

Modeling and Forecasting Mortality Rates

Daniel Mitchell*, Patrick Brockett, Rafael Mendoza-Arriaga, and Kumar Muthuraman

McCombs School of Business, University of Texas at Austin,
Austin, TX 78712, USA

November 30, 2011

Abstract

We show that by modeling the time series of mortality rate changes rather than mortality rate levels we can better model human mortality. Leveraging on this, we propose a model that expresses log mortality rate changes as an age group dependent linear transformation of a mortality index. The mortality index is modeled as a Normal Inverse Gaussian. We demonstrate, with an exhaustive set of experiments and data sets spanning 11 countries over 100 years, that the proposed model significantly out performs existing models. We further investigate the ability of multiple principal components, rather than just the first component, to capture differentiating features of different age groups and find that a two component NIG model for log mortality change best fits existing mortality rate data.

*Corresponding author, daniel.mitchell@phd.mcombs.utexas.edu

1 Introduction

Modeling and forecasting mortality rates has been an active area of research since Graunt (1662) first examined mortality in London for the purpose of trying to help create a warning system related to the onset and spread of the bubonic plague. Gaunt's work was the beginning of mortality table construction and showed that although individual life length was uncertain, there was a much more predictable pattern of longevity and mortality found in groups. From this it became clear that with an appropriately reliable mortality table life insurance rates could be determined and annuity values calculated in a rigorous manner. It fell on Edmond Halley (of Halley's comet fame) to actually construct a non-deficient mortality table using the birth and death records in the town of Breslaw, he also showed, in Halley (1693), how to perform a life annuity calculation based on this table. These were empirical tables and calculation of annuity values were very time consuming. The mathematical modeling of mortality and survival took a large step forward with the publication by de Moivre (1725) that postulated a uniform distribution of deaths model, and showed simplified annuity calculation methods. Taking a different and more biological approach to modeling and forecasting mortality, Gompertz (1825) assumed that the mortality rate, $m(t)$, represented the body's propensity to succumb to death at time t . He assumed that $w(t) = 1/m(t)$ was then the body's ability to withstand death. Assuming that the change in the body's ability to withstand death at time t is proportional to the ability it has to withstand death to begin with led him to the equation $w'(t) = c \cdot w(t)$, a differential equation whose solution is exponential. Solving this for $m(t)$, and then for the survival function that corresponds to this m yields the double exponential survival curve known as the Gompertz curve, which Brockett (1984) shows to fit mortality very well before the age of about 35, but which loses its fit in older ages. Makeham (1860) added some parameters to account for more rapid death at older ages, and provided formulae for calculating annuities from these.

Future death rates are extremely important to both national governments and insurance companies. National governments use forecasts of mortality rates to plan social security and health care programs. Gutterman, England, Parikh and Pokorski (2002) find that life expectancy in the developed world has been increasing by approximately 1.2 months every year, and globally, life expectancy at birth has increased by 4.5 months per year on average over the second half of the 20th century. It is estimated, in Couzin-Frankel (2011), that every additional year of life expectancy in the United States costs the

U.S. Social Security Administration \$50 billion. Insurance companies use these forecasts for pricing annuity contracts and to decide how much cash to keep on reserve in the case of life insurance pay outs. Similarly, corporations offering defined benefit pension plans must assure proper funding of future liabilities, however these future liabilities depend on the yet to be observed future mortality rates. A 2006 study, by Pension Capital Strategies and Jardine Lloyd Thompson (2006), in the UK found that recognizing the underestimation of expected lifetimes present in FTSE100 index companies would cause the aggregate deficit in pension reserves to more than double from £46 billion to £100 billion. In 2010 improved life expectancy added £5 billion to corporate pension obligations in the U.K. as seen in Reuters (2010). In the U.S. the level of pension contributions needed for adequate reserving will increase pension liabilities by 5-10%, as seen in Halonen (2007). If these institutions do not fully understand future trends in mortality rates, generations of people could be left without their pension plans due to lack of funds or insurance companies could go bankrupt from paying too many life insurance claims too quickly. It is in everyone's best interest for governments and insurance companies to have a better understanding of future mortality trends. In order to identify, elucidate and quantify these trends we must have a model that adequately captures the dynamics of mortality rates.

One of the first papers to model the separate effects of current age and year was Lee and Carter (1992). The authors in this paper propose a new bilinear model to better understand future trends in human mortality rates. This paper has been heavily cited and has been recommended for use by two U.S. Social Security Technical Advisory Panels, it is also used by the U.S. Bureau of the Census as a benchmark model, see for example Hollmann (2000).

1.1 Contributions and outline

In this paper we build upon the idea of bilinear modeling of age and time from Lee and Carter (1992), however we propose a model that looks at mortality data from a different perspective. We first identify and demonstrate one of the critical limitations of modeling mortality rates. We show that by first performing a simple transformation of the data the subsequent modeling vastly improves our ability to replicate the dynamics of mortality rates through time. We also propose a Normal Inverse Gaussian based mortality index model that is extremely easy to calibrate, has relatively few parameters and performs extremely well. We document this by comparing our model to several other models using several metrics found in

literature.

The rest of this paper is organized as follows. In Section 2, we describe the pitfalls of modeling the mortality rate data directly and how one can overcome this limitation. Section 3 formulates our model in more detail, describes how to fit the model and provides an overview of the data that we have used in this paper. In Section 4 we compare our model to four other models using in-sample statistical measures of goodness of fit. In Section 5 we discuss forecasting future mortality rates using our model and propose a few different ways to extrapolate into the future, and in Section 6 we conclude.

2 Modeling Mortality Levels

In Lee and Carter (1992) U.S. mortality rates are modeled for different age groups by extracting an unobserved state variable from the historical data on mortality rates. This state variable is interpreted as a single temporal mortality index applicable for each age group in entire population. Since each age group is allowed to respond to the mortality index in different ways, the mortality rate of each age group is some linear function of the mortality index. This yields a bilinear model involving the age group, x , and the year, t . The mortality index itself is modeled as a Brownian motion with a drift so future predictions of mortality rates can be made by extrapolating the index.

Making this specific, in Lee and Carter (1992) the authors model $m(x, t)$, the central mortality rate of age group x at time t , as

$$(2.1) \quad m(x, t) = e^{a_x + b_x \kappa_t + \epsilon_{x,t}}$$

a bilinear model for $\ln[m(x, t)]$. Here a_x describes the general shape of the mortality curve, κ_t is the mortality index mentioned above, and b_x describes each age group's response to the mortality index. If $b_x = 0$ or κ_t is constant, then one returns to the traditional stochastic mortality table construction.

