

Durham Research Online

Deposited in DRO:

03 January 2018

Version of attached file:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Dowd, K. and Cairns, A.J.G. and Blake, D. (2019) 'Hedging annuity risks with the age-period-cohort two-population gravity model.', *The North American actuarial journal.*, 25 (sup1). S170-S181.

Further information on publisher's website:

<https://doi.org/10.1080/10920277.2019.1652102>

Publisher's copyright statement:

This is an Accepted Manuscript of an article published by Taylor Francis in *The North American actuarial journal* on 20 November 2019 available online: <http://www.tandfonline.com/10.1080/10920277.2019.1652102>

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

Hedging Pension Risks with the Age-Period-Cohort Two-Population Gravity Model

Kevin Dowd^{*}, Andrew J.G. Cairns[#] and David Blake^{*},

This draft: February 2nd 2017

Abstract

We consider the effectiveness of an illustrative annuity hedging problem in which a forward annuity predicated on one population is hedged by a position in a forward annuity predicated on another population. Our analysis makes use of the age-period-cohort two-population gravity model that takes account of the observed inter-dependence between the two populations' mortality rates; it also considers the implications of parameter uncertainty, individual death or Poisson risk and interest-rate risk for hedge effectiveness. We consider horizons of up to 20 years. For the most part, our results are robust and indicate strong hedge effectiveness, with estimates of relative risk reduction varying from about 0.70 in the least effective case to well over 0.95 in the most effective cases.

^{*} Pensions Institute, Cass Business School, 106 Bunhill Row, London, EC1Y 8TZ, United Kingdom. Email: kevin.dowd@hotmail.co.uk. Corresponding author.

[#] Maxwell Institute for Mathematical Sciences, and Actuarial Mathematics and Statistics, Heriot-Watt University, Edinburgh, EH14 4AS, United Kingdom.

^{*} Pensions Institute, Cass Business School, 106 Bunhill Row, London, EC1Y 8TZ, United Kingdom.

Hedging Pension Risks with the Age-Period-Cohort Two-Population Gravity Model

1. Introduction

Consider the problem of how to hedge an annuity predicated on a cohort from one population (e.g., that to which a pension fund is exposed) with an annuity predicated on another population (e.g., the national population). The annuity being hedged represents a typical pension fund's exposure and the hedge instrument is a second annuity offered by a capital markets institution. Equivalently, the hedge can also be regarded as a deferred annuity swap.

More specifically, the pension fund is assumed to have a deferred annuity exposure at horizon T , and its annuitant population is considered to be a representative large sample drawn from the UK Continuous Mortality Investigation (CMI) male assured lives population. The pension fund seeks to hedge this exposure with a deferred annuity of the same horizon predicated on the England & Wales (E&W) male mortality index (i.e., on the national population of England & Wales). Both annuities involve annual payments beginning at future time T of £1 to each surviving member of a cohort who will be aged 65 at T . Using obvious notation, we denote these annuities as $a(T, CMI, 65)$ and $a(T, E \& W, 65)$ respectively.

We know that the mortality rates of the two populations are highly, but not perfectly, correlated. This suggests that a hedge might be useful in reducing the pension fund's risk exposure, but is unlikely to be perfect: some basis risk is likely to remain.

Our problem is to design and assess an appropriate hedge position with this hedge instrument. To do this, we use the framework for basis risk analysis and hedge effectiveness laid out in Coughlan *et al.* (2011). The framework has three components: (i) the development of an informed understanding of basis risk, (ii) the appropriate calibration of the hedging instrument and (iii) an evaluation of hedge effectiveness.

In applying this framework, this article is organized as follows. Section 2 discusses why basis risk and hedge calibration requires a two-population mortality model that takes account of the two populations' interdependence. Section 3 discusses basis risk and hedge effectiveness using the Age-Period-Cohort (APC) version of the two-population gravity model of Dowd *et al.* (2010)^{1,2}: we initially use the simplest version of the model which disregards parameter uncertainty and Poisson death risk.³ Section 4 then generalizes the analysis to cover these

¹ Alternatively, we could have used the MCMC two-population model of Cairns *et al.* (2011a).

² Our work also complements the earlier analyses of Coughlan *et al.* (2011) and Cairns *et al.* (2011b), both of which have considered somewhat similar hedging problems to those considered here.

³ By Poisson death risk, we mean the random variation risk associated with death rates in a small sample of the overall population. If the sample is small, as will be the case in a small pension fund, the realized pattern of


cases. Section 5 further generalizes the analysis to cover interest-rate risk. Section 6 offers some conclusions and suggestions for extensions of the present work.

Appendices *A-D* deal with various technical issues: *A* outlines the two-population gravity model used in this paper, *B* discusses parameter simulation, *C* discusses the simulation of Poisson deaths uncertainty and *D* discusses the simulation of future annuity prices.

2. Why We Need a Two-Population Mortality Model

To see why we need a two-population mortality model for evaluating hedge effectiveness, first consider the result of using two independent mortality models: one for each population. To be more precise, let us assume that we wish to simulate mortality q rates using the APC model applied to each population independently. Assume also for the time being that the parameters of the mortality model are certain and there is no Poisson death risk.

