

2012

Pricing Mortality Securities with Correlated Mortality Indexes

Yijia Lin

University of Nebraska–Lincoln, yijialin@unl.edu

Sheen Liu

Washington State University, sheen.liu@wsu.edu

Jifeng Yu

University of Nebraska–Lincoln, jifeng.yu@unl.edu

Follow this and additional works at: <http://digitalcommons.unl.edu/managementfacpub>



Part of the [Corporate Finance Commons](#), [Finance and Financial Management Commons](#), and the [Portfolio and Security Analysis Commons](#)

Lin, Yijia; Liu, Sheen; and Yu, Jifeng, "Pricing Mortality Securities with Correlated Mortality Indexes" (2012). *Management Department Faculty Publications*. 112.

<http://digitalcommons.unl.edu/managementfacpub/112>

This Article is brought to you for free and open access by the Management Department at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in Management Department Faculty Publications by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

Published in *The Journal of Risk and Insurance* 2012 (28 pp);

doi: 10.1111/j.1539-6975.2012.01481.x

Published online (before inclusion in an issue) July 20, 2012.

Copyright © 2012 The Journal of Risk and Insurance. Used by permission.



Pricing Mortality Securities with Correlated Mortality Indexes

Yijia Lin,¹ Sheen Liu,² and Jifeng Yu³

1. Department of Finance, College of Business Administration, University of Nebraska–Lincoln

2. Department of Finance, College of Business, Washington State University, Tri-Cities

3. Department of Management, College of Business Administration,
University of Nebraska–Lincoln.

Corresponding author — Y. Lin, yijialin@unl.edu

Abstract

This article proposes a stochastic model, which captures mortality correlations across countries and common mortality shocks, for analyzing catastrophe mortality contingent claims. To estimate our model, we apply particle filtering, a general technique that has wide applications in non-Gaussian and multivariate jump-diffusion models and models with nonanalytic observation equations. In addition, we illustrate how to price mortality securities with normalized multivariate exponential titling based on the estimated mortality correlations and jump parameters. Our results show the significance of modeling mortality correlations and transient jumps in mortality security pricing.

Introduction

Over the last century, populations of different countries have been increasingly linked by flows of information, goods, transportation, and communication, and as a consequence the world has become more closely connected and interdependent. While the trend of globalization has substantially driven market growth and international trade, it has also helped to spread some of the deadliest infectious diseases across borders (Daulaire, 1999). Thus, it seems improper to forecast mortality for an individual national population in isolation from others. Indeed, in practice, intercountry mortality correlation has long been a serious concern for insurers that underwrite life insurance business.

Mortality forecast that takes into account a country's linkage to others is important in the sense that not only does it facilitate better understanding of mortality risk, but it also has enormous implications for pricing mortality securities. Recent financial innova-

tion makes mortality securitization a viable option for insurers or reinsurers to transfer catastrophe mortality risk arising from the possible occurrence of pandemics or large-scale terrorist attacks. By segregating its cash flows linked to extreme mortality risk, an insurance firm is able to repackage them into securities that are traded in capital markets (Blake and Burrows, 2001; Lin and Cox, 2005; Cox and Lin, 2007). Since the first publicly traded mortality security issued by Swiss Re in 2003, almost all mortality transactions determine the coupons and principals based on three or more population mortality indexes, with the only exception—the Tartan mortality bond sold in 2006. This indicates that insurers or reinsurers are keenly interested in transferring potential country-correlated mortality risk embedded in their business. For instance, the mortality risk of the 2003 Swiss Re mortality bond was defined in terms of an index based on the weighted average annual population death rates in the United States, the United Kingdom, France, Italy, and Switzerland (Lin and Cox, 2008). As another example, the mortality bond issued by the Nathan Ltd. in 2008 depended on the annual population death rates of four countries, namely, the United States, the United Kingdom, Canada, and Germany. Given that the existing (and possible future) mortality securities bundle multination mortality risks, mortality correlation among countries merits serious consideration in mortality securitization pricing.

In the recent literature, a number of stochastic mortality models have been proposed. Despite the importance of mortality correlation, surprisingly very few papers treat correlation as an indispensable element. For example, in order to account for catastrophic mortality death shocks, Chen and Cox (2009) incorporate a jump-diffusion process into the original Lee–Carter model to forecast mortality rates and price the 2003 Swiss Re mortality bond. Yet, while the Swiss Re bond payments depended on five-country weighted mortality index, the authors price this bond only based on the U.S. mortality rates. Hence, it is not clear how to extend their method to multipopulation correlated mortality scenarios.

Beelders and Colarossi (2004) and Chen and Cummins (2010), on the other hand, use the extreme value theory to measure mortality risk of the 2003 Swiss Re bond. However, they simply model the combined index without considering the mortality correlations among different countries. So the question of how to model multipopulation mortality correlation remains open. Bauer and Kramer (2009) borrow a credit risk modeling approach introduced by Lando (1998) to describe stochastic force of mortality. To incorporate dramatic mortality changes that are crucial in measuring mortality risk (Lin and Cox, 2008), Bauer and Kramer (2009) propose a mortality model with an affine jump-diffusion process. Then they use their mortality model to price the Tartan transaction, which is the only publicly traded mortality bond to date solely based on one-country (i.e., United States) mortality experience. However, the authors acknowledge that their model misses correlations and diversification effects across genders and populations. Given the prevalence of multicountry combined mortality indexes in the mortality security markets, it is questionable whether their model is adequate for other mortality securities. Indeed, from a practical point of view, understanding the combined index as a function of various population death rates is at least as important, if not more so, than understanding it as a function of different age classes of a single country for mortality securitization.

Moreover, the existing mortality literature has demonstrated the significance of catastrophic death events in the pricing and tranche structure for a mortality risk bond (Cox,

Lin, and Wang, 2006; Bauer and Kramer, 2009; Chen and Cox, 2009). While longevity risk modeling usually simplifies the analysis by ignoring dramatic mortality changes (Cairns, Blake, and Dowd, 2006; Schrager, 2006; Kogure and Kurachi, 2010; Wills and Sherris, 2010; Yang, Yue, and Huang, 2010; Cox et al., 2012), mortality jumps must be considered in order to successfully structure and price mortality-linked securities. Thus, in this article, as the first objective, we develop a tractable mortality model, which captures the mortality correlations among countries and incorporates mortality jumps. Specifically, we extend Cox, Lin, and Wang's (2006) model to a more general setting and disentangle transient jumps from persistent volatilities. As a departure from Cox, Lin, and Wang who model unanticipated mortality jumps as permanent shocks, we model them as transient jumps with a double-jump process. Since in most cases severe short-term events such as epidemics underlie mortality risk (Cox, Lin, and Petersen, 2010), our model provides a better fit for historical data.

A prominent calibration challenge for a model accounting for correlations is that the number of pairs grows quadratically with the number of countries of interests. To address this issue, as the second objective of this article, we employ a particle filtering approach for learning about unobservable mortality shocks and states from discretely observed population mortality rates. The particle filter algorithm is easy to implement and fast to compute, requiring only simulation from a proposal distribution. It is able to handle a large number of correlations with relatively low computation costs. Moreover, particle filters are highly adaptable and easy to be adjusted for various applications. In particular, they are able to deal with nonlinear, non-Gaussian systems, and variables in either continuous or discrete states as well as their combinations (Heijden et al., 2004). This property of particle filtering is important for mortality modeling because multivariate jump-diffusion processes are an indispensable component of a comprehensive mortality model, which is typically not Gaussian. In addition, the particle filtering approach is a Bayesian approach through which we can deal with estimation error and prediction uncertainty unaddressed in some mortality models. As highlighted in Cairns, Blake, and Dowd (2006), a model that takes into account parameter uncertainty is able to generate more reliable forecasts. So far the particle filtering approach has many practical applications in engineering and finance. Despite its usefulness, this technique has received little attention in the mortality literature.¹ In this article, we show that it is useful to estimate variables related to mortality trend, volatility, correlation, and jump, which are crucial for mortality security pricing and risk management.

Aside from the question of how to appropriately model mortality stochastic process, there is an ongoing debate on how to price mortality-linked securities. If a security's payments are contingent on the correlation of mortality risks across countries, a common feature in mortality securitization, the pricing problem will become more challenging (Chen and Cox, 2009). Mortality securities, with only a limited number issued in financial markets, are not liquidly traded. As a result, to price a mortality risk linked security, the underlying mortality risk process needs to be risk adjusted (Wills and Sherris, 2010). Dif-

1. To the best of our knowledge, there is only one paper by Bauer and Kramer (2009) that applies the particle filtering in the mortality context. However, Bauer and Kramer apply the particle filtering in the mortality evolution of the U.S. population without considering mortality correlation, whereas mortality correlation is one of our main focuses.

ferent mortality pricing methods have been proposed. Dahl (2004) applies financial risk models for mortality risk modeling and then uses market data to calibrate risk-adjusted probability measures. Lin and Cox (2005), Dowd et al. (2006), and Denuit, Devolder, and Goderniaux (2007) apply the Wang transform (Wang, 2000) to physical distributions in pricing distortion distributions. Bayraktar et al. (2009) propose to compensate a mortality risk taker according to an “instantaneous Sharpe ratio,” which is defined as the additional return in excess of the risk-free rate divided by the standard deviation of a mortality portfolio after all diversifiable risk is hedged away. In line with their Bayesian mortality models, Kogure and Kurachi (2010) present a Bayesian pricing approach, the entropy maximization principle, to risk neutralize the predictive distribution of future survival rates. However, while in these articles the transformed parameters are calibrated, they are built on a univariate setting or an independent assumption, ignoring the connections among different cohorts and populations. Thus, their pricing methods have limited implications for multivariate and correlated mortality cases. To address this pricing issue, we employ normalized multivariate exponential tilting to take into account correlations across countries for mortality securitization.