To fit this model the authors use the singular value decomposition of the matrix of age specific log mortality rates through time to find the matrix of rank one that best approximates the actual log mortality rates. This is numerically equivalent to performing principal component analysis on the covariance matrix of log mortality rate levels. This can lead to a large overestimation of the dependence structure, however, since mortality rates have been trending downwards over at least the last 100 years for all age groups

and this confuses the first principal component with time.

The fact that mortality rates are trending downwards means the covariance matrix of mortality rates vastly overestimates dependence. For example, the covariance over 100 years between the log mortality rates of people aged 5 to 14 and people aged 65 to 74 is necessarily very high because early in the time series the mortality rate for both age groups was relatively high compared to their respective means, and later in the time series the mortality rates for both age groups was relatively low. This however does not necessarily mean that if we observe a better than average change in mortality for 65 - 74 year olds we should expect a better than average change for 5 - 14 year olds. A cure for cancer would certainly have a large impact on older people and a relatively lesser effect on children, due to the variability in the causes of death for different age groups, but when the cure is found both mortality rates may decline due to their longterm trends. As a result of this phenomenon the poor fit of the Lee Carter model can be explained by an exogenous variable.

Consider, for example, the consumer price index in Argentina, the miles driven in a particular car, the population of the earth and the GDP of China. These four variables have little relation to each other, but a principal component analysis would certainly show much “explained” by the first component because all these variables are highly correlated with calendar time.

As a simple experiment to illustrate this phenomenon we generated 11 *independent* Brownian motions each with a negative drift and sampled 100 points along each path. The drifts and volatilities of each Brownian motion were randomly chosen from uniform distributions. We then “demeaned” the data and performed a singular value decomposition on the data matrix. After repeating the experiment 1,000 times we find that according to this matrix decomposition the first singular value, on average, accounted for 99.2% of the variability in the data! Reminding ourselves that this implies that a model with one source of randomness explains 99.2% of the variability, it clearly does not align with the fact that the 11 Brownian motions were generated independently. The only thing these Brownian motions have in common is they all have negative drift, however, following the reasoning of Lee and Carter (1992) we might be led to infer that the first singular value is very informative and we can model all the data by simply modeling this first singular value.

Other papers since Lee and Carter (1992) have tried to improve upon the Lee Carter model by adding more principal components, or a cohort effect or any range of similar statistical quantities, but they all

suffer from the same fault: when we model the level, dependence between age groups is overestimated due to the fact that the mortality rate in 2011 is below the mean mortality rate over the last 100 years and the mortality rate in 1912 was above the mean.

Booth et al. (2002) modify the Lee Carter model by optimally choosing the time period over which to fit the model and adjust the state variable, κ_t , to fit the total number of deaths in each year. De Jong and Tickle (2006) reduce the number of parameters in Lee Carter to model mortality rates as a smoothed state space model. Yang et al. (2010) use multiple principal components to expand the Lee Carter model. Chen and Cox (2009) introduce jumps into modeling the state variable, found in Lee and Carter (1992), to increase goodness of fit measures and price insurance linked securities. Deng et al. (2012) use a more advanced jump diffusion model to fit the temporal state variable. A cohort effect, which incorporates the year of birth into the model, is added to the Lee Carter model in Renshaw and Haberman (2006). In Booth et al. (2006) the authors compare five variants of the Lee Carter model with data from several countries. Each of these models are fairly incremental, however all suffer from the same design flaw as Lee and Carter (1992), in that they all model the level of log mortality rates and hence misrepresent the dependence structure of mortality rates by age group.

This paper proposes a new model that avoids the common problem of modeling log mortality levels by first *de-trending* the data. Our model is similar to the Lee Carter model, however we model *changes* in log mortality rates rather than levels of log mortality rates themselves. By considering the *changes* we are able to more accurately capture the dependence structure between ages of mortality and use this to construct a more encompassing model.

Referring back to the independent Brownian motions experiment explained above, we performed a singular value decomposition on the matrix of differences through time of the independent Brownian motions and find that now on average the first singular value only accounts for 27.2% of the variability of the data. Theoretically this value should be about 9%, because each singular value should explain exactly the same amount of information and there are 11 singular values¹. This indicates that by de-trending the data we can learn much more about dependence in the data than if we model trending data. By de-trending the data in our simple experiment we learned that the data does not exhibit much dependence, a result which we would expect since the components were generated independently. In this

¹In our experiment, however, we only sampled the Brownian motions at 100 points in time, leading to large standard errors in estimations of covariance.

paper we show that mortality rates across age groups are indeed dependent upon each other, however by de-trending the actual data we can model this dependence better than if we consider the trending data. We show this by considering actual mortality data from 11 countries and by comparing our model to several models from literature.

We compare our model to several models existing in literature, using data from several countries, and show that in almost all cases our model has much less error than the others. Several of these competing models are multi-factor models while our model is a single factor model, indicating that our sparse model is more informative than many more complicated models. We will also show that including a second time factor to the model can increase the quality of the model, since some age groups are not adequately represented by a single factor². We also show that our model makes good predictions of the future distribution of mortality rates.

3 Model Formulation and Estimation

Following up on this insights described in Section 2, we construct a new model that is based on log mortality rate changes rather than levels. We model the mortality rate for each age group as

$$(3.1) \quad m(x, t) = m(x, t - 1)e^{\alpha_x + \beta_x k_t + \epsilon_{x,t}},$$

where α_x describes the average change in log mortality rate for each age group, k_t describes a mortality change index that is the same for all age groups, β_x tells us how each age group responds to the mortality index, and $\epsilon_{x,t}$ is the age and year specific error, which has expected value zero and which we assume to be uncorrelated across time and age group. This is analagous to how the log mortality level is a linear function of the mortality index in Lee and Carter (1992). Now however the trend in the data comes from previous years' mortality rate rather than a trend in the hidden mortality index. In this way we are able to condition on the mortality rates for the current year to make predictions about the future years' mortality. This leads to very high quality in-sample measures of fit and out-of-sample predictions.

²All other models considered have more error than our two factor model.

We can also work with the equivalent

$$(3.2) \quad \ln[m(x, t + 1)] - \ln[m(x, t)] = \alpha_x + \beta_x k_t + \epsilon_{x,t}.$$

This means that the change in log mortality rate for each age group, x , from year t to year $t + 1$ is a linear transformation of a time indexed variable k_t plus some error. The way we interpret this model is that in each year we observe a random vector of log mortality rate changes and the log mortality rate the next year for each age group is the log mortality rate from this year plus the random vector. In this way we escape the main drawback of the Lee Carter model because we have de-trended the data by considering changes rather than levels.

The values here of α_x , β_x and k_t are not unique in this representation of the model. For example, if we were to multiply every value of k_t by some constant we could then also divide every value of β_x by that same constant and the model would not change. We could also add a constant to every value of α_x and then adjust k_t in a similar way so that the model would not change. To address this identifiability problem we need to impose restrictions on the model to guarantee uniqueness. We force α_x to be the mean change in log mortality rates, which forces k_t to have expected value zero. We also force the sum of b_x across age groups to be equal to 1, as in Lee and Carter (1992).