Figure 1: Fan Chart Projections of q Rates for Age 65: Populations Treated Independently



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials. The model used is the ‘parameters certain’ version of the APC model with no allowance for individual deaths Poisson risk.

deaths could differ significantly from that of the national population. Since we assume that actual deaths follow a Poisson distribution, we will call the random variation risk Poisson risk.

Fan chart projections of the two populations' mortality or q rates for age 65 are given in Figure 1. Both mortality rates are projected to fall sharply, with the E&W mortality rates falling somewhat more rapidly than the CMI ones; the fan charts are also relatively narrow. At first sight, this looks promising as a way of capturing realistic future behaviour for the q rates for both populations.

Following Coughlan *et al.* (2011) and Cairns *et al.* (2011b), we now consider a hedging strategy in which the size of the hedge ratio (in terms of units of the E&W forward annuity) is given by the standard formula:⁴

$$h(T) = -\rho(T) \frac{\sigma_{CMI}(T)}{\sigma_{E\&W}(T)} \quad (1)$$

where $\rho(T)$ is the forward correlation between the projected values of the two annuities at horizon T and $\sigma_{CMI}(T)$ and $\sigma_{E\&W}(T)$ are the standard deviations of the projected values of the CMI and E&W annuities at horizon T .

Thus, the unhedged and hedged positions are given by (2a) and (2b) below:

$$a(T, CMI, 65) \quad (2a)$$

$$a(T, CMI, 65) + h(T) \times a(T, E \& W, 65) \quad (2b)$$

The effectiveness of the hedge can be assessed using the following formula for relative risk reduction (RRR):

$$RRR = 1 - \frac{R_{hedged}}{R_{unhedged}} \quad (3)$$

where R_{hedged} and $R_{unhedged}$ are appropriate dispersion-based risk measures for the hedged position (i.e., original liability plus hedge) and unhedged position (i.e., original liability), respectively.⁵ By this criterion, a perfect hedge would have an RRR equal to 1 and a good hedge will have an RRR 'close' to 1; an RRR equal to 0 indicates a completely ineffective hedge and a negative RRR indicates a 'hedge' that is worse than useless because it increases overall risk exposure.

The problem with modelling the two populations independently then becomes apparent: the forward correlation $\rho(T)$ is zero, and the optimal size of the hedge position is itself also

⁴ See Coughlan *et al.* (2004) for more details of optimizing hedge effectiveness in a general setting and Coughlan *et al.* (2007a,b) for the first applications to longevity hedging.

⁵ We take these to be the 95% Expected Shortfalls (ES) relative to the median, but the standard deviation or the VaR relative to the median would also do and appear to give much the same results, presumably because of the underlying assumption of bivariate Gaussianity. Similar results also hold for different prediction intervals.

zero! Very simply, working with two independent models is, at best,⁶ useless for hedging analysis; instead, we need a two-population model that takes account of how the two populations are related.

3. Basis Risk Analysis and Hedge Effectiveness Using a Two-Population Mortality Model


We now assume that mortality rates are generated by a two-population model that takes account of how the two populations are related: such a model makes intuitive sense because in practice the two populations will have common mortality drivers and therefore their evolution will be related (see Coughlan *et al.* 2011). The particular model chosen is a gravity two-population generalisation of the APC model (Dowd *et al.*, 2011b). (For further details on this model, see Appendix A.)

Figure 2 gives fan chart projections of the two populations' q rates for age 65 using this two-population model calibrated to actual historical data. At first sight, these projections are very close to those of Figure 1: the only notable difference is that the CMI fan charts are a little narrower than they were before.

Figure 3 shows the forward correlations between the E&W and CMI q rates for age 65 for horizons of up to 20 years, as generated by the parameter certain (PC) version of the gravity two-population model without taking account of individual death or Poisson risk: unlike the previous case where correlations were zero, these correlations rise from a little under 0.6 for a 1-year horizon to a little over 0.9 for a horizon of 20 years.


⁶ We say “at best” because simulation exercises suggest that the optimal hedge ratio is indeed zero: other hedge ratios appear to produce RRRs that are negative. This also makes intuitive sense: adding a substantial position to the first position which is uncorrelated with the first position merely adds noise.

Figure 2: Fan Chart Projections of the APC Two-Population Gravity Model q Rates for Age 65



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials. The model used is the 'parameters certain' version with no allowance for individual deaths Poisson risk.

Figure 3: q Forward Correlations for Age 65: APC Two-Population Gravity Model




Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials. The model used is the 'parameters certain' version of the gravity two-population model with no allowance for individual deaths Poisson risk.

The effectiveness of the hedge is shown in Figure 4. The red curve represents the probability density of possible future values of the original unhedged position, and the much narrower blue curve represents the distribution of possible future values of the hedged position. The narrowness of the latter relative to the former indicates that the hedge is highly effective, and this is borne out by the RRR which is 0.854: the hedge reduces the pension fund's exposure by just over 85%.

This is (to say the least!) a great improvement in comparison with the case where the two populations are treated as independent of each other and confirms the appropriateness of using a two-population model for hedge effectiveness analysis.

Figure 4: Hedged and Unhedged Annuity Positions for Age 65, $T=10$: APC Two-Population Gravity Model




Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials and an assumed constant interest rate of 0.04. The model used is the ‘parameters certain’ version of the model with no allowance for individual deaths Poisson risk. ‘RRR’ is relative risk reduction.

