ARC Centre of Excellence in Population Ageing Research

Working Paper 2014/09

A Multivariate Forward-Rate Mortality Framework

Daniel H. Alai[1], Katja Ignatieva[2] and Michael Sherris[3]

[1]Lecturer, School of Mathematics, Statistics and Actuarial Science
University of Kent and Associate Investigator, CEPAR email: D.H.Alai@kent.ac.uk
[2]Senior Lecturer, School of Risk and Actuarial Studies, UNSW Business School,
UNSW Australia email: k.ignatieva@unsw.edu.au
[3]Professor, School of Risk and Actuarial Studies, UNSW Business School, UNSW
Australia and Chief Investigator, CEPAR email: m.sherris@unsw.edu.au

A Multivariate Forward-Rate Mortality Framework

Daniel H. Alai[1]   Katja Ignatieva[2]   Michael Sherris[3]

School of Mathematics, Statistics and Actuarial Science
University of Kent, Canterbury, Kent CT2 7NF, UK

CEPAR, Risk and Actuarial Studies, Australian School of Business
UNSW, Sydney NSW 2052, Australia

Abstract

Stochastic mortality models have been developed for a range of applications from demographic projections to financial management. Financial risk based models build on methods used for interest rates and apply these to mortality rates. They have the advantage of being applied to financial pricing and the management of longevity risk. Olivier and Jeffery (2004) and Smith (2005) proposed a model based on a forward-rate mortality framework with stochastic factors driven by univariate gamma random variables irrespective of age or duration. We assess and further develop this model. We generalize random shocks from a univariate gamma to a univariate Tweedie distribution and allow for the distributions to vary by age. Furthermore, since dependence between ages is an observed characteristic of mortality rate improvements, we formulate a multivariate framework using copulas. We find that dependence increases with age and introduce a suitable covariance structure, one that is related to the notion of a minimum. The resulting model provides a more realistic basis for capturing the risk of mortality improvements and serves to enhance longevity risk management for pension and insurance funds.

Keywords: longevity risk, Olivier-Smith model, forward-rate mortality framework, minimum covariance pattern, copulas

---

[1]d.h.alai@kent.ac.uk
[2]k.ignatieva@unsw.edu.au
[3]m.sherris@unsw.edu.au

1 Introduction

A variety of empirical studies across many developed nations show that mortality trends have been improving stochastically; see e.g. CMI (2005), Luciano and Vigna (2005), Blake *et al.* (2006), Liu (2008), and Blackburn and Sherris (2012). Numerous stochastic mortality models proposed in the literature apply extensions of interest rate term structure modelling, known as *short-rate models.* In particular, the Cox-Ingersoll-Ross (CIR) model, Cox *et al.* (1985), has been adapted in Luciano and Vigna (2005), who use a time inhomogeneous version of the process to model the mortality dynamics. Furthermore, Biffis (2005), Russo *et al.* (2010), and Blackburn and Sherris (2012) develop a variety of affine frameworks extended from interest rate term structure modelling. Blake *et al.* (2006), Cairns *et al.* (2006), Bauer (2006) and Bauer and Ruß (2006) demonstrate that if mortality risk can be traded through securities such as longevity bonds and swaps, then the techniques developed in financial markets for pricing bonds and swaps can be adapted for mortality risk. Such models have the advantage of being able to incorporate a price of longevity risk and have applications to the valuation of mortality linked contracts, which can be used to mitigate longevity risk on life insurance products, including annuities. A typical underlying assumption in these models is that mortality rates are Gaussian.

An alternative to short-rate models is provided by forward-rate approaches that model forward forces of mortality and forward survival probabilities. Olivier and Jeffery (2004) and Smith (2005) apply a forward-rate mortality framework with stochastic factors driven by univariate gamma random variables irrespective of age or duration. Although the model is conceptually interesting, not much has been done to assess the validity of the model assumptions. In addition to restricting the stochastic factors to identical gamma distributions, another critical assumption is that of independence amongst these factors across age. We examine these assumptions using England and Wales female mortality data for 1960 to 2009; a dataset used in Cairns (2007).

Based on the empirical analysis, we generalize the Olivier-Smith model in order to provide a more realistic basis for capturing the risk of mortality improvements. Previous studies, e.g. Wills and Sherris (2010) and Jevtić *et al.* (2013), have verified that dependence plays a critical role in mortality modelling. A more realistic risk factor structure that captures such features of mortality data has the potential to provide enhanced risk management techniques for mitigating the longevity risk faced by pension funds and annuity providers. We assume random shocks are driven by non-identical univariate Tweedie random variables. We, thereby, allow the stochastic components to originate from distributions with varying levels of volatility. Our empirical investigations show this to be highly desirable. Furthermore, by specifying the Tweedie distribution, we can readily incorporate a relevant set of marginal distributions. We test both the Gaussian and gamma distributions and provide evidence that the gamma is the more suitable of the two. Since dependence between ages is an observed characteristic of mortality rate improvements, we further generalize the model by formulating a multivariate distribution that incorporates age dependence using copulas. In order to accomplish this we introduce a new covariance pattern, one that is governed by the minimum of the two age covariates. That is, the correlation between mortality rates of two ages depends only on the least of the two; this notion is applied to the age range 49-99.

Organization of the paper: In Section 2 we introduce the necessary notation for the forward-rate mortality framework and describe the Olivier-Smith model. We assess the model using data from England and Wales in Section 3. In Section 4 we introduce a univariate Tweedie generalization and in Section 5 consider a multivariate framework using copulas. Section 6 concludes the paper.

2 The Olivier-Smith Model

Let r_t denote the short rate at time t, and let B_t be the cash account defined as

$$B_t = \exp \int_0^t r_s ds.$$

Let $\mu(t, x+t)$ denote the force of mortality at time t for age $x+t$. This implies that x is the age at time zero. The age-specific survivor index is given by

$$S(t, x) = \exp \int_0^t -\mu(s, x+s) ds.$$

The survivor index $S(t, x)$ is the probability an individual aged x at time zero survives to time t. Notice the similarities of the cash account and the survivor index. Both are based on an instantaneous process. In the case of the cash account, the short rate; in the case of the survivor index, the force of mortality. However, unlike the short rate, the force of mortality also includes an age component. Clearly, $B_0 = S(0, x) = 1$; with the restriction of negative interest rates and forces of mortality, we have that the cash account increases and the survivor index decreases.

Consider an asset, call it a zero-coupon longevity bond, that pays $B_T S(T, x)$ at maturity T. Let $L(t, T, x)$ denote its price at time t. Given that the pay-off is tied to the cash account, it has no market interest rate risk. The longevity bond is strictly exposed to longevity risk. Under the traditional risk neutral measure $\mathbb{Q}$ we have that

$$L(t, T, x) = B_t E_{\mathbb{Q}} \Big[B_T^{-1} \Big(B_T S(T, x) \Big) \Big| \mathcal{F}_t \Big],$$

where $\mathcal{F}_t$ is the filtration representing the evolution of interest rates and mortality up to an including time t. It is clear that this result depends on the existence of an arbitrage-free market, which is clearly not the case. Therefore, as pointed out in Norberg (2010), the application of risk-neutral pricing could be called into question. It is not our intention to enter this debate, rather to demonstrate the results under these principles and further encourage the emergence of a longevity bond market.

Therefore, we assume there exists probability measure $\mathbb{Q}$ such that the discounted prices of longevity bonds are $\mathbb{Q}$-martingales. We do not assume a unique martingale measure. For model assessment purposes, see Section 3, we work with the historical measure. This is appropriate since we are most interested in volatility and correlation. We have

$$
\begin{aligned}
L(t, T, x) &= B_t \, E_{\mathbb{Q}}[S(T, x) | \mathcal{F}_t] \\
&= B_t \, S(t, x) \, E_{\mathbb{Q}} \Big[\exp \int_t^T -\mu(s, x+s) ds \Big| \mathcal{F}_t \Big] \\
&= B_t \, S(t, x) \, p(t, t, T, x),
\end{aligned}
$$

where, p is the mortality forward-rate defined, in general, as

$$p(t, T_1, T_2, x) = E_{\mathbb{Q}}\left[\exp\int_{T_1}^{T_2} -\mu(s, x+s)ds \Big| \mathcal{F}_t\right].$$

Therefore, $p(t, T_1, T_2, x)$ is the probability an individual age x at time zero survives to time T_2 given survival to time T_1, based on available information $\mathcal{F}_t$. For clarity, we describe the indices of p in greater detail:

- t is current time,

- T_1 is the *forward-time*, that is, the time to which survival is conditioned from time zero,

- T_2 the *maturity-time*, that is, the time to which survival is measured from time zero; a result is that $T_2 - T_1$ is the *time-horizon* of the forward mortality rate,

- and x is the age at time zero, which implies that $x + T_1$ is the *effective* age of the forward mortality rate, we refer to this age as the *forward-age*.

The martingale property implies

$$p(t, t, T, x) = E_{\mathbb{Q}}[p(t+1, t, T, x)|\mathcal{F}_t]. \tag{1}$$

The Olivier-Smith model is given below; see Olivier and Jeffery (2004), Smith (2005).