3.1 Estimating the Model

In order to fit this model we employ methods similar to the Lee Carter (1992) method. First we construct the matrix, $M_{x,t}$, of log mortality rate changes. To do this take the matrix of log mortality rates from the second year to last year under consideration and then subtract from that matrix the matrix of log mortality rates from the first year to the penultimate year under consideration, call this matrix of differences M . In this way we see that $M_{x,t} = \ln[m(x, t + 1)] - \ln[m(x, t)]$.

Then from this matrix we must find the mean change over time for each age group, call that α_x , and create a demeaned matrix \hat{M} such that $\hat{M}_{x,t} = M_{x,t} - \alpha_x$. From this matrix we can finally apply the singular value decomposition to obtain β_x and k_t . The singular value decomposition of \hat{M} finds matrices U , S and V such that $\hat{M} = USV'$, where U and V are orthogonal matrices and S is a non-negative diagonal matrix, with the same dimension as \hat{M} , and $S_{1,1}$ is the largest element of S . This means k_t is

a scalar multiple times the first column of U times $S_{1,1}$ and β_x is the first column of V divided by that same scalar. Then the model prediction of $\ln[m(x, t + 1)]$ is equal to $\ln[m(x, t)] + \alpha_x + \beta_x k_t$.

As we mentioned before this is exactly equivalent to performing principal component analysis on the covariance matrix of \hat{M} . For example, if we were to find the eigen-decomposition of the covariance matrix of \hat{M} and then multiply the first eigenvector by \hat{M} we would find the first principal component. This would be exactly the same as k_t found using the singular value decomposition, up to a scalar multiple. For details on singular value decomposition and principal component analysis see, for example, Jolliffe (2002).

3.2 Data Description

In order to fit our model and compare it to other models we use historical data from two different sources. For U.S. mortality rates we use data from the HIST290 National Center for Health Statistics³, the same data used in Chen and Cox (2009) and Deng et al. (2012). The data lists death rates per 100,000 people and separates the population into 11 different age groups: (<1), (1-4), (5-14), (15-24), then every 10 years until (75-84) and finally (>85). We combine both sex and race categories to consider all people in the United States in each age group. We use the time frame from 1900 - 2004.

In this data we have the much discussed flu pandemic of 1918 which resulted in a large jump in mortality for this flu season. Several authors, such as Chen and Cox (2009), have discussed this one year because it seems like an outlier. Some authors, such as Lee and Carter (1992), suggest ignoring it because it is an outlier, while other authors, such as Deng et al. (2012), suggest leaving it in because it has valuable information about rare events. We choose to leave it in because we also believe outliers have valuable information. If one's goal is to forecast future mortality, leaving out arguably the most significant mortality jump in recent history could render the data set more smooth than it actually is.

For international mortality rates we use data from the Human Mortality Database. From this database we select the same countries considered in Booth et al. (2006). The data from the Human Mortality Database is richer than the HIST290 data, with data available for every age from infants up to 105. It is also available separated into five year age groups. We use the grouped data because it makes comparison between models simpler and some data from the Human Mortality Database is smoothed over age groups,

³Source: www.cdc.gov/nchs/datawh/statab/unpubd/mortabs.htm

which we wish to avoid. For a few countries data goes back before 1900, in one case as far back as the 1700's, but we only consider data starting in 1900. In this international data there are more outliers than in the U.S. data. This is because there have been more extreme events in Europe than in the U.S. over the time frames considered, such as the flu pandemic and World Wars I and II. This means that in order to create a unified model for all countries we must consider a model with heavy tailed distributions.

The Human Mortality Database also has mortality data for the U.S. but the data there only goes back to 1933, whereas the HIST290 data goes back to 1900. For this reason, we only use the HIST290 for the US.

4 Performance Evaluation

Now that we know how to fit the model to data, we would like to see how well our model performs relative to other models found in literature. We will compare our model to the others using data from 11 different countries. For fairness and direct comparison to existing literature, such as in Booth et al. (2006), we judge the quality of the model using two different in-sample measures of goodness of fit. In Cairns et al. (2011) the authors also use a qualitative comparison of model forecasts. We will also discuss this in Section 5. Most of the models we compare in this section are two factor models while our model is a single factor model. Even with the extra factor considered in these competing models, our model outperforms them in most countries using both measures of fit. We then see that if we include a second time component our model's performance is improved such that we have less error than *every* other model for *all* 11 countries.

4.1 Other Models

The models we consider in this section are the Lee Carter (LC) model, the model of Renshaw and Haberman (2006), which includes a cohort effect (RH), and two sub-models of this⁴, denoted H₁ and M3. We choose these models because in Cairns et al. (2011) and Haberman and Renshaw (2011) the author find that these models perform well with both in-sample measures and qualitatively out of sample. The

⁴Labeled as in Cairns et al. (2011) and Haberman and Renshaw (2011)

RH model states that

$$(4.1) \quad \ln[m(x, t)] = a_x + b_x^{(1)}k_t + b_x^{(2)}\gamma_{t-x} + \epsilon_{x,t},$$

where γ_{t-x} is the cohort effect. In the model H₁ we set $b_x^{(2)} = 1$ for all x and in the model M3 we set $b_x^{(1)} = 1$ and $b_x^{(2)} = 1$ for all x . We also observe that this is the LC model when $b_x^{(2)} = 0$.

Each of these models can be fit using numerical optimization, where we try to maximize the likelihood function of the data. In Renshaw and Haberman (2006) the authors describe how this can be done using Newton’s method. The authors present a table with formulae to update each variable iteratively in order to find the maximum for the RH model. In order to fit the other models we can use these same formulae and add the applicable constraints to the optimization routine by assigning specific values to the appropriate parameters.

Figure 1 shows the actual log U.S. mortality rate for two different age groups and the Lee Carter prediction of the log mortality rate and our model’s prediction of the log mortality rate. We see that our model predicts next year’s mortality rate much better than the Lee Carter model. We show age group 15-24 and >85 because they are among the best and worst levels of accuracy respectively for the Lee Carter model.

