policy, China has made great strides in reducing mortality. According to the World Bank, the life expectancy at birth for the unisex population of China has increased from 59.085 years in 1970 to 76.47 years in 2017. The reduction in Chinese pension insurers' mortality is even more remarkable. A comparison between the 1990-1993 and 2010-2013 insurance life tables for Chinese pension insurers shows an increase in life expectancy of approximately 8.7 years over the 20-year period, which is much faster than the typical rate of increase in life expectancy (3 years per decade) in the developed world. The gain in life expectancy is certainly an important social achievement, but it also poses a threat to the population's retirement income security. The problem is exacerbated by China's infamous one-child policy, which leaves the older generations increasingly dependent on retirement funds.

The current urban pension system in China is based on the 1997 State Council Document No. 26. In line with the recommendations provided by the World Bank (1994), it is a three-pillar system comprised of a mandatory publicly managed pillar, a mandatory privately managed pillar, and a voluntary pillar. The first pillar is a defined-benefit public plan that includes a pay-as-you-go portion financed by employer contributions equal to 20% of wages plus a funded portion supported by employee contributions equal to 8% of wages. The second pillar is formed by defined-contribution occupational plans that are known as Enterprise Annuities. According to the Ministry of Labour and Social Security Statement No. 20 published in 2004, Enterprise Annuities are funded by employer and employee contributions of up to one sixth of the total salary roll. All of the contributions go directly to Enterprise Annuity accounts, which are decumulated upon retirement. We refer interested readers to Cai and Cheng (2014) and Dorfman et al. (2012) for further information about pension systems in China.

Although the first pillar of the urban pension system offers individuals some protection against longevity risk by providing a modest lifetime income, the second pillar provides no protection against longevity risk as the benefit from an Enterprise Annuity must be received as either a lump sum or instalments with a fixed term. For additional protection against the risk of outliving financial resources, individuals' only option is to purchase life annuities from insurance companies. For this reason, the life annuity market in China has a massive growth potential as the coverage of the second pillar broadens. In 2018, the total amount of funds accumulated in Enterprise Annuity accounts was 1.477 trillion yuan, which is almost eight times the amount that was in these funds in 2008 (Ministry of Human Resources and Social Security, 2019). Adding to the 1.477 trillion yuan are the funds from the third (voluntary) pillar. Chen and Zhu (2009) found that the total amount of life annuities purchased with individual savings in 2006 was 62.6 billion yuan. In 2016, this number increased to 150.0 billion yuan according to China Insurance Regulatory Commission. The rising demand for life annuities can already be seen in the total annuity benefit payout of the Chinese insurance industry, which has risen significantly from 14.0 billion yuan in 2010 to 85.3 billion yuan in 2017 (China Insurance Regulatory Commission, 2011, 2018).

A large part of the longevity risk entailed in life annuities is "trend risk," which arises from the uncertainty surrounding the trend in Chinese mortality over time. Because the risk systematically affects all of the annuitants in the Chinese insurance industry, the more life annuities an insurer sells, the more the insurer is exposed to the risk. Although the offsetting exposure in the insurer's life insurance book may naturally hedge the trend risk acquired from life annuity sales, the effect of such a natural hedge may be limited for reasons such as differences in underwriting, duration, and age profile (Zhu and Bauer, 2014), and the (limited) potential of natural hedging may eventually be exhausted as the life annuity book grows. Furthermore, the newly introduced China Risk Oriented Solvency System (C-ROSS) specifically requires insurers operating in China to hold longevity risk solvency capital. As a result, the C-ROSS may possibly compress the ability of the Chinese insurance industry to offer life annuities at affordable prices.

The question for the insurance industry is who can share the trend risk. One possible candidate is the Chinese government, which could take longevity trend risk exposures from the insurance market explicitly by issuing longevity bonds (Blake et al., 2013) or implicitly by "bailing out" one or more insurance companies in the case of a systemic failure in the insurance industry due to longevity risk (Basel Committee on Banking Supervision, 2013). However, as the Chinese government is already assuming huge longevity trend risk due to its public pension plan, which had assets of 2826.9 billion yuan at 2013 year-end, it is not positioned to accept further risk of the same kind. A more promising candidate is the capital markets in China. Capital market investors may be interested in taking longevity trend risk exposures, because of the risk premium and diversification benefits they offer. In 2014, the total market capitalization of the equity markets in China was USD8.3 trillion, and the total notional amount of derivatives traded in Chinese exchanges was USD271 trillion.[1] These figures suggest that the capital markets in China, in theory, can absorb at least some of the longevity trend risk exposures from the insurance industry.

A recent OECD (2014) report discussed the potential of capital markets to assist the life insurance industry in continuing to provide longevity protection to individuals. The report further recommended that financial institutions create standardized index-based mortality derivatives, which could resolve the misalignment of incentives between annuity providers, who wish to mitigate their longevity trend risk exposures, and capital market investors, who demand liquidity and are likely to be discouraged by the information asymmetry arising from the fact that insurers have better knowledge of the mortality experience of their annuitants. To follow this recommendation, an indispensable prerequisite is the creation of standardized mortality indexes, upon which derivative securities like swaps and forwards can be written. Although tradable mortality indexes such as the LifeMetrics index provided by the Life and Longevity Markets Association (LLMA) already exist, they are based on mortality experience in the Western world and may therefore be unsuitable for use in China. We believe that with a population of over 1.35 billion, China requires its own standardized mortality index.

To render a standardized mortality index useful, an appropriate hedging strategy is needed. A number of longevity hedging strategies have recently been introduced by researchers including Cairns (2011, 2013), Cairns et al. (2014), Coughlan et al. (2011), Dahl et al. (2008), Li and Luo (2012), Luciano et al. (2012), Tan et al. (2014), and Zhou and Li (2017). The first objective of this paper is to, by adapting the work of Zhou and Li (2017), produce a dynamic hedging strategy that is compatible with the demographic situation in China. We demonstrate that this strategy can offload a meaningful portion of longevity risk

[1] Sources: Hong Kong Security and Futures Commission, China Security Regulation Commission, World Federation of Exchanges, and the authors' own calculations.

from insurers' annuity books. We also study how the different sources of longevity risk facing insurers operating in China can be quantified using a stochastic mortality model. Our modeling approach stems from our parallel study (Li et al., 2019), which attempts to overcome the challenge of inadequate mortality data by the Bayesian methods (Czado et al., 2006; Pedroza, 2006).

