



Pensions
Institute

DISCUSSION PAPER PI-1205

A Piecewise Linear Cohort Extension to the Cairns-Blake-Dowd Model

Paul Sweeting & Yiwei Zhao

November 2016

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University London
106 Bunhill Row
London EC1Y 8TZ
UNITED KINGDOM

<http://www.pensions-institute.org/>

A Piecewise Linear Cohort Extension to the Cairns-Blake-Dowd Model

Paul Sweeting (University of Kent)
Yiwei Zhao (University of Kent, Scottish Widows)

ABSTRACT

Age-Period-Cohort (“APC”) models have been criticised on a number of grounds. One area of concern is in relation to projecting future cohorts. However, we would argue that such projection is unnecessary in some key cases, such as for closed defined benefit pension schemes.

More fundamental issues relate to the fit itself. APC models typically use at least one parameter for each cohort, in addition to those used for parameters age and period. This leads to a large number of parameters which are not necessarily independent.

However, the model we propose here uses a potentially far smaller number of parameters that essentially describe times where a new type of cohort emerges. This is similar to the trend-change models of mortality improvement discussed by as described by Sweeting (2011), Coelho and Nunes (2011), and van Berkum et al (2014). Because this cohort approach identifies a small number of changes in cohort rather than imposing a new cohort parameter for each year of birth, this reduces the risk of interdependence.

1. INTRODUCTION

Despite fitting mortality data well (Cairns et al 2009), Age-Period-Cohort (“APC”) models have been criticised on a number of grounds. The first area of concern is in relation to projecting future cohorts.

As Booth and Tickle (2008) remark: “The APC model has been usefully applied in describing the past, but has been considered less useful in forecasting”. This is because it is difficult to find any discernible pattern for cohort factors, meaning that the projection of the adjustment for future cohorts is no straightforward.

However, for a closed portfolio of lives – such as that which would be found in a pension scheme closed to future accrual or even just to new members – there would be no future cohorts. As such, the projection of future cohorts would not be necessary. Such schemes are increasingly common. According to the Pension Protection Fund (2015), 51% of defined benefit schemes are closed to new members whilst a further 34% are closed to future accrual. In both cases, the population is already known, and no future cohorts will be added.

A more fundamental issue with APC models relates to the fit itself. In a typical APC model, the mortality rate at each age is determined by a combination of parameters relating to the age, period and cohort of rate in question. Currie (2012) computes the canonical correlations between the estimates for the three variables of the APC model and found for the data in that paper and

finds that all lie between 0.4 and 0.7, significantly different from zero, quantifying concerns around APC models first raised by Clayton and Schifflers (1987).

Furthermore, the number of variables in most APC models significantly increases the risk of over-fitting. In Cairns et al (2009), the number of variables increases from 102 to 144 when the Lee-Carter model (Lee and Carter, 1992) is expanded by Currie (2006) to allow for a cohort effect, and to 203 if the Renshaw and Haberman (2006) alternative is used. Similarly, the Cairns-Blake-Dowd or CBD model (Cairns et al, 2006) has only 88 variables, but adding a cohort effect increases this to 159.

We believe that this over-parameterisation can be avoided by viewing cohort effects differently. Starting with the CBD model, we propose an extension that views the logit-mortality surface as being described adequately by age and period factors except for those instances when a cohort “event” occurs. This then change the direction of the cohort’s mortality until another “event” occurs. Because the location of these events is carried out independently from the fitting of the model, this process also removes one element of the interdependence between the age, period and cohort factors. As such, it offers a parsimonious and robust approach to APC modelling. We call the new model a Piecewise Linear Cohort-based (PLC) mortality model.