4.2 In-Sample Fit

With the data described above we can now compare our model to the others using in-sample measures of fit. We first consider in-sample tests to stay consistent with existing literature. The first measure of fit we use is the root sum of squared errors, or RSSE. This measures the square root of the sum of squared differences between the actual log mortality rate and the model prediction. We define the RSSE as

$$(4.2) \quad \text{RSSE} = \sqrt{\sum_{x,t} \epsilon_{x,t}^2}.$$

In Table 1 we show the results of all 5 models in each country. As each country has different amounts of data we include columns to describe the years considered and the number of age groups for each country, this can explain some of the difference between errors in different countries. In all tables, the model with

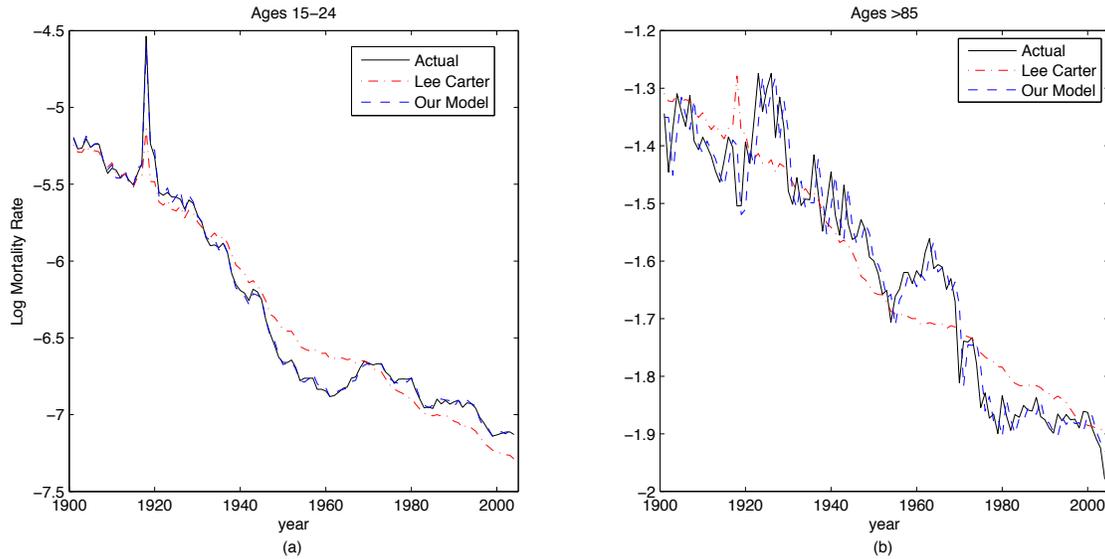


Figure 1: Comparison of log mortality rate predictions

the lowest error is highlighted.

We see that only in Denmark and Norway does the RH model have less error than our model. For all other countries our model has the least error of all models considered. The RH model is a two factor model, because it includes a factor for time and cohort, while our model is only a single factor model. As we will show in the next section, if we include a second time factor, such as in Yang et al. (2010), then the error of our model is greatly reduced and no competing model (including the RH model) outperforms our model in any country.

The next measure of fit we consider is the one considered in Lee and Carter (1992), which measures the unexplained variance (UV_x) of the each age group in the model. For each age group we find the variance of the time series of log mortality rates and the variance of the error between the model predicted log mortality rate and the actual log mortality rate. The ratio of these two variances is then called the unexplained variance, as in

$$(4.3) \quad UV_x = \frac{\text{Var}(\epsilon_{x,t})}{\text{Var}(\log[m(x,t)])},$$

where the variance is taken through time. Table 2 lists the unexplained variance for our model and the

Country	Years	# Ages	Ours	LC	RH	H ₁	M3
# Factors	-	-	1	1	2	2	2
Australia	1921 - 2007	22	2.20	5.98	3.25	3.96	7.39
Canada	1921 - 2007	22	2.25	6.06	2.49	3.00	8.38
Denmark	1900 - 2009	21	4.02	6.54	3.77	4.68	9.85
Finland	1900 - 2009	22	5.23	10.39	6.23	6.98	13.19
France	1900 - 2009	22	3.79	8.92	6.35	7.11	12.24
Italy	1900 - 2008	22	4.29	9.04	5.69	6.93	13.58
Norway	1900 - 2009	22	4.50	9.26	3.85	5.66	11.41
Sweden	1900 - 2010	21	3.41	8.66	3.73	4.69	10.23
Switzerland	1900 - 2009	21	3.47	5.94	3.54	4.10	7.90
U.K.	1922 - 2009	22	2.77	7.65	2.95	3.55	9.08
U.S.A.	1900 - 2004	11	1.10	3.99	2.04	2.54	6.39

Table 1: RSSE for several models and countries

other models in each age group, using only U.S. data. Using this measure it is not possible to get an overall measure of the model because there is no way to combine the variance of one age group with the variance of another age group.

Model	<1	1-4	5-14	15-24	25-34	35-44	45-54	55-64	65-74	75-84	>85
LC	0.025	0.0028	0.0013	0.041	0.040	0.010	0.039	0.075	0.062	0.086	0.12
Ours	0.0023	0.0013	0.0008	0.0009	0.0014	0.0012	0.0027	0.0063	0.0094	0.016	0.054
RH	0.0093	0.0015	0.0022	0.015	0.010	0.013	0.0089	0.018	0.016	0.018	0.083
H ₁	0.017	0.0029	0.0029	0.018	0.019	0.011	0.041	0.030	0.013	0.025	0.069
M3	0.0096	0.077	0.011	0.081	0.046	0.027	0.094	0.14	0.099	0.24	0.54

Table 2: Unexplained variance of models using U.S. data

Here all of the models do relatively poorly in explaining the variance of older age groups, with the 85 and older age group performing the worst. This is best explained by the fact that in the current state of the world people routinely live to be much older than 85 and thus grouping everyone older than 85 together fails to differentiate people that live to be quite old. Neither our model nor the Lee Carter model are limited to these age groups, however we selected this data for ease of comparison. We will later see that this is also due to the fact that a single temporal component may not adequately represent the older age groups, and we will discuss a way to correct this.

We can also examine how well the other models predict the changes by considering $(a_x + b_x^{(1)}k_t + \beta_x^{(2)}\gamma_{t-x}) - (a_x + b_x^{(1)}k_{t-1} + \beta_x^{(2)}\gamma_{t-1-x})$, the model predictions of the mortality rate in year t minus the

prediction in year $t - 1$. This is attained by calibrating the each model to the levels, as before, and then calculating the difference in model predictions between years. The RSSE of this compared to the actual change in mortality rates displayed in Table 3. For our model the RSSE is the same for levels and changes. We see that even though some models do not predict levels well, they do a relatively better job at predicting changes. In the next section we show that, as before, if we include a second time component to our model we have less error than every other model in every country.