We are also interested in the effectiveness of hedges over a range of possible horizons. Accordingly, Figure 5 shows the fan chart projections of the future values of the hedged and unhedged positions:⁷ we see that the hedged position has projections that are much narrower than those of the unhedged position, and the narrowness of the former relative to the latter is particularly pronounced for longer horizons. This indicates that the hedge is very effective over a range of horizons, and that its effectiveness increases with the length of the hedge horizon.

⁷ More precisely, whereas Figure shows pdfs for unhedged vs. hedged positions for $T=10$, Figure 4 shows fan chart representations of the pdfs of unhedged vs. hedged positions for T values going from 1 to 20 years.


Figure 5: Fan Chart Projections of the Values of Hedged and Unhedged Positions for Age 65: APC Two-Population Gravity Model



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005, 1000 simulation trials and an assumed constant interest rate of 0.04. The model used is the ‘parameters certain’ version of the gravity model with no allowance for individual deaths Poisson risk. The black line to year 2005 is a plot of model-based annuity prices.

This conclusion is confirmed by Figure 6 which shows that the RRR is high and rises with T : the RRR rises from an initial value of almost 80% to nearly 90% for $T = 20$.

Figure 6: Relative Risk Reduction for Age 65: APC Two-Population Gravity Model




Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005, 1000 simulation trials and an assumed constant interest rate of 0.04. The model used is the ‘parameters certain’ version of the gravity model with no allowance for individual deaths Poisson risk.

4. Allowing for Parameter Uncertainty and Poisson Risk

We now generalize our analysis to incorporate the possibilities of parameter uncertainty and Poisson risk. (Appendix B gives further details on the parameter uncertainty simulations and Appendix C gives further details on the Poisson risk simulations.) Figure 7 gives the q forward correlations for age 65 for versions of the model with and without parameter uncertainty, and with and without Poisson deaths risk. We see that the PC and parameter uncertain (PU) correlations are close, but the correlations incorporating Poisson risk are notably lower than those leaving it out: incorporating Poisson risk leads to correlations of about 0.70 rather than about 0.90 for a horizon of 20 years.


Figure 7: q Forward Correlations for Age 65: APC Two-Population Gravity Model



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials. 'PC' means 'parameters certain' and 'PU' means 'parameters uncertain'. 'Poisson Risk' means that the simulations take account of individual deaths Poisson risk; 'No Poisson Risk' indicates the opposite.

Figure 8 gives the corresponding fan chart projections of the two populations' q rates. We see that adding each extra source of risk – adding parameter uncertainty or adding Poisson risk – serves to widen the fan charts somewhat.


Figure 8: Fan Chart Projections of the q Rates for Age 65: APC Two-Population Gravity Model



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials. ‘PC’ means ‘parameters certain’ and ‘PU’ means ‘parameters uncertain’. ‘Poisson Risk’ means that the simulations take account of individual deaths Poisson risk; ‘No Poisson Risk’ indicates the opposite.

Figure 9 presents the pdf charts of hedged vs. unhedged annuity positions for all four cases, for the illustrative case of a hedge horizon of $T = 10$. We see that the RRR is fairly high, ranging from 0.703 for the PU case with Poisson risk to the 0.854 for the PC case without Poisson risk.


Figure 9: Hedged and Unhedged Annuity Positions for Age 65, $T=10$: APC Two-Population Gravity Model



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials, and assuming a constant interest rate of 0.04. 'PC' means 'parameters certain' and 'PU' means 'parameters uncertain'. 'RRR' is relative risk reduction. 'Poisson Risk' means that the simulations take account of individual deaths Poisson risk; 'No Poisson Risk' indicates the opposite.

Figure 10 gives the corresponding fan chart projections for horizons of up to 20 years. As with those in Figure 5, these fan charts indicate a high degree of hedge effectiveness which increases with the length of the horizon period T .


Figure 10: Fan Chart Projections of the Values of Hedged and Unhedged Positions for Age 65: APC Two-Population Gravity Model



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials, and assuming a constant interest rate of 0.04. 'PC' means 'parameters certain' and 'PU' means 'parameters uncertain'. The black lines to year 2005 are plots of model-based annuity prices. 'Poisson Risk' means that the simulations take account of individual deaths Poisson risk; 'No Poisson Risk' indicates the opposite.

Figure 11 shows the corresponding plots of relative risk reduction for horizons of up to 20 years: we see that the additional each risk factors serve to lower the RRRs somewhat, although they are all still quite high.

Figure 11: Relative Risk Reduction for Age 65: APC Two-Population Gravity Model



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials, and assuming a constant interest rate of 0.04. 'PC' means 'parameters certain' and 'PU' means 'parameters uncertain'. 'Poisson Risk' means that the simulations take account of individual deaths Poisson risk; 'No Poisson Risk' indicates the opposite.


5. Allowing for Interest-Rate Risk

A final extension is to incorporate interest-rate risk. We model interest-rate risk by assuming that the spot interest-rate process is governed by an illustrative Cox-Ingersoll-Ross (CIR, 1985) process with a mean reversion parameter 0.25, a mean interest rate of 0.04 and a standard deviation of 0.01.⁸ The term structure of interest rates is assumed to be flat.

⁸ These values are taken from Dowd *et al.* (2011b).