Exponential tilting is an incomplete market pricing method that neutralizes statistical distributions, which is consistent with the literature on nonarbitrage pricing of contingent claims (see Buhlmann, 1980; Gerber and Shiu, 1996; Madan and Unal, 2004; Kijima, 2006; Wang, 2007; and others). It can be applied in pricing risks embedded in loan defaults, mortgage refinancing, electricity trading, weather derivatives, catastrophic insurance, and insurance-linked securities (Duffie, 1992; Karatzas and Shreve, 1992; Heston, 1993; Gerber and Shiu, 1996; Cox, Lin, and Wang, 2006; Milidonis, Lin, and Cox, 2011). Kijima (2006) and Wang (2007) extend univariate exponential tilting to multivariate cases. The need for changing multivariate probability measures arises from pricing contingent claims on multiple underlying assets or liabilities. As noted earlier, the payoffs of the existing mortality securities are contingent on several population mortality indexes. To properly compensate investors for risk arising from mortality correlations among countries, as the last objective of this article, we apply the normalized multivariate exponential tilting to price mortality securities. Specifically, we first introduce the concept of normalized exponential tilting and then formulate probability distortions for the multivariate case. For demonstration, we utilize the normalized multivariate exponential tilting to make inference about the likely market prices of risk from the pricing information contained in the existing mortality securities. To the best of our knowledge, this is the first article that not only accounts for mortality correlations but also calibrates different market prices of risk for various mortality diffusions and jumps.

The article is organized as follows. In the next section, we introduce our proposed stochastic model for a set of correlated population mortality indexes with jumps. The “Estimation With Particle Filtering” section provides an overview on the particle filtering approach. The multivariate exponential tilting as an incomplete market pricing method is introduced in “Pricing Mortality Securities With Multivariate Exponential Tilting.” In “Empirical Applications,” we first estimate our proposed model by applying the particle filtering technique based on historical data. We show that mortality correlation and jump process play important roles in mortality securitization modeling. Then we price some mortality bonds issued in 2006 by applying normalized multivariate exponential tilting and compare our estimated market prices of risk to those when we do not consider mor-

tality correlations and jumps. After a discussion of our findings, “Conclusions” concludes and provides an outlook on future research.

Model Specification

Our approach combines Brownian motions and compound Poisson processes. As a distinctive feature of our model, we simultaneously consider n -country population mortality death rates and their correlations. Let $q_t^{(i)}$ denote the observed population mortality index for country i at time t . The logarithm of $q_t^{(i)}$ equals

$$y_{i,t} = \ln q_t^{(i)}$$

We assume that the n population mortality indexes jointly solve:

$$\begin{aligned} dy_{1,t} &= \alpha_1 + \sigma_1 dB_t^{(1)} + \sum_{v=1}^2 (m_{c1}^{(v)} + s_{c1}^{(v)} Z_1^{(v)}) dN_t^{(v)}, \\ dy_{2,t} &= \alpha_2 + \sigma_2 dB_t^{(2)} + \sum_{v=1}^2 (m_{c2}^{(v)} + s_{c2}^{(v)} Z_2^{(v)}) dN_t^{(v)}, \\ &\vdots \\ dy_{n,t} &= \alpha_n + \sigma_n dB_t^{(n)} + \sum_{v=1}^2 (m_{cn}^{(v)} + s_{cn}^{(v)} Z_n^{(v)}) dN_t^{(v)}, \end{aligned} \quad (1)$$

where α_i is country i 's instantaneous expected force of population mortality index; σ_i is country i 's instantaneous volatility of population mortality index, conditional on no jumps. The factor $B_t^{(i)}$ captures the correlation of $y_{i,t}$'s. We denote the correlation coefficient matrix of n countries, $\Sigma_{WW} = \{r_{ij}\}$, as

$$\Sigma_{WW} = \begin{pmatrix} 1 & r_{12} & \dots & r_{1n} \\ r_{12} & 1 & \dots & r_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ r_{1n} & r_{2n} & \dots & 1 \end{pmatrix}$$

Assume the matrix, Σ_{WW} is symmetric and positive definite. Then, there is a matrix A such that $\Sigma_{WW} = AA'$ with

$$A = \begin{pmatrix} A_{11} & A_{12} & \dots & A_{1n} \\ A_{12} & A_{22} & \dots & A_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ A_{1n} & A_{2n} & \dots & A_{nn} \end{pmatrix}$$

where each element $\{A_{ij}\}$ of A is a function of $\{r_{ij}\}$, $i, j = 1, \dots, n$.

Under this specification, the diffusion process, $B_t^{(i)}$ of country i , is a linear combination of A_{ij} 's (Shrieve, 2000):

$$B_t^{(i)} = \sum_{j=1}^n A_{ij} W_t^{(j)}$$

where $W_t^{(1)}, W_t^{(2)}, \dots, W_t^{(n)}$ are n independent Brownian motions with mean 0 and variance t .

All countries are subject to common random surprise events or shocks such as pandemics, which follow a Poisson process $N_t^{(v)}$ with intensity k_c . The common jumps (e.g., the 1918 worldwide flu) act as a fundamental shock to mortality curves of all countries and decrease the expected lifetime of the population across countries. To model those jumps, we assume both positive and negative movements by applying a double-jump process. As customers in queues come and go, so does a pandemic. As such, the common double-jump process has two jump size parameters $m_{ci}^{(v)} = 1$ or 2 with opposite signs and two jump size volatilities $s_{ci}^{(v)} = 1$ or 2 . $Z_i^{(v)}$ is the standard normal random variable:

$$Z_i^{(v)} \sim N(0,1).$$

From model (1), we can derive the mortality index $y_{i,t+h}$ given $y_{i,t}$ as follows:

$$y_{i,t+h} = y_{i,t} + \alpha_i h + \sigma_i \sum_{j=1}^n A_{ij} W_h^{(j)} + \sum_{v=1}^2 \sum_{l > N_t^{(v)}}^{N_{t+h}^{(v)}} (m_{ci}^{(v)} + s_{ci}^{(v)} Z_{il}^{(v)}) \quad (2)$$

For simplicity, for the double-jump process, we assume the same magnitude of positive and negative jumps:

$$m_{ci}^{(1)} = -m_{ci}^{(2)} = m_{ci}, \quad s_{ci}^{(1)} = s_{ci}^{(2)} = s_{ci}.$$

Estimation With Particle Filtering

Why Use Particle Filtering?

The particle filtering approach has several advantages over other estimation methods. First, it is rooted in Bayesian statistics. In the spirit of recursive Bayesian estimation, particle filtering is an iterative application of Bayes rule to sequential parameter estimation. In each step, the particle filtering approach first uses all information up to the step to predict future values. Then, based on the differences between predicted values and realized values, the parameter estimates are modified and updated. This procedure, to some extent, is similar to our learning process as it does not assume stationary model parameters. So it is superior to non-Bayesian approaches, such as the maximum likelihood method, that require stationary model parameters.

Second, the particle filtering is more adaptable than other filtering methods like the Kalman filter, which is built on the linear and Gaussian assumptions. Instead, as mentioned earlier, the particle filtering approach combined with sequential Monte Carlo method is more flexible and provides a way to estimate nonlinear, non-Gaussian, and nonstationary models. To be shown later, the particle filter generates a set of “particles,” which are samples of the distribution from Monte Carlo simulations. Those particles are then used to forecast the future and are updated when new observations are available. Mortality shocks are rare events so typically the distribution of mortality rates is non-Gaussian. To properly handle this situation, we use the particle filtering to estimate our model.

Particle Filtering Estimation Procedure

The objective of our estimation is to filter out the non-Gaussian state variables from unwanted noises, and this is where particle filters step in. Specifically, our estimation is to infer from a Bayesian perspective the probability function of the latent state, $S_t = \{x_{1,0:t}, x_{2,0:t}, \dots, x_{n,0:t}\}$, and the parameter vector, Θ_ν given the observed mortality time sequences, $Y_t = \{y_{1,0:t}, y_{2,0:t}, \dots, y_{n,0:t}\}$. The basis for the estimation is to specify state and measurement equations. Given the time-discrete observations of mortality indexes, we discretize Equation (1) using the Euler method:

$$x_{i,t+h} = x_{i,t} + \alpha_i h + \sigma_i \sum_{j=1}^n A_{ij} \Delta W^{(j)}(h) + \sum_{v=1}^2 \sum_{l > N_l^{(v)}}^{N_{t+h}^{(v)}} (m_{ci}^{(v)} + s_{ci}^{(v)} Z_{il}^{(v)}) \quad (3)$$

for $i = 1, 2, \dots, n$, where $\Delta W(h)$ is a standard normal random variable, and h is the time discretization interval (1 year in this article, i.e., $h = 1$), during which there are $N_{t+h}^{(v)}$ positive or negative jumps ($v = 1$ or 2). That is, our latent state S_t is composed of unobservable mortality Brownian motions and double jumps. Equation (3) serves as the state equation in the particle filtering. The jump size of $m_{ci}^{(v)} + s_{ci}^{(v)} Z_{il}^{(v)}$ retains its normal distributional structure, with the mean $m_{ci}^{(v)}$ and the standard deviation $s_{ci}^{(v)}$.

The measurement equation is given by

$$y_{i,t} = x_{i,t} + u_{i,t} \quad (4)$$

in which the measurement error $u_{i,t}$ is independent and identically, normally distributed

$$u_{i,t} \sim N(0, \sigma_i^2)$$

where σ_i^2 is the variance of $u_{i,t}$.