Country	Ours	LC	RH	H ₁	M3
Australia	2.20	2.56	2.28	2.29	2.25
Canada	2.25	3.25	1.82	1.79	1.91
Denmark	4.02	4.29	3.94	3.91	4.16
Finland	5.23	6.46	5.72	5.97	6.05
France	3.79	5.67	4.67	5.33	5.73
Italy	4.29	4.85	4.00	4.37	4.51
Norway	4.50	4.85	3.80	4.25	4.55
Sweden	3.41	3.78	3.30	3.39	3.60
Switzerland	3.47	4.02	3.75	3.71	3.91
U.K.	2.77	4.81	1.96	1.98	1.97
U.S.A.	1.10	1.66	1.40	1.46	1.62

Table 3: RSSE of predicted changes for several models and countries

4.3 Two Factors

In the model described above we have restricted ourselves to only include one time component for predicting mortality rates for all age groups. The response of each age group to this time component is completely determined by β_x . Figure 2 shows the value for β_x , labeled $\beta^{(1)}$ in the figure, for each age group in the U.S. data and we can see that for the older age groups β_x gets very close to zero. This tells us that for older age groups, the age groups most important for pension plans, the k_t index has little effect on the change in mortality rate. This suggests that it may be prudent to add a second time component, as in Yang et al. (2010), in order to account for the declining importance of k_t . Also, when we consider the eigenvalues of the covariance matrix, we observe that the first singular value accounts for 80% of the variation in the U.S. data and the first two account for 90% of the variation in the data. If we were to include a third singular value we could explain 95% of the variation but we stop at two to

keep the model parsimonious. In the international data, the first singular value accounts for 50.5% of variation on average, the first and second account for 65.5% on average and the first three singular values account for 75.1% of variation on average. The first three singular values account for less variation in international data than in the U.S. data because there are more age groups in the international data than in the U.S. data. There are as many singular values as there are age groups so we cannot expect the fit to be as good for this many age groups. Despite this however, with only two factors we can outperform every other model.

In order to add a second time component we propose a second model where

$$(4.4) \quad \ln[m(x, t + 1)] - \ln[m(x, t)] = \alpha_x + \beta_x^{(1)} k_t^{(1)} + \beta_x^{(2)} k_t^{(2)} + \epsilon_{x,t}.$$

This model is calibrated exactly the same way that the first model, except that we find $\beta_x^{(2)}$ and $k_t^{(2)}$ from the second rows and columns of the singular value decomposition of the data matrix.

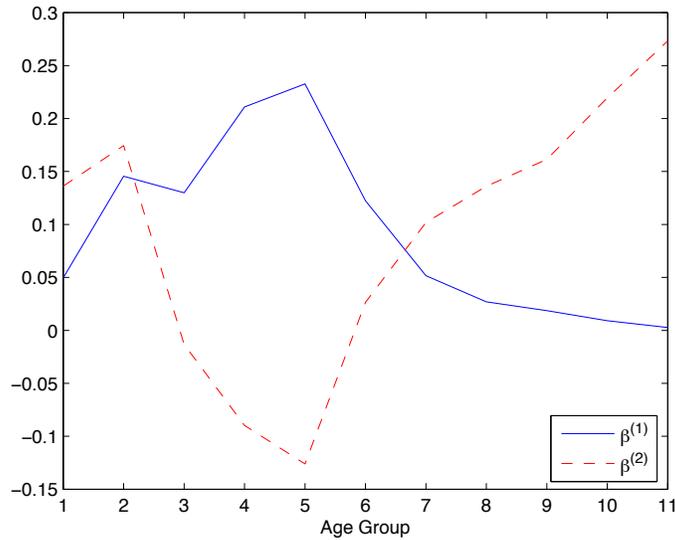


Figure 2: $\beta_x^{(1)}$ and $\beta_x^{(2)}$ for each age group in the US

In Figure 2 we plot the values of $\beta_x^{(1)}$ and $\beta_x^{(2)}$, using U.S. data, and see that for the older age groups the second time component is much more important than the first. Also, for the middle age groups

the second component is negative, indicating that perhaps the first component over states correlation between these ages and the others, because in years in which $k_t^{(2)}$ is positive the predicted log mortality change for the middle age groups decreases, where as it increases for the younger and older age groups, and vice versa for years in which $k_t^{(2)}$ is negative.

With the addition of this second time component we also improve the goodness of fit of the model. When we used only one time component there were some cases when the two factor RH model had less error than our one factor model. This, however, is not the case when we expand our model to include two time factors. In this case our model has less error than every other model considered in every country. In Table 4 we report the RSSE of the different models with respect to changes. This is essentially a copy of Table 3 except we add the column for the two factor model. In this table we see that no other model, in any country, has less error than our two factor model. We can also compare the second column of Table 4 to the columns in Table 1 because the error of changes in our models is the same as the error of levels in our models.

Country	1 Factor	2 Factor	LC	RH	H ₁	M3
Australia	2.20	1.88	2.56	2.28	2.29	2.25
Canada	2.25	1.63	3.25	1.82	1.79	1.91
Denmark	4.02	3.41	4.29	3.94	3.91	4.16
Finland	5.23	4.39	6.46	5.72	5.97	6.05
France	3.79	3.05	5.67	4.67	5.33	5.73
Italy	4.29	3.38	4.85	4.00	4.37	4.51
Norway	4.50	3.76	4.85	3.80	4.25	4.55
Sweden	3.41	2.88	3.78	3.30	3.39	3.60
Switzerland	3.47	3.03	4.02	3.75	3.71	3.91
U.K.	2.77	1.49	4.81	1.96	1.98	1.97
U.S.A.	1.10	0.82	1.66	1.40	1.46	1.62

Table 4: RSSE of predicted changes for several models and countries

Using two factors also decreases the error substantially for older age groups. When we consider the RSSE calculated over just the oldest age group in the U.S. using our one factor model the error is 0.22, but when we add the second factor the error is reduced to 0.063, which is a 72% reduction in error for this age group.

5 Forecasting

After we have fit this model we would also like to make forecasts of mortality rates in the future and evaluate the quality of these predictions. In order to do this we must model the future distribution of mortality rates. In this section we propose a few different models to fit the k_t series and use them with a normal approximation of $\epsilon_{x,t}$ to make forecasts of future mortality rates. We compare these different models for k_t using the Bayesian Information Criterion (BIC). We first consider forecasts using our one factor model and then expand to two factors.

Figure 3 shows the time series of k_t for our model. We see that this does not seem to follow any particular process, so we interpret this to mean that k_t are independent observations of a random variable. This does however pose a problem around 1918 because right after a big jump up we observe a big jump down. This could be indicative of a state dependent process, but this problem is also observed in the Lee Carter model when calibrating $\kappa_{t+1} - \kappa_t$ as independent observations of a Gaussian random variable. We ignore this problem and simply interpret k_t as independent observations of a random variable.

