

Kent Academic Repository

Full text document (pdf)

Citation for published version

Sweeting, Paul (2010) A Trend-Change Extension of the Cairns-Blake-Dowd Model. In: International Congress of Actuaries, 7-12 March 2010, Cape Town, South Africa.

DOI

Link to record in KAR

<http://kar.kent.ac.uk/47886/>

Document Version

Publisher pdf

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

A Trend-Change Extension of the Cairns-Blake-Dowd Model

P.J.Sweeting

February 3, 2010

Abstract

This paper builds on the two-factor model developed by Cairns et al. (2006) and updated by Cairns et al. (2009) for projecting future mortality. It is shown that these two factors do not follow a random walk, as proposed by Cairns et al. (2006), but should instead be modelled as a random fluctuation around a trend, the trend changing periodically. Projecting mortality rates in this way suggests much greater uncertainty over future mortality improvements.

p.j.sweeting@kent.ac.uk; School of Mathematics, Statistics and Actuarial Science, University of Kent, Canterbury, Kent, CT2 7NZ

1 Introduction

In recent years, there have been a number of models developed to try to ascertain future changes in mortality rates. The most well-known of these is the model developed by Lee and Carter (1992), but this has seen many developments, such as the more robust treatment of errors by Brouhns et al. (2002) and the addition of cohort effects by Renshaw and Haberman (2003, 2006). There are also continuous time variants such as those developed by Milevsky and Promislow (2001), Dahl (2004), Dahl and Møller (2005), Milterson and Persson (2005), Biffis (2005) and Schrage (2006). The version I consider, however, has its roots in a discrete time model developed by Cairns et al. (2006). A more robust version of this model is given as model M5 in Cairns et al. (2009). This re-bases the age used to calculate the age-related component of the model, centring it around the average age of the data. Cairns et al. (2009) also describe a number of more complex versions of the original formulation. However, the simplicity of the original model, retained in model M5, is attractive and better allows the demonstration of the approach in this paper. I therefore use the M5 formulation, which I call the CBD model.

The CBD model is an innovative two-factor model. It assumes that each of the two parameters follows a random walk with drift, the rate of drift being

constant and changes in the parameters being correlated. This approach is well-suited for the pricing of mortality-related derivatives with a term of a few years; however, when considering the very long-term, it becomes clear that the patterns for these two factors do not necessarily resemble a random walk with drift process; for most periods each of the factors can be modelled as a random fluctuation around a trend, the trend changing periodically.

This is not the first research to question the use of a random walk with a single rate of drift as the appropriate method of mortality projection. Booth et al. (2002) find that the level of drift for Australian data is not constant, and allow for this by fitting their model only for the period where a linear trend in mortality is apparent. Using the same dataset, De Jong and Tickle (2006) allow for varying rates of drift over time, and suggest improvements such as the addition of an autoregressive process to the parameters used; however, like Booth et al. (2002), mortality rates are assumed to follow a random walk.

The CBD model describes the logit of the initial mortality rate with a slope term and an intercept term, allowing for the number of deaths to follow a Poisson distribution. Future stochastic simulations are then obtained by projecting these two terms as following correlated random walks. In other words:


$$\log\left(\frac{q_{x,y}}{1-q_{x,y}}\right) = k_{1,y} + [(x - \bar{x}) \times k_{2,y}] + e_{x,y} \quad (1)$$

where $q_{x,y}$ is the initial mortality rate for a life aged x in year y ; \bar{x} is the average of the ages x used in the analysis; $k_{1,y}$ is the intercept term in year y ; $k_{2,y}$ is the slope term in year y ; and e_y is an error term. The parameters $k_{1,y}$ and $k_{2,y}$ are members of sets described by $k_1 = \{k_{1,y} : y = 1 \dots Y\}$ and $k_2 = \{k_{2,y} : y = 1 \dots Y\}$, and k_1 and k_2 are equivalent to the parameters described in the CBD model as kappa 1 and kappa 2 respectively.

The CBD model is calibrated using data for England and Wales males from age 60 to 89 over the period 1961 to 2002. However, this is a relatively short time scale which has seen a generally steady pattern of improvements in life expectancy. It is not reasonable to expect this constancy to continue indefinitely into the future, something which can be seen by considering past mortality rate improvements over a longer period. One way of appreciating these past changes is to fit the CBD model over a larger dataset, covering the period from 1841 to 2005, still using England and Wales males age 60 to 89. Here, clear patterns can be seen in $k_{1,y}$ and $k_{2,y}$, as shown in (1). The data used for these calculations is obtained from the Human Mortality Database (2008). Figure 1 suggests that both $k_{1,y}$ and $k_{2,y}$ exhibit variation around a trend rather than following random walks. They also suggest that these trends change suddenly and definitely, that changes in the trend for $k_{1,y}$ and $k_{2,y}$ tend to happen at the same time, and that there is strong negative correlation between the direction of these trend changes.

These sudden changes in the trend are what one would expect, considering changes that have happened over the last 160 years or so. Cutler et al. (2006) describe three phases of change in mortality rates. From the mid eigh-

Figure 1: $k_{1,y}$ and $k_{2,y}$ for the CBD Model, 1841-2005



Source: Human Mortality Database (2008), Author's Calculations

teenth century to the mid nineteenth century, they see mortality rates falling as a result of improved nutrition and economic growth; in the late nineteenth and early twentieth centuries, improvement is seen as a result of clean water, better treatment of effluent and better health information, all of which became necessary as a result of urbanisation; and the 1930's onwards is characterised as the decade of big medicine with immunisation, antibiotics, and ever increasing medical technology. Another phenomenon that continues to have a major effect on mortality rates is the effect of the reduction in smoking in the last quarter of the twentieth century, as described by Humble and Wilson (2008).

The negative correlation is also understandable. From (1) it can be seen that at time t , $k_{1,y}$ represents the level of the logit of mortality, and $k_{2,y}$ the extent to which the logit of mortality varies with age. The second variable therefore reflects changes in the relative levels of mortality at different ages, but the first variable determines the underlying change in mortality around which the relative changes are based.

The organisation of this paper is, then, as follows. First, I consider ways in which change-points in the trend can be determined. I then fit trends to different sections of the data. Within each of these sections, I then consider whether the observations are trend- or difference-stationary. Next, I propose an approach for projection mortality rates based on this structure, and finally I show the results of some of these projections.

2 Determining the Change-Points in the Trend

Looking at (1), it is clear that there are breaks in the trend for both k_1 and k_2 . However, it is not necessarily clear exactly where these breaks occur, and whether any changes are significant enough to be considered changes in trend.