We use $p(S_t | Y_t)$ to denote the likelihood, that is, the distribution of latent states conditional on the observed vector Y_t . Given an initial probability distribution function

$p(S_0|Y_0) = p(S_0)$ and the time step $h=1$, the update $p(S_t|Y_{t-1})$ can be obtained sequentially through prediction:

$$p(S_t|Y_{t-1}) = \int p(S_t|S_{t-1}) p(S_{t-1}|Y_{t-1}) dS_{t-1} \quad (5)$$

where $p(S_t|S_{t-1})$ can be estimated using the state Equation (3). The probability $p(S_t|Y_t)$ is updated by the Bayes rule,

$$p(S_t|Y_t) = \frac{p(Y_t|S_t)p(S_t|Y_{t-1})}{p(Y_t|Y_{t-1})} \quad (6)$$

In general, however, the posterior probability $p(S_t|Y_t)$ of our model cannot be determined analytically because of the mixture of diffusions and jumps in the model. Mixing diffusions with jumps makes analytic treatment of the filtering problem difficult. The particle filtering provides a solution to the problem in this setting. This technique makes two assumptions (Doucet et al., 2001):

1. The state evolution can be accurately simulated. In other words, accurate samples can be drawn from $p(S_t|S_{t-1})$;
2. The likelihood can be accurately evaluated as a function of state variables and observations.

With these two assumptions, we can estimate our model based on particle filtering with the Sampling-Importance-Resampling (SIR) algorithm. The key idea of particle filtering with the SIR algorithm is summarized in Appendix A.

Pricing Mortality Securities With Multivariate Exponential Tilting

A complete market is a market where any derivative security can be replicated and hedged by holding cash and the underlying asset. Mortality securities, contingent on mortality catastrophe events, are derivatives with mortality rates as their underlying indexes. Currently, there are only a limited number of mortality securities, which are illiquid. Furthermore, a catastrophic event can happen at a random time with a random size, which makes it impossible to create a hedge portfolio to replicate mortality jumps. This implies that mortality security markets are incomplete and contingent claims cannot be replicated by trading. Thus, it leaves open the problem of pricing securities in this situation.

One possibility of valuing mortality securities in an incomplete market is to apply exponential tilting. Given a statistical distribution of potential outcomes and their respective likelihoods from historical data, the exponential tilting risk neutralizes the statistical distribution. This technique, consistent with the existing literature on no-arbitrage pricing of contingent claims (Duffie, 1992; Heston, 1993; Karatzas and Shreve, 1992; Gerber and Shiu, 1996), has some popularity in pricing insurance-related risks (e.g., catastrophe property and mortality risks).

Another reason for us to apply the exponential tilting method in mortality securitization is that its multivariate version can account for multiple underlying assets or liabilities.

ties of a contingent claim and price the effect of their correlations accordingly (Cox, Lin, and Wang, 2006). This is important because the payoffs of existing mortality securities are typically determined by multiple population mortality indexes. By changing multivariate probability measures, the multivariate exponential tilting can capture the correlations of these indexes so as to properly disentangle different components of the risk premium and price such securities.

In the following discussion, we first present the basics of the exponential tilting approach. Then we adopt a normalization procedure proposed by Cox, Lin, and Wang (2006) and Wang (2007) to ensure a consistent interpretation of the risk-adjustment parameter. Next we formulate normalized multivariate exponential tilting that provides a general framework for pricing risks with respect to multiple reference risks. Later in “Empirical Applications,” we provide an example to show how to apply this approach to price mortality securities.

Normalized Univariate Exponential Tilting

Given a probability space (Ω, P) , consider two risks X and V . If X is continuous with respect to V , the exponential tilting of X with respect to V is defined as

$$f_X^*(x) = f_X(x) \frac{E[\exp(\lambda V) | X = x]}{E[\exp(\lambda V)]} \quad (7)$$

where λ is the risk adjustment parameter. $f_X(x)$ is the statistical or physical probability density function and $f_X^*(x)$ is the transformed probability density function after applying the exponential tilting. The ratio

$$RN(x) = \frac{E[\exp(\lambda V) | X = x]}{E[\exp(\lambda V)]} \quad (8)$$

is called the Radon-Nikodym derivative of f_X^* with respect to f_X .

If X uses itself as the reference, that is, $X = V$, Equation (7) becomes

$$f_X^*(x) = f_X(x) \frac{\exp(\lambda x)}{E[\exp(\lambda X)]} \quad (9)$$

This is the well-known Esscher transform, which has been implemented in option pricing, insurance-linked security pricing, and contingent claims valuation (Gerber and Shiu, 1996; Cox, Lin, and Petersen, 2010).

Except that V is normally distributed, the risk adjustment parameter λ does not have a consistent interpretation because “the scale and shape of the reference variable V can significantly impact the result of exponential tilting” (Wang, 2007). As such, we adopt a normalization procedure of the reference variable V through percentile mapping to a standard normal variable Z following Cox, Lin, and Wang (2006) and Wang (2007). Mathematically,

$$V = F_V^{-1}(\Phi(Z)) \quad (10)$$

where Φ is the cumulative distribution function of Z . After normalization, we replace V with Z as the reference variable. That is, we implement “normalized” exponential tilting of X with respect to V by using Z . Formally,

Definition 1: Let Z a normalized variable of the reference V . The normalized exponential tilting of X with respect to Y is defined as:

$$f_X^*(x) = f_X(x) \frac{E[\exp(\lambda Z)|X = x]}{E[\exp(\lambda Z)]} \quad (11)$$

For more details, see Wang (2007).

Normalized Multivariate Exponential Tilting

Consider n variables X_1, X_2, \dots, X_n with k references V_1, V_2, \dots, V_k in a probability space (Ω, P) .

Definition 2: Given a probability space (Ω, P) , the exponential tilting of X_1, X_2, \dots, X_n with respect to references V_1, V_2, \dots, V_k is defined as the following probability density function:

$$f_X^*(x) = f_X(x) \frac{E[\exp(\lambda Z)|X = x]}{E[\exp(\lambda Z)]} \quad (12)$$

The references V_1, V_2, \dots, V_k can be flexibly specified to fit a particular situation. For instance, the references can be the risks themselves, that is, $(V_1, V_2, \dots, V_k) = (X_1, X_2, \dots, X_k)$ given $n = k$. In other cases, the references V_1, V_2, \dots, V_k represent a firm’s specific financial variables or some industry aggregate indexes. In the context of this article, V_1, V_2, \dots, V_k refer to population mortality indexes of different countries.

Similar to the univariate case in “Normalized Univariate Exponential Tilting,” in order to get a meaningful interpretation of the parameters λ_j in Equation (12) where $j = 1, \dots, k$, we need to apply a normalization procedure to all references V_1, V_2, \dots, V_k . Specifically, we calculate

$$V_1 = F_{V_1}^{-1}(\Phi(Z_1)), V_2 = F_{V_2}^{-1}(\Phi(Z_2)), \dots, V_k = F_{V_k}^{-1}(\Phi(Z_k)) \quad (13)$$

by percentile mapping where Z_1, Z_2, \dots, Z_k are standard normal variables.

Then, we can define the normalized multivariate exponential tilting of X_1, X_2, \dots, X_n with respect to references V_1, V_2, \dots, V_k as follows:

$$\frac{f^*(x_1, x_2, \dots, x_n)}{f(x_1, x_2, \dots, x_n)} = a E[\exp(\sum_{j=1}^k \lambda_j Z_j) | X_1 = x_1, X_2 = x_2, \dots, X_n = x_n] \quad (14)$$

With this setup, Wang (2007) proves the following theorem:

Theorem 1: Assume that $\{X_1, X_2, \dots, X_n; V_1, V_2, \dots, V_k\}$ follow a normal copula with correlation matrix:

$$\Sigma = \begin{pmatrix} \Sigma_{xx} & \Sigma_{xv} \\ \Sigma_{xv} & \Sigma_{vv} \end{pmatrix} \quad (15)$$

The normalized multivariate exponential tilting in Equation (14) of X_1, X_2, \dots, X_n w.r.t. V_1, V_2, \dots, V_k is equivalent to applying the multivariate Wang transforms to X_i with:

$$F_{X_i}^*(x_i) = \Phi[\Phi^{-1}(F_{X_i}(x_i)) + \beta_i], \text{ and } \beta_i = \sum_{j=1}^k \rho_{X_i, V_j} \cdot \lambda_j, (\text{for } i = 1, 2, \dots, n) \quad (16)$$

In addition, Wang (2007) shows that the correlation matrix between X_1, X_2, \dots, X_n is unchanged after the normalized multivariate exponential tilting, $\Sigma_{XX}^* = \Sigma_{XX}$. This is a desirable property because it simplifies the pricing procedure.

Empirical Applications

In this section, an example is provided to illustrate how to apply the proposed methods. We first use particle filtering to estimate mortality stochastic of six countries ($n = 6$): United States, United Kingdom, France, Germany, Japan, and Canada. Then, based on the mortality rates forecasted from model (1) and the normalized multivariate exponential tilting introduced in “Pricing Mortality Securities With Multivariate Exponential Tilting,” we derive the market prices of risk for different mortality diffusion and jump processes. The same technique can be easily extended to more complicated situations where a greater number of countries are involved.