Figure 12 gives the fan chart projects of the values of hedged and unhedged positions in the presence of interest-rate risk so modelled.⁹ The presence of interest-rate risk makes the fan charts somewhat wider than they were in Figure 10.

Figure 12: Fan Chart Projections of the Values of Hedged and Unhedged Positions with Interest-Rate Risk for Age 65: APC Two-Population Gravity Model



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005, 1000 simulation trials and an assumed initial interest rate of 0.04. The spot interest rate is assumed to be governed by a CIR process with a mean reversion parameter 0.25, a mean interest rate of 0.04 and a standard deviation of 0.01, and the spot term structure is assumed to be flat. The black lines to year 2005 are plots of model-based annuity prices. 'PC' means 'parameters certain' and 'PU' means 'parameters uncertain'. 'Poisson Risk' means that the simulations take account of individual deaths Poisson risk; 'No Poisson Risk' indicates the opposite.


The corresponding RRR plots are shown in Figure 13. If we compare these to the earlier ones of Figure 11, we draw the following conclusions:

⁹ We recognize, of course, that practitioners would usually hedge interest rate and longevity risk exposures separately. It is, however, a useful exercise to see how the presence of interest rate risk affects the effectiveness of the longevity hedge.

- Adding parameter uncertainty leads to slight falls in RRRs.
- Adding Poisson risk usual (in 3 cases out of 4) leads to a reduction in RRRs.
- Adding interest rate risk leads to an increase in RRRs.

Of course, we should be careful about generalising from these results. In the worst case (PU, Poisson risk, no IRR) RRRs vary from just under 0.6 to about 0.75, but this is something of an outlier. In the best case (PC, Poisson risk, IRR) RRRs are very close to 0.9. However, in most cases, RRRs are in the region of about 0.8. Thus, our results suggest that a high degree of hedging effectiveness.

Figure 13: Hedge Effectiveness with Interest-Rate Risk for Age 65: APC Two-Population Gravity Model



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005, 1000 simulation trials and an assumed initial interest rate of 0.04. ‘PC’ means ‘parameters certain’ and ‘PU’ means ‘parameters uncertain’. ‘Poisson Risk’ means that the simulations take account of individual deaths Poisson risk; ‘No Poisson Risk’ indicates the opposite. The spot interest rate is assumed to be governed by a CIR process with a mean reversion parameter 0.25, a mean 0.04 and a standard deviation of 0.01, and the spot term structure is assumed to be flat.

6. Conclusions and Possible Extensions

This article has considered the effectiveness of an illustrative annuity hedging problem in which a forward annuity predicated on one population is hedged by a position in a forward annuity predicated on another population. Our analysis makes use of a two-population gravity model that takes account of the observed inter-dependence between the two populations' mortality rates; it also considers the implications of parameter uncertainty, individual death or Poisson risk and interest-rate risk for hedge effectiveness. We consider horizons of up to 20 years.

For the most part, our results are robust and indicate strong hedge effectiveness, with estimates of relative risk reduction varying from about 0.70 in the least effective case to well over 0.95 in the most effective cases. This has important implications for defined benefit pension plans, annuity insurers and reinsurers. It suggests that hedges based on longevity indices based on national populations (e.g., the LifeMetrics Index) can with appropriate calibration be highly effective in reducing longevity risk.

There are many possible extensions of this work, including:

- For the simple position considered here, we can try alternative hedging strategies (alternative hedge sizes, hedges based on alternative ages, more sophisticated hedges such as 'gamma' hedges instead of the simple 'delta' hedges considered here, etc), and we can examine cash-flow rather than value-hedges.
- We can also consider the hedging of more complicated positions (e.g., portfolios of annuities predicated on different ages, genders or countries) and, besides the above extensions, consider the additional scope for diversification across ages, genders and countries.
- Where interest-rate risk is present, we might consider hedging strategies that separate out the hedging of interest-rate risk from the hedging of longevity risk, and so bring to bear the existing corpus of interest-rate risk management tools.¹⁰
- Given the similarities between mortality and interest-rate risks – both being positive random variables with term structures, etc – it would also be useful to explore how existing interest-rate risk management tools and strategies might be adapted to manage longevity risk exposures.
- We can (and, indeed, should) extend our analysis to take account of credit risk considerations, e.g., such as the risk of counterparty default, which is not an inconsiderable risk especially over longer horizons.

¹⁰ For example, we might hedge each of the two annuities using bonds of the same value and duration. For more on fixed-interest hedging, see, e.g., Fabozzi (2000).

Despite a robust framework being available, the analysis of the hedge effectiveness of longevity-related securities is, thus, clearly still at a very early stage and a great deal remains to be done.

We end with two caveats regarding the interpretation of our results:

First, we have only considered the effectiveness of a single longevity hedge. It is therefore quite possible, and indeed likely, that additional hedges would improve RRRs further: for example, a hedging strategy might also include an interest-rate hedge to supplement the longevity hedge we have considered here, and would likely produce higher RRRs than those we have reported. In short, we should not ignore the possible scope for additional hedges to generate further improvements in hedging effectiveness.