The second objective of this paper is to study the longevity risk component of C-ROSS, a new solvency system that has drawn considerable attention from both domestic and foreign insurers in recent years (see Zhao, 2014). To this end, we first illustrate how the C-ROSS longevity solvency risk capital is calculated with the prescribed adverse scenario factors. We then demonstrate the benefit of index-based longevity hedges to insurers by estimating how much C-ROSS solvency capital such hedges can release.

The rest of this paper is organized as follows. Section 2 describes the mortality data used in this study. Section 3 details the Bayesian stochastic mortality model that is built specifically for assessing longevity risk in China. Section 4 presents the dynamic longevity hedging strategy we consider. Section 5 describes how we estimate the C-ROSS capital relief from an index-based longevity hedge. Finally, Section 6 summarizes the contributions of this study.

2 Data

Our research objectives require historical mortality data for the entire population of mainland China. In what follows, we briefly describe the relevant data that are available to us. We refer interested readers to Li et al. (2019) for further information about Chinese mortality data.

The China Knowledge Resource Integrated Database provides age- and gender-specific death and mid-year population counts for the entire population of mainland China. These mortality data are extracted from the 1988 to 2015 China Population and Employment Statistics Yearbooks, issued by the National Bureau of Statistics of China. The data span an age range of 0-99 and a time period of 1981-2014, covering 3,400 age-time cells. However, as described below, the data for some of the age-time cells within this age range and time period are not available.

(i) No mortality data are available for 1982-1985, 1987, 1988, and 1990-1993.

(ii) For 1989, 1994, 1997-1999, 2001-2004 and 2006-2009, the mortality data above age 89 are unavailable; for 1996, the mortality data above age 85 are unavailable.

(iii) No death count is reported for a number of individual age-time cells (e.g., age 7, 2008, males).

Overall, there are 1,197 (1,229) age-time cells in the data set for males (females) with no death and/or population count, representing 35.21% (36.15%) of the total number of age-time cells.

The available mortality data are also not collected from the same source. For 1981, 1989, 2000 and 2010, the data are based on national censuses. For 1986, 1995 and 2005, the data are based on survey sampling of 1% of the national population. For the rest of years, the data are based on survey sampling of 0.1% of the national population. A mortality model for China needs to take this situation into account.

3 The Bayesian Stochastic Mortality Model for China

In this section, we describe the stochastic mortality model used in this paper. The model captures the movement of the mortality trends of the Chinese national population, and quantities the extent of different sources of longevity risk. The materials in this section draw heavily from our parallel study (Li et al., 2019), which is devoted to a Bayesian approach to model the evolution of Chinese mortality over time, taking into account all of the problems associated with the data set. In what follows, we first explain the specifications of the model for China. We then describe how it can be estimated given the limited available data.

3.1 Model Specification

The stochastic mortality model we consider is an adapted version of the classical Lee-Carter model (Lee and Carter, 1992). The model is specified as follows.

$$\ln m(x,t) = a(x) + b(x)k(t) + \epsilon(x,t), \tag{1}$$

where $m(x,t)$ denotes the central death rate at age x and in year t, $a(x)$ is an age-specific parameter representing the average level of mortality at age x over time, $k(t)$ is a time-varying parameter, $b(x)$ is an age-specific parameter indicating the sensitivity of $\ln(m(x,t))$ to $k(t)$, and $\epsilon(x,t)$ is the error term. We can interpret $k(t)$ to mean the overall level of mortality in year t. A reduction in $k(t)$ implies a parallel downward shift of the log-transformed curve of central mortality rates.

It is assumed that $\epsilon(x,t)$ is normally distributed with zero means and constant variances of $\sigma_\epsilon^2(t)$, where $\sigma_\epsilon^2(t)$ takes three distinct values, reflecting the nature of the data for year t. To ensure that the resulting mortality forecasts are demographically reasonable, a cubic B-splines method that smooths the pattern of β_x across ages is applied.

As in the classical Lee-Carter model, the evolution of $k(t)$ over time is modelled by a random walk with drift:

$$k(t) = c + k(t-1) + \zeta(t), \tag{2}$$

where c is a constant and $\zeta(t)$ follows a normal distribution with a zero mean and a constant variance of σ_ζ^2.

3.2 Model Estimation

Because of the missing data values, the model cannot be estimated with simple methods such as singular value decomposition. We overcome the estimation challenge by following the Bayesian method of Pedroza (2006), in which the model is formulated jointly as a Gaussian state-space model. The time-varying factor $k(t)$ is treated as hidden states, whereas $a(x)$, $b(x)$, c, $\sigma_\epsilon^2(t)$, and σ_ζ^2 are considered as model parameters that are assumed to be random themselves.

The iterative estimation procedure consists of the following major components.

Gibbs sampling

It is assumed that $\ln(m(x,t))$ is normally distributed. Under this assumption, the conditional posterior distribution of each parameter can be analytically obtained by using an appropriate conjugate prior of a normal distribution. The conjugate priors we use include normal (for parameters $a(x)$, $b(x)$, and c) and inverse-gamma (for parameters $\sigma_\epsilon^2(t)$ and σ_ζ^2). From the conditional posterior distributions, we can readily draw samples of the model parameters.

Sequential Kalman filtering and smoothing

Given a Gaussian state-space formulation, the hidden states ($k(t)$ for all t in the calibration window) can be retrieved using a sequential Kalman updating algorithm (to incorporate the information up to and including time t) and a sequential Kalman smoothing algorithm (to incorporate information beyond time t).

Imputation of missing data

On the basis of the sample of parameters drawn and the hidden states retrieved in the most recent iteration, we simulate the single cell values of $\ln(m(x,t))$ that are missing in the data set. The imputed data and the observed data are combined to form a 'new' data sample for the Gibbs sampling and the sequential Kalman filtering and smoothing in the next iteration. For the rest of missing data, we treat their missingness as Missing Completely at Random, and thus do not simulate them.

Enforcement of identifiability constraints

It is well-known that the Lee-Carter model and its variants are subject to the identifiability problem. To stipulate parameter uniqueness, the following constraints are used:

$$\sum_x b(x) = 1 \quad \text{and} \quad \sum_t k(t) = 0.$$

The identifiability constraints are applied at the end of each iteration.

We estimate the model to the entire available data described in Section 2. As usual in Bayesian methods, the first batch of 500 samples are regarded as burn-in and therefore discarded. To mitigate auto-correlation, only one sample is recorded in every 100 values drawn after the burn-in period. 5,000 samples are drawn and used to form the joint empirical posterior distribution of the model parameters. We refer interested readers to Li et al. (2019) for further details about the algorithms for Gibbs sampling and sequential Kalman filtering and smoothing used in the estimation procedure.