One way of investigating possible changes-points is to fit lines to sections of k_1 and k_2 and to calculate the Durbin-Watson (DW) statistic, using the test designed by Durbin and Watson (1950, 1951). Under the null hypothesis that a fitted line covers data within a single trend, then the DW statistic should show no evidence of significant serial correlation between subsequent values of $k_{1,y}$ or $k_{2,y}$. However, if a line covers, say two periods where the rate of change is lower in the second period than the first, then the first and last sections of the data would lie below the line whilst the middle section would lie above it. This would lead to the DW statistic showing significant positive serial correlation in these variables, suggesting that the null hypothesis could be rejected and that there was evidence for a break in the trend. The DW test also incorporates a middle ground of uncertainty, where the test is inconclusive. This approach is not exact, but gives some clues as to where breaks might occur.

The methodology used for the above approach is to fit a line for each year y where $y = 1841$ to 2003 (so using at least three years of data to calculate the DW statistic) and, within each year for each period p where $p = 3$ to 2005 $-y$. This approach can point to broad areas where changes in trend might occur; however, when fitting a series of lines to this data, some trial and error

is required to get a good fit.

3 Fitting a Model

Having identified potential breakpoints both visually and using the DW test, the results are verified using more exact approaches. The broad approach used to fit lines to $k_{1,y}$ and $k_{2,y}$ is a weighted least squares approach. This is necessary because of the heteroskedasticity present in both $k_{1,y}$ and $k_{2,y}$ the variance decreases substantially over time. There are several possible causes for this. The first is simply that the quality of data is likely to have improved over the period of observation. However, as life expectancies have improved, so have the number of lives over age 60. This leads to less stochastic variation in the mortality rates calculated, and thus less variation in the estimates of $k_{1,y}$ and $k_{2,y}$ from year to year. A third possibility is that improved understanding and treatment of infectious diseases together with improved air and water quality reduced the incidence of various epidemics that also led to significant year-to-year variation in mortality.

The weights used are the reciprocals of the variances for the seven years centred on the observation in question. The weights for the first three observations are set equal to the fourth observation; the weights for the last three observations are set equal to the fourth from last observation.

I first consider k_1 . Let the final year of a trend (and the first year of the next) within k_1 be identified as $b_{k_1,n(k_1)}$ where $n(k_1) = 1 \dots N(k_1) - 1$. This means that the number of break points is $N(k_1) - 1$ and the number of lines is $N(k_1)$. Each line is expressed as a constant, $\alpha_{k_1,n(k_1)}$ plus a slope, $\beta_{k_1,n(k_1)}$, the latter being multiplied by the year. This means that the estimate of $k_{1,y}$ is described as follows:

$$\hat{k}_{1,y} = \begin{cases} \alpha_{k_1,1} + \beta_{k_1,1}y & \text{if } y \leq b_{k_1,1} \\ \alpha_{k_1,1} + \beta_{k_1,1}y & \text{if } b_{k_1,1} < y \leq b_{k_1,2} \\ \vdots & \\ \alpha_{k_1,N(k_1)} + \beta_{k_1,N(k_1)}y & \text{if } y > b_{k_1,N(k_1)-1} \end{cases} \quad (2)$$

Similarly, the estimate of $k_{2,y}$ is described as follows:

$$\hat{k}_{2,y} = \begin{cases} \alpha_{k_2,1} + \beta_{k_2,1}y & \text{if } y \leq b_{k_2,1} \\ \alpha_{k_2,1} + \beta_{k_2,1}y & \text{if } b_{k_2,1} < y \leq b_{k_2,2} \\ \vdots & \\ \alpha_{k_2,N(k_2)} + \beta_{k_2,N(k_2)}y & \text{if } y > b_{k_2,N(k_2)-1} \end{cases} \quad (3)$$

It can also be said that:

$$k_{1,y} = \hat{k}_{1,y} + \epsilon_{k_1,y} \quad (4)$$

and

$$k_{2,y} = \hat{k}_{2,y} + \epsilon_{k2,y} \quad (5)$$

where $\epsilon_{k1,y}$ and $\epsilon_{k2,y}$ are residuals. Initially, $N(k1)$ is set equal to $N(k2)$. Furthermore, all $b_{k1,n(k1)}$ are set equal to $b_{k2,n(k2)}$. This second point seems sensible, since a visual inspection of $k_{1,y}$ and $k_{2,y}$ suggests that changes in trends occur at about the same time in one as in the other, although there are perhaps some instances when a significant change occurs in only one series. There is also a rationale for expecting changes in $k_{1,y}$ and $k_{2,y}$ to occur at the same time, since a change in both determines the age around which relative changes in mortality pivot. This is investigated below.

When fitting the lines, all potential changes in trend are considered initially. The lines are then fitted by minimising the sum of squared errors subject to the restriction that:

$$\alpha_{k1,n(k1)} + \beta_{k1,n(k1)} b_{k1,n(k1)} = \alpha_{k1,n(k1)+1} + \beta_{k1,n(k1)+1} b_{k1,n(k1)} \quad (6)$$

and:

$$\alpha_{k2,n(k2)} + \beta_{k2,n(k2)} b_{k2,n(k2)} = \alpha_{k2,n(k2)+1} + \beta_{k2,n(k2)+1} b_{k2,n(k2)} \quad (7)$$

for $n(k1) = 1 \dots N(k1) - 1$ and $n(k2) = 1 \dots N(k2) - 1$. The first test I carry out is a Chow test, as described by Chow (1960). Under the Chow test, the null hypothesis is that $\alpha_{k1,n(k1)} = \alpha_{k1,n(k1)+1}$ and $\beta_{k1,n(k1)} = \beta_{k1,n(k1)+1}$ for consecutive groups of data within $k1$, the same being true for $k2$. To calculate the test statistic, a single line is fitted covering both groups, and the test statistic is given as:

$$CT = \frac{(SSR_{All} - (SSR_1 + SSR_2)) / v}{(SSR_1 + SSR_2) / (N_1 + N_2 - 2v)} \sim F_{N_1 + N_2 - 2v}^v \quad (8)$$

where SSR_1 , SSR_2 and SSR_{All} are the sums of squared residuals from the first section of data, the second section of data and the combined dataset respectively. The number of variables, given by v , is two, and the numbers of observations in the first and second groups of data are given by N_1 and N_2 . The Chow test is helpful in examining structural breaks since it considers not the extent to which individual observations represent a change from an established trend, but rather the extent to which two groups of variables display different patterns.