2. THE COHORT EFFECT

The last quarter of a century has seen great improvements in mortality modelling, from the Lee-Carter model (Lee and Carter 1992) and its extensions (Brouhns et al, 2002; Renshaw and Haberman, 2003a, b, c, 2006; Continuous Mortality Investigation Bureau, 2005) to the CBD model (Cairns et al, 2006) and its extensions (Cairns, et al, 2009). Cairns, et al (2009) comprehensively compare and rank these models in terms of the Bayesian Information Criterion (BIC).

Cairns et al (2009) consider three families of model. The first is the Lee-Carter – or L-C – model family. The L-C model, which is a single-factor model that assumes no smoothness across ages or years; the second is the P-spline model which uses penalised basis splines to impose smoothness across years and ages; and the third is the CBD model family which assumes smoothness across ages in same year, but makes no assumption of smoothness between different years. It is the third family that we use in this paper.

According to Cairns, et al. (2009), the cohort factor representing the effect of the birth year improves the fit of mortality models considerably. This is due to the fact that it takes into account the cohort effect.

The starting point for this analysis is the CBD model, given in Cairns et al (2006), and referred to as model M5 in Cairns et al (2009). This is defined in Equation 1.

$$(Eq. 1) \quad \text{logit } q(t, x) = \kappa_t^{(1)} + \kappa_t^{(2)}(x - \bar{x}) + \varepsilon_{t,x}$$

where $q(t, x)$ is the initial mortality rate for a life aged x in year t ; \bar{x} is the average of the ages x ; $\kappa_t^{(1)}$ denotes the intercept parameter in year t ; $\kappa_t^{(2)}$ denotes the slope parameter in year t ; and $\varepsilon_{t,x}$ is an error term. The function $\text{logit } q(t, x)$ is calculated as:

$$(Eq. 2) \quad \text{logit } q(t, x) = \frac{q(t, x)}{1 - q(t, x)}$$

The logit function is used because the initial mortality rate, $q(t, x)$, cannot be allowed to exceed one – the initial population must always be at least as great as the number of deaths in the period. As such, although $\text{logit } q(t, x)$ can take values from zero to infinity, its inverse is bounded by zero and one.

Model M5 is extended by Cairns et al. (2009) with the addition of a series of cohort factors, γ_{t-x} , to give model M6, shown in Equation 3:

$$(Eq. 3) \quad \text{logit } q(t, x) = \kappa_t^{(1)} + \kappa_t^{(2)}(x - \bar{x}) + \gamma_{t-x}^{(3)} + \varepsilon_{t,x}$$

In all our analysis, we fit data from the Human Mortality Database (HMD) for England and Wales using $t = 1961, 1962, \dots, 2013$ and $x = 60, 61, \dots, 89$. Cairns et al (2009) use the same age range, but their observations cease in 2004; we are able to extend our analysis to 2013. Following Cairns et al, we exclude cohorts with fewer than five observations. We then convert the initial rate of mortality, $q(t, x)$, to the central rate, $m(t, x)$, using the relation in Equation 4:

$$(Eq. 4) \quad m(t, x) = -\ln[1 - q(t, x)]$$

The models are fitted by assuming that deaths $D(t, x)$ follow a Poisson distribution with a mean of $E(t, x)m(t, x)$, as described by Brouhns et al (2002). This means that the likelihood, L , is given by:

$$(Eq. 5) \quad L = \prod_{t,x} \frac{e^{-E(t,x)\hat{m}(t,x)}(E(t,x)\hat{m}(t,x))^{D(t,x)}}{D(t,x)!},$$


and the log likelihood, $\ln L$, is given by:

$$(Eq. 6) \quad \ln L = \sum_{t,x} [D(t, x) \ln E(t, x)\hat{m}(t, x) - E(t, x)\hat{m}(t, x) - \ln D(t, x)!],$$

where $\hat{m}(t, x)$ is the estimated central mortality rate derived from the mortality model.