Model. *For all ages x and forward-times $T = t, t+1, \ldots$*

$$p(t+1, T, T+1, x) = p(t, T, T+1, x)^{b(t,T,T+1,x)G(t+1)},$$

where $G(1), G(2), \ldots$ are independent and identically distributed gamma random variables with shape and rate parameter α, such that $E_{\mathbb{Q}}[G(t)] = 1$ and $Var_{\mathbb{Q}}(G(t)) = 1/\alpha$. Furthermore, the $b(t, T, T+1, x)$ are $\mathcal{F}_t$-measurable bias correction functions given by

$$b(t, T, T+1, x) = -\frac{\alpha p(t, t, T, x)^{-1/\alpha}(p(t, T, T+1, x)^{-1/\alpha} - 1)}{\ln p(t, T, T+1, x)}.$$

For an arbitrage-free market, we require the martingale property, that is, equation (1) to be satisfied. Consequently, we obtain

$$p(t, t, T, x) = \frac{\alpha^\alpha}{(\alpha - \sum_{u=t}^{T-1} b(t, u, u+1, x) \ln p(t, u, u+1, x))^\alpha};$$

see e.g. Cairns (2007) for the details of this derivation, which rests on standard properties of the gamma distribution and the use of moment generating functions.

The model restricts the stochastic nature of mortality evolution by imposing gamma, $\Gamma(\alpha, \alpha)$, random variables. Henceforth, we generalize this stochastic component and make use of the notation $Z(t+1, T, T+1, x)$, where

$$Z(t+1, T, T+1, x) = \frac{\ln p(t+1, T, T+1, x)}{\ln p(t, T, T+1, x)}.$$

4

3 Model Assessment

The Olivier-Smith model makes three assumptions regarding the stochastic mortality component. The first is that of a gamma distribution. The second is that the distribution that generates the components does not depend on time t, forward-time T, or age at time zero x; that is, that mortality components are identically distributed. Lastly, that the components are mutually independent. We presently investigate the validity of these assumptions using real data. We assume $b = 1$ and treat population mortality rates as risk neutral.

In a risk neutral framework with constant annual forces of mortality and one-year forward rates, we have, for all t, all x, and $T = t - 1, \ldots, \omega - 1 - x$,

$$p(t, T, T + 1, x) = \exp\{-\mu(t, x + T)\},$$

where ω is the maximum attainable age. Notice that the above formulation of the forward rate does not include any imposed improvement trend. That is, the forward mortality probability for an individual aged x in T years is taken from an individual presently aged $x + T$. The reason for this is that any inclusion of trend must necessarily be a part of bias correction functions b. Since it is our aim to isolate the random component in this empirical analysis, considering an improvement trend would have no added benefit, it would merely augment b, leaving the stochastic mortality component Z unchanged.

We define for all t, all x, and $T = t, \ldots, \omega - 1 - x$,

$$Z(t + 1, T, T + 1, x) = \frac{\ln p(t + 1, T, T + 1, x)}{\ln p(t, T, T + 1, x)} = \frac{\mu(t + 1, x + T)}{\mu(t, x + T)}. \tag{2}$$

We use female central mortality rates from the Human Mortality Database (2013) for *England and Wales Total Population* for the period 1960-2009, and ages 0-105 ($\omega = 106$). We extract observations of $Z(t + 1, T, T + 1, x)$ from the data. England and Wales data was selected since it is a commonly used benchmark dataset.
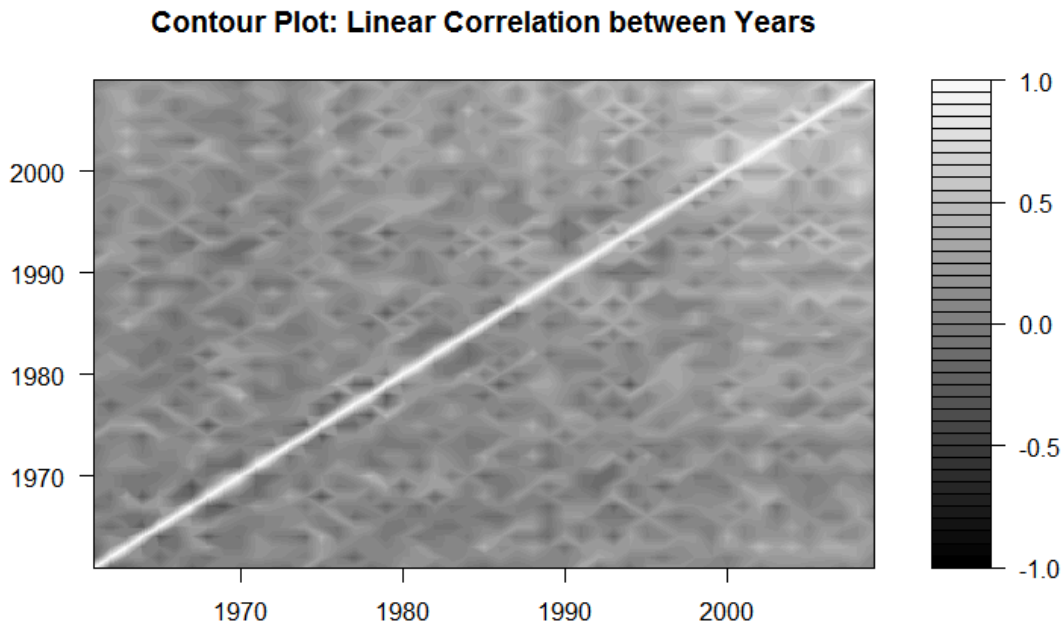
3.1 Independent Distributions

In this section, we investigate whether the stochastic components are independent. We address this in two ways: by considering the linear correlation structure between appropriately selected sub-samples, and applying a principal component analysis.

Correlation Analysis

We explore the correlation structure between the stochastic mortality components over time and forward-age. Figure 1 suggests no particular dependence pattern over time; roughly 90% of the correlations have absolute value less than 0.35. That is, the collection of observations of the stochastic component from calendar year s appear uncorrelated with the collection from calendar year $t \neq s$. This is evidence of independence over the time dimension.

Figure 2 provides contour plots of correlation between the stochastic components over forward-age. In contrast to the contour plots over time, we do notice varying levels of correlation over forward-age, especially amongst the older ages.

Figure 1: Correlations of $Z(s+1, s, s+1, x)$ and $Z(t+1, t, t+1, x)$

Contour Plot: Linear Correlation between Years

The correlation appears very volatile for the younger forward-ages. This is a technical consequence of having comparatively smaller sample sizes for forward-ages 0-48. Furthermore, correlation patterns are difficult to discern when centenarians are included. This is related to the high volatility of mortality data for centenarians. Consequently, we provide the contour plot for ages 49-99 in the second panel of Figure 2. Lighter shades of gray indicate higher positive correlation; the plot provides clear evidence that correlation increases with forward-age.

Principal Component Analysis

In the previous section we found evidence that the stochastic mortality components are correlated by forward-age, especially pronounced for the older ages. In order to further investigate this form of dependence we perform a principal component analysis (PCA) for the random components taken with respect to different forward-ages. We focus on forward-ages 49-99; represented in the second panel of Figure 2. Therefore, with 49 years of observations, we obtain a 51 by 49 matrix of stochastic components by forward-age and time.

PCA is applied to multivariate data in order to take advantage of possible redundancy due to dependence. By using linear combinations of the original variables, we transform the 51 correlated variables (stochastic components by forward-ages) into uncorrelated ones of a lower dimension. In general, PCA may either be applied to the covariance, or correlation matrix. PCA applied to the correlation matrix is typically used when variables are measured in different units. Since this is not the case in our data, we elect to apply PCA to the covariance matrix. In addition, standardising produces data with the same variation. This is not desirable since we want to preserve differences in variation by forward-age.