In Figure 1 we show the estimates of $a(x)$, $b(x)$, $k(t)$, c, $\sigma_\epsilon^2(t)$ and σ_ζ^2 parameters for $x = 0, \ldots 99$ and $t = 1981, \ldots, 2014$. The fan chart in each panel of the top row shows the central 10% prediction interval for $a(x)$, $b(x)$ and $k(t)$ with the heaviest shading, surrounded by the 20%, 30%, ..., 90% prediction intervals with progressively lighter shading. The downward trend in $k(t)$ indicates a steady reduction in the overall level of mortality over the past couple of decades. The histogram in each panel of the bottom row shows the empirical posterior distributions of $\sigma_\epsilon^2(t)$, σ_ζ^2 and c. Note that, for $\sigma_\epsilon^2(t)$, we show three distinct histograms for $t = 2010$ (census year), $t = 2005$ (1% survey sampling) and $t = 2014$ (0.1% survey sampling).

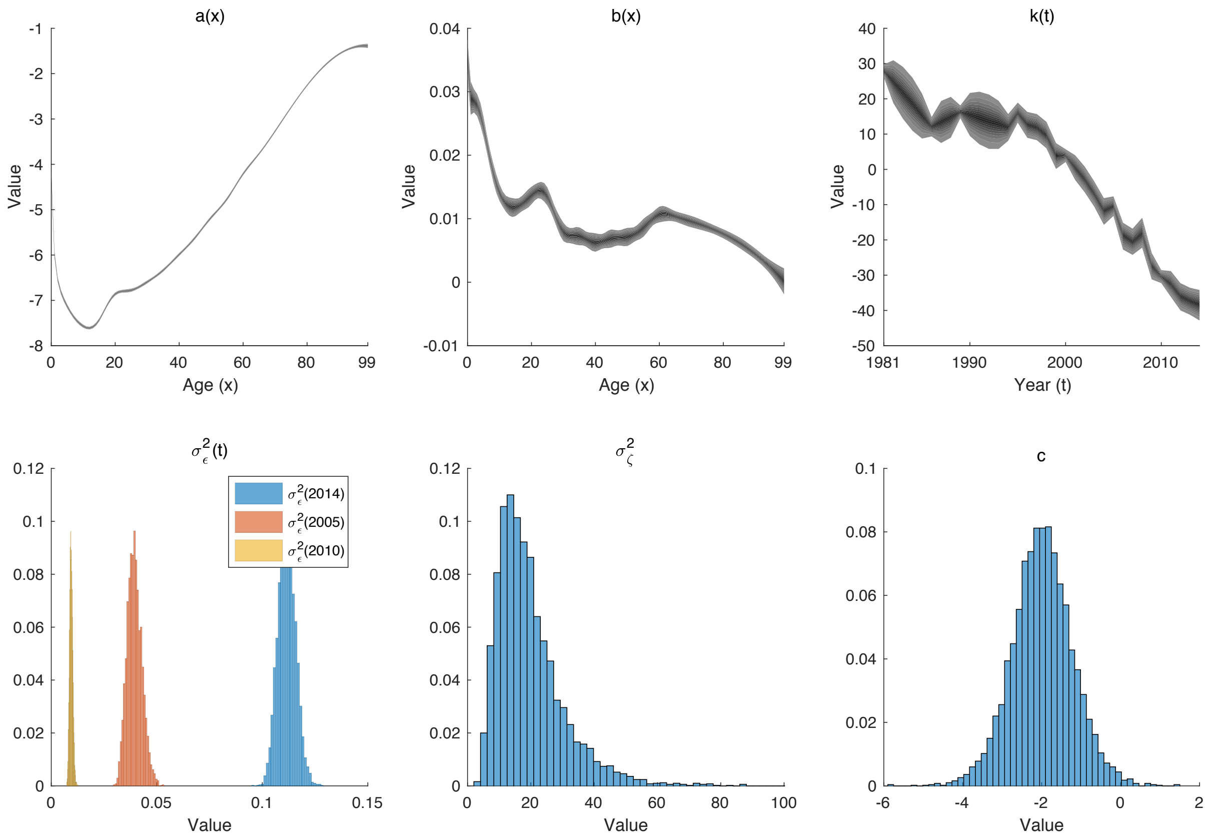

Figure 1: The estimates of $a(x)$, $b(x)$, $k(t)$ and $\sigma_\epsilon^2(t)$ in equation (1) for $x = 0, \ldots 99$ and $t = 1981, \ldots, 2014$, and c and σ_ζ^2 in equation (2), Chinese males.

4 Hedging Strategies

In this section, we investigate how Chinese insurers can use a national mortality index to offload longevity risk from their balance sheets. We adapt the work of Zhou and Li (2017) to form a dynamic delta hedging strategy, in which the hedge parameters (the deltas) of the insurer's portfolio and the portfolio of hedging instruments are matched.

We begin this section with a description of the liability being hedged, followed by explanations about how the longevity risk involved in the liability can be mitigated by using instruments written on an age-specific mortality index. We then detail how hedge effectiveness may be measured, and estimate the degrees of hedge effectiveness that longevity hedges for annuity liabilities can achieve. We conclude this section with an analysis of various factors that may affect the performance of a longevity hedge for China. Throughout this section, the Bayesian mortality model presented in Section 3 is assumed.

4.1 The Set-up

Let us first define several notations. We let

$$S_{x,t}(T) = \prod_{s=1}^{T}(1 - q(x + s - 1, t + s))$$

be the *ex post* probability that an individual who is aged x at time t (the end of year t) would have survived to time $t + T$, where $q(x, t)$ denotes the probability that an individual dies between time $t - 1$ and t (during year t), provided that he/she has survived to age x at time $t - 1$. It is clear from the definitions that $S_{x,t}(T)$ is not known prior to time $t + T$, whereas $q(x, t)$ is not known prior to time t. We also let

$$p_{x,u}(T, \mathcal{F}_t) = \mathrm{E}(S_{x,u}(T)|\mathcal{F}_t),$$

where $u \geq t$ and $\mathcal{F}_t$ represents the information about the evolution of mortality up to and including time t. Because the assumed mortality model is based on central death rates, we need to approximate $q(x, t)$ from $m(x, t)$. We use the relation $q(x, t) = 1 - \exp(-m(x, t))$, which holds exact if the force of mortality between two consecutive integer ages is constant.

Let us suppose that the liability being hedged is a portfolio of life annuities, which are associated with the cohort of individuals who are aged x_0 at time t_h when the longevity hedge is established. We further assume that the each life annuity pays $1 at the end of each year until death. It follows that the time-t value of the insurer's future liabilities (per policyholder at time t) is

$$FL_t = \sum_{s=1}^{\infty}(1 + r)^{-s}\, p_{x_0+t-t_h,t}(s, \mathcal{F}_t), \quad t \geq t_h,$$

where r is the interest rate for discounting purposes.