Second, our results are projections based on assumptions and are only as good as those assumptions might turn out to be. Amongst the most important of these are that the underlying ‘laws of motion’ remain stable over the long horizons we are considering, and that the economy does not experience shocks that might blow our projections well off course.¹¹ These assumptions are self-evidently heroic; in any case, no projections can take account of the ‘unknown unknowns’, to quote Donald Rumsfeld’s famous phrase. Thus, our projections are best interpreted as a form of stochastic scenario analysis and should *not* be interpreted as forecasts.

¹¹ Such possibilities include rapid inflation, which would undermine our interest rate assumptions, or mass defaults, which might cause the institution that issued the hedge to itself default.

References

- Cairns, A. J. G. (2000) A discussion of parameter and model uncertainty in insurance, *Insurance: Mathematics and Economics*, 27: 313-330.
- Cairns, A. J. G. (2004) *Interest Rate Models: An Introduction*. Princeton University Press, Princeton NJ.
- Cairns, A. J. G. (2007) The CBD 2-factor model: A note on approximations to survivor and financial functions. Mimeo. Heriot-Watt University.
- Cairns, A. J. G., Blake, D., and Dowd, K. (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., Cairns, A.J.G., Blake, D., Dowd, K., Coughlan, G.D., and Khalaf-Allah, M. (2011) Bayesian Stochastic Mortality Modelling for Two Populations. To appear in *ASTIN Bulletin* 41(1).
- Cairns, A.J.G., Dowd, K., Blake, D., and Coughlan, G.D. (2011b) Longevity hedge effectiveness: A decomposition. Pensions Institute Discussion Paper, forthcoming.
- Coughlan, G. D., Emery, S. and Kolb, J. (2004). HEAT (Hedge Effectiveness Analysis Toolkit): A consistent framework for assessing hedge effectiveness under IAS 39 and FAS 133, *Journal of Derivatives Accounting*, 1 (2): 221-272.
- Coughlan, G. D., Epstein, D., Ong, A., Sinha, A., Balevich, I., Hevia-Portocarrera, J., Gingrich, E., Khalaf-Allah, M., and Joseph, P. (2007a) LifeMetrics: A toolkit for measuring and managing longevity and mortality risks, Technical Document, JPMorgan: London, 13 March [www.lifemetrics.com].
- Coughlan, G. D., Epstein, D., Sinha, A. and Honig, P. (2007b) q-forwards: Derivatives for transferring longevity and mortality risk, JPMorgan: London, 2 July [www.lifemetrics.com].
- Coughlan, G. D., Khalaf-Allah, M., Yijing Ye, Kumar, S., Cairns, A. J. G., Blake, D., and Dowd, K. (2011) Longevity hedging 101: A framework for longevity basis risk analysis and hedge effectiveness. To appear in *North American Actuarial Journal*.
- Dowd, K., Cairns, A.J.G., Blake, D., Coughlan, G.D., and Khalaf-Allah, M. (2011a) A gravity model of mortality rates for two related populations. To appear in *North American Actuarial Journal*.
- Dowd, K., Blake, D., and Cairns, A.J.G. (2011b) A computationally efficient algorithm for estimating the distribution of future annuity values under interest-rate and longevity risks. To appear in *North American Actuarial Journal*.
- Fabozzi, F. J. (2000) *Bond Markets, Analysis and Strategies*. Prentice Hall.
- Gunawardena, D., C. Hicks and D. O'Neill. (2008). Pension annuities: pension annuities and the open market solution. Association of British Insurers *Research Paper No. 8*.
- Osmond, C. (1985) Using age, period and cohort models to estimate future mortality rates. *International Journal of Epidemiology*, 14: 124–29.

Appendix A: The Age-Period-Cohort Two-Population Gravity Model

A.1. Introduction

The gravity model chosen is a two-population version of the APC model, alternatively known as M3B (see, e.g., Osmond (1985), Jacobsen *et al.* (2002) and Cairns *et al.* (2009a, Table 1)). The APC model is a useful one for our purposes because it is relatively tractable and because it has a cohort effect, which evidence suggests is an important feature of the behaviour exhibited in E&W mortality data and in a number of other important countries' datasets. However, the general gravity approach can also be applied to most other stochastic mortality models.

The successful implementation of any two-population gravity model also requires a solution to the difficult problem of obtaining consistent estimates of both the (unobservable) state variables of the model and of the parameters governing the dynamics of those state variables. In the case of the APC model, the state variables are the period effects (denoted κ_t below) and the cohort effects (denoted γ_c below).

In the standard one-population context, it is relatively straightforward to obtain estimates of both the state variables and their parameters.

However, in the two-population case, we also have to allow for the fact that the state variables themselves cannot be estimated without taking account of the two-population inter-dynamics and hence their parameters. The upshot is that estimation of the state variables requires estimates of the parameters that govern them, and yet we can only obtain estimates of those parameters if we have estimates of the state variables to start with. The solution to this 'chicken and egg' problem is an iterative approach in which we start with preliminary estimates of the state variables based on the assumption of population independence. We then estimate the parameters. After this, we revise the estimates of the state variables and then re-estimate the parameters again based on the new estimates of the state variables. We go on to repeat this process over and over again until estimates of both state variables and parameters have converged.