Figure 2: Correlations of $Z(t+1, t, t+1, x-t)$ and $Z(t+1, t, t+1, y-t)$

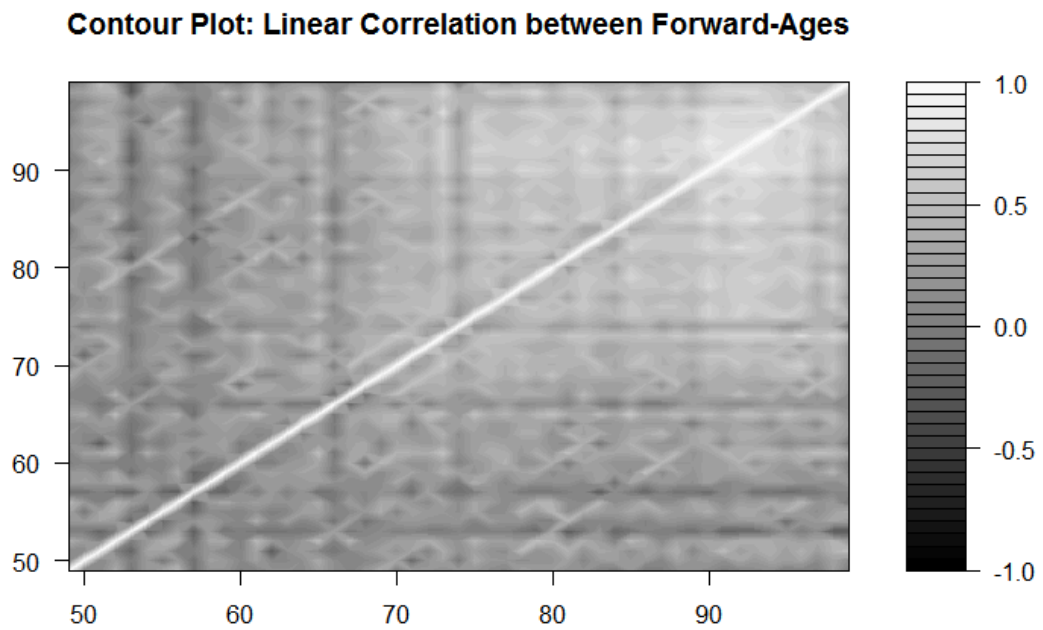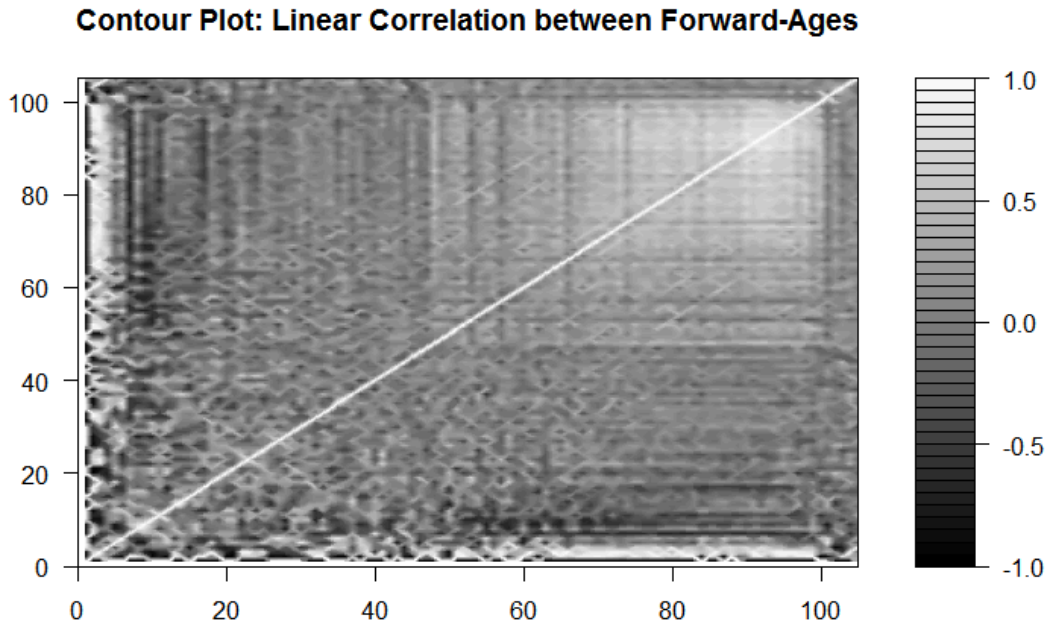**Contour Plot: Linear Correlation between Forward-Ages**

Contour Plot: Linear Correlation between Forward-Ages

Table 1: PCA Results

Eigenvalue	Individual Variance	Cumulative Variance
λ_1	0.2818	0.2818
λ_2	0.2096	0.4915
λ_3	0.1416	0.6331
λ_4	0.1028	0.7359
λ_5	0.0434	0.7793
λ_6	0.0312	0.8104
λ_7	0.0214	0.8319
λ_8	0.0167	0.8485
λ_9	0.0156	0.8641
λ_{10}	0.0138	0.8779
λ_{11}	0.0122	0.8901
λ_{12}	0.0112	0.9013
λ_{13}	0.0103	0.9116
λ_{14}	0.0096	0.9213
λ_{15}	0.0074	0.9287
λ_{15}	0.0068	0.9355

Table 1 summarises the results of the PCA applied to the covariance matrix. Note that as a criteria for deciding on the number of components, one typically chooses the number of components that cumulatively account for a meaningful percentage of the variance (typically around 70%). We select the first 4 components since they account for approximately 74% of the variance; see Table 1.

Figure 3 shows the top 4 principal components. They represent key movement in the random component by forward-age. We observe that some factors affect the stochastic components irrespective of forward-age, while others affect stochastic components from younger and older forward-ages quite differently. For example, the first principal component, which explains approximately 28% of the variation, exhibits a clear upward trend. This indicates that stochastic components from older forward-ages are more influenced than those from lower forward-ages. This confirms the result from Figure 2 that correlation increases by forward-age. Although the third and fourth principal components exhibit a downward trend, they are of less significance since they collectively account for less than 25% of the variance.

Finally, the biplots in Figure 4 show the magnitude and sign of each variable's contribution to the first two (upper panel) or three (lower panel) principal components, and how each observation is represented in terms of those components. The axes in the plot represent the principal components, the coefficients of the variables (forward-ages) on the principal components are represented as vectors, and the points are the scores of the observations (years) on the principal components.

Both the empirical investigation and the PCA provide evidence of correlation between the stochastic mortality components, especially among the older ages. The assumption of independent distributions for the stochastic mortality component is therefore violated.

Figure 3: Top 4 Principal Components

3.2 Identical Distributions

In this section, we investigate whether the stochastic mortality components come from identical distributions. Although we have identified potential dependence over forward-age, it is not conclusive, and therefore, we also investigate plots of the stochastic component over this dimension.

Figure 5 shows the observations of Z over time t for various fixed combinations of T and x. Over time t, the Z resemble a random sample. However, the underlying distribution changes with respect to the *forward-age*, $x + T$, of the mortality components. The scale of the first plot is roughly five times larger than the others. This implies that the first plot demonstrates very high volatility, which is not surprising given that this plot represents a forward-age of 105.

Figure 6 shows the observations of Z over age x for various fixed t. Recall that we have no observations of Z for $T < t$; for $T > t$, we obtain a similar plot as that of $T = t$, where the only difference is a shift in the x-axis, representative of a translation of the forward-age. We present the plots for $T = t$ as they are most informative. Figure 6 reinforces what is observed in Figure 5. The volatility of the stochastic component clearly varies with age. The forward-ages 0-40 and 90 plus generally exhibiting high volatility, with the remaining forward-ages, 40-90, exhibiting comparatively lower volatility.

In Figure 7 we fix the year-of-birth and plot over time, which is representative of tracking a single cohort. We notice that volatility varies with forward-age, which is visible to a greater or lesser extent depending on which forward-age-range the cohort traverses; notice the different scales of the plots in this figure. However, given the contour plots in Section 3.1, the samples in Figure 7 most likely violate the independence assumption. Therefore, caution must be exercised when interpret-
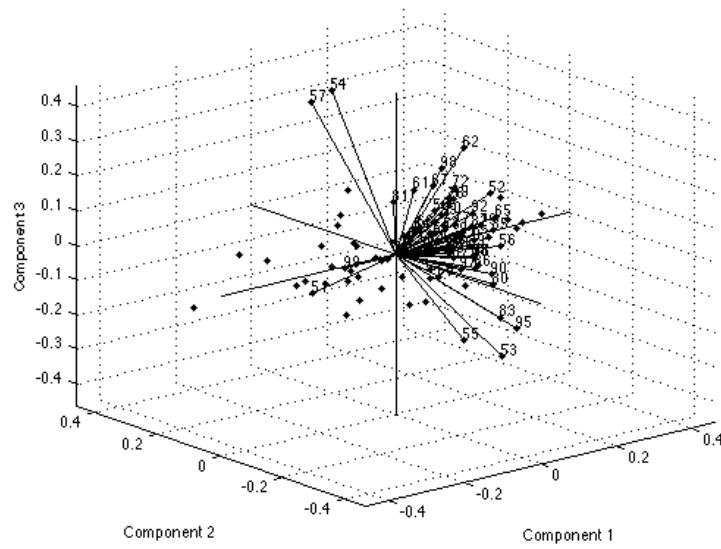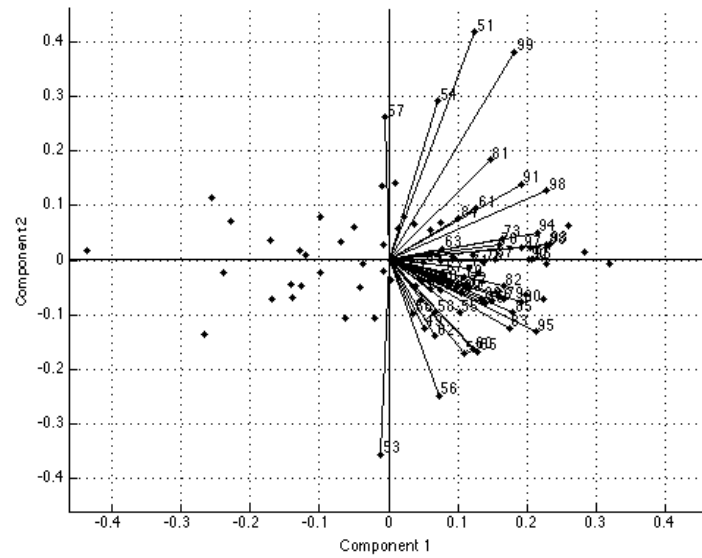
Figure 4: Biplots of the Principal Components