The results for our fit of M5 are shown in Figure 1, whilst our fit of M6 is shown in Figure 2. In Figure 1. The first panel row shows the intercept parameter $\kappa_t^{(1)}$ whilst the second shows the slope parameter $\kappa_t^{(2)}$. In Figure 2, these two parameters are shown again in the first row, with the cohort parameter, γ_{t-x} , shown in the second row.

Figure 1: Estimated values of parameters $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$ in M5 (Equation 1), England and Wales, males, age 60-89, year 1961-2013



Source: Human Mortality Database; authors' calculations

Figure 2: Estimated values of parameters $\kappa_t^{(1)}$, $\kappa_t^{(2)}$ and γ_{t-x} in M6 (Equation 3), England and Wales, males, age 60-89, year 1961-2013



Source: Human Mortality Database; authors' calculations

The second row in Figure 2 shows the impact of the cohort effect. As mentioned by Willets (2004) and Renshaw and Haberman (2006), several noticeable features should be taken into account, the most noteworthy of which is the steep relative fall in mortality rates for those born from the mid 1920's. It is also interesting to note the discontinuity in γ_{t-x} in 1919 and 1920. There is a similar but much smaller discontinuity in 1947-9. The fact that these discontinuities are around the ends of the two world wars suggests that the end of the world wars coincided with data anomalies. This is important, as it suggests that any parameterisations involving changes around these points should be viewed with suspicion. Rather, it suggests that an M6-like adjustment may be appropriate for anomalous cohorts only. We investigate this later.

Renshaw and Haberman (2006) point out that there is another discontinuity in 1887, which can be traced to a set of outliers, and that this is possibly due to mis-stated exposures for this particular cohort. Because this is near to the start of our dataset, it is difficult to see that anomaly here. More obvious is the move from cohorts with declining mortality to improving mortality in 1900. A detailed analysis of the cohort effect in the UK is given by Willets (2004).

The first two parameters in Figure 2 give the trend of mortality improvements ($\kappa_t^{(1)}$) and the impact of these improvements at various ages ($\kappa_t^{(2)}$). As Cairns et al (2006) note for the basic CBD model, these parameters suggest that mortality rates have been falling – that is, improving – over the entire period, and that rates at higher ages are decreasing at a lower rate.

3. PROPOSED MODELS

In each calendar year, mortality rates can be fitted as a smooth linear function using the CBD model. However, because of the cohort effect, there are diagonal breaks across the entire surface. As discussed above, model M6 deals with this by adding a separate cohort factor to the logit of mortality rates for each year of birth (Eq 3). However, an alternative is to use two or more linear functions to fit logit q_{tx} in each calendar year, with the break occurring between the same cohort or cohorts across the years. The principal here is similar to that used to model changes in the period parameter, as described by Sweeting (2011), Coelho and Nunes (2011), and van Berkum et al (2014). To do this, we need to know the cohorts for which there is a break. We do this by finding each break in turn using a sup-LR approach, motivated by Andrews (1993, 2003). The sup-LR test identifies a break by identifying the point where a break is most statistically significant. It does through a likelihood ratio test. The test statistic here is $-2 \ln(L_0/L_{1,t-x})$, where $L_{1,t-x}$ is the PLC model with a break at year of birth $t-x$, and L_0 is model M5. The test statistic has a chi-squared distribution where the number of degree of freedom is the difference between the number of parameters used. Since L_0 is the same in each case, the breakpoint can be found by looking for the highest value of $\ln(L_{1,t-x})$. This does assume that the number of parameters is the same for all values of $\ln(L_{1,t-x})$, which is true only for the middle cohorts; however, the difference in the values of the log likelihood function is such that the order of significance would not change.


Having found the first break point, the procedure is then repeated with that breakpoint already in place to find a second breakpoint, and then a third. At this point, the number of variables is

similar to that under model M6. We take this to be the maximum number of desirable break points for the dataset that we have.