Suppose that the hedging horizon is Y years and that the hedge portfolio is adjusted annually. Due to the dynamic nature of the hedge, the value of FL_t at the beginning of each of the Y years has to be computed. As FL_t takes no analytical form, evaluating the hedge over the hedging horizon requires nested simulations. To reduce the computation burden, an approximation formula is used to compute each value of FL_t. The approximation formula is derived by applying a second order Taylor expansion on the probit transformation of $p_{x_0+t-t_h,t}(s, \mathcal{F}_t)$ about the best estimates of $k(t)$. We refer readers to Cairns (2011) and Zhou and Li (2017) for a detailed discussion of the approximation method.

4.2 Hedging with an Age-Specific Mortality Index

We suppose here that the hedging instruments used are q-forwards written on age-specific death probabilities of a national population. A q-forward is a zero-coupon swap with a floating leg proportional to the realized death probability at a certain reference age during the year immediately prior to maturity and a fixed leg proportional to the corresponding forward mortality rate that is fixed at inception. To hedge the longevity risk involved in the life annuity portfolio, the hedger should participate in the q-forwards as the fixed-rate receiver, so that he/she will receive a net payment from the counterparty when mortality turns out to be lower than expected.

Let us consider a q-forward that is linked to the national population of China and a reference age x_f. Assume that the q-forward is issued at time t_0 and matures at time t_0+T^*. By definition, the payoff from the q-forward depends on the realized value of $q(x_f, t_0+T^*)$. Let $q^f(x_f, t_0+T^*)$ be the corresponding forward mortality rate, which is fixed at $t=t_0$ when the q-forward is first launched. At $t=t_0,\ldots,t_0+T^*-1$, the value of the q-forward (per \$1 notional) from the perspective of the hedger (fixed-rate receiver) is given by

$$\begin{aligned}Q_t(t_0) &= (1+r)^{-(t_0+T^*-t)}(q^f(x_f,t_0+T^*) - \mathrm{E}(q(x_f,t_0+T^*)|\mathcal{F}_t))\\ &= (1+r)^{-(t_0+T^*-t)}(q^f(x_f,t_0+T^*) - (1-\mathrm{E}(S_{x_f,t_0+T^*-1}(1)|\mathcal{F}_t)))\\ &= (1+r)^{-(t_0+T^*-t)}(q^f(x_f,t_0+T^*) - (1-p_{x_f,t_0+T^*-1}(1,\mathcal{F}_t)).\end{aligned}$$

Suppose that at time t during the hedging horizon, the hedger uses the aforementioned q-forward (with $t_0\le t$) as the only hedging instrument. The main idea behind the delta hedging strategy is to ensure that the annuity portfolio and the q-forward portfolio have similar sensitivities to changes in $k(t)$. To achieve this goal, the hedge ratio h_t (i.e., the notional amount of the q-forward) is chosen in such a way that

$$\frac{\partial FL_t}{\partial k(t)} = h_t\frac{\partial Q_t(t_0)}{\partial k(t)},$$

where $\partial FL_t/\partial k(t)$ and $h_t\partial Q_t(t_0)/\partial k(t)$ represent the time-t deltas of the annuity portfolio and the (calibrated) q-forward portfolio, respectively.

The hedge portfolio has a value of $h_tQ_t(t_0)$ at time t and a value of $h_tQ_{t+1}(t_0)$ at time $t+1$. At time $t+1$, the q-forward written at time t is closed out, and another q-forward portfolio is constructed. The process repeats from the beginning to the end of the hedging horizon.

When evaluating such a hedge, we need to compute the value of $Q_t(t_0)$ for every t over the hedging horizon, but $Q_t(t_0)$ cannot be analytically calculated. To avoid the need for nested simulations, an approximation formula is used to calculate $Q_t(t_0)$. The approximation is based on a first order Taylor's expansion of the probit transformation of $p_{x_f,t_0+T^*-1}(1,\mathcal{F}_t)$ about the best estimate of $k(t)$. We refer readers to Cairns (2011) and Zhou and Li (2017) for further details about the approximation of $Q_t(t_0)$. The values of $\partial FL_t/\partial k(t)$ and $\partial Q_t(t_0)/\partial k(t)$ are calculated on the basis of the approximation formulas for FL_t and $Q_t(t_0)$, respectively.

4.3 Measuring Hedge Effectiveness

We can evaluate the effectiveness of a dynamic longevity hedge by simulating a large number of mortality scenarios from the estimated Bayesian mortality model.

We let $PL_{t_h}=FL_{t_h}$ and

$$PL_t = \sum_{s=1}^{t-t_h}(1+r)^{-s}S_{x_0,t_h}(s) + (1+r)^{-(t-t_h)}S_{x_0,t_h}(t)FL_t,\quad t=t_h+1,\ldots,t_h+Y.$$

We can interpret PL_t to mean the value of all annuity payments at time t_h when the hedge is established, given the information up to and including time t. For $t>t_h$, the value of $PL_t|\mathcal{F}_{t_h}$ is random in part because the value of $S_{x_0,t_h}(s)$ depends on the realizations of $k(t_h+1),\ldots,k(t_h+s)$, and in part because the value of FL_t depends on the realizations of $k(t)$.

Define by PA_t the time-t_h value of the assets backing the pension plan at time t, where $t\ge t_h$. We assume that the asset value equals the liability value when the hedge is established; i.e., $PA_{t_h}=PL_{t_h}$. To simplify exposition, we assume that all of the q-forwards used have the same maturity T^* and reference age x_f. We also assume that at every time point t when the hedge portfolio is adjusted, a freshly launched q-forward is written (i.e., $t_0 = t$ for $t = t_h, \ldots, t_h + Y - 1$). Under these assumptions, we have

$$PA_t = PA_{t-1} + (1+r)^{-(t-t_h)} h_{t-1} Q_t(t-1)$$

for $t = t_h + 1, \ldots, t_h + Y$.

The potential deviation between PA_t and PL_t is the residual risk that is not eliminated by the longevity hedge. Hence, we may measure hedge effectiveness by the following metric:

$$HE_u = 1 - \frac{\text{Var}(PA_{t_h+u} - PL_{t_h+u}|\mathcal{F}_{t_h})}{\text{Var}(PL_{t_h+u}|\mathcal{F}_{t_h})}, \quad u = 1, \ldots, Y,$$

which is close to 1 if the hedge is effective and 0 if it is not.