A.2. Single-population APC model

For the single-population case, the APC model postulates that the true underlying death rate, $m_{t,x}$, satisfies:

$$\log m_{t,x} = \beta_x + n_a^{-1} \kappa_t + n_a^{-1} \gamma_c \quad (\text{A1})$$

where the β_x are age-dependent parameters, κ_t is a time-dependent state variable that represents the period effect and γ_c is a year-of-birth-dependent state variable that represents the cohort effect, n_a is the number of ages in the sample data used to estimate the parameters,

and the variables x , t and $c = t - x$ represent the age, current year and year of birth, respectively. We assume κ_t follows a one-dimensional random walk with drift:

$$\kappa_t = \kappa_{t-1} + \mu + CZ_t \quad (\text{A2})$$

in which μ is a constant drift term, C a constant volatility and Z_t a one-dimensional iid $N(0,1)$ error. We follow Cairns *et al.* (2008) and model γ_c as an ARIMA(1,1,0) process:

$$\Delta\gamma_c = (1 - \alpha^{(\gamma)})\mu^{(\gamma)} + \alpha^{(\gamma)}\Delta\gamma_{c-1} + C^{(\gamma)}Z_c^{(\gamma)} \quad (\text{A3})$$

where $\mu^{(\gamma)}$, $\alpha^{(\gamma)}$ and $C^{(\gamma)}$ are given parameters and $Z_c^{(\gamma)}$ is iid $N(0,1)$. The parameters of this model can be estimated using MLE.

This model requires the imposition of identifiability constraints. The first two are:

$$\sum_t \kappa_t = 0 \text{ and } \sum_c \gamma_c = 0 \quad (\text{A4})$$

We also need a third identifiability constraint based on a tilting parameter δ . This is chosen within an iterative scheme to minimize:

$$S(\delta) = \sum_x (\beta_x + \delta(x - \bar{x}) - \bar{\beta}_x)^2$$

where $\bar{\beta}_x = n_y^{-1} \sum_t \log m(t, x)$. This then implies that:

$$\delta = -\frac{\sum_x (x - \bar{x})(\beta_x - \bar{\beta}_x)}{\sum_x (x - \bar{x})^2} \quad (\text{A5})$$

Given that the κ_t and γ_c already satisfy the first two constraints, we now apply the third constraint which requires us to revise our parameter estimates according to the following formulae:

$$\begin{aligned} \tilde{\kappa}_t &= \kappa_t - n_a \delta(t - \bar{t}) \\ \tilde{\gamma}_c &= \gamma_c + n_a \delta((t - \bar{t}) - (x - \bar{x})) \\ \tilde{\beta}_x &= \beta_x + \delta(x - \bar{x}) \end{aligned} \quad (\text{A6})$$

A.3. Two-population APC model

We now wish to model q rates for two related populations, focussing on the case where one population is much larger than the other. Let us denote the large population using the

superscript ‘(1)’ and the small population using the superscript ‘(2)’. The gravity model then involves the following 2-population κ_t process:

$$\kappa_t^{(1)} = \kappa_{t-1}^{(1)} + \mu^{(1)} + C^{(11)}Z_t^{(1)} \quad (\text{A6})$$

$$\kappa_t^{(2)} = \kappa_{t-1}^{(2)} + \phi^{(\kappa)}(\kappa_{t-1}^{(1)} - \kappa_{t-1}^{(2)}) + \mu^{(2)} + C^{(21)}Z_t^{(1)} + C^{(22)}Z_t^{(22)} \quad (\text{A7})$$

where $0 \leq \phi^{(\kappa)} < 1$ is the ‘gravity parameter’ that pulls the $\kappa_t^{(2)}$ state variable towards the $\kappa_t^{(1)}$ state variable. Now let $\mu^{(\gamma_1)}$ be the mean reversion level for $\Delta\gamma_c^{(1)}$, $\mu^{(\gamma_2)}$ be the mean reversion level for $\Delta\gamma_c^{(2)}$, α^{γ_1} be the AR(1) parameter for $\Delta\gamma_c^{(1)}$, α^{γ_2} be the AR(1) parameter for $\Delta\gamma_c^{(2)}$ and $0 \leq \phi^{(\gamma)} < 1$ be the gravity parameter for the 2-population γ_c process. The 2-population AR(1) γ_c processes is then:

$$\gamma_c^{(1)} = (1 + \alpha^{(\gamma_1)})\gamma_{c-1}^{(1)} - \alpha^{(\gamma_1)}\gamma_{c-2}^{(1)} + \mu^{(\gamma_1)}(1 - \alpha^{(\gamma_1)}) + C^{(\gamma_11)}Z_c^{(\gamma_1)} \quad (\text{A8})$$

$$\begin{aligned} \gamma_c^{(2)} = & (1 + \alpha^{(\gamma_2)} - \phi^{(\gamma)})\gamma_{c-1}^{(2)} - \alpha^{(\gamma_2)}\gamma_{c-2}^{(2)} + \phi^{(\gamma)}\gamma_{c-1}^{(1)} + \mu^{(\gamma_2)}(1 - \alpha^{(\gamma_2)}) + \\ & C^{(\gamma_21)}Z_c^{(\gamma_1)} + C^{(\gamma_22)}Z_c^{(\gamma_2)} \end{aligned} \quad (\text{A9})$$

Further details on the model and its calibration are given in Dowd *et al.* (2011a).