An attraction of the sup-LR approach is that it gives an indication of the relative importance of the breaks. However, there is always the possibility that other combinations of breaks – combinations that might not be found using the above approach – might give a higher log likelihood. We therefore also implement an Iterative Grid Search (IGS) method (Thisted, 1988) for the case where we have three breaks. In our implementation of this approach, we first look for the combination of dates giving the highest log likelihood with the dates being spaced by sixteen years. Within the cube of years specified by the sixteen years either side of this point for each break, we then look for the highest log likelihood with the dates being spaced by four years. The cube is then re-defined and the log likelihood-maximising combination of years is determined.

The mechanism of the breakpoints can be shown as follows. Supposing there is a single kink in the logit of mortality rates in a particular year t occurring for lives born before and after year c . Suppose also that this occurs in the years when the people who were born exactly in year c were between ages 60 to 89 (Figure 3). In this case, two linear functions can be used to fit the logit of mortality rates, instead of the one linear function in model M5. The model can be fitted using the Poisson maximum likelihood approach described by Brouhns et al (2002).

Figure 3: Method of modelling logit mortality using two linear functions and a continuity constraint



Source: authors

We assume that there is a constraint that the two linear functions connect to each other. The constraint ensures that there is no first order discontinuity between the two linear functions, and also reduces the number of parameters required. This model (denoted as PLC1, where) can be described as:

$$(Eq. 4) \quad \text{logit } q(t, x) = \begin{cases} \kappa_t^{(11)} + \kappa_t^{(12)}(x - \bar{x}) + \varepsilon_{tx}, & t - x \geq c \\ \kappa_t^{(21)} + \kappa_t^{(22)}(x - \bar{x}) + \varepsilon_{tx}, & t - x \leq c \end{cases}$$


$$\text{Constraint: } \kappa_t^{(11)} + \kappa_t^{(12)}(t - c - \bar{x}) = \kappa_t^{(21)} + \kappa_t^{(22)}(t - c - \bar{x}),$$

where year c is a kink point of birth year; $\kappa_t^{(i1)}$, $i = 1, 2$ denotes intercept factor in year t : $i = 1$, fits the equation for ages less than or equal to $t + c$, while $i = 2$ fits the equation to those whose

ages are more than or equal to $t - c$; $\kappa_t^{(i2)}$, $i = 1, 2$ denotes slope factor in year t , with i having the same meaning as before; and ε_{tx} is an error term. The term c remains constant across the years investigated. Once a value for c has been found that maximizes the log likelihood, the cohort birth year is fixed at year $t - c$.

PLC Model Two (PLC2) allows for two kink points rather than one, which improves the goodness-of-fit. The mechanism using two kink points is similar to PLC1, breaking the original smoothly linear function of CBD model into a piecewise linear function with the two kink points for given birth years as shown in Figure 4. This shows an instance when both kink points are located between age 60 and 89, although the kink points can also be outside the span of ages shown here. This model is also fitted under the assumption that the number of deaths follows a Poisson distribution.

Figure 4: Method of modelling logit mortality using three linear functions and a continuity constraint



Source: authors

This model (denoted as PLC2) is described as:

$$(Eq. 5) \quad \text{logit } q(t, x) = \begin{cases} \kappa_t^{(11)} + \kappa_t^{(12)}(x - \bar{x}) + \varepsilon_{tx}, & t - x \geq c_2 \\ \kappa_t^{(21)} + \kappa_t^{(22)}(x - \bar{x}) + \varepsilon_{tx}, & c_1 \leq t - x \leq c_2 \\ \kappa_t^{(31)} + \kappa_t^{(32)}(x - \bar{x}) + \varepsilon_{tx}, & t - x \leq c_1 \end{cases}$$