4.4 An Illustration

In this subsection, we illustrate the use of the dynamic delta hedging strategy to mitigate the longevity risk associated with annuity portfolios in China. The following assumptions are made in the illustration.

1. The liability being hedged is a portfolio of life annuities that are sold to individuals who are aged 60 at the end of 2014. Each annuity pays \$1 at the end of each year until the annuitant dies or reaches age 90, whichever is the earliest.
2. The mortality experience of the annuitants is the same as that of the male population of China.
3. The hedge begins at the end of 2014 and the hedging horizon is 30 years. The hedge portfolio is adjusted annually.
4. The hedging instruments used are q-forwards that are also linked to the male population of China. They all have a time-to-maturity of 10 years and a reference age of 75.
5. All of the q-forwards have a zero risk premium, which means $q^f(x_f, t_0 + T^*) = \text{E}(q(x_f, t_0 + T^*))$. This working assumption has no effect on the resulting hedge effectiveness.
6. The market for q-forwards is liquid and no transaction cost is required.
7. The interest rate for all durations is $r = 4\%$ per annum and remains constant over time. The hedger can invest or borrow at this rate.
8. The evaluation of hedge effectiveness is based on 5,000 mortality scenarios that are generated from the Bayesian stochastic mortality model presented in Section 3.

The hedging results are presented in Figure 2. In this figure, the grey (larger) fan chart shows the distributions of $PL_{t_h+u}|\mathcal{F}_{t_h}$ for $u = 1, \ldots, Y$, whereas the green (smaller) fan chart depicts the distributions of $PA_{t_h+u} - PL_{t_h+u}|\mathcal{F}_{t_h}$ for $u = 1, \ldots, Y$. The difference between the widths of the two fan charts reflects the amount of longevity risk that is removed from the dynamic q-forward hedge. The value of HE_{30} is 0.7895, indicating that 78.95% of the variance in the liability being hedged have been removed by the q-forward hedge.

4.5 A Decomposition of Risks

The full model used for generating the results in the previous subsection incorporates various sources of risk, including trend risk (the uncertainty arising from $\zeta(t)$), error risk (the uncertainty arising from $\epsilon(x,t)$), and parameter risk (the uncertainty in estimating the parameters in the Lee-Carter structure and the time-series processes). To better understand how these risks contribute to the erosion in hedge effectiveness, we now re-evaluate the longevity hedge using restricted models in which some of the stochastic components are switched off. In particular, we consider the following four scenarios, which have different levels of conservatism.

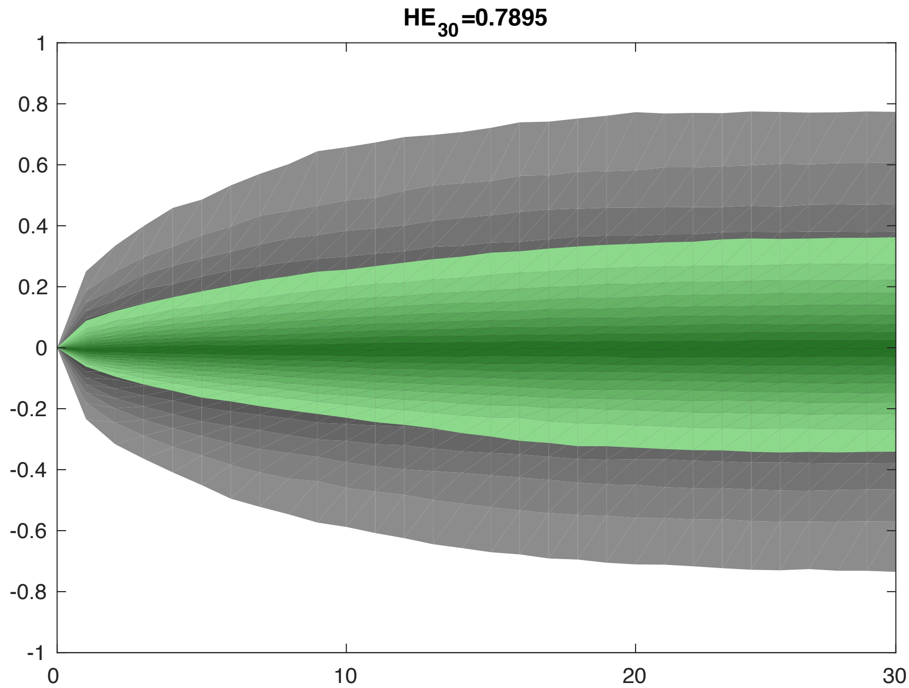

Figure 2: Simulated distributions of $PL_{t_h+u}|\mathcal{F}_{t_h}$ (the grey fan charts) and $PA_{t_h+u} - PL_{t_h+u}|\mathcal{F}_{t_h}$ (the green fan charts), $u = 1, \ldots, Y$, for the annuity liabilities that are associated with males in China.

Scenario 1: Trend risk only

This scenario in the most optimistic. In this scenario, it is assumed that the Lee-Carter structure captures mortality patterns perfectly and that all of the parameters are accurately and precisely estimated. The simulation procedure incorporates only the randomness arising from $\zeta(t)$.

Scenario 2: Trend risk and trend-related parameter risk only

This scenario is less optimistic than Scenario 1. In addition to the randomness arising from $\zeta(t)$, we also consider in this scenario the uncertainty about the parameters that are associated with the time-series processes (i.e., parameters c and σ^2_ζ).

Scenario 3: All but error risk

It is assumed in this scenario that the Lee-Carter structure describes mortality patterns perfectly. The hedging results under this scenario are obtained by assuming $\sigma^2_\epsilon(t) = 0$ in the simulation procedure. All of the other sources of risk are retained.

Scenario 4: All sources of uncertainty

This scenario is the most conservative. The full simulation model that incorporates all sources of uncertainty is used. The results under this scenario are identical to those obtained in the previous subsection.

We display the hedging results under the four scenarios in Figure 3. Let us begin with the most optimistic view. If parameter and error risks are assumed to be non-existent, then there is less uncertainty associated with the future mortality rates, leading to narrower fan charts for both the hedged and unhedged liabilities. Overall, the hedge effectiveness becomes significantly higher. The value of HE_{30} under this scenario is close to 100%, which means the hedge is almost perfect. The results generated under this scenario are the most comparable to those produced by Cairns (2011) and Zhou and Li (2017), who did not consider parameter and error risks.