Data

Our U.S. data from 1900 to 1998 are obtained from the Vital Statistics of the United States (VSUS).² The VSUS reports the U.S. age-adjusted death rates per 100,000 standard million population (2000 standard) for selected causes of death. Our mortality data for United States from 1999 to 2005 and United Kingdom and France from 1900 to 2005 are drawn from the Human Mortality Database.³ The Human Mortality Database reports the death and population size of country i for different ages. We divide the total number of deaths in different ages by the total population to get each population mortality index $q_t^{(i)}$. Following the same procedure, we calculate the population mortality index for Canada using the cohort data from 1900 to 1920 and the period data from 1921 to 2005 from the

2. Data source: <http://www.cdc.gov>

3. Data source: Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Available at <http://www.mortality.org> or <http://www.humanmortality.de> (the data were downloaded on June 8, 2008).

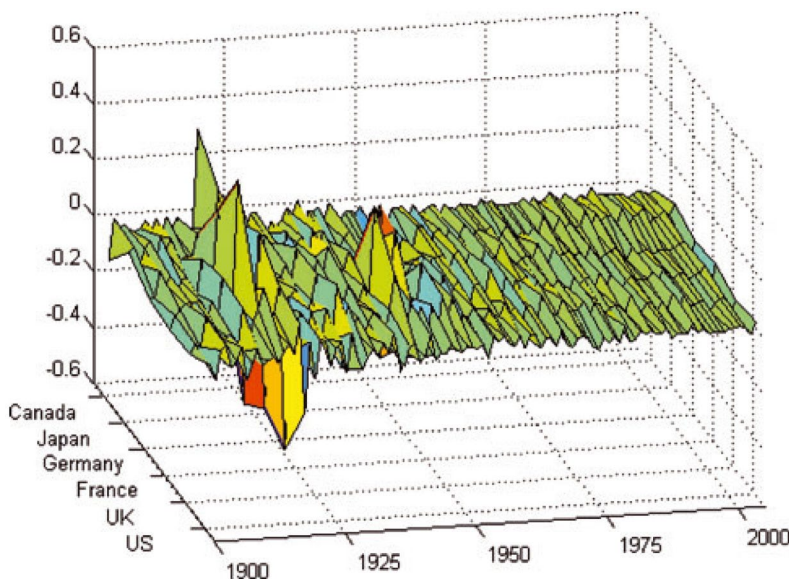


Figure 1. The Changes of Log-Mortality Rates, 1900 Through 2005. The figure presents the change of log-mortality rate, $\Delta \ln q_t^{(i)}$, for country i ($i = 1, 2, \dots, 6$) from 1900 to 2005. Country 1 is United States, country 2 is United Kingdom, country 3 is France, country 4 is Germany, country 5 is Japan, and country 6 is Canada.

Human Mortality Database.⁴ Our Japanese population mortality data are obtained from the Vital Statistics of the Ministry of Health, Labor and Welfare in Japan.⁵ The Japanese Vital Statistics database does not include the mortality data from 1944 to 1946. So we use the linear interpolation to generate the population mortality indexes for these 3 years. The German population mortality data from 1900 to 2005 are provided by the German Federal Statistical Office.⁶

We plot our data as the change of log-mortality rate, $\Delta \ln q_t^{(i)}$, from 1900 to 2005 in Figure 1, which shows a remarkable mortality fluctuation caused by the influenza breakout in 1918. Another visible fluctuation caused by influenza and pneumonia was around 1950 but its magnitude was much smaller as compared to that of the 1918 worldwide flu.

4. The Canadian cohort data do not capture the impact of the 1918 flu epidemic on the whole population. As such, we adjust the Canadian population mortality index in 1918 upward by 0.0045 to account for the excess death rate due to the flu based on the 1918 pandemic information from the Canadian Encyclopedia at <http://www.thecanadianencyclopedia.com>

5. <http://www.mhlw.go.jp>

6. <http://www.destatis.de>

Table 1. Particle Filtering Parameter Estimates Based on the United States, United Kingdom, France, Germany, Japan, and Canada Annual Population Mortality Indexes $q_t^{(i)}$ ($t = 1900, \dots, 2005$).

	i	α_i	σ_i	m_{ci}	s_{ci}
United States	1	-0.0091 (0.0013)	0.0404 (0.0085)	0.1314 (0.0276)	0.0602 (0.0178)
United Kingdom	2	-0.0032 (0.0006)	0.0810 (0.0237)	0.0001 (1.5×10^{-5})	0.1638 (0.0394)
France	3	-0.0077 (0.0013)	0.0433 (0.0094)	0.1947 (0.0514)	0.1317 (0.0227)
Germany	4	-0.0058 (0.0011)	0.0671 (0.0129)	0.0001 (3.1×10^{-5})	0.1249 (0.0347)
Japan	5	-0.0055 (0.0018)	0.0218 (0.0034)	0.1888 (0.0325)	0.0562 (0.0142)
Canada	6	-0.0060 (0.0008)	0.0336 (0.0076)	0.1218 (0.0247)	0.1711 (0.0265)

The standard errors are reported in parentheses.

Parameter Estimation

We need to estimate the parameters $\Theta = \{\theta_k, k = 1, 2, \dots, M\}$ besides the state vector. One can achieve, in particle filtering, the joint estimation of states and parameters by the augmented state space $\{S_t; \Theta_t\}$ where each parameter is specified as a Gaussian random walk

$$\theta_{k,t} = \theta_{k,t-1} + \varsigma_{k,t}, \quad \varsigma_{k,t} \sim N(0, \sigma_{\theta_k})$$

Then the same steps described in “Estimation With Particle Filtering” can be followed to estimate the augmented state variables. It is worth noting that the use of random walk may lead to estimated posterior distributions far more diffuse than the actual posteriors. To remedy this problem, we adopt the kernel smoothing approach (Doucet et al., 2001), which is described in Appendix B.

Estimation Results. Table 1 presents our parameter estimates after we apply the particle filtering approach. On average, the population mortality of the United States, United Kingdom, France, Germany, Japan, and Canada improves 0.91 percent, 0.32 percent, 0.77 percent, 0.58 percent, 0.55 percent, and 0.60 percent, per year, respectively. The six countries’ mortality volatilities σ_i are comparable, with United Kingdom having the highest value of 0.0810. The common jump size volatility parameter s_{ci} for each country i is much larger than that of the Brownian motion volatility σ_i . It coincides with what we observe from Figure 1: mortality jumps have a much larger magnitude than mortality normal fluctuations. Moreover, our estimate of the Poisson parameter equals $k_c = 0.0120$, implying that the worldwide mortality jump is approximately a one-in-eighty-three-year ($1/k_c \approx 83$) event.

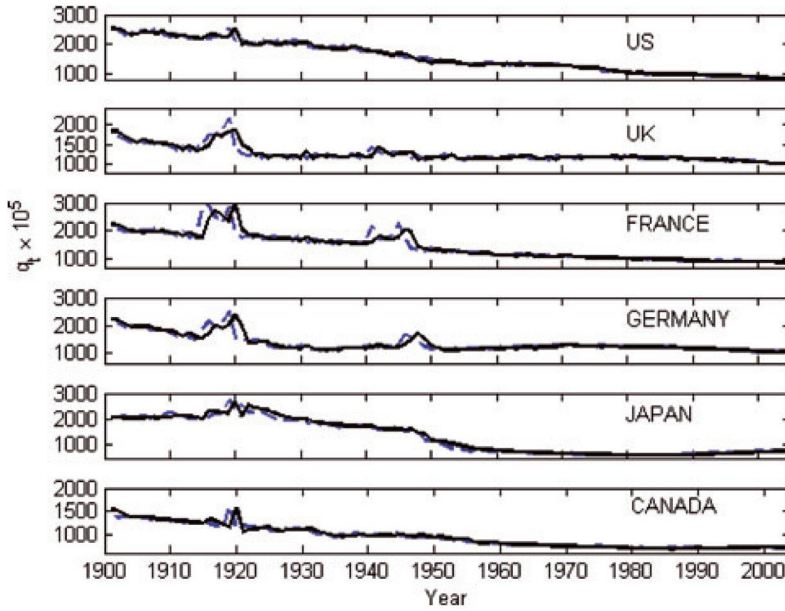


Figure 2. Filtering Results and Data. The figure compares the actual and posterior means of the filtered population death rate per 100,000 ($= 100,000 q_t^{(i)}$) distribution for the United States ($i= 1$), United Kingdom ($i= 2$), France ($i= 3$), Germany ($i= 4$), Japan ($i= 5$), and Canada ($i= 6$) in years $t= 1900, 1901, \dots, 2005$. The dash line represents the historical data and the solid line plots the death rates estimated from the particle filtering.

Our estimated correlation matrix of the six countries' instantaneous mortality indexes, Σ_{WW} is shown as follows:

$$\Sigma_{WW} = \begin{pmatrix} 1 & r_{12} & r_{13} & r_{14} & r_{15} & r_{16} \\ r_{12} & 1 & r_{23} & r_{24} & r_{25} & r_{26} \\ r_{13} & r_{23} & 1 & r_{34} & r_{35} & r_{36} \\ r_{14} & r_{24} & r_{34} & 1 & r_{45} & r_{46} \\ r_{15} & r_{25} & r_{35} & r_{45} & 1 & r_{56} \\ r_{16} & r_{26} & r_{36} & r_{46} & r_{56} & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0.11 & 0.01 & 0.11 & 0.28 & 0.35 \\ 0.11 & 1 & 0.43 & 0.56 & 0.05 & 0.02 \\ 0.01 & 0.43 & 1 & 0.80 & 0.18 & 0.16 \\ 0.11 & 0.56 & 0.80 & 1 & 0.22 & 0.17 \\ 0.28 & 0.05 & 0.18 & 0.22 & 1 & 0.21 \\ 0.35 & 0.02 & 0.16 & 0.17 & 0.21 & 1 \end{pmatrix} \quad (17)$$

From the correlation matrix (17), we can see that the Brownian motions of the six countries are all positively correlated. Among all of these correlations, the correlation between France and Germany is the highest ($r_{34}= 0.80$). This can be explained by the fact that both of them are in continental Europe and they share a common border. Their geographical connection leads to a high level of mortality comovements.