The restrictions in (6) and (7) effectively mean that the slope parameters are optimised subject to restrictions on the level parameters. I therefore look at the difference between successive slope parameters to identify where any change in slope is not statistically significant, calculating a t-statistic using the joint standard error calculated assuming unequal sample sizes and unequal variances in each sample, the null hypothesis being that there is no change in slope.

I also consider the DW statistic for the various sections. If separate trends are combined, then this may be highlighted by significant positive serial correlation, suggesting that a change in trend has been missed.

Looking first at k_1 , a number of break points appear to be less than convincing, suggesting the need for further investigation; however, the only break point not strongly suggested by the Chow Tests in k_2 is the first.

Trying various combinations of the trends and examining the various statistics suggests that the breakpoints identified in Tables 3 and 4, below, better describe the data.

Table 1: Parameters for Fitted Values of k_1

First Year of Trend	k_1 Intercept (SE)	k_1 Slope (SE)	DW Statistic	DF Statistic	Change in k_1 Slope	Chow Test Statistic
1841	-2.966830 (4.994174)	0.000168 (0.002700)	1.701532	-3.518565 *	0.006502 ** (0.002856)	3.694137 **
1860	-15.054700 (5.491361)	0.006671 (0.002943)	1.219745 +	-2.279956	-0.007657 *** (0.002069)	2.668865 *
1873	-0.720550 (2.678539)	-0.000986 (0.001419)	1.496792	-4.129292 **	-0.005125 *** (0.001407)	13.189975 ***
1902	9.021928 (2.679251)	-0.006111 (0.001399)	1.256072 ++	-3.055781	-0.001805 (0.002875)	0.355020
1930	12.504380 (7.156921)	-0.007917 (0.003695)	1.042455 ++	-2.349303	0.008425 (0.006203)	5.760253 **
1945	-3.874190 (15.239244)	0.000509 (0.007817)	1.974323	-2.828847	-0.003804 (0.005403)	3.221651 *
1955	3.558806 (2.530577)	-0.003295 (0.001289)	1.950550	-4.133678 **	-0.011955 *** (0.001123)	26.251961 ***
1973	27.133480 (1.875906)	-0.015250 (0.000948)	1.784664	-3.667513 **	-0.009425 ** (0.003682)	23.103952 ***
1988	45.851490 (10.024917)	-0.024675 (0.005033)	2.584364	-5.415041 ***	-0.008391 (0.005222)	48.665600 ***
1998	62.607600 (10.800210)	-0.033066 (0.005396)	1.951402	-2.331239		

Significance level: *** 1%; ** 5%; * 10%.

Serial correlation: ++ evidence of significant positive serial correlation (statistic < dL); + evidence of possible positive serial correlation (dL < statistic < dU).

Table 2: Parameters for Fitted Values of k_2

First Year of Trend	k_2 Intercept (SE)	k_2 Slope (SE)	DW Statistic	DF Statistic	Change in k_2 Slope	Chow Test Statistic
1841	0.019917 (0.150937)	0.000031 (0.000082)	2.503566	-5.360380 ***	-0.000344 ** (0.000126)	1.727529
1860	0.658869 (0.272553)	-0.000313 (0.000146)	1.148115 +	-2.600994	0.000385 *** (0.000092)	12.562046 ***
1873	-0.060940 (0.080100)	0.000072 (0.000042)	1.535067	-4.872952 ***	0.000574 *** (0.000041)	94.695468 ***
1902	-1.152990 (0.078094)	0.000646 (0.000041)	2.402800	-6.538536 ***	-0.000800 *** (0.000073)	76.076018 ***
1930	0.389713 (0.178596)	-0.000154 (0.000092)	1.348259 +	-4.161198 **	0.000512 *** (0.000141)	40.616383 ***
1945	-0.605620 (0.339579)	0.000358 (0.000174)	2.369726	-3.530067 *	-0.000430 *** (0.000125)	85.953783 ***
1955	0.234761 (0.108824)	-0.000072 (0.000055)	2.138283	-4.409199 ***	0.000430 *** (0.000056)	29.847510 ***
1973	-0.613360 (0.110733)	0.000358 (0.000056)	1.178452 +	-2.773581	0.000141 ** (0.000064)	11.426100 ***
1988	-0.892990 (0.142363)	0.000499 (0.000071)	2.166787	-4.649186 ***	-0.000346 *** (0.000063)	69.007492 ***
1998	-0.202820 (0.107361)	0.000154 (0.000054)	1.330504 +	-2.244186		

Significance level: *** 1%; ** 5%; * 10%.

Serial correlation: ++ evidence of significant positive serial correlation (statistic < dL); + evidence of possible positive serial correlation (dL < statistic < dU).

Table 3: Revised Parameters for Fitted Values of k_1

First Year of Trend	k_1 Intercept (SE)	k_1 Slope (SE)	DW Statistic	DF Statistic	Change in k_1 Slope	Chow Test Statistic
1841	-2.966830 (4.994174)	0.000168 (0.002700)	1.696690	-3.518565 *	0.006537 ** (0.002856)	5.133120 ***
1860	-15.120000 (5.491361)	0.006706 (0.002943)	1.220154 +	-2.279956	-0.007700 *** (0.002069)	8.041144 ***
1873	-0.705050 (2.678539)	-0.000995 (0.001419)	1.495064	-4.129292 **	-0.005182 *** (0.001046)	16.802856 ***
1902	9.145043 (1.434794)	-0.006176 (0.000746)	1.447660 ++	-4.727278 ***	0.003530 *** (0.000809)	6.216409 ***
1945	2.282530 (1.673403)	-0.002646 (0.000854)	1.673548	-4.705058 ***	-0.012966 *** (0.000992)	33.455055 ***
1973	27.852100 (2.180101)	-0.015612 (0.001101)	1.885997	-3.667513 **	-0.008897 ** (0.003612)	100.169916 ***
1988	45.522080 (9.768575)	-0.024510 (0.004904)	2.629184	-5.415041 ***	-0.008834 (0.005106)	32.807038 ***
1998	63.162860 (10.588903)	-0.033343 (0.005290)	2.030812	-2.331239		

Significance level: *** 1%; ** 5%; * 10%.

Serial correlation: ++ evidence of significant positive serial correlation (statistic<dL); + evidence of possible positive serial correlation (dL<statistic<dU).