Constraints:

$$\kappa_t^{(11)} + \kappa_t^{(12)}(t - c_2 - \bar{x}) = \kappa_t^{(21)} + \kappa_t^{(22)}(t - c_2 - \bar{x})$$

$$\kappa_t^{(21)} + \kappa_t^{(22)}(t - c_1 - \bar{x}) = \kappa_t^{(31)} + \kappa_t^{(32)}(t - c_1 - \bar{x})$$

where, year c_1 and c_2 are two kink points of birth year, and $c_1 < c_2$; $\kappa_t^{(i1)}$, $i = 1, 2, 3$ denotes intercept factor in year t : $i = 1$ fits the equation for ages less than and equal to $t - c_2$, $i = 2$ fits the equation for ages greater than or equal to $t - c_2$ but less than or equal to $t - c_1$, and $i = 3$

fits the equation for ages greater than or equal to $t - c_1$; $\kappa_t^{(i2)}$, $i = 1, 2, 3$ denotes slope factor in year t , with i having the same meaning as before; ε_{tx} is an error term.

By the same reasoning, more kink points can be added. Therefore, a comprehensive model with continuity constraints can be given by:

$$(Eq. 6) \quad \text{logit } q(t, x) = \begin{cases} \kappa_t^{(11)} + \kappa_t^{(12)}(x - \bar{x}) + \varepsilon_{tx}, & t - x \geq c_n \\ \kappa_t^{(21)} + \kappa_t^{(22)}(x - \bar{x}) + \varepsilon_{tx}, & c_{n-1} \leq t - x \leq c_n \\ \dots & \dots \\ \kappa_t^{((n+1)1)} + \kappa_t^{((n+1)2)}(x - \bar{x}) + \varepsilon_{tx}, & t - x \leq c_1 \end{cases}$$

Constraints:

$$\begin{aligned} \kappa_t^{(11)} + \kappa_t^{(12)}(t - c_n - \bar{x}) &= \kappa_t^{(21)} + \kappa_t^{(22)}(t - c_n - \bar{x}) \\ &\dots \\ \kappa_t^{(n1)} + \kappa_t^{(n2)}(t - c_1 - \bar{x}) &= \kappa_t^{((n+1)1)} + \kappa_t^{((n+1)2)}(t - c_1 - \bar{x}) \end{aligned}$$


where the kink points are c_i , $i = 1, 2, \dots, n$; $c_1 < c_2 < \dots < c_{n-1} < c_n$.

We denote a PLC model as $PLCk(c_1, c_2, \dots)$, in which k represents the number of kinks used in estimation, and (c_1, c_2, \dots) are cohort years (cohort kinks).

4. EVALUATION

As for the analysis of M5 and M6, we fit data from the Human Mortality Database (HMD) for England and Wales using ages 60 to 89 and years 1961 to 2013. Following Cairns et al, we exclude cohorts with fewer than five observations. Again, we calculate the central mortality rate, $m(t, x)$, from the initial mortality rate given by the model, and fit the model using the Poisson approach of Brouhns et al (2002). Using a sup-LR approach, we calculate the year of birth for which the first break is likely. Including this break we then look for the second and then the third. The traces of the log likelihoods are shown in Figure 5, together with those for models M5 and M6. Summary information for the fitted models is then given in Table 1. This includes the Bayesian Information Criterion (BIC) and Akaike Information Criterion (AIC). Then, in Figure 6, I give $\kappa_t^{(i1)}$ and $\kappa_t^{(i2)}$ for $i = 1$ to 3 before showing the fit of these functions for the calendar years 1985 in Figure 7. We also use IGS to derive the break points, and this approach gives the same results. This is likely to be because the number of breaks is small, and because the final break is between the other two.