Figure 2 provides a graphical depiction of the filter's performance with yearly sampling frequency. It displays the actual population death rate per 100,000 $= 100,000 q_t^{(i)}$

and the filtered posterior mean for country i from 1900 to 2005. On average, the posterior mean of each country tracks the actual observation nicely.

Out-of-Sample Performance. Most of the existing papers on mortality risk modeling do not consider mortality correlations and temporary mortality common jumps in a single model. In “Model Specification,” we use the combination of Brownian motions and a double jump process to describe the dynamics of correlated mortality indexes of several countries. To justify this, we compare the prediction standard errors of our model, based on different prediction time horizons, with those of (1) a model with jumps but without correlations and (2) a model with neither common jumps nor correlations.

We first estimate these three models over the periods 1900 to 1950, 1900 to 1965, and 1900 to 1980, respectively.⁷ Then we conduct out-of-sample forecasting for each of them. In each sample period, we simulate 10,000 paths of the mortality indexes and measure prediction standard error across all six countries, prediction years, and paths. Mathematically, the prediction standard error is the square root of the prediction variance. In our case, the prediction variance is the variation associated with the difference between the true and predicted log-mortality rates, $\ln q_t^{(i)}$. A lower prediction standard error suggests a good prediction power. We compare the 3-, 4-, and 5-year prediction standard errors of the three models and the results are reported in Table 2.⁸ Table 2 shows that the model with correlations and common jumps consistently has the lowest prediction standard error in all specifications, which suggests that it has a better forecasting power than the other two models with less parameters.

Pricing Mortality Securities

To illustrate how to apply the normalized multivariate exponential tilting on the multidimension mortality dynamics described by our proposed mortality model (1), we will use market data to calibrate risk-adjusted probability measures.⁹ In “Estimation Results,” we have obtained the parameters for six Brownian motions for United States, United Kingdom, France, Germany, Japan, and Canada, and a common jump process. To esti-

7. To conserve space, we do not report the parameter estimates for these models in different sample periods. The results are available upon request.

8. Pricing mortality securities relies on a good mortality model. Based on seven mortality securities issued in 2006, in “Pricing Mortality Securities” we calibrate market prices of risk for different diffusion and common jump processes with the parameters estimated from our proposed model with correlations and common jumps. These mortality securities have a maturity of 3, 4, or 5 years. To show the performance of our model in the life of these securities, we present the prediction standard errors of 3, 4, and 5 years. We reach the same conclusions for longer prediction periods and different sample periods while not presented here to conserve space.

9. The risk adjustment factor in the Wang transform is usually called the market price of risk (Lin and Cox, 2005; Dowd et al., 2006; Denuit, Devolder, and Goderniaux, 2007; Chen and Cox, 2009; and others). Following this line of research, due to the equivalence of the normalized multivariate exponential tilting and the multivariate Wang transform, we also call the risk adjustment factor λ in the normalized multivariate exponential tilting “market price of risk.”

Table 2. Three-, Four-, and Five-Year Prediction Standard Errors of Models With and Without Correlations and Common Jumps Based on Sample Periods of 1900 Through 1950, 1900 Through 1965, and 1900 Through 1980.

Sample Period 1900 through 1950				
Correlation	Common Jump	Prediction Period		
		3-Year	4-Year	5-Year
No	No	0.1246	0.1415	0.1572
No	Yes	0.0889	0.0947	0.0975
Yes	Yes	0.0837	0.0893	0.0924

Sample Period 1900 Through 1965				
Correlation	Common Jump	Prediction Period		
		3-Year	4-Year	5-Year
No	No	0.1208	0.1363	0.1470
No	Yes	0.0837	0.0863	0.0899
Yes	Yes	0.0748	0.0791	0.0821

Sample Period 1900 Through 1980				
Correlation	Common Jump	Prediction Period		
		3-Year	4-Year	5-Year
No	No	0.1162	0.1284	0.1420
No	Yes	0.0825	0.0832	0.0851
Yes	Yes	0.0730	0.0770	0.0799

mate the market prices of risk for these seven processes, we select the following seven mortality securities issued in 2006 that meet our needs: Vita II mortality bond Tranche B, Vita II mortality bond Tranche C, Vita II mortality bond Tranche D, Tartan mortality bond Tranche B, Osiris mortality bond Tranche B2, Osiris mortality bond Tranche C, and Osiris mortality bond Tranche D.

In the following discussion, we first briefly introduce these mortality securities and then show how to estimate the market prices of risk for different processes underlying these securities.

Mortality Securities Issued in 2006 for Our Calibration. Capital markets have grown increasingly important in recent years as a solution to deal with mortality risk and expand the capacity of paying catastrophic mortality losses. For instance, to transfer catastrophic death risk from events such as epidemics, earthquakes, hurricanes, and other natural or man-made disasters, Swiss Re launched its second mortality transaction in 2006, the Vita

Capital II Ltd. mortality bond (Vita II) with three tranches (i.e., Tranches B, C, and D) after it issued the first pure mortality bond in December 2003.

Tranches B, C, and D of this deal provide Swiss Re total coverage of \$362 million from investors against extreme mortality risk scenarios. This bond is a 5-year deal, with the principal subject to mortality risk. Tranches B, C, and D pay a par spread of 0.90 percent, 1.4 percent, and 1.9 percent, respectively. The mortality risk is defined in terms of a combined mortality index (CMI), \bar{q}_t , which is the average of two annual index values over the period concerned:

$$\bar{q}_t = (q_{t-1} + q_t) \div 2, \quad t = 2007, \dots, 2011$$

The annual index q_t is a weighted average annual population death rates of five countries: 62.5 percent of q_t is based on the U.S. mortality index, 17.5 percent on the U.K., 7.5 percent on Germany, 7.5 percent on Japan, and 5 percent on Canada:

$$q_t = 0.65 q_t^{(1)} + 0.175 q_t^{(2)} + 0.075 q_t^{(3)} + 0.075 q_t^{(4)} + 0.005 q_t^{(5)}$$

Both the trigger level and exhaustion level for observed mortality in the risk period will be measured against 2002/2003 index value, q_0 :

$$q_0 = (q_{2002} + q_{2003}) \div 2$$

For each tranche of Vita II, a trigger event is deemed to have occurred when the CMI exceeds i_d of the base level. Take Tranche B as an example: If \bar{q}_t ($t = 2007, \dots, 2011$) exceeds $i_d = 120$ percent of the 2002/2003 index value, q_0 , the principal payment will be reduced. The principal will be forfeited when \bar{q}_t equals or goes above $i_u = 125$ percent of q_0 . Specifically, the principal loss percentage, L_t in year t , is determined as follows:

$$L_t = \begin{cases} 0 & \text{if } \bar{q}_t \leq 1.2q_0 \\ \frac{\bar{q}_t - 1.2q_0}{0.05q_0} & \text{if } 1.2q_0 < \bar{q}_t \leq 1.25q_0 \\ 1 & \text{if } \bar{q}_t > 1.25q_0. \end{cases}$$

However, a year can be counted only once toward a loss event. For example, if year 2007 resulted in a loss in 2006 through 2007, then the next earliest possible loss would be in 2008 through 2009. No loss of the principal would be possible in 2007 through 2008 because 2007 was already considered.

Therefore, given the size of \$62 million, the principal of Tranche B to be paid at maturity is equal to

Tranche B's Payment at Maturity

$$= 62,000,000 \times \begin{cases} 100\% - \sum_{t=2007}^{2011} L_t & \text{if } \sum_{t=2007}^{2011} L_t < 100\% \\ 0 & \text{if } \sum_{t=2007}^{2011} L_t \geq 100\% \end{cases} \quad (18)$$

The payment of Tranche C (Tranche D) of Vita II at maturity is similarly determined except that its size is \$200 million (\$100 million) and its attachment point i_d and detachment point i_u are 1.15 and 1.2 (1.1 and 1.15), respectively. The above discussion on the three tranches of Vita II is summarized in Table 3. Table 3 also provides an overview on the other four 2006 mortality securities we use for calibrating market prices of risk.

Pricing. In the previous discussion, we have shown how to describe the U.S., U.K., France, Germany, Japan, and Canada population mortality dynamics by using Brownian motions

$$W_t^{(1)}, W_t^{(2)}, W_t^{(3)}, W_t^{(4)}, W_t^{(5)}, W_t^{(6)}$$

and jump sizes

$$J^{(c1)}, J^{(c2)}, J^{(c3)}, J^{(c4)}, J^{(c5)}, J^{(c6)}$$

where $J^{(ci)} = \exp(m_{ci} + \sigma_{ci} Z_i)$. The correlations between $W_t^{(j)}$'s are Σ_{WW} as shown in (17). The common jump sizes $J^{(ci)}$'s are assumed to be independent of $W_t^{(j)}$'s. Further assume that

$$W_t^{(1)}, W_t^{(2)}, W_t^{(3)}, W_t^{(4)}, W_t^{(5)}, W_t^{(6)}, J^{(c1)}, J^{(c2)}, J^{(c3)}, J^{(c4)}, J^{(c5)}, J^{(c6)}$$

use themselves as references (Cox, Lin, and Wang, 2006).