Table 4: Revised Parameters for Fitted Values of k_2

First Year of Trend	k_2 Intercept (SE)	k_2 Slope (SE)	DW Statistic	DF Statistic	Change in k_2 Slope	Chow Test Statistic
1841	0.255659 (0.069743)	-0.000097 (0.000038)	1.863737	-5.240416 ***	0.000127 (0.000038) ***	3.351952 **
1873	0.017483 (0.071755)	0.000030 (0.000038)	1.903972	-6.725793 ***	0.000619 (0.000040) ***	109.848615 ***
1902	-1.159045 (0.078098)	0.000649 (0.000041)	2.419991	-2.656571 ***	-0.000803 (0.000074) ***	71.745047 ***
1930	0.389737 (0.181001)	-0.000154 (0.000093)	1.331124 +	-3.071266 **	0.000512 (0.000141) ***	40.299153 ***
1945	-0.605591 (0.338518)	0.000358 (0.000174)	2.377781	-3.264308 *	-0.000430 (0.000125) ***	79.795420 ***
1955	0.234785 (0.109126)	-0.000072 (0.000056)	2.133573	-3.613035 ***	0.000430 (0.000055) ***	30.105090 ***
1973	-0.613332 (0.107336)	0.000358 (0.000054)	1.215524 +	-3.819187	0.000141 (0.000065) **	7.237369 ***
1988	-0.892961 (0.146787)	0.000499 (0.000074)	2.105288	-3.912770 ***	-0.000346 (0.000063) ***	58.677171 ***
1998	-0.202798 (0.101564)	0.000154 (0.000051)	1.414360	-3.978984		

Significance level: *** 1%; ** 5%; * 10%.

Serial correlation: ++ evidence of significant positive serial correlation (statistic < dL); + evidence of possible positive serial correlation (dL < statistic < dU).

Again, looking first at k_1 , the Chow test statistic strongly suggests breaks at all the points remaining. However, the change-in-trend test also suggests changes in all but one of these instances, and further combining the trends in the one instance where no change is suggested results in significant positive serial correlation. In fact, one of the combinations carried out has also resulted in a single case of positive serial correlation; however, looking at the results in Figure 2 suggest that this is more likely to be a result of random clustering than two separate trends having been combined.

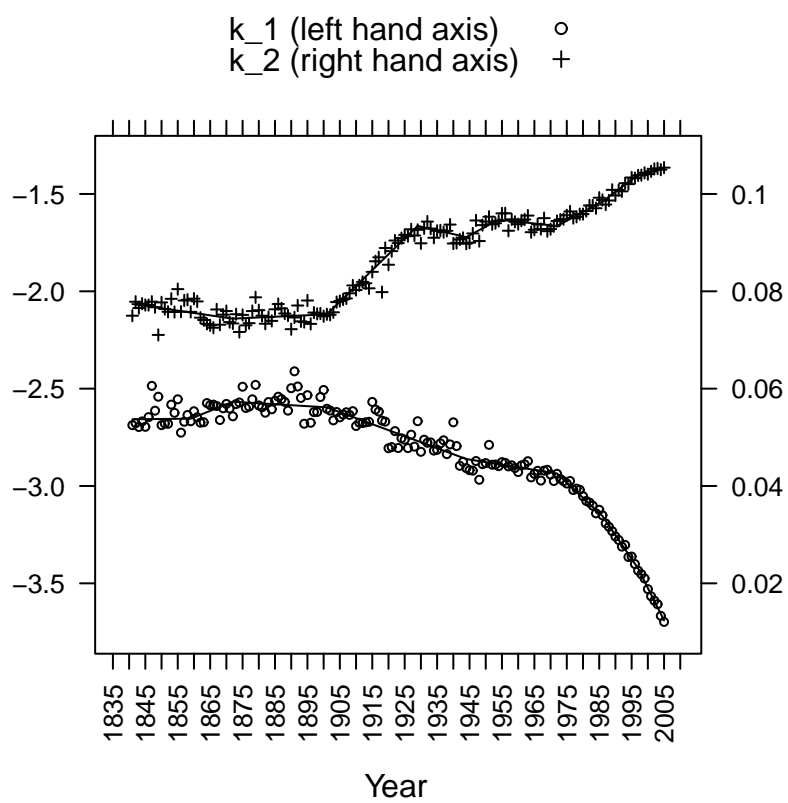
The single change in k_1 is less controversial the Chow test and the change in trend analysis both suggest that no further changes are needed, and the DW statistics do not suggest that any trends have been combined incorrectly.

Once reported statistic has not been discussed in detail: the Dickey-Fuller (DF) statistic. Under the null hypothesis of the test designed by Dickey and Fuller (1979) the time series under investigation follows a random walk with drift in other words, it has a unit root. The alternative hypothesis is that the time series consists of random deviations from a trend. For k_1 , the DF test suggests that the data follows random deviations around a trend in six out of eight periods; for k_2 , the figure is seven out of nine periods. In most case, the level of significance is 1%.

Ideally, the DF statistic would be calculated for the whole period. However, since the calculation of the statistic involves regression against a time trend and it has already been shown that this trend changes over time, such a test is not straightforward. In particular, changes in the direction of the trend of either the level or the slope parameter will result in unreliable results. The change-of-trend issues is also why this test is considered now, rather than before the investigations into the break-points in the trend. One possible modification to the DF test in the presence of a changing trend is to substitute the linear time trend with a series constructed from the slope parameters. If the slope in year y for parameter k_1 is ${}_y\beta_{k_1}$, then this replacement trend for k_1 can be calculated as ${}_yz_{k_1} = {}_y\beta_{k_1} + {}_{y-1}\beta_{k_1}$, where ${}_1z_{k_1} = {}_1\beta_{k_1}$. Note that ${}_y\beta_{k_1}$, the slope parameter in year y , should not be confused with $\beta_{k_1, N(k_1)}$, the slope parameter for all y falling in trend $N(k_1)$. Similarly, the replacement trend for k_2 can be calculated as ${}_yz_{k_2} = {}_y\beta_{k_2} + {}_{y-1}\beta_{k_2}$, where ${}_1z_{k_2} = {}_1\beta_{k_2}$. Using this approach, the null hypothesis for the data series as a whole is rejected at the 1% level of significance, with the test statistics for k_1 and k_2 being -10.317059 and -13.064194 respectively. This provides strong evidence that each parameter follows a trend-stationary process, albeit with a trend that changes periodically.


Figure 2 shows how the fitted lines described in Table 3 and Table 4 relate to the parameters shown in Figure 1. However, the scale of these diagrams is such that it is difficult to see the extent to which the trend does actually change. Figure 3 and Figure 4 are therefore given. These show the values of $k_{1,y}$ and $k_{2,y}$ for each consecutive pairs of trends, as described in Table 3 and Table 4. They also show the fitted line for the first trend extended to the period covered by the second trend. This is intended to show the extent to which a change in trend occurs. In most cases the change is marked; in others it is less definitive but nonetheless apparent.