Appendix B: Parameter Simulation

The simulation of the parameters is as follows:

- The μ , V , $\mu^{(\gamma)}$ and $V^{(\gamma)}$ parameters are simulated using the algorithms set out in Cairns et al (2006) and later set out for the gravity two-population APC model in Dowd *et al.* (2011a).
- The $\alpha^{(\gamma)}$ parameters are simulated using the approach of Cairns (2000) and Dowd (2011b).
- The method used to simulate $\phi^{(k)}$ and $\phi^{(\gamma)}$ parameters is based on the assumptions that (a) the relevant ϕ comes from single-peaked beta(ν, ω) distribution (b) whose mean matches the empirical estimate of ϕ , $\hat{\phi}$, and (c) which has the maximum dispersion, so the prior is as uninformative as possible. Combining these requirements means that ω is the lowest possible value no less than 2 such that $\nu = \omega \hat{\phi} / (1 - \hat{\phi})$.

Appendix C: Simulation of the Poisson Process

Where we wish to incorporate the Poisson (or individual deaths) process, we proceed as follows:

For any given population and age x at t , we simulate the number of deaths $D(t, x)$ from a Poisson process with ‘arrival rate’ $m(t, x) \times E(t, x)$, where $m(t, x)$ is the crude central death rate from our mortality tables and $E(t, x)$ is the number of exposures. We then update the number of exposures next period using

$$E(t+1, x+1) = E(t, x) \cdot e^{-m(t, x)} \quad (\text{C1})$$

based on an initial (i.e., $t=0$) exposure equal to the 2005 value for the appropriate age. (C1) ensures that the number of exposures is always positive, which is helpful computationally.¹²

¹² Another consideration is that having extra Poisson risk in the $E(t, x)$ as well as $D(t, x)$ would not have a significant impact on the $D(t, x)$ Poisson randomness, except perhaps at the high ages and as $E(t, x)$ gets very small.

Appendix D: Simulation of Future Annuity Prices

D.1. Annuity price simulation

The model used to value the annuity at future time T involves the estimation of the expected survivor rates as of time T for the relevant cohort: these can then be treated as the expected annuity payments, and the annuity price is the present value of these payments. For any given version of the mortality model (i.e., with or without parameter uncertainty, and with or without Poisson risk), let $S^{(i)}(t, x)$ be the survivor index at time t of a cohort from population i (where $i=1,2$) aged x in year 0: that is, $S^{(i)}(t, x)$ is the probability, measured retrospectively, that an individual aged x at time 0 survives to time t . For any given x , $S^{(i)}(0, x)=1$ and $S^{(i)}(t, x)$ will decrease as t gets bigger and eventually approach 0 as t gets large. Given any path of $q^{(i)}(t, x)$, we then obtain a corresponding path of $S^{(i)}(t, x)$ from the relationship between the survivor index and q rates:

$$(D1) \quad S^{(i)}(t+1, x) = (1 - q^{(i)}(t+1, x))S^{(i)}(t, x)$$

These survivor rates are driven off the state variables $\kappa_t^{(1)}$ and $\kappa_t^{(2)}$. Hence, for our purposes, we wish to simulate sets of state variables out to a future date T and then estimate the expectations of (D1), conditional on surviving to the specified future date and conditional on the future values of the state variables, i.e., $\kappa_T^{(1)}$ and $\kappa_T^{(2)}$.

We take j simulation paths of each state variable out to period T , and let $\kappa_T^{(j)} = [\kappa_T^{(1j)}, \kappa_T^{(2j)}]$ be the j^{th} set of simulated state variables for period T . Assuming for the moment that the interest rate R is constant throughout, then the fair value of the annuity for population i at time T , conditional on the time- T simulated state variables under simulation path j , is

$$(D2) \quad a^{(ij)}(T, \kappa_T^{(j)}) = (1 + \psi) \sum_{h=0}^{115-x-T} \exp(-hR) E[S^{(i)}(T+h, x) / S^{(i)}(T, x) | \kappa_T^{(j)}]$$

where ψ is the loading factor built into the annuity value (assumed to be 0) and we assume that no-one lives beyond age 115. The term $E[S^{(i)}(T+h, x) / S^{(i)}(T, x) | \kappa_T]$ is to be interpreted as the expected probability that an individual aged x will survive to year $T+h$, conditional on their surviving to T and conditional on the mortality state parameters $\kappa_T^{(j)}$ at T .

However, we cannot compute (D2) directly because there is no simple formula for $E[S^{(i)}(T+h, x) / S^{(i)}(T, x) | \kappa_T]$ in terms of the mortality state variables κ_T^j . Nor is it practically feasible to use stochastic simulation to estimate $E[S^{(i)}(T+h, x) / S^{(i)}(T, x) | \kappa_T]$ for

each set of $\kappa_T^{(j)}$, as this would require a simulation tree within a simulation tree and would be computationally prohibitively expensive.¹³