Equation (16) suggests that the risk-neutralized distributions are equivalent to changing measure on the Brownian motions and jump sizes. That is, $W_t^{(j)*} = W_t^{(j)} + \beta_{W^{(j)}} t$ for $j=1, \dots, 6$, and $m_{ci}^* = m_{ci} + \beta_J \sigma_{ci}$ for $i=1, \dots, 6$. Accordingly, we define a vector \mathbf{B} in such a way that it is composed of parameters for risk-adjusting physical distributions of various mortality Brownian motions and jump sizes:

$$\mathbf{B}' = (\beta_{W^{(1)}} \beta_{W^{(2)}} \beta_{W^{(3)}} \beta_{W^{(4)}} \beta_{W^{(5)}} \beta_{W^{(6)}} \beta_J) \quad (19)$$

Table 3. Mortality Securities for Calibrating Market Prices of Risk

	Vita II Tranche B	Vita II Tranche C	Vita II Tranche D
Size	\$62M	\$200M	\$100M
Issue date	April 2006	April 2006	April 2006
Maturity	5 years	5 years	5 years
Coupon (bps)	LIBOR+90	LIBOR+140	LIBOR+190
Attachment	$1.2 \times [(q_{2002}+q_{2003})/2]$	$1.15 \times [(q_{2002}+q_{2003})/2]$	$1.1 \times [(q_{2002}+q_{2003})/2]$
Detachment	$1.25 \times [(q_{2002}+q_{2003})/2]$	$1.2 \times [(q_{2002}+q_{2003})/2]$	$1.15 \times [(q_{2002}+q_{2003})/2]$
Index	United Sates 62.5%, United Kingdom 17.5%, Germany 7.5%, Japan 7.5%, Canada 5%		
	Tartan Tranche B		
Size	\$80M		
Issue date	May 2006		
Maturity	3 years		
Coupon (bps)	LIBOR+300		
Attachment	$1.1 \times [(q_{2004}+q_{2005})/2]$		
Detachment	$1.15 \times [(q_{2004}+q_{2005})/2]$		
Index	United States 100%		
	Osiris Tranche B2	Osiris Tranche C	Osiris Tranche D
Size	Euro 50M	\$150M	\$100M
Issue date	November 2006	November 2006	November 2006
Maturity	4 years	4 years	4 years
Coupon (bps)	LIBOR+120	LIBOR+285	LIBOR+500
Attachment	$1.14 \times [(q_{2004}+q_{2005})/2]$	$1.10 \times [(q_{2004}+q_{2005})/2]$	$1.06 \times [(q_{2004}+q_{2005})/2]$
Detachment	$1.19 \times [(q_{2004}+q_{2005})/2]$	$1.14 \times [(q_{2004}+q_{2005})/2]$	$1.10 \times [(q_{2004}+q_{2005})/2]$
Index	United States 15%, Japan 25%, France 60%		

Based on Equation (16), \mathbf{B} is given by

$$\mathbf{B} = \Sigma \Lambda \quad (20)$$

where Σ is the correlation matrix of different countries' mortality Brownian motions and common jump,

$$\Sigma = \begin{pmatrix} \Sigma_{WW} & \Sigma_{JW} \\ \Sigma_{WJ} & \Sigma_{JJ} \end{pmatrix}$$

where Σ_{WW} has the estimated values shown in (17). Given no correlation between the Brownian motions and the jump process, it is easy to obtain

$$\Sigma = \begin{pmatrix} 1 & 0.11 & 0.01 & 0.11 & 0.28 & 0.35 & 0 \\ 0.11 & 1 & 0.43 & 0.56 & 0.05 & 0.02 & 0 \\ 0.01 & 0.43 & 1 & 0.80 & 0.18 & 0.16 & 0 \\ 0.11 & 0.56 & 0.80 & 1 & 0.22 & 0.17 & 0 \\ 0.28 & 0.05 & 0.18 & 0.22 & 1 & 0.21 & 0 \\ 0.35 & 0.02 & 0.16 & 0.17 & 0.21 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \quad (21)$$

The vector Λ' includes the parameters that control the magnitude of risk adjustments,

$$\Lambda' = (\lambda_{W(1)} \lambda_{W(2)} \lambda_{W(3)} \lambda_{W(4)} \lambda_{W(5)} \lambda_{W(6)} \lambda_J) \quad (22)$$

Then we can calculate all elements in the vector \mathbf{B} as follows:

$$\begin{aligned} \beta_{W(1)} &= \lambda_{W(1)} + r_{12}\lambda_{W(2)} + r_{13}\lambda_{W(3)} + r_{14}\lambda_{W(4)} + r_{15}\lambda_{W(5)} + r_{16}\lambda_{W(6)}, \\ \beta_{W(2)} &= r_{12}\lambda_{W(1)} + \lambda_{W(2)} + r_{23}\lambda_{W(3)} + r_{24}\lambda_{W(4)} + r_{25}\lambda_{W(5)} + r_{26}\lambda_{W(6)}, \\ \beta_{W(3)} &= r_{13}\lambda_{W(1)} + r_{23}\lambda_{W(2)} + \lambda_{W(3)} + r_{34}\lambda_{W(4)} + r_{35}\lambda_{W(5)} + r_{36}\lambda_{W(6)}, \\ \beta_{W(4)} &= r_{14}\lambda_{W(1)} + r_{24}\lambda_{W(2)} + r_{34}\lambda_{W(3)} + \lambda_{W(4)} + r_{45}\lambda_{W(5)} + r_{46}\lambda_{W(6)}, \\ \beta_{W(5)} &= r_{15}\lambda_{W(1)} + r_{25}\lambda_{W(2)} + r_{35}\lambda_{W(3)} + r_{45}\lambda_{W(4)} + \lambda_{W(5)} + r_{56}\lambda_{W(6)}, \\ \beta_{W(6)} &= r_{16}\lambda_{W(1)} + r_{26}\lambda_{W(2)} + r_{36}\lambda_{W(3)} + r_{46}\lambda_{W(4)} + r_{56}\lambda_{W(5)} + \lambda_{W(6)}, \\ \beta_J &= \lambda_J. \end{aligned} \quad (23)$$

Based on Equations (1) and (20), the six-country population mortality indexes forecasted from the estimated parameters in Table 1 and the correlation matrix Σ in (21), the U.S. Treasury rates on December 29, 2006, as well as the spreads of the seven mortality securities in Table 3, our estimated market prices of risk for different diffusion and jump processes equal

$$\begin{aligned} \Lambda' &= (\lambda_{W(1)} \lambda_{W(2)} \lambda_{W(3)} \lambda_{W(4)} \lambda_{W(5)} \lambda_{W(6)} \lambda_J) \\ &= (0.0726 \ 0.3329 \ 0.1062 \ 0.1212 \ 0.0277 \ 0.0562 \ 3.2885) \end{aligned} \quad (24)$$

In (24), the market price of risk for the common jump ($\lambda_J = 3.2885$) is much higher than those of the country-specific Brownian motions ($\lambda_{W(i)}$, $i = 1, \dots, 6$). The high-risk premium of the common jump reflects the risk aversion of insurers or investors when they face unhedgeable catastrophe mortality risks (Froot and Stein, 1998; Lin and Cox, 2008).¹⁰ However, insurers pay a low premium to transfer their mortality diffusion risk.

10. Most of the mortality securities were issued to protect insurance and reinsurance companies. Here we use "insurers" to refer to both insurers and reinsurers.




Figure 3. Normalized Multivariate Exponential Tilting Transformed Probability Distribution of $(q_{2009}^* + q_{2010}^*)/2$ (Shown as Broken Line) and Physical Probability Distribution of $(q_{2009} + q_{2010})/2$ (Shown as Solid Line) Based on 10,000 Simulation Trials. The horizontal axis is the average of the weighted population mortality indexes in 2009 and 2010 defined in the Vita II bond (62.5 percent on U.S. population death rate, 17.5 percent on the U.K., 7.5 percent on Germany, 7.5 percent on Japan, and 5 percent on Canada) and the vertical axis stands for the probability.

This can be explained by the fact that the insurers can predict mortality normal deviations from the trend to some extent. The regular mortality fluctuations are less likely to impose them a high financial distress cost so the risk premium is lower. Among all market prices of risk for various diffusions, the United Kingdom's Brownian motion has the highest market price of risk ($\lambda_{W(2)} = 0.3329$). This is consistent with our parameter estimations in Table 1: the United Kingdom has a much higher volatility $\sigma_2 = 0.0810$ than other sample countries. A higher volatility implies a higher risk so it requires a higher market price of risk.

We can also obtain the transformed probability distribution of the weighted average population mortality rate q_t^* for different mortality securities. For example, the transformed annual CMI of Vita II equals

$$q_t^* = 0.625 q_t^{*(1)} + 0.175 q_t^{*(2)} + 0.075 q_t^{*(4)} + 0.075 q_t^{*(5)} + 0.05 q_t^{*(6)}$$

where $q_t^{*(i)}$ is the transformed country i 's population mortality rate after applying normalized multivariate exponential tilting.

As noted earlier, the payoffs of the Vita II mortality bond depend on $(q_{t-1} + q_t)/2$ ($t = 2007, \dots, 2011$) where q_t is a weighted average of the population mortality indexes of United States (62.5 percent), United Kingdom (17.5 percent), Germany (7.5 percent), Japan (7.5 percent), and Canada (5 percent). Figure 3 plots its transformed probability distribution of $(q_{2009}^* + q_{2010}^*)/2$ with the market prices of risk in (24). Figure 3 also compares the

transformed distribution to its physical distribution $(q_{2009}+q_{2010})/2$.¹¹ As what we expect, the transformed distribution lies on the right of the physical distribution. After transformation, we put more weight on catastrophic mortality events (i.e., the right tail). It implies that investors expect a higher probability of extreme mortality losses than what the physical distribution suggests.