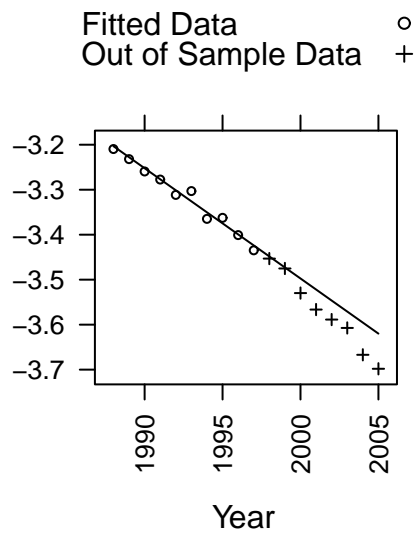
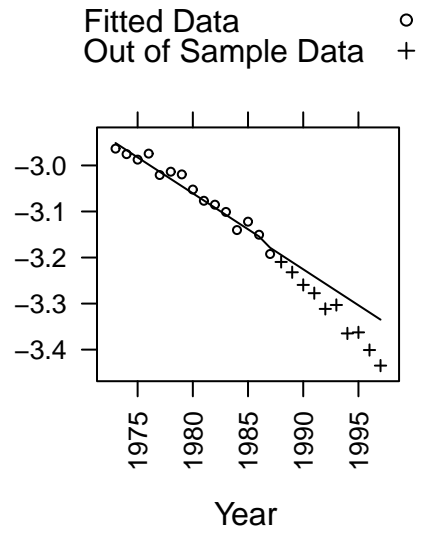
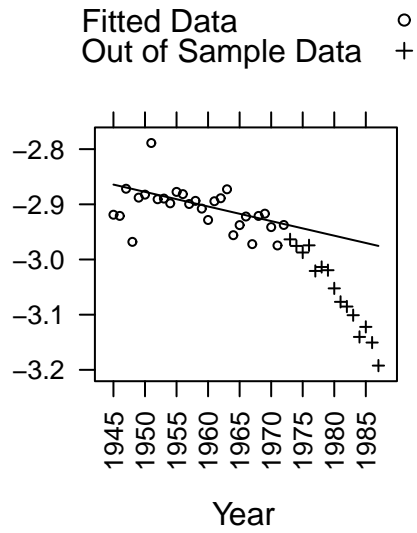
Figure 2: $k_{1,y}$ and $k_{2,y}$ for the CBD Model, 1841-2005



Source: Human Mortality Database (2008), Author's Calculations

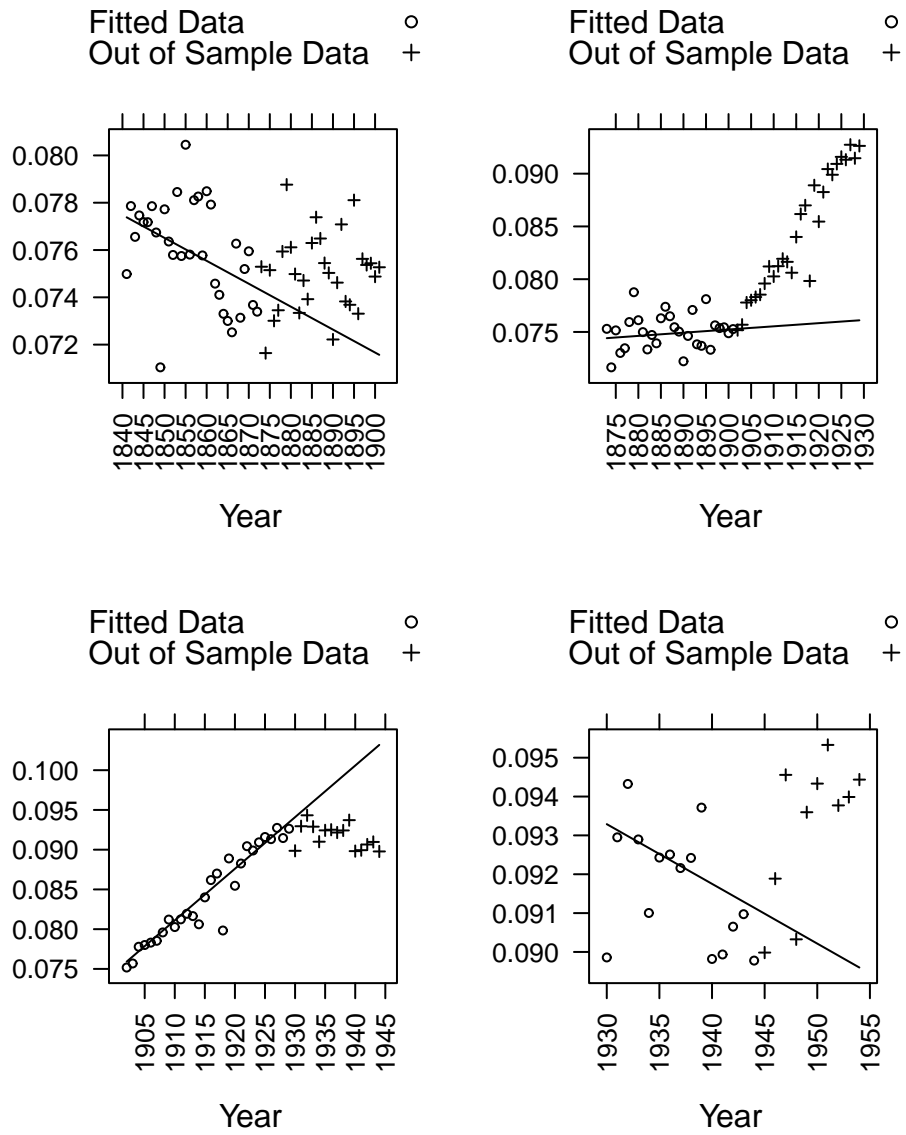
Figure 3: Changes in Trend for $k_{1,y}$ Slope – Post-Change Observations vs Pre-Change Fitted Line