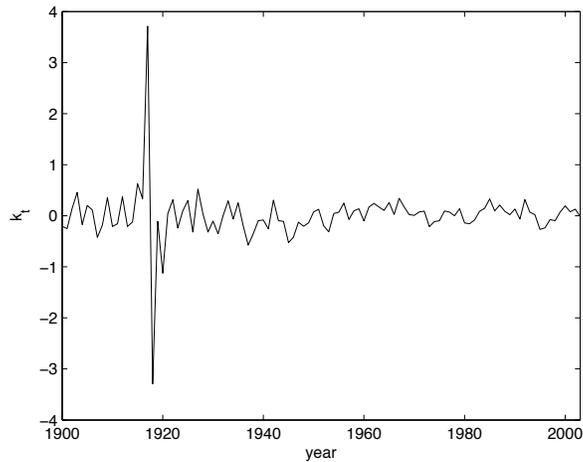


Figure 3: The k_t time series

5.1 Modelling k_t

There are several different ways we can model k_t , each having its own merits. The first model to consider is that the k_t are independent observations of a Gaussian random variable with a certain mean and variance. This model is the easiest to calibrate as the maximum likelihood estimation of the parameters is the usual average and population variance, requiring no numerical optimization. This distribution is simple, however events like the 1918 flu and World War II are extremely unlikely under this model and therefore the fit to actual data is quite bad.

We propose another model to fit the k_t series, called the Normal Inverse Gaussian (NIG). This distribution comes from the area of stochastic processes called subordination and is seen applied to mortality modeling in Wang et al. (2011). The distribution describes the value of a Brownian motion, with some drift and unit variance, evaluated at a random time described by an Inverse Gaussian random variable. An Inverse Gaussian random variable describes the distribution of time it takes for a second, independent, Brownian motion with drift to reach a fixed level. This leads to a very rich distribution that can incorporate skewness and excess kurtosis, both of which are seen in the k_t series. The pdf of a random variable, X , that is distributed like a NIG is

$$(5.1) \quad f_X(x) = e^{\lambda/\theta + \mu(x-\delta)} \sqrt{\frac{\lambda(\lambda + \mu^2\theta^2)}{\pi^2\theta^2(\lambda + (x-\delta)^2)}} K_1\left(\frac{1}{\theta} \sqrt{(\lambda + \mu^2\theta^2)(\lambda + (x-\delta)^2)}\right),$$

where $\mu \in \mathbb{R}$ represents the drift of the Brownian motion, $\delta \in \mathbb{R}$ represents the starting value of the Brownian Motion, $\theta > 0$ is the mean time at which the Brownian motion is evaluated and $\lambda > 0$ is related to the volatility of the second Brownian motion used to construct the Inverse Gaussian. Also, K_ν is the modified Bessel function of the third kind with $\nu = 1$ in our setting. Although the modified Bessel function is a special function, most numerical packages, including R and Matlab, contain built in routines to approximate it numerically. Using this pdf we can easily find a closed form for the log likelihood function.

Figure 4 shows an example of four Brownian motions starting at δ . Here each Brownian motion runs for a different amount of time, determined by an Inverse Gaussian random variable. At the end of each Brownian motion there is a dot, and the y coordinate here represents an observation of an NIG random variable.

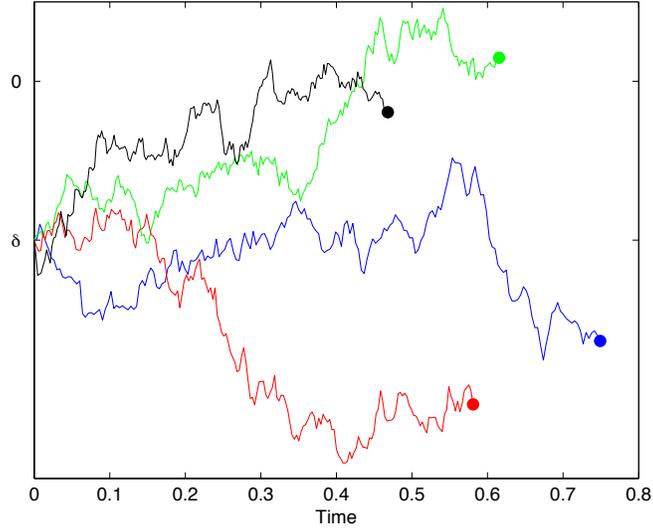


Figure 4: Four Brownian motions with drift evaluated at random times

We must use numerical optimization to find the optimal values of $(\mu, \delta, \theta, \lambda)$ but since we have the likelihood function in closed form we can analytically calculate the first and second derivatives to facilitate optimization. Using standard optimization routines included in Matlab finding the optimal $(\mu, \delta, \theta, \lambda)$, from a random initial guess, takes less than one second.

Table 5 reports the optimal log likelihood values of these three models, along with the Bayesian Information Criterion (BIC) for each model using U.S. data. The likelihood function measures the probability of observing the given data as a function of the parameters of the model. We want to maximize the log likelihood value and minimize the BIC. The BIC is used in statistics for model selection and it discounts the likelihood function by the number of parameters in the model and the number of data points. If L is the optimal value the log likelihood function, n is the number of parameters in the model and q is the number of data points then the BIC is $-2L + n \ln(q)$. The BIC is just one measure of goodness of fit of a probability distribution function to data.

We see that with the additional two parameters of the NIG, over the Gaussian model, the BIC is reduced by about 65%, which is a significant gain. The Double Exponential Jump Diffusion seen, in Deng et al. (2012), also fits the data well but we do not consider it here because the computational complexity

Model	Num. Parameters	Likelihood	BIC
Gaussian	2	-84.62	178.52
NIG	4	-22.42	63.41

Table 5: Log likelihood value and BIC for different models of k_t

is not necessary to adequately fit the data. We will consider the NIG model for forecasting in the next sections because of the computational simplicity and the improved fit over the Normal model.

5.2 Forecasting mortality rates

Now that we have decided on a model for k_t we would like to make forecasts of the log mortality rates for each age group. To do this we use Monte Carlo simulation because we need to include the error term in future predictions. We assume that the error for each age group, $\epsilon_{x,t}$, is Gaussian with zero mean and standard deviation σ_x . This means the distribution of future log mortality rate is the sum of an NIG and a Normal, which would need to be computed using numerical convolution. But we can easily generate NIG random variables and Gaussian random variables so we can combine them to simulate the distribution of future log mortality rate changes. Then using the Monte Carlo results we can obtain the expected value of future mortality rates along with error boundaries. For example, we can get a 95% predictive interval for future log mortality rates by considering the 2.5% and 97.5% values from the Monte Carlo.

Although the NIG may seem like a complicated distribution, because of the simple way that it is constructed it is easy to simulate random variables from the distribution. As explained before, an NIG is simply a Brownian motion, starting at δ , evaluated at a random time. This means that if T is distributed like an Inverse Gaussian with parameters (θ, λ) and $(X|T)$ is distributed like a normal with mean $\mu T + \delta$ and variance T , then the marginal distribution of X is a NIG with parameters $(\mu, \delta, \theta, \lambda)$. So to generate a NIG first generate an Inverse Gaussian random variable, and then generate a normal with the mean and variance determined by T , repeating these two steps in each of Monte Carlo draw. Each normal random variable will have a different mean and variance and this collection of normal random variables, each with random means and variances, are NIG random variables. See Michael et al. (1976) for information on generating Inverse Gaussian random variables.