Impact of Correlations and Jumps on Pricing. Do the correlations and common jumps play an important role in mortality security pricing? To answer this question, we compare the market prices of risk in (24) from the model with correlations and common jumps with those of (1) the model with common jumps but without correlations and (2) the model without common jumps and correlations.

To obtain the market prices of risk for the model with common jumps but without correlations, we first run 10,000 trials of simulation based on the parameters in Table C1 in Appendix C where mortality correlations among six sample countries are not considered. Then we apply the normalized multivariate exponential tilting and obtain the following market prices of risk:

$$\begin{aligned}\Lambda'_{nc} &= (\lambda_{W(1)} \lambda_{W(2)} \lambda_{W(3)} \lambda_{W(4)} \lambda_{W(5)} \lambda_{W(6)} \lambda_J) \\ &= (0.2628 \ 0.8366 \ 0.8056 \ 0.2591 \ 0.0695 \ 0.0759 \ 3.6322)\end{aligned}\quad (25)$$

Assuming zero correlations, the market prices of risk in (25) are all higher than those in (24) when we account for the correlations of six-country population death rates. According to our proposed model (1) with correlations and jumps, all six countries have positive correlations with each other (see the correlation matrix (17)). That is, an increase in a country's mortality rate will be coincident with mortality rate increases of the other five countries. As such, a CMI based on the model with positive correlations has a higher mortality risk since it is more likely to exceed the trigger level above which the principal will be reduced. Given the fixed par spreads of the seven mortality securities, the higher mortality risk from the real world model implies the lower market prices of risk when the positive mortality correlations are incorporated in pricing. In contrast, failing to model mortality correlations would overestimate market prices of risk and overprice a "new" mortality security based on these estimates, thus discouraging insurers from transferring catastrophe mortality risk to capital markets. High-risk premia at least partially explain the failure of some insurance-linked securities, such as the EIB longevity bond issued in 2004 (Lin and Cox, 2008).

When we ignore both mortality correlations and common jumps, based on the parameters in Table C2 in Appendix C and 10,000 trials of simulation, we get the following mar-

11. We obtain similar graphs for various population mortality indexes in different years underlying the mortality securities in Table 3. To conserve space, we do not show them here.

ket prices of risk that deviate dramatically from the values in (24):

$$\begin{aligned}\Lambda'_{ncj} &= (\lambda_{W(1)} \lambda_{W(2)} \lambda_{W(3)} \lambda_{W(4)} \lambda_{W(5)} \lambda_{W(6)} \lambda_J) \\ &= (0 \quad 4.6414 \quad 0 \quad 1.1936 \quad 0 \quad 0.6426 \quad 0)\end{aligned}\tag{26}$$

In (26), $\lambda_J = 0$ because the model excludes mortality common jumps. As the pricing formula is based on tilting the real-world model, the adjustment parameter in terms of market price of risk is meaningful only if the baseline historical model is reliably specified. When the mortality common jumps are estimated as part of the diffusions, the model has a poor out-of-sample performance (see “Out-of-Sample Performance”) and thus leads to unreliable market prices of risk for different diffusion processes. From the above-mentioned analysis, we conclude that both mortality correlations and common jumps are indispensable components for mortality security pricing.

Conclusions

Mortality correlations become an increasingly important consideration in life insurance business and mortality security pricing, because populations of different countries and cohorts are more interdependent than ever. Life insurers that operate in multiple countries have to evaluate the effect of massive catastrophic mortality events such as worldwide flu epidemics. In the mortality security markets, most transactions, albeit their difference in structure, have one thing in common: Based on mortality data reported by official agencies, they tend to calculate a weighted average mortality index of multiple underlying reference populations. Therefore, the payoffs of these mortality securities are dependent on the correlations among these population mortality indexes.

In this article, we propose a time-continuous model for describing mortality correlations among different countries. For each included country, the mortality model is composed of two parts: (1) a baseline component describes a country’s mortality normal deviations from the trend and its correlation with other countries’ mortality rates and (2) a common jump component is governed by a double-jump process. The double-jump process aims at capturing transient catastrophic death events.

Many existing papers ignore mortality correlations to simplify their calibration procedures because the correlations to be estimated grow quadratically with the number of countries of interests. However, as shown in the article, mortality correlation has a significant impact on parameter estimation and mortality security pricing, so it should not be ignored. To tackle this calibration challenge in correlation estimation, we apply the particle filtering approach. We describe its procedures and resulting parameterizations in the context of mortality stochastic modeling.

Finally, we utilize the normalized multivariate exponential tilting method to price some mortality securities issued in 2006 and calibrate market prices of risk for different mortality diffusion and jump processes. The market prices of risk derived from the normalized multivariate exponential tilting depend on intercountry mortality correlations and jumps. Our numerical example confirms that both mortality correlations and common jumps are important in pricing those securities.

In sum, this article contributes to the mortality literature in three ways. First, we develop a mortality dynamic model that captures both mortality correlation among countries and common jumps that affect all countries. Second, we propose to use a double-jump process to model transient mortality shocks. Third, we apply the particle filtering approach to estimate multicountry mortality models with correlations. While we apply the particle filtering approach in the context of multicountry population mortality indexes, we can extend the application of this technique to investigate cohort mortality dependency within a country or among countries. We leave this for future work.

Acknowledgments — The authors thank two anonymous referees from *The Journal of Risk and Insurance* for their very helpful suggestions and comments during the revision process. Yijia Lin gratefully acknowledges financial support from the Research Council Grand-In-Aid Award—Maude Hammond Fling Faculty Research Fellowship of the University of Nebraska–Lincoln.

Appendix A

The key idea of particle filtering with the SIR algorithm is briefly described in the following recursive steps.

Step 1. This step generates random samples from previous state variables and observations. Specifically, given the observation Y_{t-1} and the estimated S_{t-1} at time $t-1$, we simulate shocks $\{\Delta W^j, \Delta N^j, Z^j\}$ and draw x_t^j ($j=1, \dots, L$) from $p(S_t|S_{t-1})$ by simulating the latent state vector forward using the state Equation (3). We draw L samples of x_t^j from the prior probability density of (3) and approximate $p(S_t|Y_{t-1})$ by using the set of L particles.

Step 2. We estimate the likelihood function $p(Y_t|S_t)$. The measurement Equation (4) shows that $p(Y_t|S_t)$ is normally distributed so it can be evaluated as a function of S_t and Y_t . For each sample x_t^j , set the importance weight equal to $p(Y_t|x_t^j)$. Then we calculate the normalized importance weight w_t^j ,

$$w_t^j = \frac{p(Y_t|x_t^j)}{\sum_{j=1}^L p(Y_t|x_t^j)}$$

where $\sum_{j=1}^L w_t^j = 1$. The set of $\{w_t^j\}$ can be viewed as a discrete probability distribution of x_t^j ($j=1, \dots, L$), that is, an approximate density $p^L(S_t|Y_t)$ to the true density $p(S_t|Y_t)$.

Step 3. In this step the particle filter estimates an approximation, $\hat{p}^L(S_t|Y_t)$, using L unweighted particles, of the true filtering density $p(S_t|Y_t)$. The resampling step modifies the weighted density $p^L(S_t|Y_{t-1})$, estimated previously, to an unweighted density $\hat{p}^L(S_t|Y_{t-1})$ by eliminating particles with low importance weights and by multiplying particles with high importance weights. That is, the resampling procedure eliminates selectively over time those particles with very lower probability, replaces them with particles with higher probability, and at the same time, keeps the number of particles unchanged.

12. There are a number of resampling algorithms. Three of them, multinomial resampling, systematic resampling, and residual resampling, are used in this study.

This step generates a new set of L particles, \hat{x}_t^j ($j = 1, \dots, L$).¹² As the number of particles, L , increases, the approximation converges to the true filtering distribution.

Appendix B

The process of the kernel smoothing approach (Doucet et al., 2001) is briefly described as follows:

Step 1. Given the observation Y_{t-1} and $L = 5,000$ samples of $\{x_{t-1}^j; \theta_{t-1}^j\}$, first simulate shocks $\{\Delta W^j, \Delta N^j, Z^j\}$ for the time interval from $t-1$ to t . Then sample the state variable x_{t-1}^j , $j = 1, \dots, L$, using (3) and compute the mean of x_{t-1}^j

$$\mu_t = E(S_t | S_{t-1}, \Theta_{t-1})$$

and the kernel location

$$\tau_t = \sqrt{1 - \xi^2} \Theta_{t-1} + \left(1 - \sqrt{1 - \xi^2}\right) \bar{\Theta}_{t-1}$$

where $\xi > 0$ is a smooth parameter and $\bar{\Theta}_{t-1}$ is the mean of the prior $p(\Theta_{t-1} | Y_{t-1})$.

Step 2. For $j = 1, 2, \dots, L$, sample a parameter vector θ_t^j from the j -th normal component of the kernel density, $\theta_t^j \sim \mathcal{N}(\cdot | \tau_t, \xi^2 V_t)$, where V_t is the variance/covariance matrix of $p(\Theta_t | Y_{t-1})$.

Step 3. Sample the values of the current state vector $\{x_t^j\}$ again. This time use the sampled parameter vector θ_t^j estimated in step 2. That is, step 3 draws samples $\{x_t^j\}$ from the distribution $p(S_t | S_{t-1}, \Theta_t)$ compared to step 2 where we draw samples $\{x_t^j\}$ from the distribution $p(S_t | S_{t-1}, \Theta_{t-1})$.