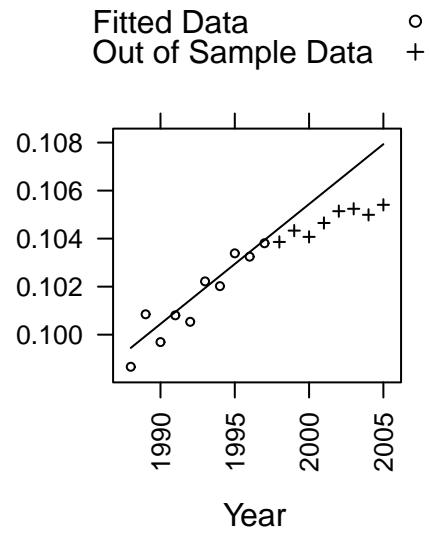
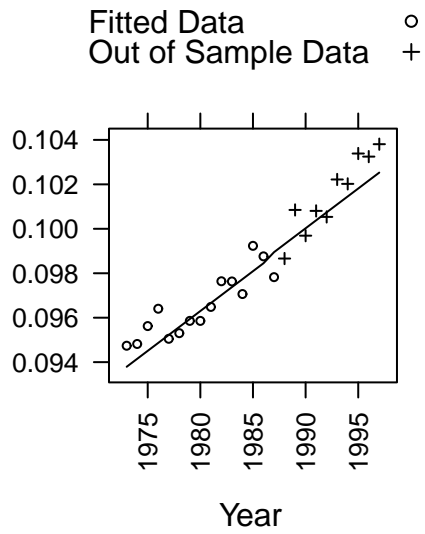
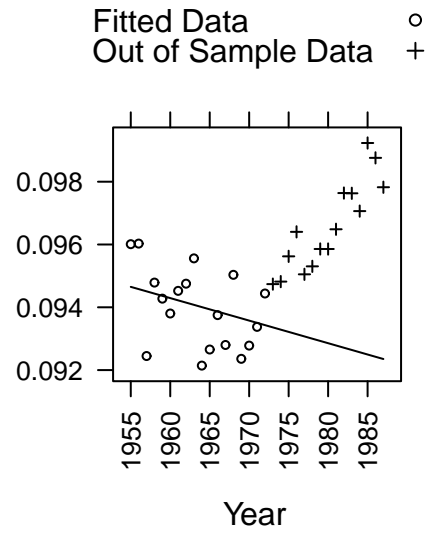
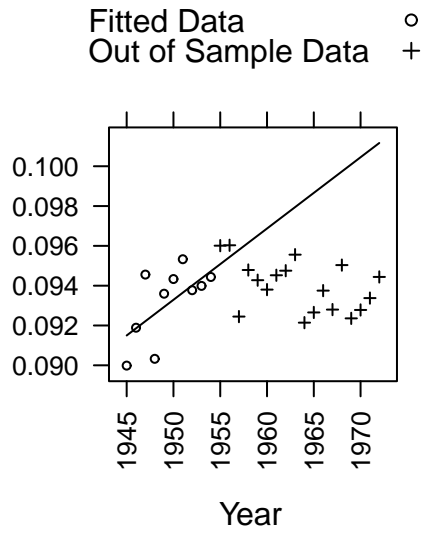




Source: Human Mortality Database (2008), Author's Calculations

Figure 4: Changes in Trend for $k_{2,y}$ Slope – Post-Change Observations vs Pre-Change Fitted Line





Source: Human Mortality Database (2008), Author's Calculations

4 A Projection Approach

Having fitted a model to the data, I next look at an approach to projecting mortality rates. The first stage here is to consider how the parameters behave in the historical data.

Since (6) and (7) mean that the changes in $\alpha_{k_1,n(k_1)}$ and $\alpha_{k_2,n(k_2)}$ are defined by the changes in $\beta_{k_1,n(k_1)}$ and $\beta_{k_2,n(k_2)}$, it is only the latter that we need to consider when looking at changes of the trends in $k_{1,y}$ and $k_{2,y}$. The trend in $k_{1,y}$ changes 7 times out of a possible 164 times, a probability of 0.036585; the trend in $k_{2,y}$ changes 8 times in the same period, a probability of 0.054878. However, $\beta_{k_1,n(k_1)}$ and $\beta_{k_2,n(k_2)}$ change together only six of these times. This means that:

- the probability of *only* $\beta_{k_1,n(k_1)}$ changing is 1 in 164, or 0.006098;
- the probability of *only* $\beta_{k_2,n(k_2)}$ changing is 2 in 164, or 0.012195;
- the probability of *both* $\beta_{k_1,n(k_1)}$ and $\beta_{k_2,n(k_2)}$ changing is 6 in 164, or 0.036585.

These changes can therefore be modelled by simulating a uniform random variable, $0 \leq \Delta\beta < 1$, such that:

- if $0 \leq \Delta\beta < 0.006098$ then $\beta_{k_1,n(k_1)}$ only changes;
- if $0.006098 \leq \Delta\beta < 0.042683$ then $\beta_{k_1,n(k_1)}$ and $\beta_{k_2,n(k_2)}$ change (since $0.042683 - 0.006098 = 0.036585$);
- if $0.042683 \leq \Delta\beta < 0.054878$ then $\beta_{k_2,n(k_2)}$ only changes (since $0.054878 - 0.042683 = 0.012195$; and
- if $\Delta\beta \geq 0.054878$ then neither $\beta_{k_1,n(k_1)}$ nor $\beta_{k_2,n(k_2)}$ change (since $0.006098 + 0.012195 + 0.036585 = 0.054878$, or $8/164$).

The next stage is to look at the variation in $\beta_{k_1,n(k_1)}$ and $\beta_{k_2,n(k_2)}$ when they change. One approach would be to consider the root mean square (RMS) of the deviation of each variable from its previous value and use this to define the volatility. This would mean essentially assuming that the expectation was for the current trends in $k_{1,y}$ and $k_{2,y}$ to continue indefinitely, with each being as likely to accelerate as to decelerate. The alternative is to use the standard deviation. This would give a lower level of volatility, but would require an assumption that changes in $k_{1,y}$ and $k_{2,y}$ would be expected to accelerate or decelerate (depending on whether the average was positive or negative). This latter view seems unrealistic, as it suggests that both $k_{1,y}$ and $k_{2,y}$ would eventually tend towards positive or negative infinity, leading to biologically unrealistic mortality rates. I therefore use approach based on the RMS, describe above. I also calculate a measure of correlation between $\beta_{k_1,n(k_1)}$ and $\beta_{k_2,n(k_2)}$ when both change using a similar approach, calculating a measure of covariance based

on deviations from zero rather than from the mean, then dividing the result by the RMS of $\beta_{k1,n(k1)}$ and $\beta_{k2,n(k2)}$. The correlation calculated by this method – effectively an “RMS correlation coefficient” – is -0.301142, compared with the “true” correlation coefficient – calculated as the actual covariance of $\beta_{k1,n(k1)}$ and $\beta_{k2,n(k2)}$ divided by the standard deviation of each – of 0.380513. As discussed above, there is a good reason to use the RMS instead of the standard deviation when measuring volatility. This reason – that using the standard deviation would lead to unrealistic mortality rates – also applies to the measure of correlation. This suggests that the RMS correlation coefficient is more appropriate than the true one. This view is strengthened by the observation that the two lines in Figure 2 diverge in a reasonably consistent fashion over the full period of investigation. Divergence and convergence are both consistent with a negative correlation coefficient; a positive coefficient would suggest that the two series should move in broadly the same direction.