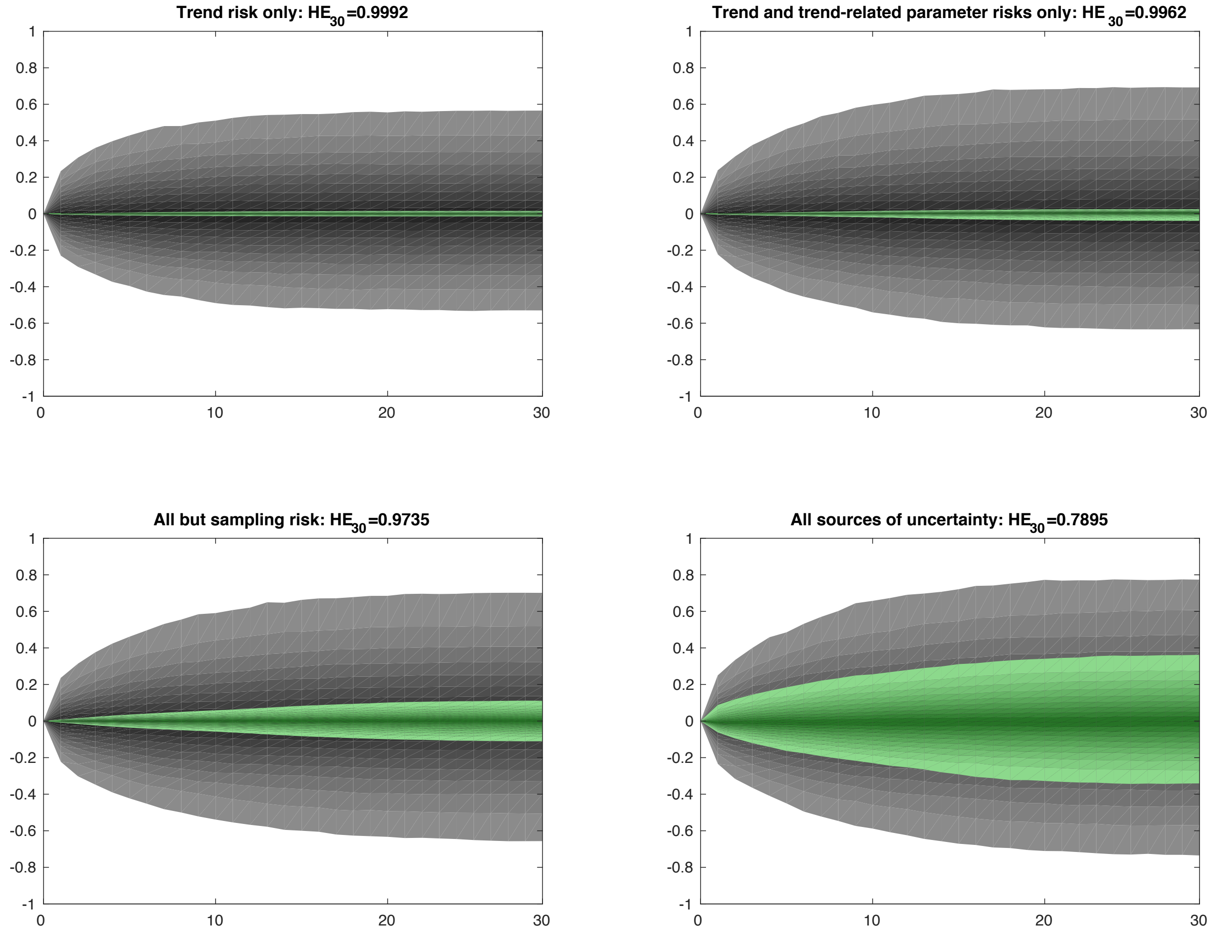

Figure 3: Simulated distributions of $PL_{t_h+u}|\mathcal{F}_{t_h}$ (the grey fan charts) and $PA_{t_h+u} - PL_{t_h+u}|\mathcal{F}_{t_h}$ (the green fan charts), $u = 1, \ldots, Y$, under four scenarios: (1) trend risk only, (2) trend and trend-related parameter risks only, (3) all but error risk, and (4) all sources of uncertainty.

We then move on to Scenarios 2 and 3. If we take a slightly less optimistic view by incorporating the portion of parameter risk that is related to the time-series processes for $k(t)$, then the value of HE_{30} is reduced to 99.62%, which is still substantially higher than that in the most conservative scenario. If we also incorporate the portion of parameter risk that is related to the rest of the model, then the value of HE_{30} drops to about 97.35%. We can interpret the difference between the values of HE_{30} in Scenarios 1 and 3 as the effect of parameter risk on the effectiveness of the longevity hedge.

The only difference between Scenarios 3 and 4 is the incorporation of error risk, so the difference between the values of HE_{30} in Scenarios 3 and 4 can be regarded as the erosion of hedge effectiveness due to error risk. The effect of error risk is significant in this illustration because we use $\sigma_\epsilon^2(t) = \sigma_\epsilon^2(2014)$, the value for 0.1% survey sampling and also the largest one of the three distinct values of $\sigma_\epsilon^2(t)$.

5 Longevity Risk Solvency Capital under C-ROSS

The C-ROSS was introduced by the China Insurance Regulatory Commission (CIRC) in 2012 to supersede the Insurance Company Solvency Regulations (ICSR) established in 2008. The C-ROSS can be seen as the Chinese version of Europe's Solvency II, in which regulations and capital requirements are emphasized on a risk-oriented system rather than on a factor-based system. Some domestic and foreign insurance companies in China tried implementing C-ROSS when preparing their reserve calculations. In February 2015, the CIRC released the official version of C-ROSS.

Similar to Solvency II, C-ROSS adopts a regulatory framework with three pillars: Quantitative Capital Requirements, Qualitative Supervisory and Market Discipline Mechanism. In the first pillar, the calculation of the minimum capital requirement (MCR) for insurers' quantifiable risks, namely insurance risk, market risk, and credit risk, is explicitly specified using actual versus minimum capital assessment standards. In the second pillar, the Solvency Aligned Risk Management Requirements and Assessment (SARMRA) is introduced to evaluate insurers' overall solvency level through an integrated risk rating system for qualitative risks including operational risk, strategic risk, reputational risk, and liquidity risk. The third pillar imposes supervision of insurance companies from rating agencies, financial reports, media, and the general public by enforcing risk disclosure, risk transparency, and market disciplines.

In terms of longevity risk management, the C-ROSS classifies mortality and longevity risks as part of insurance risk, and explicitly specifies the calculation of the MCRs for these risks. For simplicity, in what follows we ignore insurance risks other than mortality and longevity risks. We use $MCR^{(M)}$ and $MCR^{(L)}$ to represent the C-ROSS MCRs for mortality and longevity risks, respectively.