Step 4. We estimate the likelihood function $p(Y_t | S_{t-1}, \Theta_t)$ and evaluate the normalized importance weight

$$w_t^j = \frac{p(Y_t | x_t^j, \theta_t^j)}{p(Y_t | \mu_t, \tau_t)}$$

The set of w_t^j can be viewed as a discrete probability distribution of $\{x_t^j, \theta_t^j\}$, that is, an approximation of the true density $p(S_{t-1}, \Theta_t | Y_t)$.

Step 5. The resampling step modifies the weighted approximate density $p^L(S_{t-1}, \Theta_t | Y_{t-1})$ to an unweighted density $\hat{p}^L(S_{t-1}, \Theta_t | Y_{t-1})$ as in step 3 of Appendix A. The particles with low importance weights are eliminated and the particles with high importance weights are propagated so that the total number of particles remains constant. This step generates $\{\hat{S}_t, \hat{\Theta}_t\}$ where \hat{S}_t and $\hat{\Theta}_t$ are the new set of unweighted particles. Their density approximates to the true filtering density $p(S_{t-1}, \Theta_t | Y_t)$.

Repeat the above steps toward a final posterior estimation of states and parameters.

Appendix C

Based on the same data shown in the "Data" section, Table C1 shows the estimates when we assume zero mortality correlations between the six countries but with common jumps. Our estimate of the Poisson intensity in this model equals 0.0125, close to $k_c = 0.0120$ when we model mortality correlations and common jumps. Table C2 shows the estimation results for the model without mortality correlations and common jumps.

Table C1. Particle Filtering Parameter Estimates Based on the Model With Common Jumps but Without Correlations (1900 Through 2005)

	i	α_i	σ_i	m_{ci}	s_{ci}
United States	1	-0.0104 (0.0010)	0.0386 (0.0037)	0.1360 (0.0136)	0.0576 (0.0057)
United Kingdom	2	-0.0035 (0.0004)	0.0857 (0.0067)	0.0001 (8.6×10^{-6})	0.1451 (0.0147)
France	3	-0.0077 (0.0005)	0.0421 (0.0091)	0.2220 (0.0210)	0.1194 (0.0128)
Germany	4	-0.0060 (0.0005)	0.0689 (0.0070)	0.0001 (9.9×10^{-6})	0.1417 (0.0131)
Japan	5	-0.0061 (0.0007)	0.0207 (0.0050)	0.1784 (0.0174)	0.0616 (0.0059)
Canada	6	-0.0059 (0.0006)	0.0326 (0.0054)	0.1105 (0.0112)	0.1937 (0.0192)

The standard errors are reported in parentheses.

Table C2. Particle Filtering Parameter Estimates Based on the Model Without Correlations and Common Jumps (1900 Through 2005)

	i	α_i	σ_i
United States	1	-0.0100 (0.0010)	0.0373 (0.0039)
United Kingdom	2	-0.0038 (0.0004)	0.0682 (0.0070)
France	3	-0.0046 (0.0005)	0.0982 (0.0094)
Germany	4	-0.0045 (0.0005)	0.0726 (0.0071)
Japan	5	-0.0071 (0.0007)	0.0531 (0.0049)
Canada	6	-0.0058 (0.0006)	0.0526 (0.0054)

The standard errors are reported in parentheses.

References

- Bauer, D., and F.W. Kramer, 2009, Risk and Valuation of Mortality Contingent Catastrophe Bonds, Working Paper, Georgia State University and Allianz Investment Management SE.
- Bayraktar, E., M. A. Milevsky, S. D. Promislow, and V. R. Young, 2009, Valuation of Mortality Risk via the Instantaneous Sharpe Ratio: Applications to Life Annuities, *Journal of Economic Dynamics and Control*, 33(3): 676-691.
- Beelders, O., and D. Colarossi, 2004, Modelling Mortality Risk With Extreme Value Theory: The Case of Swiss Re's Mortality-Indexed Bond, *Global Association of Risk Professionals*, 4: 26-30.
- Blake, D., and W. Burrows, 2001, Survivor Bonds: Helping to Hedge Mortality Risk, *Journal of Risk and Insurance*, 68(2): 339-348.
- Buhlmann, H., 1980, An Economic Premium Principle, *ASTIN Bulletin*, 11(1): 52-60.
- Cairns, A. J. G., D. Blake, and K. Dowd, 2006, A Two-Factor Model for Stochastic Mortality With Parameter Uncertainty: Theory and Calibration, *Journal of Risk and Insurance*, 73(4): 687-718.
- Chen, H., and S. H. Cox, 2009, Modeling Mortality With Jumps: Applications to Mortality Securitization, *Journal of Risk and Insurance*, 76(3): 727-751.
- Chen, H., and J. D. Cummins, 2010, Longevity Bond Premiums: The Extreme Value Approach and Risk Cubic Pricing, *Insurance: Mathematics and Economics*, 46(1):150-161.
- Cox, S. H., and Y. Lin, 2007, Natural Hedging of Life and Annuity Mortality Risks, *North American Actuarial Journal*, 11(3): 1-15.
- Cox, S. H., Y. Lin, and H. Petersen, 2010, Mortality Risk Modeling: Application to Insurance Securitization, *Insurance: Mathematics and Economics*, 46(1):242-253.
- Cox, S. H., Y. Lin, R. Tian, and L. F. Zuluaga, 2012, Mortality Portfolio Risk Management, *Journal of Risk and Insurance*; doi: 10.1111/j.1539-6975.2012.01469.x
- Cox, S. H., Y. Lin, and S. Wang, 2006, Multivariate Exponential Tilting and Pricing Implications for Mortality Securitization, *Journal of Risk and Insurance*, 73(4):719-736.
- Dahl, M., 2004, Stochastic Mortality in Life Insurance: Market Reserves and Mortality-Linked Insurance Contracts, *Insurance: Mathematics and Economics*, 35(1):113-136.
- Denuit, M., P. Devolder, and A. Goderniaux, 2007, Securitization of Longevity Risk: Pricing Survivor Bonds With Wang Transform in the Lee-Carter Framework, *Journal of Risk and Insurance*, 74(1): 87-113.
- Doucet, A., N. de Freitas, N. Gordon, and A. Smith, 2001, *Sequential Monte Carlo Methods in Practice* (New York: Springer).
- Dowd, K., D. Blake, A. J. G. Cairns, and P. Dawson, 2006, Survivor Swaps, *Journal of Risk and Insurance*, 73(1): 1-17.
- Duffie, D., 1992, *Dynamic Asset Pricing Theory* (Princeton, NJ: Princeton University Press).
- Froot, K. A., and J. C. Stein, 1998, Risk Management, Capital Budgeting, and Capital Structure Policy for Financial Institutions: An Integrated Approach, *Journal of Financial Economics*, 47(1): 55-82.

- Gerber, H. U., and E. S. W. Shiu, 1996, Martingale Approach to Pricing Perpetual American Options on Two Stocks, *Mathematical Finance*, 6(3): 303-322.
- Heijden, F., R. Duin, D. Ridder, and D. M. J. Tax, 2004, *Classification, Parameter Estimation and State Estimation: An Engineering Approach Using MATLAB* (New York: Wiley).
- Heston, S. L., 1993, A Closed-Form Solution for Options With Stochastic Volatility With Applications to Bond and Currency Options, *Review of Financial Studies*, 6(2):327-343.
- Karatzas, I., and S. E. Shreve, 1992, *Brownian Motion and Stochastic Calculus*, 2nd edition (New York: Springer).
- Kijima, M., 2006, A Multivariate Extension of Equilibrium Pricing Transforms: The Multivariate Esscher and Wang Transforms for Pricing Financial and Insurance Risks, *ASTIN Bulletin*, 36(1): 269-284.
- Kogure, A., and Y. Kurachi, 2010, A Bayesian Approach to Pricing Longevity Risk Based on Risk-Neutral Predictive Distributions, *Insurance: Mathematics and Economics*, 46(1): 162-172.
- Lando, D., 1998, On Cox Processes and Credit Risky Securities, *Review of Derivatives Research*, 2: 99-120.
- Lin, Y., and S. H. Cox, 2005, Securitization of Mortality Risks in Life Annuities, *Journal of Risk and Insurance*, 72(2): 227-252.
- Lin, Y., and S. H. Cox, 2008, Securitization of Catastrophe Mortality Risks, *Insurance: Mathematics and Economics*, 42(2): 628-637.
- Madan, D., and H. Unal, 2004, Risk-Neutralizing Statistical Distributions: With an Application to Pricing Reinsurance Contracts on FDIC Losses, FDIC Center for Financial Research Working Paper No. 2004-01.
- Milidonis, A., Y. Lin, and S. H. Cox, 2011, Mortality Regimes and Pricing, *North American Actuarial Journal*, 15(2): 266-289.
- Schrager, D. F., 2006, Affine Stochastic Mortality, *Insurance: Mathematics and Economics*, 38(1): 81-97.
- Shrieve, S. E., 2000, *Stochastic Calculus for Finance II: Continuous Time Model* (New York: Springer).
- Wang, S. S., 2000, A Class of Distortion Operations for Pricing Financial and Insurance Risks, *Journal of Risk and Insurance*, 67(1): 15-36.
- Wang, S. S., 2007, Normalized Exponential Tilting: Pricing and Measuring Multivariate Risks, *North American Actuarial Journal*, 11(3): 89-99.
- Wills, S., and M. Sherris, 2010, Securitization, Structuring and Pricing of Longevity Risk, *Insurance: Mathematics and Economics*, 46(1): 173-185.
- Yang, S. S., J. C. Yue, and H.-C. Huang, 2010, Modeling Longevity Risks Using a Principal Component Approach: A Comparison With Existing Stochastic Mortality Models, *Insurance: Mathematics and Economics*, 46(1): 254-270.