When both $\beta_{k1,n(k1)}$ and $\beta_{k2,n(k2)}$ change, the RMS of the former is 0.008406 (compared with a standard deviation of 0.005597) and of the latter is 0.000405 (compared with a standard deviation of 0.000352). However, the RMS for $\beta_{k1,n(k1)}$ using all changes is 0.015998 (compared with a standard deviation of 0.013719), and for $\beta_{k2,n(k2)}$ is 0.000332 (compared with a standard deviation of 0.000287). An F-test suggests that the volatility differs for $k_{1,y}$, but only at the 10% level; there is no significant difference for $\beta_{k2,n(k2)}$. However, the small number of observations means that it is difficult to draw any strong conclusions. In the analysis below, I use the same volatility when only one variable changes or both change.

Finally, given that the data is to be modelled assuming random volatility around a trend, the nature of this volatility needs to be investigated. Analysis of the standard deviation of the errors, $\epsilon_{k1,y}$ and $\epsilon_{k2,y}$, shows that the volatility decreases successively in each period, but there is a particularly large and sustained fall for both $\beta_{k1,n(k1)}$ and $\beta_{k2,n(k2)}$ in 1973. I therefore assume that the volatility around the trend is given by the volatility calculated using data since this date, 0.012098 for $\epsilon_{k1,y}$ and 0.000627 for $\epsilon_{k2,y}$.

Having determined the parameters for the projection of mortality, I then project mortality forward stochastically. This is done by projecting values for $\beta_{k1,n(k1)}$ and $\beta_{k2,n(k2)}$. A uniform random variable taking values between 0 and 1 is generated in each period. The value of the random variable determines whether either or both of $\beta_{k1,n(k1)}$ and $\beta_{k2,n(k2)}$ change, as described earlier. If either variable changes, then another random variable is generated with a mean of zero and a standard deviation taken as the RMS of past changes in $\beta_{k1,n(k1)}$ or $\beta_{k2,n(k2)}$. If both change, then the random variables are generated assuming an RMS correlation as given above. The correlated random variables are generated using Cholesky decomposition, as described in Wilmott (2006) and elsewhere. If $\beta_{k1,n(k1)}$ and/or $\beta_{k2,n(k2)}$ change, then $\alpha_{k1,n(k1)}$ and/or $\alpha_{k2,n(k2)}$ also change as described in (6) and (7) – in other words, they change in such a way as to avoid there being a discontinuity at the point that the trend changes. These values are used to calculate projected values of $\hat{k}_{1,y}$ and $\hat{k}_{2,y}$, extrapolating the relationship in (2) and (3).

As discussed earlier, $\hat{k}_{1,y}$ and $\hat{k}_{2,y}$ are only estimators, and there is uncertainty in these estimates. This is given by $\epsilon_{k1,y}$ and $\epsilon_{k2,y}$, as described in (4) and (5). The projected results for $\hat{k}_{1,y}$ and $\hat{k}_{2,y}$ are therefore modified by the addition of further random variables, $\epsilon_{k1,y}$ and $\epsilon_{k2,y}$, which are independent and normally distributed with means of zero and standard deviations of 0.012098 and 0.000627 respectively, reflecting the volatility in recent years.

5 Projection Results

Using the method and data above, I carry out 1,000 simulations of $k_{1,y}$ and $k_{2,y}$. I then use these to calculate the period life expectancy of a 60-year old male for the fifty year period from 2006 to 2056. I show the results in Figure 5. Displayed are the median and various percentile limits, together with three sample paths.

This shows that the range of results in early years is relatively narrow, but uncertainty does grow rapidly. To illustrate, the 90% confidence interval – calculated as the difference between the 5th and the 95th centiles – for period life expectancy in 2056 is 18.7 years (37.6 years for the 95th centile less 18.9 years for the 5th centile). Dowd et al. (2008) perform similar calculations to arrive at cohort (rather than period) life expectancies for 65 year old males for the same projection period (2006 to 2056), using a version of the CBD model that allows for cohort effects. If parameter uncertainty is ignored, the 90% confidence interval in 2056 spans 3.6 years; even if parameter uncertainty is allowed for, the range rises to only 7.6 years. In other words, this trend-change model suggests more than twice as much uncertainty as the cohort-adjusted CBD model with parameter uncertainty over a fifty-year time horizon, and the shape of the funnel suggests that the difference in uncertainty continues to increase.

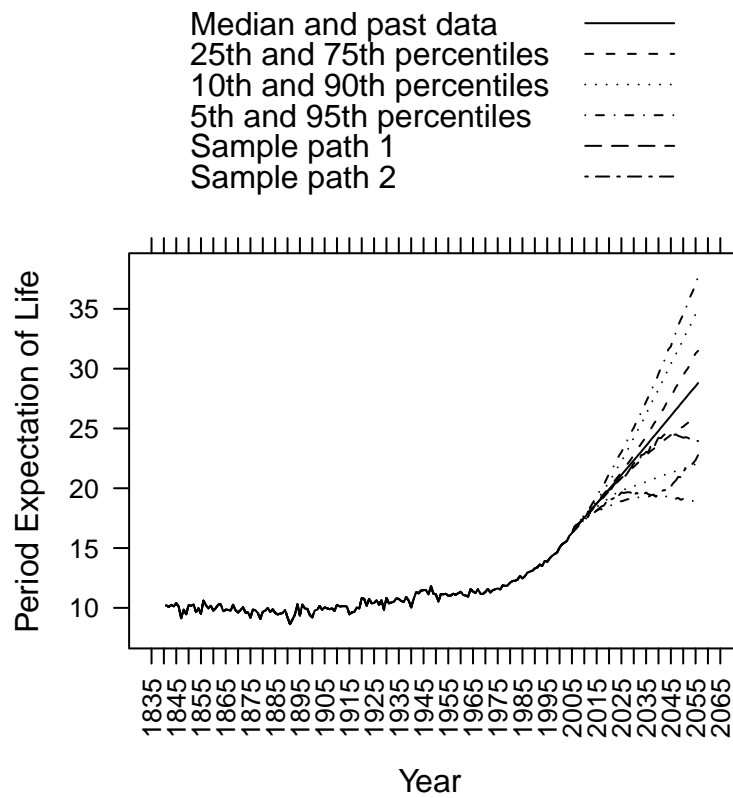
6 Conclusion

If a two parameter model of the type described by Cairns et al. (2006) is used to model mortality, then the parameters follow clear trends that change periodically. The changes in the parameters frequently occur at the same time and are negatively correlated. Within each trend, the parameters do seem to be random fluctuations around a trend rather than random walks with drift. Modelling mortality into the future this way suggests a much greater degree of uncertainty than that implied by random walks with drift.