Figure 5: Log likelihood for model M5, PLC1, PLC2 and PLC3, England and Wales, males, age 60-89, year 1961-2013




Source: Human Mortality Database; authors' calculations

Table 1: Summary data for models M5, M6, PLC1, PLC2 and PLC3, England and Wales, males, age 60-89, year 1961-2013

Model	Parameters	Log likelihood	AIC	BIC	Break Years
M5	106	-18,454.0	37,119.9	37,688.0	N/A
M6	180	-9,999.1	20,358.1	21,322.7	N/A
PLC1	136	-16,401.0	33,074.1	33,802.9	1900
PLC2	158	-15,109.0	30,533.9	31,380.6	1900, 1932
PLC3	188	-14,601.0	29,578.0	30,585.5	1900, 1920, 1932

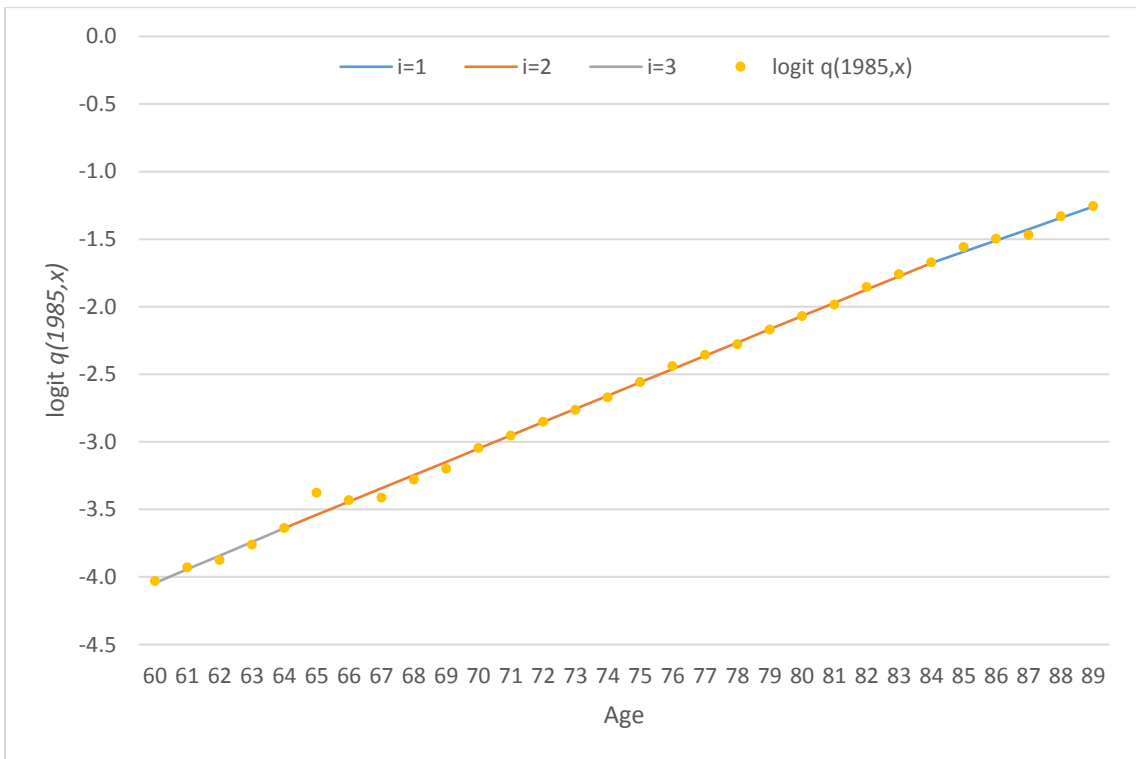
Source: Human Mortality Database; authors' calculations

Figure 6: Estimated values of parameters $\kappa_t^{(i1)}$ and $\kappa_t^{(i2)}$ for $i = 1$ to 3 in PLC models, England and Wales, males, age 60-89, year 1961-2013



Source: Human Mortality Database; authors' calculations

Figure 7: Logit $q(1985,x)$ for PLC3, England and Wales, males, age 60-89




Source: Human Mortality Database; authors' calculations

As Table 1 shows, each successive version of PLC model is a significant improvement on model M5 according to both the AIC and the BIC. However, by the time there are three breaks, the number of parameters exceeds that of model M6, which still clearly outperforms all PLC models.