A property of the NIG distribution that is useful for forecasting is that if $\{X_i\}_{i=1}^n$ are independent random variables distributed like NIG's with parameters $(\mu, \delta, \theta, \lambda)$ then $\sum_{i=1}^n X_i$ is also a NIG, but with parameters $(\mu, n\delta, n\theta, n^2\lambda)$. So if the change in the log mortality index each year is an NIG then the change over n years is the sum of n changes. This means we can easily forecast log mortality rates n years into the future without having to rely on sums between years in the Monte Carlo simulation.

We have considered two sample periods to forecast future log mortality rates for U.S. data. First we considered 1900 - 1980 as the training set and 1981 - 2004 as the testing set, then we considered 1900 - 1960 as the training set and 1961 - 2004 as the testing set. Using the training data we calculated α_x, β_x, k_t and the variance of each $\epsilon_{x,t}$, and then fit an NIG distribution to k_t . For the forecast, in each future year we simulated 100,000 random draws from the appropriate distribution, using the parameters found in the training set, and found the average value along with the 95% predictive interval. The results for 1981 - 2004 are presented in Figure 5 for a few age groups and the results for 1961 - 2004 are presented in Figure 6. In these two figures we compare the 95% predictive interval for our method and the Lee Carter model; we do not consider forecasts of the other models because there is no general consensus on how to best model the cohort effect. The predictive intervals for the Lee Carter model are the red dashed lines, our model's predictive interval is shown in the blue line and the green line represents the mean predicted log mortality rate. In general the predictive interval in our model is tighter than that of the Lee Carter model. This is true in all but three cases; in all but two of the 22 cases our model's predictive interval is inside that of the Lee Carter model and in one the predictive intervals overlap. In each age group and in each year the future log mortality rate never left the 95% predictive interval for our model or the Lee Carter model.

We have also run these experiments for all the other countries we considered. We use different training and testing periods than the U.S. data because data availability is different for each country. This results in many graphs similar to those presented in Figures 5 and 6 and therefore we do not display them in this paper. In the international data we have considered many age groups, many countries and several years of forecasting leading to about 4000 predictions and find that for our model the data escapes the 95% predictive intervals about 3% of the time. For the LC model the data never leaves the 95% predictive intervals. The fact that the data never escapes the 95% predictive interval for the LC model over this many trials leads us to believe that their model over predicts variation in future mortality rates.

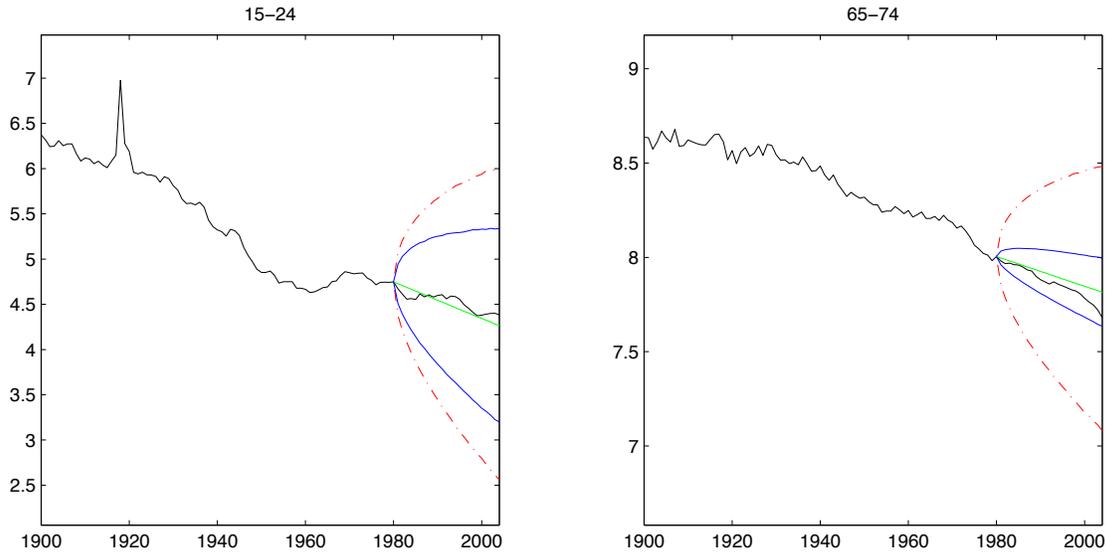


Figure 5: Forecast log mortality rates along with 95% predictive intervals

5.3 Two Factors

Making forecasts of future mortality rates using the two factor model is very similar to using the one factor model. In the case of two factors we need a model to fit $k_t^{(1)}$ and $k_t^{(2)}$. Since $k_t^{(1)}$ and $k_t^{(2)}$ are obtained using the singular value decomposition they are necessarily uncorrelated, therefore we assume that $k_t^{(1)}$ and $k_t^{(2)}$ are independent. This means that they may be fit independently of each other. Here we assume that $k_t^{(1)}$ is an NIG random variable and so is $k_t^{(2)}$. Each k must be fit using the methods described above, then to forecast we follow the same procedure as the one factor model, except we also include $k_t^{(2)}$ in the Monte Carlo simulation of future mortality rates, which changes the variance of the error term from the one factor model. For the younger age groups the forecasts of future mortality rates do not change much using the two factor model, but for older age groups the 95% predictive interval is quite different for some.

In Table 6 we have reported a selection of predictive intervals for different age groups from different countries in the last year of available data, using different models. In each test the training set was selected to include every year of available data, except the last 20 years, and we compute the predictive