A more practical approach is to use a Taylor series approximation, as first suggested by Cairns (2007) and developed further by Dowd *et al.* (2011b). Define $f^{(i)}(h, x, \kappa) = \Phi^{-1}\left(E[S^{(i)}(T+h, x)/S^{(i)}(T, x) | \kappa_T]\right)$ as the probit transformation of $E[S^{(i)}(T+h, x)/S(T, x) | \kappa_T]$, where $\Phi(\cdot)$ is the standard normal distribution function. Let $\hat{\kappa}_T = E[\kappa_T]$ be the expectation of the mortality state variables at T . We then take the following first-order Taylor series expansion of $f^{(i)}(T, h, x, \kappa_T)$ around $\hat{\kappa}_T$:¹⁴

$$(D3) \quad f^{(i)}(T, h, x, \kappa_T) \approx \Delta_0^{(i)}(T, h, x) + \Delta_1^{(i)}(T, h, x)'(\kappa_T - \hat{\kappa}_T)$$

where $\Delta_0^{(i)}(T, h, x)$ is a scalar function of h and x , and $\Delta_1^{(i)}(T, h, x) = [\Delta_{11}^{(i)}(T, h, x), \Delta_{12}^{(i)}(T, h, x)]'$ is a 2×1 vector of first derivatives.

For any given T , h and x , these ‘ Δ ’ terms are parameters that are easily computed by Monte Carlo simulation. The simulated expected survivor rates out through to the time when the cohort has died out can then be recovered from

$$(D4) \quad E[S^{(i)}(T+h, x)/S^{(i)}(T, x) | \kappa_T] \approx \Phi\left(f^{(i)}(T, h, x, \kappa_T)\right)$$

Finally, each simulated $E[S^{(i)}(T+h, x)/S^{(i)}(T, x) | \kappa_T]$ can be plugged into (D2) to give the corresponding simulated future annuity value we are seeking.

D.2. Interest-rate simulation

We assumed earlier that the interest rate R was fixed throughout. Where we wish to randomize the interest rate, we simulate a random interest rate using the Cox-Ingersoll-Ross (CIR) model (Cox *et al.*, 1985). This model postulates that the instantaneous spot interest rate R obeys the continuous-time process:

$$(D5) \quad dR(t) = \alpha(\bar{R} - R(t))dt + \sigma\sqrt{R(t)}dW(t)$$

¹³ For example, with 1000 simulation paths in each stage for each of the mortality state variables, this would require 2 million simulation paths for the mortality state variables alone; combined with all the other calculations required, this implies a computational burden that is not practically feasible under real-time constraints.

¹⁴ Cairns (2007) actually suggests a second-order Taylor series expansion, but we use a first-order expansion here because the ‘ Δ ’ parameters of the second population are sometimes very unstable when we allow for Poisson risk.

where α determines the strength of the mean-reversion process governing R , \bar{R} is the long-term mean instantaneous spot interest rate, σ is the interest-rate standard deviation and $dW(t)$ is a standard geometric Brownian motion.

We can simulate values of $R(T)$ directly from their exact distribution using the CIR parameters and the current instantaneous spot rate $R(0)$ as inputs.¹⁵ In our random interest-rate scenarios, we then assume for convenience that the term structure at T is flat at $R(T)$.

D.3. Implementation

In principle, we can estimate the ‘ Δ ’ terms for each horizon h , but this is computationally expensive.

However, we can exploit the structure of the κ process to economize on calculation time, i.e., in particular, we can assume that the slope or $\Delta_1^{(i)}$ terms are constant across T . (By contrast, the $\Delta_0^{(i)}$ intercept terms are not constant; they reflect the expected survivorship rates for $h=0, \dots, 115-T-x$ starting as of T , and these will generally increase with T .)

We now estimate (D3) for $T=0$ and $h=0, \dots, 115-T-x$.¹⁶

We then retain the ‘ $\Delta_1^{(i)}$ ’ values so obtained throughout for all T simulations and re-estimate the ‘ $\Delta_0^{(i)}$ ’ for each T .


D.4. Estimates of the ‘ Δ ’ terms

Figure D1 shows estimates of the ‘ $\Delta_1^{(i)}$ ’ terms. For age 65 and assuming that everyone dies by age 115, there are 50 values for each type of ‘ $\Delta_1^{(i)}$ ’ term, i.e., one value of $\Delta_1^{(i)}$ for each of up to $h=50$. One will note how the estimates with Poisson risk (i.e., the green ‘-’ and red ‘- -’ plots) are somewhat less stable than the others, especially for population 2.

¹⁵ To be precise, if $R(T)$ follows a CIR process, then $(4\alpha R(T))/\{\sigma^2(1-\exp[-\alpha T])\}$ has a non-central chi-squared distribution with $4\alpha\bar{R}/\sigma^2$ degrees of freedom and a non-centrality parameter equal to $(4\alpha R(0))/\{\sigma^2(1-\exp[-\alpha T])\}$ (Cairns, 2004, Theorem 4.8 (c)).

¹⁶ Estimating (D3) for $T=0$ gives plausible ‘ $\Delta_1^{(i)}$ ’ results for our data set, but in principle any T value would do, provided the ‘ $\Delta_1^{(i)}$ ’ values are plausible and fairly stable.

Figure D1: Estimates of the ‘ $\Delta_1^{(i)}$ ’ Terms: Gravity Two-Population Model



Notes: Based on E&W and CMI male deaths and exposures data over ages 60:84 and years 1961:2005 and 1000 simulation trials. Black continuous refers to the PC case with no deaths Poisson risk, blue ‘-.’ refers to the PU case with no Poisson risk, green ‘-.’ refers to the PC case with Poisson risk, and red ‘- -’ refers to the PU case with Poisson risk.