A key reason for the underperformance of the PLC models relative to model M6 can be seen in Figure 3, and that is the outlier cohorts of 1919 and 1920. This can also be seen as the data point for age 65 in 1985 in Figure 7. In order to allow for this distortion, we use a hybrid model. This involves fitting model M6 to the data, but using only the cohort parameters for years of birth 1919 and 1920. These parameters are then subtracted from the original data before fitting the PLC models once again. These models are denoted PLC(h) models, the “h” standing for hybrid. Results of this analysis are shown in Figure 8 and Table 2. In Figure 9, I give κ_t^{i1} and κ_t^{i1} for $i = 1$ to 3 before showing the fit of these functions for the calendar years 1985 in Figure 10. IGS again yields the same results.

Figure 8: Log likelihood for model M5, PLC(h)1, PLC(h)2 and PLC(h)3, England and Wales, males, age 60-89, year 1961-2013



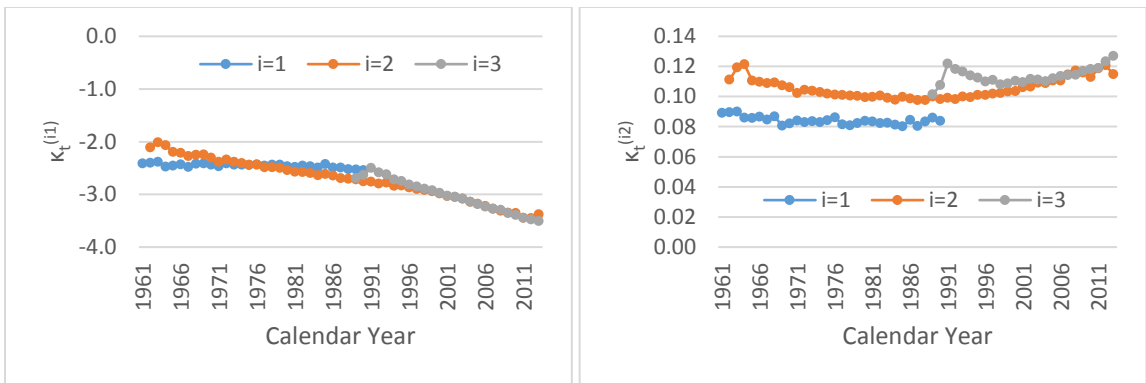
Source: Human Mortality Database; authors' calculations

Table 2: Summary data for models M5, M6, PLC(h)1, PLC(h)2 and PLC(h)3, England and Wales, males, age 60-89, year 1961-2013

Model	Parameters	Log likelihood	AIC	BIC	Break Years
M5	106	-18,454.0	37,119.9	37,688.0	N/A
M6	180	-9,999.1	20,358.1	21,322.7	N/A
PLC1	138	-11,793.5	23,862.9	24,602.4	1900
PLC(h)2	160	-10,805.6	21,931.1	22,788.5	1900, 1934
PLC(h)3	187	-10,664.5	21,703.0	22,705.1	1900, 1927, 1934


Source: Human Mortality Database; authors' calculations

Figure 9: Estimated values of parameters $\kappa_t^{(i1)}$ and $\kappa_t^{(i2)}$ for $i = 1$ to 3 in PLC(h) models, England and Wales, males, age 60-89, year 1961-2013



Source: Human Mortality Database; authors' calculations

Figure 10: Logit $q(1985,x)$ for PLC3, England and Wales, males, age 60-89



Source: Human Mortality Database; authors' calculations

As Table 2 shows, controlling for the two outlier cohorts not only significantly improves the fit of the model – for an additional penalty of only two parameters – but it also changes the position of the break years slightly in PLC(h)2 and PLC(h)3. In terms of the fit, PLC(h)3 and even PLC(h)2 are close in performance to M6 as measured by the AIC and BIC. However, because of the way the PLC and PLC(h) models deal with cohorts, the models remain more parsimonious than M6.