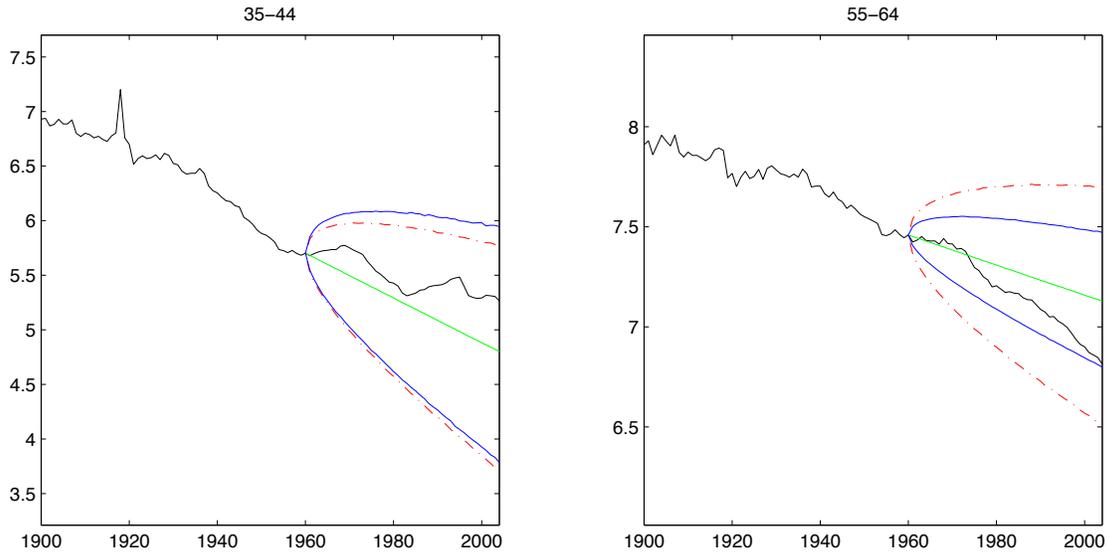


Figure 6: Forecast log mortality rates along with 95% predictive intervals

interval of the log mortality rate in the last year of available data. We see that mostly the Lee Carter model has much wider predictive intervals than our one factor model and two factor model, and the two factor model is wider than the one factor model. This is because we include the variation in the second time component, which is not included in the one factor model. In each of the tests presented in Table 6 the true value of the mortality rate lies inside each predictive interval.

Country	Age Group	Lee Carter	One Factor	Two Factor
Australia	25 - 29	(-8.94,-5.78)	(-7.99,-6.90)	(-8.04,-6.84)
Canada	≥ 100	(-1.74,-0.14)	(-1.59,-0.29)	(-2.04,0.15)
Denmark	55 - 59	(-5.34,-3.98)	(-4.98,-4.40)	(-5.04,-4.35)
Italy	85 - 89	(-2.87,-1.25)	(-2.60,-1.48)	(-2.76,-1.32)
Sweden	65 - 69	(-5.19,-3.29)	(-4.44,-4.00)	(-4.53,-3.91)

Table 6: A selection of predictive Intervals for different age groups in different countries

6 Concluding Remarks

The central idea that has been demonstrated in this paper is that by using a bilinear model to model mortality rate changes rather than the mortality rates directly, we can fit and forecast mortality rates significantly better. Apart from describing the intuitive reasoning for such a model, we have also shown its benefits with an exhaustive array of computational experiments. Almost all variants of the bilinear Lee Cartel model perform better when used on mortality rate changes rather than the mortality rate data. We have also demonstrated that a Normal Inverse Gaussian (NIG) model for the mortality index component of the model, significantly out performs simple random walk models. The NIG model has only four parameters and is extremely easy to calibrate as compared to other stochastic models like the double exponential jump distribution model. We have also studied the first two principal components of the rate change data and their differentiating abilities in capturing features.

Our exhaustive computational results show that our double component NIG model for log mortality change stands out in its ability to provide much better performance measures for all data sets and is still extremely easy to calibrate.

References

- Booth, H., Hyndman, R., Tickle, L., Jong, P. D., 2006. Lee-carter mortality forecasting: a multi-country comparison of variants and extensions. *Demographic Research* 15, 289–310.
- Brockett, P., 1984. General bivariate makeham laws. *Scandinavian Actuarial Journal*, 150–156.
- Cairns, A., Blake, D., Dowd, K., Coughlan, G., Epstein, D., Khalaf-Allah, M., 2011. Mortality density forecasts: An analysis of six stochastic mortality models. *Insurance: Mathematics and Economics* 48, 355–367.
- Chen, H., Cox, S., 2009. Modeling mortality with jumps: Applications to mortality securitization. *The Journal of Risk and Insurance* 76 (3), 727–751.
- Couzin-Frankel, J., 2011. A pitched battle over life span. *Science* 29, 549–550.
- de Moivre, A., 1725. *Annuities upon Lives*.

- DeJong, P., Tickle, L., 2006. Extending Lee-Carter mortality forecasting. *Mathematical Population Studies* 13 (1), 1–18.
- Deng, Y., Brockett, P., MacMinn, R., 2012. Longevity/mortality risk modeling and securities pricing. To appear *Journal of Risk and Insurance*.
- Gompertz, B., 1825. On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. *Philosophical Transactions of the Royal Society of London* 115, 513–585.
- Graunt, J., 1662. *Natural and Political Observations Made upon the Bills of Mortality*.
- Gutterman, S., England, C., Parikh, A., Pokorski, R., 2002. Living to 100 and beyond: Implications for retirement. *Record of the Society of Actuaries* 28.
- Haberman, S., Renshaw, A., 2011. A comparative study of parametric mortality projection models. *Insurance: Mathematics and Economics* 48, 35–55.
- Halley, E., 1693. An estimate of the degrees of the mortality of mankind, drawn from curious tables of the births and funerals at the city of Breslaw; with an attempt to ascertain the price of annuities upon lives. *Philosophical Transactions of the Royal Society of London* 17, 596–610.
- Halonen, D., January 22, 2007. IRS sets costly table. *Pension and Investments*.
- Hollmann, F., Mulder, T., Kallan, J., 2000. Methodology and assumptions for the population projections of the United States: 1999 to 2100. *Population Division, U.S. Bureau of Census* 38.
- Human Mortality Database, 2011. University of California, Berkeley (USA), and Max Plank Institute for Demographic Research (Germany).
URL www.mortality.org
- Jolliffe, I., 2002. *Principal Component Analysis*. Springer, New York.
- Lee, R., Carter, L., 1992. Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association* 87 (419), 659–671.

Makeham, W., 1860. On the law of mortality and the construction of annuity tables. *J. Inst. Actuaries and Assur. Mag.* 8, 301–310.

Pension Capital Strategies, Jardine Lloyd Thompson, 2006. Pension capital strategies releases report on FTSE100 pension disclosures.

URL www.jltgroup.com/content/UK/employee_benefits/pressreleases/070117PCSQuaReport.pdf

Renshaw, A., Haberman, S., 2006. A cohort-based extension to the lee-carter model for mortality reduction factors. *Insurance: Mathematics and Economics* 38, 556–570.

Reuters News Service, November 23, 2010. Uk pension liabilities spike as people live longer.

URL www.reuters.com/article/idUSLNE6A001120101125

Wang, C., Huang, H., Liu, I., 2011. A quantitative comparison of the Lee-Carter model under different types of non-Gaussian innovations. *The Geneva Papers* 36, 675–696.

Yang, S., Yue, J., Huang, H., 2010. Modeling longevity risks using a principal component approach: A comparison with existing stochastic mortality models. *Insurance: Mathematics and Economics* 46, 254–270.