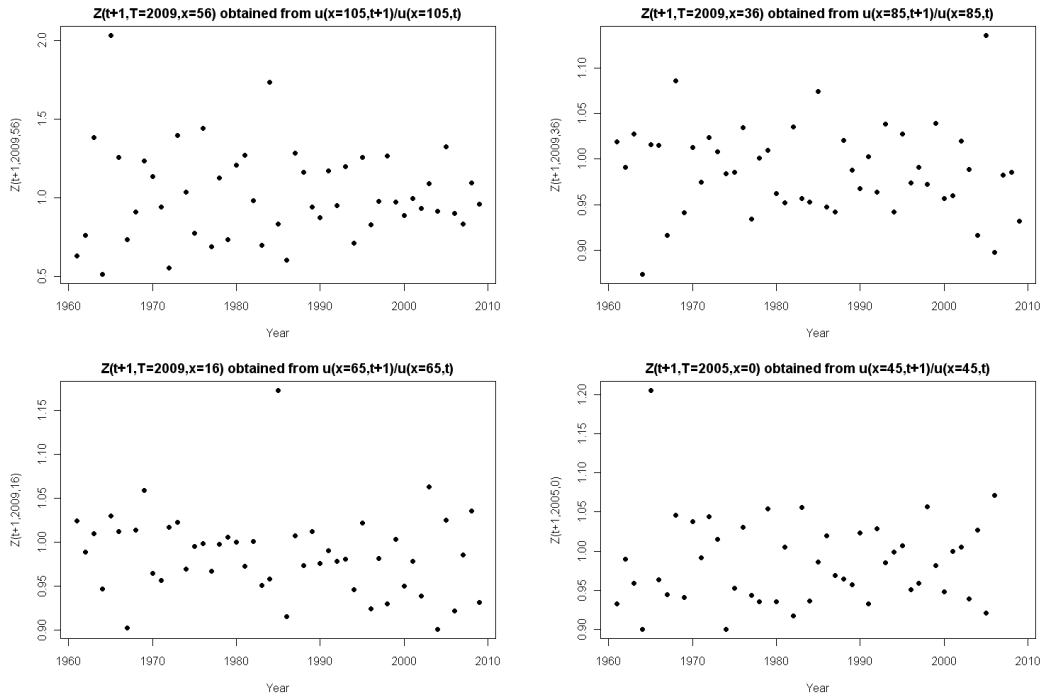Figure 5: The Stochastic Mortality Component over Time t for Fixed x

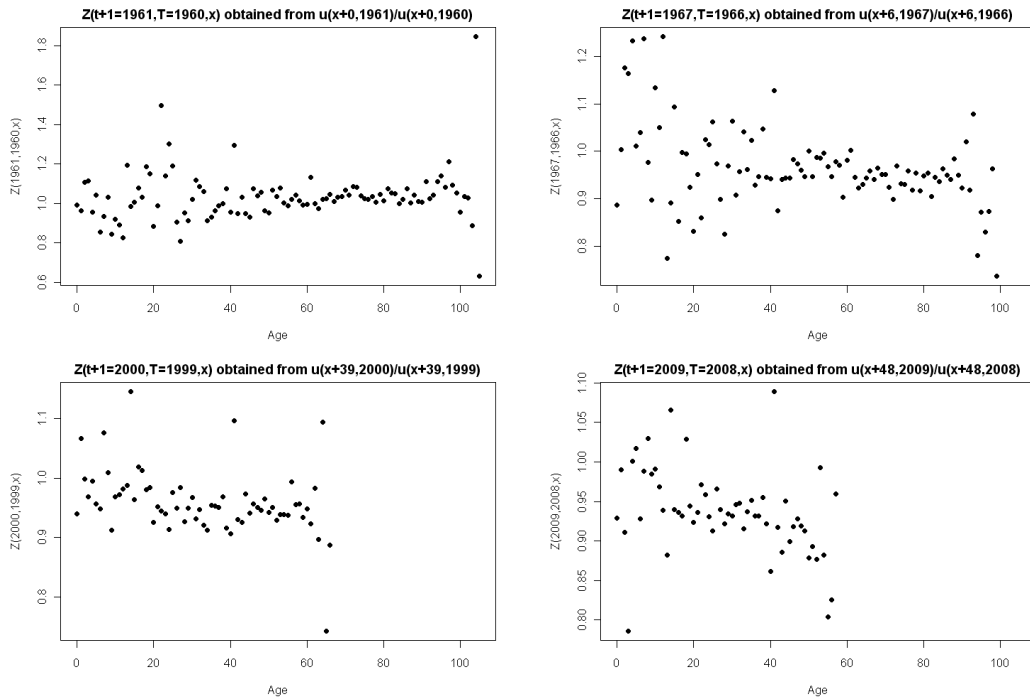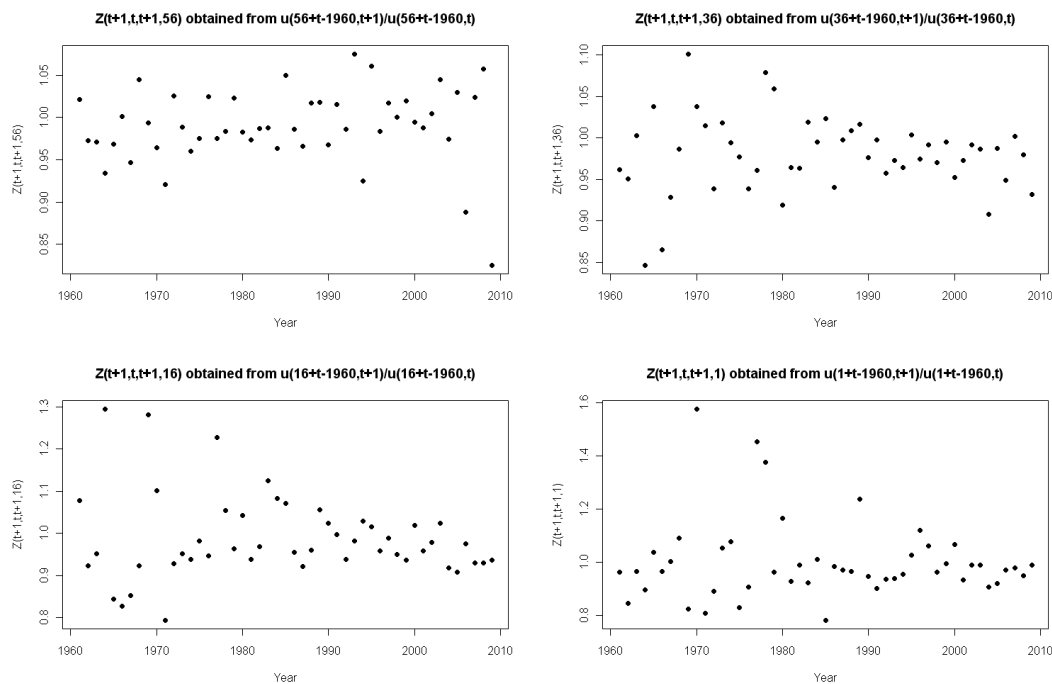Figure 6: The Stochastic Mortality Component over Age x for Fixed t

Figure 7: The Stochastic Mortality Component for Fixed Year-of-Birth

ing these plots. Figures 6-7 provide evidence that the volatility of the stochastic component are not homogeneous over forward-ages. The assumption of identical distributions for the stochastic mortality component is therefore violated.

3.3   Appropriate Parametric Distribution

From the data shown in Figures 5, 6, and 7, it is difficult to conclude whether the gamma distribution provides a good fit to the data. Fitting the data presumes independent and identically distributed observations, both of which may be violated.

However, we present quantile-to-quantile (QQ) plots and formal distributional (goodness-of-fit) tests to investigate the suitability of the gamma distribution. Figures 8, 9, and 10 show the QQ plots corresponding to Figures 5, 6, and 7, respectively. They contrast the quantiles of the empirical distribution with those from the estimated gamma distribution. A deviation from the 45 degree line indicates a departure from the assumed distribution. The results indicate the gamma distribution provides a reasonable fit for the stochastic mortality component over time t; see Figure 8. 9 and 10 are less relevant, since they represent fitting the distribution over age and year-of-birth, respectively, where we have evidence that the independent and identical distribution assumption is violated. This means the lack in fit is not very informative.