5. CONCLUSIONS

Starting with the CBD model M5, the PLC and PLC(h) models allow the inclusion of cohort effects with in a more parsimonious way than model M6. The PLC(h) model includes two additional parameters to deal with the anomalous cohorts seen for birth years 1919 and 1920. These two parameters significantly improve the fit of the model. More important than the fit, the PLC and PLC(h) models better reflect the way in which mortality rates develop over time than models which use a separate parameter or parameters for each year of birth.

REFERENCE

- Andrews, D. (1993). *Tests for Parameter Instability and Structural Change with Unknown Change Point*. *Econometrica*, 61(4), 821-856.
- Andrews, D. (2003). *Tests for Parameter Instability and Structural Change with Unknown Change Point: A Corrigendum*. *Econometrica*, 71(1), 395-397.
- van Berkum, F., K. Antonio, and M. Vellekoop (2014). *The impact of multiple structural changes on mortality predictions*. *Scandinavian Actuarial Journal*, 16, 1-23.
- Brouhns, N., M. Denuit and J.K. Vermunt (2002). *A Poisson Log-Bilinear Regression Approach to the Construction of Projected Life Tables*. *Insurance: Mathematics and Economics*, 31: 373–93.
- Cairns, A.J.G., D. Blake and K. Dowd (2006). *A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration*. *Journal of Risk and Insurance*, 73: 687–718.
- Cairns, A.J.G., D. Blake, K. Dowd, G.D. Coughlan, D. Epstein, A. Ong and I. Balevich (2009). *A Quantitative Comparison of Stochastic Mortality Models Using Data from England and Wales and the United States*, *North American Actuarial Journal*, 13(1), 1-35.
- Continuous Mortality Investigation (CMI) (2005). *Projecting Future Mortality: Towards a Proposal for a Stochastic Methodology*. Working paper 15.
- Coelho, E. and L. Nunes (2011), *Forecasting mortality in the event of a structural change*, *Journal of the Royal Statistical Society* 174(3), 713-736.
- Clayton, D. and E. Schifflers (1987). *Models for temporal variation in cancer rates. II: age-period-cohort models*. *Statistics in Medicine*, 6, 469-481.
- Currie, I.D. (2006). *Smoothing and Forecasting Mortality Rates with P-Splines*. Paper given at the Institute of Actuaries, June 2006.
- Currie, I.D. (2012). *Forecasting with the age-period-cohort model?* Proceedings of 27th International Workshop on Statistical Modelling, Prague, 87-92.
- Lee, R.D. and L.R. Carter (1992). *Modeling and Forecasting U.S. Mortality*. *Journal of the American Statistical Association*, 87: 659–75.
- Renshaw, A.E. and S. Haberman (2003a). *Lee-Carter mortality forecasting with age-specific enhancement*. *Insurance: Mathematics and Economics*, 33, 255-272.
- Renshaw, A.E. and S. Haberman (2003b). *Lee-Carter mortality forecasting: a parallel generalized linear modelling approach for England and Wales mortality projections*. *Applied Statistics*, 51(1), 119-137.
- Renshaw, A.E. and S. Haberman (2003c). *On the forecasting of mortality reduction factors*. *Insurance: Mathematics and Economics*, 32, 379-401.
- Renshaw, A.E. and S. Haberman (2006). *A cohort-based extension of the Lee-Carter model for mortality reduction factors*. *Insurance: Mathematics and Economics*, 38, 556-570.

Sweeting, P.J. (2011). *A trend-change extension of the Cairns-Blake-Dowd model*. Annals of Actuarial Science 5(2) , 143-162.

Thisted, R.A. (1988). Elements of Statistical Computing. Chapman and Hall, London.

Willets, R.C. (2004). *The cohort effect: insights and explanations*. British Actuarial Journal, 10(4), pp.833-877.