This has important consequences for risk management and pricing, in particular when considering longer term insurance products such as annuities. The greater uncertainty over mortality rates in the distant future suggests that higher reserves are needed than those suggested by a random walk with drift approach inherent in the original CBD model. The need for higher reserves should also be reflected in pricing.

Another consequence of this pattern of changing mortality rates is that a low level of volatility in mortality rates over a given period should not be interpreted

Figure 5: Historical and Projected Period Life Expectancy for 65 Year Old Males, England and Wales, 1841-2056



Source: Human Mortality Database (2008), Author's Calculations

as a fall in the level of risk – a change in the trend rate of mortality improvement could occur at any point in time, potentially causing a large change in predicted mortality rates.

The potential application of this trend-change approach is not confined to CBD-type models. In particular, it could also be applied to the model proposed by Lee and Carter (1992) as well as those based on this approach. In this way, it would act as an additional extension to those suggested by De Jong and Tickle (2006). As discussed earlier, this paper has several suggestions for dealing with unstable trends, including taking care over the choice of parameterisation period and adding an autoregressive element to the description of the parameters. However, allowing a random change in the trend rates of mortality improvement combined with volatility around the trend rate would give greater long-term variation in mortality rates than any of the approaches suggested by De Jong and Tickle (2006).

In conclusion, then, the approach I propose suggests that there should be much greater allowance for uncertainty in mortality rates in the distant future.

References

- BIFFIS, E.: *Affine Processes for Dynamic Mortality and Actuarial Valuations*, Insurance: Mathematics and Economics, **37**, (2005), 443-468.
- BOOTH, H., J. MAINDONALD AND L. SMITH: *Applying Lee-Carter under conditions of mortality decline*, Population Studies, **56**, (2002), 325-326.
- BROUHNS, N., M. DENUIT AND J. K. VERMUNT: *A Poisson Log-Bilinear Regression Approach to the Construction of Projected Life Tables*, Insurance: Mathematics and Economics, **31**, (2002), 373-393.
- CAIRNS, A. J. G., D. BLAKE AND K. DOWD: *A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration*, Insurance: Mathematics and Economics, **73**, (2006), 687-718.
- CAIRNS, A. J. G., D. BLAKE, K. DOWD, G. D. COUGHLAN, D. EPSTEIN, A. ONG AND I. BALEVICH: *A Quantitative Comparison of Stochastic Mortality Models Using Data from England & Wales and the United States*, North American Actuarial Journal, **13**, (2009), 1-35.
- CUTLER, D., A. DEATON AND A. LLERAS-MUNEY: *The Determinants of Mortality*, The Journal of Economic Perspectives, **20**, (2006), 97-120.
- CHOW, G.C.: *Tests of Equality Between Sets of Coefficients in Two Linear Regressions*, Econometrica, **28**, (1960), 591-605.
- DAHL, M.: *Stochastic Mortality in Life Insurance: Market Reserves and Mortality-Linked Insurance Contracts*, Insurance: Mathematics and Economics, **35**, (2004), 113-136.

- DAHL, M. AND T. MØLLER: *Valuation and Hedging of Life Insurance Liabilities with Systematic Mortality Risk*, in the Proceedings of the 15th International AFIR Colloquium, Zurich, (2005), available online at www.afir2005.ch.
- DE JONG, P. AND L. TICKLE: *Extending the Lee-Carter Model of Mortality Projection*, Mathematical Population Studies, **13**, (2006), 1-18.
- DICKEY, D.A. AND W.A. FULLER: *Distribution of the Estimators for Autoregressive Time Series with a Unit Root*, Journal of the American Statistical Association, **74**,(1979), 427-431.
- DOWD, K., D. BLAKE AND A. J. G. CAIRNS: *Facing Up to Uncertain Life Expectancy: The Longevity Fan Charts*, Pensions Institute Discussion Paper **PI-0703**, (2008), available online at www.pensions-institute.org.
- DURBIN, J. AND G.S. WATSON: *Testing for Serial Correlation in Least Squares Regression, I.*, Biometrika, **37**, (1950), 409-428.
- DURBIN, J. AND G.S. WATSON: *Testing for Serial Correlation in Least Squares Regression, II.*, Biometrika, **38**, (1951), 159-179.
- HUMAN MORTALITY DATABASE: UNIVERSITY OF CALIFORNIA, BERKELEY (USA) AND MAX PLANCK INSTITUTE FOR DEMOGRAPHIC RESEARCH (GERMANY): downloaded data, available at www.mortality.org or www.humanmortality.de, data downloaded on 1 August 2008.
- HUMBLE R. AND B. WILSON: *Drivers of Longevity Projections in Particular with Reference to Smoking*, paper presented to the Staple Inn Actuarial Society, available online at www.sias.org.uk.
- LEE, R. D. AND L. R. CARTER: *Modeling and Forecasting U.S. Mortality*, Journal of the American Statistical Association, **87**, (1992), 659-675.
- MILEVSKY, M. A. AND S. D. PROMISLOW: *Mortality Derivatives and the Option to Annuitize*, Insurance: Mathematics and Economics, **29**, (2001), 299-318.
- MILTERSEN, K. R. AND S.-A. PERSSON: *Is Mortality Dead? Stochastic Forward Force of Mortality Determined by No Arbitrage*, (2005), working paper, University of Bergen.
- RENSHAW, A. E. AND S. HABERMAN: *Lee-Carter Mortality Forecasting with Age-specific Enhancement*, Insurance: Mathematics and Economics, **33**, (2003), 255-272.
- RENSHAW, A. E., AND S. HABERMAN *A Cohort-Based Extension to the Lee-Carter Model for Mortality Reduction Factors*, Insurance: Mathematics and Economics, **38**, (2006), 556-570.
- SCHRAGER, D. F.: *Affine Stochastic Mortality*, Insurance: Mathematics and Economics, **38**, (2006), 81-97.

WILMOTT, P.: *Paul Wilmott on Quantitative Finance, 2nd Edition*, (2006),
John Wiley & Sons.