We apply the goodness-of-fit tests suggested in D'Agostino and Stephens (1986) and Stephens (1974). We test the null hypothesis that the data belongs to the theoretical (hypothesized) gamma distribution, that is, $H_0 : \hat{F}_n(x) = F_n(x)$, where $\hat{F}_n(x)$ represents the empirical cumulative distribution function (cdf) and $F_n(x)$, the theoretical cdf with estimated parameters obtained from maximum likelihood. Let the x_i be ordered observations, and $z_i = F_n(x_i)$, the resulting estimated quantiles. We consider the following test statistics:

Figure 8: The Stochastic Mortality Component over Time t for Fixed x;
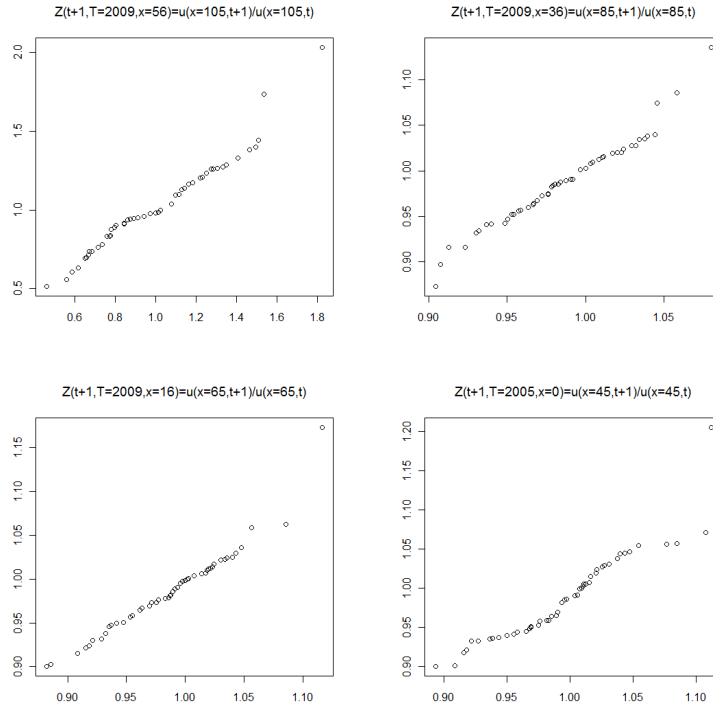QQ plot corresponding to Figure 5

Figure 9: The Stochastic Mortality Component over Age x for Fixed t;
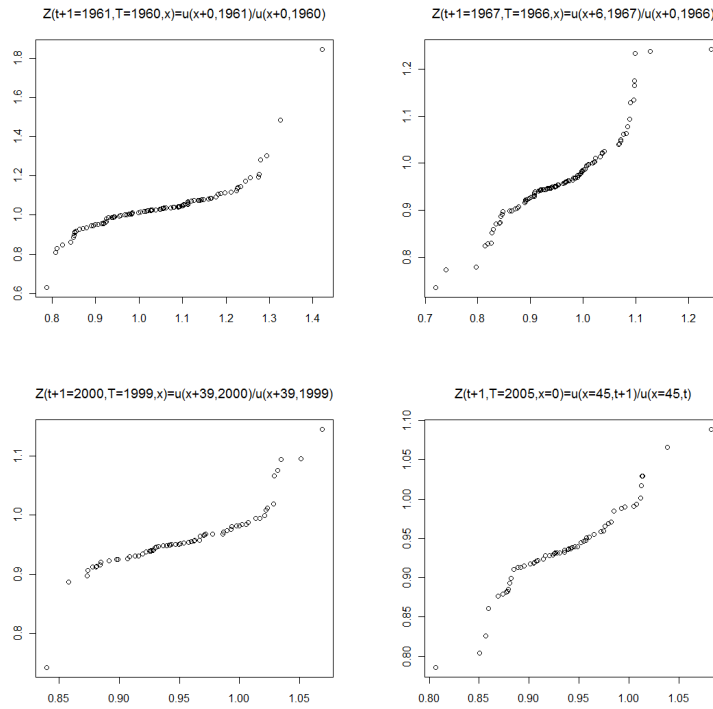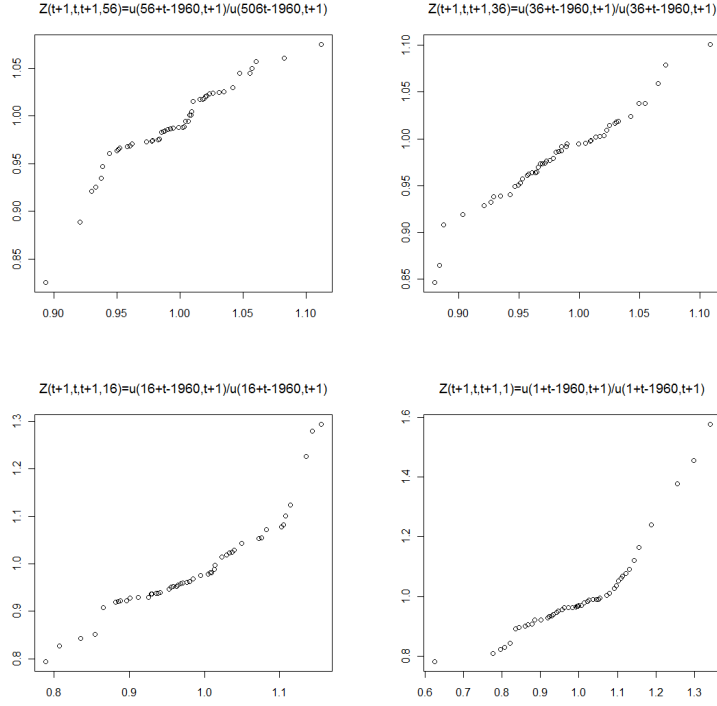QQ plot corresponding to Figure 6

Figure 10: The Stochastic Mortality Component for Fixed Year-of-Birth;
QQ plot corresponding to Figure 7

1. The Anderson-Darling statistic A^2, given by

$$A^2 = -n - 1/n \sum_{i=1}^{n} (2i-1)(\ln(z_i) + \ln(1 - z_{n+1-i})).$$

2. The Kolmogorov statistic D, given by $D = \max(D^+, D^-)$, where

$$
\begin{aligned}
D^+ &= \max_{1 \leq i \leq n} [(i/n) - z_i], \\
D^- &= \max_{1 \leq i \leq n} [z_i - (i-1)/n].
\end{aligned}
$$

The modified form statistic proposed in Stephens (1974) together with the critical values corresponds to $D(\sqrt{n} + 0.12 + 0.11/\sqrt{n})$.

3. The Cramér-von Mises statistic W^2, given by

$$W^2 = 1/12n + \sum_{i=1}^{n} (z_i - (2i-1)/2n)^2.$$

The modified form statistic reported in Stephens (1974) is given by $(W^2 - 0.4/n + 0.6/n^2)(1 + 1/n)$.

The critical values for the above tests are specified in Table 2, as provided in D'Agostino and Stephens (1986) and Stephens (1974).

Table 2: Critical Values for the Goodness-of-Fit Tests

Stat.	Crit.5%	Crit.1%
A^2	2.492	3.857
D	1.358	1.628
W^2	0.461	0.743

Table 3: Distributional Test Results

	α	β	A^2	D	W^2
Panel A: The Stochastic Mortality Component over Time t for Fixed x					
Z(t+1,T=2009,x=56)	12.943646	12.624253	0.2181661**	0.5312711**	0.1075524**
Z(t+1,T=2009,x=36)	415.72498	421.08874	0.2475575**	0.5700078**	0.1057159**
Z(t+1,T=2009,x=16)	452.03843	458.56811	0.4509184**	0.6866471**	0.1270422**
Z(t+1,T=2005,x=0)	331.04846	335.43804	0.6111859**	0.7313687**	0.1552846**
Panel B: The Stochastic Mortality Component over Age x for Fixed t					
Z(t+1=1961,T=1960,x)	71.337103	69.005897	3.843054*	1.571849*	0.7661192
Z(t+1=1967,T=1966,x)	127.56511	132.22064	2.624833*	1.256487**	0.5699055**
Z(t+1=2000,T=1999,x)	312.29324	325.31700	2.905791*	1.167076**	0.5625603*
Z(t+1=2009,T=2008,x)	297.92308	317.92973	1.075193**	1.010486**	0.2789330**
Panel C: The Stochastic Mortality Component for Fixed Year-of-Birth					
Z(t+1,t,t+1,56)	497.3803	502.4013	0.8524845**	0.9784892**	0.2081623**
Z(t+1,t,t+1,36)	444.92502	454.13746	0.5299603**	0.5954966**	0.1529977**
Z(t+1,t,t+1,16)	107.80345	109.21021	1.463862**	1.020509**	0.332074**
Z(t+1,t,t+1,1)	53.05649	53.04923	2.440254*	1.390677*	0.520064*

Table 3 summarizes the results for the goodness-of-fit tests; * and ** indicate the gamma assumption is not rejected at a 5% and 1% significance level, respectively, where the gamma density is given by

$$f(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}.$$

The distributional tests confirm the results observed from the QQ plots. The gamma distribution is not rejected at either significance level for the observations of stochastic mortality component over time (Panel A). The results for stochastic mortality component over age for fixed calendar year (Panel B) and over time for fixed year-of-birth (Panel C) are mixed. Again, the mixed results in Panel B and C are likely due to the violation of the independent and identical distribution assumption of these samples. Finally, note from Table 3 that the estimates of the shape and rate parameters, α and β, are approximately equal, which confirms the assumption that the expected value of the stochastic component is one.

4   Univariate Tweedie Generalization

We briefly introduce the Tweedie distribution, first formulated in Tweedie (1984); also see McCullagh and Nelder (1989). For applications to actuarial science, see e.g. Aalen (1992), Jørgensen and De Souza (1994), Smyth and Jørgensen (2002), Furman and Landsman (2010).