For an unhedged insurance/annuity liability, we have

$$MCR^{(i)} = \max(V((1 + SF^{(i)})\boldsymbol{m})) - V(\boldsymbol{m}), 0), \quad i = M, L,$$

where $V(\cdot)$ is the present value of all of the cash flows from the insurance/annuity liability evaluated at a certain mortality curve, $\boldsymbol{m}$ is the best-estimate mortality curve for the duration of the liability, and SF is the adverse scenario factor. In C-ROSS, $SF^{(M)}$ is a parallel shock to the mortality curve reflecting the Value-at-Risk at a certain conservative confidence level:

$$SF^{(M)} = \begin{cases} 10\%, & N > 200, \\ 15\%, & 100 < N \le 200, \\ 20\%, & N \le 100, \end{cases}$$

where N denotes the number of contracts; and $SF^{(L)}$ is specified as follows:

$$SF^{(L)} = \begin{cases} (1 - 3\%)^t - 1, & 0 < t \le 5, \\ (1 - 3\%)^5(1 - 2\%)^{t-5} - 1, & 5 < t \le 10, \\ (1 - 3\%)^5(1 - 2\%)^5(1 - 1\%)^{t-10} - 1, & 10 < t \le 20, \\ (1 - 3\%)^5(1 - 2\%)^5(1 - 1\%)^{10} - 1, & t > 20, \end{cases}$$

where t is the number of years after the assessment date. Finally, to calculate the overall MCR for both mortality and longevity risks, the following formula is used:

$$MCR = \sqrt{\boldsymbol{M}\Sigma\boldsymbol{M}'},$$

where

$$\boldsymbol{M} = (MCR^{(M)}, MCR^{(L)})$$

and

$$\Sigma = \begin{pmatrix} 1 & -0.25 \\ -0.25 & 1 \end{pmatrix},$$

which is predetermined by the CIRC.

The column labeled "unhedged liability" in Table 1 displays the time-0 values of $MCR^{(M)}$ and $MCR^{(L)}$ for the annuity liability defined in Section 4. The value of $\boldsymbol{m}$ is taken as the best-estimate mortality curve implied by the assumed mortality model. Given how the liability is structured, the calculated value of $MCR^{(M)}$ is always 0, no matter what value of N is assumed. Because $MCR^{(M)} = 0$, $MCR = MCR^{(L)}$ and the correlation matrix Σ is not involved in the calculations. In

Unhedged liability		q-forward portfolio		Hedged liability	
$MCR^{(M)}$	$MCR^{(L)}$	$MCR^{(M)}$	$MCR^{(L)}$	$MCR^{(M)}$	$MCR^{(L)}$
0	0.7037	0.6370	0	0.0679	0

Table 1: The values of $MCR^{(M)}$ and $MCR^{(L)}$ for the unhedged liability, the q-forward portfolio, and the hedged liability, on the basis of projected mortality rates.

general, when the insurer's portfolio contains both insurance and annuity liabilities, the value of $MCR^{(M)}$ is not necessarily 0.

Similarly, the values of $MCR^{(M)}$ and $MCR^{(L)}$ for a portfolio of hedging instruments is given by

$$MCR^{(i)} = \max(H((1 + SF^{(i)})\boldsymbol{m})) - H(\boldsymbol{m}), 0), \quad i = M, L,$$

where $H(\cdot)$ is the present value of all of the cash flows from the portfolio of hedging instruments, evaluated at the best-estimate mortality curve $\boldsymbol{m}$. Finally, for a hedged portfolio, we have

$$MCR^{(i)} = \max(V((1 + SF^{(i)})\boldsymbol{m}) - V(\boldsymbol{m}) - H((1 + SF^{(i)})\boldsymbol{m}) + H(\boldsymbol{m}), 0), \quad i = M, L.$$

The specifications of $SF^{(M)}$, $SF^{(L)}$, and Σ also apply to the calculations of MCRs for a portfolio of hedging instruments and a hedged portfolio.

The calculated time-0 values of $MCR^{(M)}$ and $MCR^{(L)}$ for the calibrated q-forward portfolio and the hedged annuity liability are shown in Table 1. Note that the q-forward portfolio incurs a loss only if future mortality turns out to be higher than expected, so the value of its $MCR^{(L)}$ is always 0. Overall, the longevity hedge removes all of the $MCR^{(L)}$ from the annuity portfolio, but introduces some $MCR^{(M)}$.

The percentage of MCR reduced by the longevity hedge (calculated as one minus the ratio of the hedged MCR over the unhedged MCR) is 90.35%, indicating that a standardized longevity hedge can significantly reduce a Chinese insurer's required capital.

6 Concluding Remarks

China's rapidly shifting demographics have created a huge potential demand for life annuities, which protect individuals against longevity risk. However, due to the systematic nature of longevity risk, there is a limit to the amount of longevity risk an insurer can accept. To maintain insurers' ability to sell life annuities at affordable prices and to reduce the risk of a systemic failure in the insurance industry due to an excessive exposure to longevity risk, we need to develop markets for standardized mortality-linked securities, through which longevity risk can be transferred from insurers to capital market investors.

In this paper, we study how the longevity risk facing insurers operating in China can be quantified using stochastic mortality models that are developed to suit the demographic situation in China. By adapting the work of Zhou and Li (2017), we formulate a dynamic hedging strategy and examine how much of a stylized life annuity portfolio's variance can be eliminated when the hedging strategy is implemented. The resulting hedge effectiveness (measured by the proportion of variance eliminated) indicate that a meaningful portion of longevity risk can be removed by a dynamic q-forward hedge, even when trend risk, parameter risk and error risk are taken into account.

To understand the amount of capital relief that can potentially be obtained from an index-based longevity hedge under C-ROSS, we compare the C-ROSS MCRs for insurance risks when a longevity hedge is absent and when our longevity hedging strategy is implemented. Our strategy reduces the MCR by approximately 90%. Thus, we conclude that the longevity hedge created from a Chinese national mortality index yield significant reductions in C-ROSS MCR, suggesting that the development of a market for standardized longevity risk transfers is worth Chinese insurers' attention.

In terms of hedging strategies, future research could extend the work of Cairns (2011) and Zhou and Li (2017) by incorporating not only life annuities but also life insurances. Such an extension is important, as at present Chinese insurers generally run both life insurance and life annuity lines. With this extension, it would be possible to gauge the effect of natural hedging, and hence determine the need for transferring the residual longevity risk to capital markets. Another direction for future research is the development of more sophisticated hedging strategies, including "delta-gamma" hedging, which incorporates not only the first but also the second order partial derivatives with respect to $k(t)$.