Let $X \sim ED(\theta, \lambda)$ denote a random variable belonging to the additive exponential dispersion family (EDF) with probability measure $P_{\theta,\lambda}$, absolutely continuous with respect to Q_λ,

$$dP_{\theta,\lambda}(x) = e^{[\theta x - \lambda \kappa(\theta)]} dQ_\lambda(x),$$

where $\kappa(\theta)$ is called the cumulant; see Jørgensen (1997), Section 3.1. The parameters θ and λ are called the canonical and dispersion parameter, respectively, $\kappa(\theta)$ is a twice differentiable function, and the expectation of X is given by

$$\mu = \lambda\kappa'(\theta).$$

Furthermore, $\kappa'(\theta)$ is a one-to-one map and there exists an inverse function

$$\theta = \theta(\mu) = (\kappa')^{-1}(\mu).$$

The function $V(\mu) = \kappa''(\theta(\mu))$ is called the unit variance function. The Tweedie subclass is the class of EDF with power unit variance function

$$V(\mu) = \mu^p,$$

where p is called the power parameter. The values $p = 0, 1, 2, 3$ correspond to the normal, over-dispersed Poisson, gamma, and inverse Gaussian distributions, respectively. The cumulant $\kappa_p(\theta) = \kappa(\theta)$ for the Tweedie subclass has the form

$$\kappa(\theta) = \begin{cases} e^\theta, & p = 1, \\ -\ln(-\theta), & p = 2, \\ \frac{\alpha-1}{\alpha}(\frac{\theta}{\alpha-1})^\alpha, & p \neq 1, 2, \end{cases}$$

where $\alpha = (p-2)/(p-1)$. Let $X \sim Tw_p(\theta, \lambda)$ denote a random variable belonging to the additive Tweedie family.

Model. *For all ages x and forward-times $T = t, t+1, \ldots$*

$$p(t+1, T, T+1, x) = p(t, T, T+1, x)^{b(t,T,T+1,x)Z(t+1)},$$

where $Z(1), Z(2), \ldots$ are independent and identically distributed Tweedie random variables, $Tw_p(\theta, \lambda)$ with $E_\mathbb{Q}[Z] = \lambda\kappa'_p(\theta)$ and $Var_\mathbb{Q}(Z) = \lambda\kappa''_p(\theta)$. Furthermore, the $b(t, T, T+1, x)$ are $\mathcal{F}_t$-measurable bias correction functions given by

$$
\begin{aligned}
&b(t, T, T+1, x) \\
&= \frac{\kappa_p^{-1}(\ln p(t, t, T+1, x)/\lambda + \kappa_p(\theta)) - \kappa_p^{-1}(\ln p(t, t, T, x)/\lambda + \kappa_p(\theta))}{\ln p(t, T, T+1, x)}.
\end{aligned}
$$

For an arbitrage-free market, we require the martingale property, that is, equation (1) to be satisfied. Consequently, we obtain

$$
\begin{aligned}
p(t, t, T, x) &= E_\mathbb{Q}[p(t+1, t, T, x)|\mathcal{F}_t] \\
&= E_\mathbb{Q}\left[\prod_{u=t}^{T-1} p(t+1, u, u+1, x)\Big|\mathcal{F}_t\right] \\
&= E_\mathbb{Q}\left[\prod_{u=t}^{T-1} p(t, u, u+1, x)^{b(t,u,u+1,x)Z(t+1)}\Big|\mathcal{F}_t\right] \\
&= E_\mathbb{Q}\left[\exp\left\{Z(t+1)\sum_{u=t}^{T-1} b(t, u, u+1, x)\ln p(t, u, u+1, x)\right\}\Big|\mathcal{F}_t\right] \\
&= M_{Z|\mathcal{F}_t}\left(\sum_{u=t}^{T-1} b(t, u, u+1, x)\ln p(t, u, u+1, x)\right),
\end{aligned}
$$

16

where

$$M_{Z|\mathcal{F}_t}(y) = E_{\mathbb{Q}}[e^{Zy}|\mathcal{F}_t] = \exp\{\lambda(\kappa_p(\theta + y) - \kappa_p(\theta))\},$$

is the moment generating function of $Z \sim Tw_p(\theta, \lambda)$. This reduces to the Olivier-Smith model if we select $p = 2$, $\lambda = \alpha$, and $\theta = -\alpha$.

The b functions are derived following a recursive procedure as shown in Cairns (2007). First, we investigate the case with a maturity of $t + 1$. From the above, we have

$$p(t, t, t + 1, x) = M_{Z|\mathcal{F}_t}(b(t, t, t + 1, x) \ln p(t, t, t + 1, x)),$$

which yields

$$b(t, t, t + 1, x) = \frac{\kappa_p^{-1}(\ln p(t, t, t + 1, x)/\lambda + \kappa_p(\theta)) - \theta}{\ln p(t, t, t + 1, x)}.$$

From this it is clear that

$$\sum_{u=t}^{T} b(t, u, u + 1, x) \ln p(t, u, u + 1, x) = \kappa_p^{-1}(\ln p(t, t, T + 1, x)/\lambda + \kappa_p(\theta)) - \theta$$

$$\sum_{u=t}^{T-1} b(t, u, u + 1, x) \ln p(t, u, u + 1, x) = \kappa_p^{-1}(\ln p(t, t, T, x)/\lambda + \kappa_p(\theta)) - \theta.$$

Substracting the two equations from one another leaves

$$b(t, T, T + 1, x)$$
$$= \frac{\kappa_p^{-1}(\ln p(t, t, T + 1, x)/\lambda + \kappa_p(\theta)) - \kappa_p^{-1}(\ln p(t, t, T, x)/\lambda + \kappa_p(\theta))}{\ln p(t, T, T + 1, x)}.$$

Note that $\ln p(t, t, t, x) = 0$, consequently, the expression above holds for the case with maturity $t + 1$. Therefore, this is the general expression and we no longer require the recursive argument.

The variance of the forward rates is given by

$$Var_{\mathbb{Q}}(p(t + 1, T, T + 1, x)|\mathcal{F}_t)$$
$$= M_{Z|\mathcal{F}_t}(2b(t, T, T + 1, x) \ln p(t, T, T + 1, x))$$
$$- M_{Z|\mathcal{F}_t}(b(t, T, T + 1, x) \ln p(t, T, T + 1, x))^2.$$

5 Multivariate Generalization

We begin by formulating the most general model.

Model. *For all ages x and forward-times $T = t, t + 1, \dots$*

$$p(t + 1, T, T + 1, x) = p(t, T, T + 1, x)^{b(t,T,T+1,x)Z(t+1,T,T+1,x)},$$

where the Z follow a multivariate distribution. The $b(t, T, T + 1, x)$ are some $\mathcal{F}_t$-measurable bias correction functions.

To preserve the martingale property, we have

$$p(t, t, T, x) = E_{\mathbb{Q}}[p(t + 1, t, T, x)|\mathcal{F}_t]$$
$$= E_{\mathbb{Q}}\left[\exp\left\{\sum_{u=t}^{T-1} Z(t + 1, u, u + 1, x)b(t, u, u + 1, x) \ln p(t, u, u + 1, x)\right\}\bigg|\mathcal{F}_t\right].$$

To proceed further, a multivariate distribution must be specified for the Z. We do so using copulas. Copulas provide the means to effectively study any form of dependence; see e.g. Embrechts *et al.* (2002), Breymann *et al.* (2003), and McNeil *et al.* (2005). We focus on the use of the elliptical family. Previous work modelling dependence using the elliptical family include Hult and Lindskog (2002), Fang *et al.* (2002), and Frahm *et al.* (2003). Figure 2 suggests that the correlation between any pair of stochastic components depends on the smallest of the two associated forward-ages. This produces the "L" shapes intersecting the diagonal of the plot; see Figure 2 above and Figures 12 and 13 below. To accommodate this form of dependence, we introduce the *minimum covariance pattern*.