Acknowledgments

This work is financially supported by the Global Risk Institute, the Society of Actuaries Center of Actuarial Excellence Programme, and the Natural Sciences and Engineering Research Council of Canada.

References

Basel Committee on Banking Supervision. (2013). Joint Forum - Longevity Risk Transfer Market: Market Structure, Growth Drivers and Impediments, and Potential Risks. Bank of International Settlements. Available at www.bis.org.

Blake, D., Boardman, T. and Cairns, A. (2013). Sharing Longevity Risk: Why Governments Should Issue Longevity Bonds. Pension Institute Discussion Paper PI-1002.

Cai, Y. and Cheng, Y. (2014). Pension Reform in China: Challenges and Opportunities. *Journal of Economic Surveys*, 28: 636-651.

Cairns, A.J.G. (2011). Modelling and Management of Longevity Risk: Approximations to Survival Functions and Dynamic Hedging. *Insurance: Mathematics and Economics*, 49: 438-453.

Cairns, A.J.G. (2013). Robust Hedging of Longevity Risk. *Journal of Risk and Insurance*, 80: 621-648.

Cairns, A.J.G., Blake, D., Dowd, K. and Coughlan, G.D. (2014). Longevity Hedge Effectiveness: A Decomposition. *Quantitative Finance*, 14: 217-235.

Chen, B. and Zhu, W. (2009). A Review of Longevity Risk Management. China Center for Insurance and Risk Management of Tsinghua SEM.

China Insurance Regulatory Commission (2011, 2018). *Yearbook of China's Insurance* China Insurance Yearbook Publishing House.

Czado, C., Delwarde, A. and Denuit, M. (2005). Bayesian Poisson Log-Bilinear Mortality Projections. *Insurance: Mathematics and Economics* 36: 260-284.

Dahl, M., Melchior, M. and Møller, T. (2008). On Systematic Mortality Risk and Risk Minimization with Mortality Swaps. *Scandinavian Actuarial Journal*, 108: 114-146.

Dorfman, M.C., Wang, D, O'Keefe, P. and Cheng. J. (2012). China's Pension Schemes for Rural and Urban Residents. In: R. Hinz, R. Holzmann and D. Tuesta, eds. *Matching Contributions for Pensions: A Review of International Experience.* World Bank Publications. 2012.

Lee, R. and Carter, L. (1992). Modeling and Forecasting U.S. Mortality. *Journal of the American Statistical Association*, 87: 659-671.

Li, J.S.-H. and Luo, A. (2012). Key q-duration: A Framework for Hedging Longevity Risk. *ASTIN Bulletin*, 42: 413-452.

Li, J.S.-H., Zhou, K.Q., Zhu, X., Chan, W.-S., and Chan, F.W.-H. (2019). A Bayesian approach to developing a stochastic mortality model for China. Journal of the Royal Statistical Society: Series A (Statistics in Society), 182: 1523-1560.

Luciano, E., Regis, L. and Vigna, E. (2012). Delta-Gamma Hedging of Mortality and Interest Rate Risk. *Insurance: Mathematics and Economics*, 50: 402-412.

Ministry of Human Resources and Social Security. (2019). Nationwide Enterprise Annuity Funds Data Summary 2018. Available at http://www.mohrss.gov.cn/shbxjjjds/SHBXJDSgongzuodongtai/201904/W020190402665288131918.pdf.

OECD. (2014). *Mortality Assumptions and Longevity Risk: Implications for Pension Funds and Annuity Providers*. OECD Publishing, Paris.
DOI: http://dx.doi.org/10.1787/9789264222748-en.

Pedroza, C. (2006). A Bayesian Forecasting Model: Predicting U.S. Male Mortality. *Biostatistics*, 7: 530-550.

Tan, C.I., Li, J, Li, J.S.-H. and Balasooriya, U. (2014). Parametric Mortality Indexes: From Index Construction to Hedging Strategies. *Insurance: Mathematics and Economics*, 59: 285-299.

The World Bank. (1994). *Averting the Old Age Crisis: Policies to Protect the Old and Promote Growth*. Oxford University Press.

Zhao, Y. (2014). China's C-ROSS: A New Solvency System Down the Road. *The Actuary Magazine*, 11: 16-18.

Zhou, K.Q. and Li, J.S.-H. (2017). Dynamic Longevity Hedging in the Presence of Population Basis Risk: A Feasibility Analysis from Technical and Economic Perspectives. *The Journal of Risk and Insurance*, 84: 417-437.

Zhu, N. and Bauer, D. (2014). A Cautionary Note on Natural Hedging of Longevity Risk. *North American Actuarial Journal*, 18: 104-115.

About The Society of Actuaries

With roots dating back to 1889, the Society of Actuaries (SOA) is the world's largest actuarial professional organization with more than 31,000 members. Through research and education, the SOA's mission is to advance actuarial knowledge and to enhance the ability of actuaries to provide expert advice and relevant solutions for financial, business and societal challenges. The SOA's vision is for actuaries to be the leading professionals in the measurement and management of risk.

The SOA supports actuaries and advances knowledge through research and education. As part of its work, the SOA seeks to inform public policy development and public understanding through research. The SOA aspires to be a trusted source of objective, data-driven research and analysis with an actuarial perspective for its members, industry, policymakers and the public. This distinct perspective comes from the SOA as an association of actuaries, who have a rigorous formal education and direct experience as practitioners as they perform applied research. The SOA also welcomes the opportunity to partner with other organizations in our work where appropriate.

The SOA has a history of working with public policymakers and regulators in developing historical experience studies and projection techniques as well as individual reports on health care, retirement and other topics. The SOA's research is intended to aid the work of policymakers and regulators and follow certain core principles:

Objectivity: The SOA's research informs and provides analysis that can be relied upon by other individuals or organizations involved in public policy discussions. The SOA does not take advocacy positions or lobby specific policy proposals.

Quality: The SOA aspires to the highest ethical and quality standards in all of its research and analysis. Our research process is overseen by experienced actuaries and nonactuaries from a range of industry sectors and organizations. A rigorous peer-review process ensures the quality and integrity of our work.

Relevance: The SOA provides timely research on public policy issues. Our research advances actuarial knowledge while providing critical insights on key policy issues, and thereby provides value to stakeholders and decision makers.

Quantification: The SOA leverages the diverse skill sets of actuaries to provide research and findings that are driven by the best available data and methods. Actuaries use detailed modeling to analyze financial risk and provide distinct insight and quantification. Further, actuarial standards require transparency and the disclosure of the assumptions and analytic approach underlying the work.

Society of Actuaries
475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
www.SOA.org