Minimum Covariance Pattern

The minimum covariance pattern makes the assumption that the correlation between any pair of responses is determined solely by the minimum value of the corresponding covariate; in this case, forward-age. That is, $\text{Corr}(Z_{i,j}, Z_{i,k}) = \rho_{min(j,k)}$. Note that this results in $n-1$ correlation parameters.

$$\text{Corr}(Z_i) = \begin{pmatrix} 1 & \rho_1 & \rho_1 & \cdots & \rho_1 & \rho_1 \\ \rho_1 & 1 & \rho_2 & \cdots & \rho_2 & \rho_2 \\ \rho_1 & \rho_2 & 1 & \cdots & \rho_3 & \rho_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \rho_1 & \rho_2 & \rho_3 & \cdots & 1 & \rho_{n-1} \\ \rho_1 & \rho_2 & \rho_3 & \cdots & \rho_{n-1} & 1 \end{pmatrix}.$$

5.1   Capturing Dependence with Copulas

Copulas are multivariate distribution functions relating d one-dimensional standard uniform marginals to a joint distribution. If F is a d-dimensional distribution function with marginals $F_1 \ldots, F_d$, then, for every $x_1, \ldots, x_d \in \overline{\mathbb{R}}$, there exists a copula C with

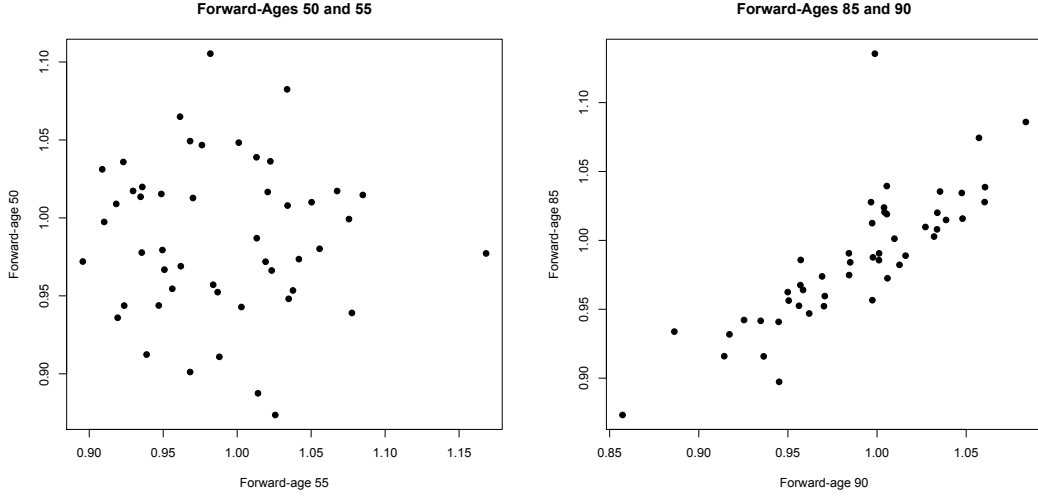$$F(x_1, \ldots, x_d) = C\{F_1(x_1), \ldots, F_d(x_d)\}.$$

Equivalently, if C is a copula and $F_1, \ldots, F_d$ are distribution functions, then the function F defined above is a joint distribution function with marginals $F_1, \ldots, F_d$. If $F_1, \ldots, F_d$ are continuous, then C is unique. Thus, if $X = (X_1, \ldots, X_d)^\top$ is a random vector with distribution $X \sim F_X$ and continuous marginals $X_j \sim F_{X_j}$ ($j = 1, \ldots d$), then the copula of X is the distribution function C_X of $u = (u_1, \ldots, u_d)^\top \in [0,1]^d$, where $u_j = F_{X_j}(x_j)$:

$$C_X(u_1, \ldots, u_d) = F_X\{F_{X_1}^{-1}(u_1), \ldots, F_{X_d}^{-1}(u_d)\}.$$

The above formulation is found in *Sklar's Theorem*; see Joe (1997) for a proof.

Two scatter-plots shown in Figure 11 highlight the different levels of dependence between observations from younger versus older forward-ages. We investigated the scatter-plots originating from observations of all possible pairs of forward-ages and found no evidence of non-standard types of dependence. Therefore, we consider the Gaussian copula. In order to allow for the presence of tail-dependence, we also consider the Student's t copula.

Figure 11: Scatter Plots for Observations by Forward-Age

With respect to the marginals, we consider both the gamma and the Gaussian distributions. The gamma distribution because it was originally considered in the Olivier-Smith model; we choose to compare it with the Gaussian distribution since it is the most widely used member of the Tweedie family. The Gaussian distribution was also considered in Olivier and Jeffery (2004). They claimed results similar to the gamma distribution, although no formal analysis was provided.

The literature suggests various approaches to estimate copulas; see e.g. Joe (1997) and Cherubini *et al.* (2004). We make use of the so-called inference for marginals (IFM) method, which is a sequential two-step maximum likelihood approach. First, one estimates the marginal parameters, subsequently, the copula to obtain the pseudo log-likelihood function, which is then maximized with respect to the copula dependence parameter. For details on the IFM and alternative estimation techniques see McLeish and Small (1988) and Joe (1997).

The minimum covariance pattern described above is not implemented in any statistical packages we are aware of. Therefore, we fit the Gaussian and Student's t copulas with an unstructured correlation matrix; given below.

$$\text{Corr}(Z_i) = \begin{pmatrix} 1 & \rho_{1,2} & \rho_{1,3} & \cdots & \rho_{1,n-1} & \rho_{1,n} \\ \rho_{1,2} & 1 & \rho_{2,3} & \cdots & \rho_{2,n-1} & \rho_{2,n} \\ \rho_{1,3} & \rho_{2,3} & 1 & \cdots & \rho_{3,n-1} & \rho_{3,n} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \rho_{1,n-1} & \rho_{2,n-1} & \rho_{3,n-1} & \cdots & 1 & \rho_{n-1,n} \\ \rho_{1,n} & \rho_{2,n} & \rho_{3,n} & \cdots & \rho_{n-1,n} & 1 \end{pmatrix}.$$

Using the output, $\widehat{\rho}_{j,k}$, we construct the parameter estimates from the minimum covariance pattern as follows:

$$\widehat{\rho}_j = \frac{1}{n-j} \sum_{k=j+1}^{n} \widehat{\rho}_{j,k}, \qquad j \in \{1, n-1\}.$$

The ability to specify the minimum covariance pattern would enhance estimation. Figure 12 shows the estimation of the correlation parameters for the Gaussian copula

19

with gamma marginal distributions in the unstructured and minimum covariance patterns. The similarities between the two contour plots confirms that the minimum covariance pattern is suitable in this case. Figure 13 shows the equivalent using the Student's t copula.
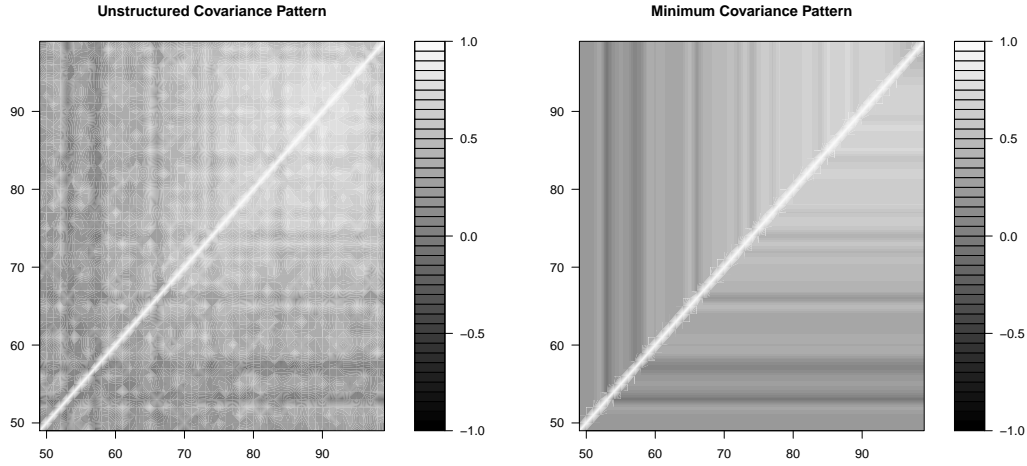
Figure 12: Contour Plots for the Gaussian Copula

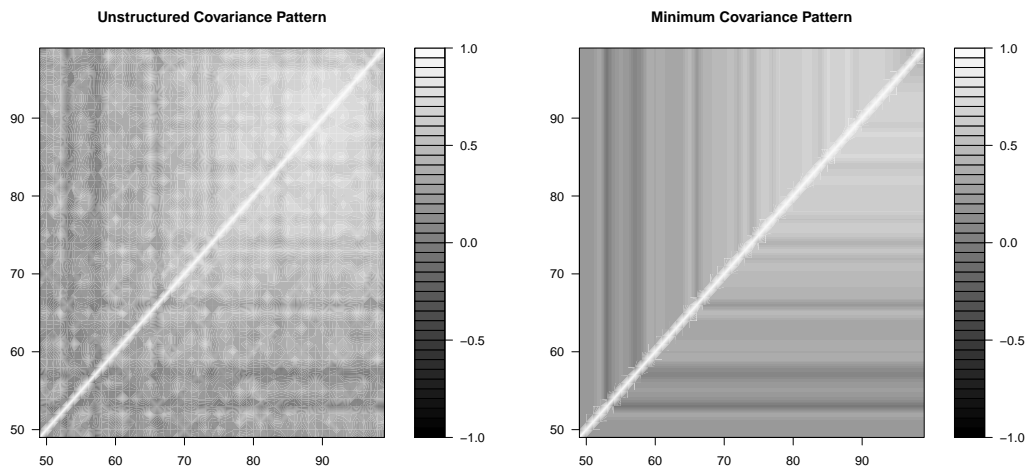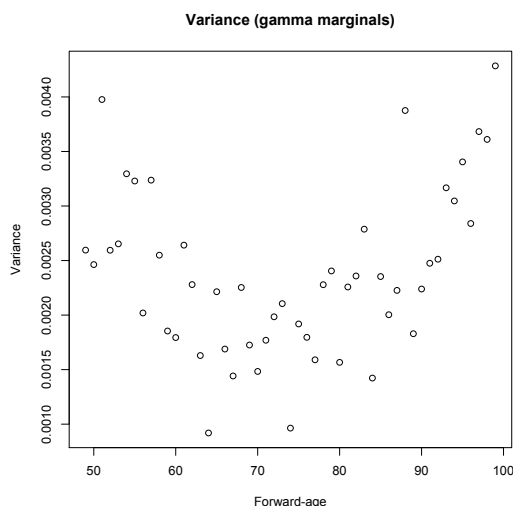Figure 13: Contour Plots for the Student's t Copula

Figure 14 shows the estimated variance of the gamma marginal distributions by forward-age. A very similar result is produced under the assumption of Gaussian marginals. The fluctuations in the variance clearly demonstrates that the marginal distributions vary by forward-age.

To produce the contour plots and the variance plot (Figures 12–14), we make use of the inversion of Kendall's tau method developed in Joe (1990). This was done because likelihood-based methods cannot readily estimate 51-dimensional copulas.

Finally, to examine the performance of the considered copula models, we use the *Akaike Information Criterion* (AIC) introduced in Akaike (1974):

$$AIC = -2l(\alpha; x_1, \ldots, x_T) + 2q,$$

Figure 14: Variance Estimation

where $l(\alpha; x_1, \ldots, x_T)$ denotes the maximized value of the log-likelihood and q the number of estimated parameters. Smaller values of AIC indicate a better fit.

We compute the AIC for groups of forward-ages, as summarized in the first column of Table 4. We split the sample into 5 sub-samples of 11 forward-ages in order to overcome the dimensionality issue arising when fitting a 51-dimensional copula using the likelihood-based IFM method. The Gaussian copula outperforms the Student's t copula for forward-ages up to 79, and under-performs for forward-ages above 79. This provides evidence of tail-dependence at the older forward-ages. With respect to the marginal distributions, the gamma consistently outperforms the Gaussian distribution. Since many models rely on Gaussian marginals, the result that the gamma outperforms the Gaussian distribution is rather significant. The fact that the modelling framework advocated in this paper allows for this added flexibility in the choice of marginal distributions is shown to be a very desirable attribute.

Table 4: Model Comparison (AIC)

Forward-Ages	Gaussian Copula		Student's t Copula	
	Gaussian	Gamma	Gaussian	Gamma
49-59	-5.95	**-7.29**	-3.93	-5.26
59-69	-23.45	**-27.46**	-21.44	-25.43
69-79	-145.26	**-154.93**	-143.23	-152.91
79-89	-223.68	-241.29	-246.88	**-266.26**
89-99	-403.88	-410.02	-405.27	**-410.61**

6 Conclusion

We investigate the Olivier-Smith model and show that, using population mortality data for England and Wales, the model requires a more general framework, with additional emphasis on age dependence and marginal distributions that vary by

forward-age. The gamma distribution provides a reasonable fit, but is a restrictive assumption. We improve the model by specifying a more general distribution, namely the Tweedie class of the exponential dispersion family. In addition to allowing the distribution to vary by forward-age, we also incorporate dependence between forward-ages using copulas. To accommodate the nature of the dependence, we specify a new covariance pattern related to the notion of *minimum*.

Acknowledgements

References

Aalen, O. O. (1992). Modelling heterogeneity in survival analysis by the compound Poisson distribution. *Annals of Applied Probability*, **2**(4), 951–972.

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transaction on Automatic Control*, **19**(6), 716–723.

Bauer, D. (2006). An arbitrage-free family of longevity bonds. University of Ulm.

Bauer, D. and Ruß, M. (2006). Pricing longevity bonds using implied survival probabilities. Meeting of the American Risk and Insurance Association (ARIA).

Biffis, E. (2005). Affine processes for dynamic mortality and actuarial valuations. *Insurance: Mathematics and Economics*, **37**(3), 443–468.

Blackburn, C. and Sherris, M. (2012). Consistent dynamic affine mortality models for longevity risk applications. UNSW Working Paper Series. Available at http://papers.ssrn.com/abstract=1832014.

Blake, D., Cairns, A. J. G., and Dowd, K. (2006). Living with mortality: Longevity bonds and other mortality-linked securities. *British Actuarial J.*, **12**(1), 153–197.

Breymann, W., Dias, A., and Embrechts, P. (2003). Dependence structures for multivariate high-frequency data in finance. *Quantitative Finance*, **3**(1), 1–14.

Cairns, A. J. G. (2007). A multifactor generalisation of the Olivier-Smith model for stochastic mortality. In Proceedings of the 1st IAA Life Colloquium, Stockholm, 2007.

Cairns, A. J. G., Blake, D., and Dowd, K. (2006). Pricing death: Frameworks for the valuation and securitization of mortality risk. *ASTIN Bulletin*, **36**(1), 79–120.

Cherubini, U., Luciano, E., and Vecchiato, W. (2004). *Copula Methods in Finance*. Wiley Finance Series.

CMI (2005). Projecting future mortality: Towards a proposal for a stochastic methodology. Working Paper 15 of the Continuous Mortality Investigation. The Faculty of Actuaries and Institute of Actuaries.

Cox, J. C., Ingersoll Jr, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica: Journal of the Econometric Society*, **53**(2), 385–407.

D'Agostino, R. B. and Stephens, M. A. (1986). *Goodness-of-fit techniques*. Statistics: Textbooks and Monographs, Dekker, New York.

Embrechts, P., McNeil, A. J., and Straumann, D. (2002). Correlation and dependence in risk management: properties and pitfalls. *Risk management: value at risk and beyond*, pages 176–223.

Fang, H. B., Fang, K. T., and Kotz, S. (2002). The meta-elliptical distributions with given marginals. *Journal of Multivariate Analysis*, **82**(1), 1–16.

Frahm, G., Junker, M., and Szimayer, A. (2003). Elliptical copulas: applicability and limitations. *Statistics and Probability Letters*, **63**(3), 275–286.

Furman, E. and Landsman, Z. (2010). Multivariate Tweedie distributions and some related capital-at-risk analysis. *Insurance: Mathematics and Economics*, **46**(2), 351–361.

Hult, H. and Lindskog, F. (2002). Multivariate extremes, aggregation and dependence in elliptical distributions. *Advances in Applied Probability*, **34**(3), 587–608.

Human Mortality Database (2013). University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org or www.humanmortality.de (data downloaded on March 8, 2013).

Jevtić, P., Luciano, E., and Vigna, E. (2013). Mortality surface by means of continuous time cohort models. *Insurance: Mathematics and Economics*, **53**(1), 122–133.

Joe, H. (1990). Multivariate concordance. *Journal of Multivariate Analysis*, **35**, 12–30.

Joe, H. (1997). *Multivariate Models and Dependence Concepts*. Chapman & Hall, London.

Jørgensen, B. (1997). *The Theory of Dispersion Models*. Chapman & Hall, London.

Jørgensen, B. and De Souza, M. C. P. (1994). Fitting Tweedie's compound Poisson model to insurance claims data. *Scandinavian Actuarial J.*, **1**, 69–93.

Liu, X. (2008). Stochastic mortality modelling. PhD Thesis available at https://tspace.library.utoronto.ca/.

Luciano, E. and Vigna, E. (2005). Non mean reverting affine processes for stochastic mortality. International Centre for Economic Research.

McCullagh, P. and Nelder, J. A. (1989). *Generalized Linear Models*. Chapman & Hall, London, 2nd edition.

McLeish, D. L. and Small, C. G. (1988). *The Theory and Applications of Statistical Inference Functions*. Springer-Verlag, New York.

McNeil, A. J., Frey, R., and Embrechts, P. (2005). *Quantitative Risk Management*. Princeton Series in Finance, Princeton.

Norberg, R. (2010). Forward mortality and other vital rates – are they the way forward? *Insurance: Mathematics and Economics*, **47**(2), 105–112.

Olivier, P. and Jeffery, T. (2004). Stochastic mortality models. Presentation to the Society of Actuaries of Ireland. See www.actuaries.ie/Resources/eventspapers/PastCalendarListing.htm.

Russo, V., Giacometti, R., Rachev, S., and Fabozzi, F. J. (2010). Calibrating affine stochastic mortality models using insurance contracts premiums. Technical report, University of Karlsruhe.

Smith, A. D. (2005). Stochastic mortality modelling. Talk at Workshop on the Interface between Quantitative Finance and Insurance, International Centre for the Mathematical Sciences, Edinburgh. See www.icms.org.uk/archive/meetings/2005/quantfinance/.

Smyth, G. K. and Jørgensen, B. (2002). Fitting Tweedie's compound Poisson model to insurance claims data: dispersion modelling. *ASTIN Bulletin*, **32**(1), 143–157.

Stephens, M. A. (1974). EDF statistics for goodness of fit and some comparisons. *Journal of the American Statistical Association*, **69**, 730–737.

Tweedie, M. C. K. (1984). An index which distinguishes some important exponential families. in *Statistics: Applications and New Directions*, J. K. Ghosh and J. Roy (eds), Indian Statistical Institute, Calcutta, 579–604.

Wills, S. and Sherris, M. (2010). Securitization, structuring and pricing of longevity risk. *Insurance: Mathematics and Economics*, **46**(1), 173–185.