Institute and Faculty of Actuaries

Understanding and Comparing Longevity Projections

Dave Grimshaw and Jon Palin

06 November 2015

Aims of this talk

Understanding and Comparing Longevity Projections

- Many techniques are used to project longevity. Here, we will:
  - Describe and compare different longevity models
  - Examine the principles and assumptions of different models
  - Consider the benefits of particular models for different purposes
  - Highlight the merits of using multiple models.
- This session is aimed at delegates without specialist longevity expertise, who need to assess proposed assumptions.

Broad types of model

- Different models can be put into broad categories:
  - Deterministic (best-estimate or scenarios) or stochastic (distribution)
  - Modelling all-cause mortality or split by cause-of-death (e.g. circulatory, cancer, other; or more detail)
  - Extrapolation (past trends continue) or opinion (perhaps relating to drivers such as smoking).

Specific types of model

- There are lots of variations on a theme, but key distinctions between:
  - CMI Model
  - Cause of death models
  - Stochastic models
  - P-spline (not covered in this talk).

CMI Model

- Deterministic
- All-cause
- Opinion – multiple parameters and requires a view on the long-term rate
- Basic idea – mortality improvements are blended between current rates and a long-term assumption, considering age-period and cohort effects separately – but complex implementation
- Under review – consultation in March 2016.

CMI Model: age-period terms

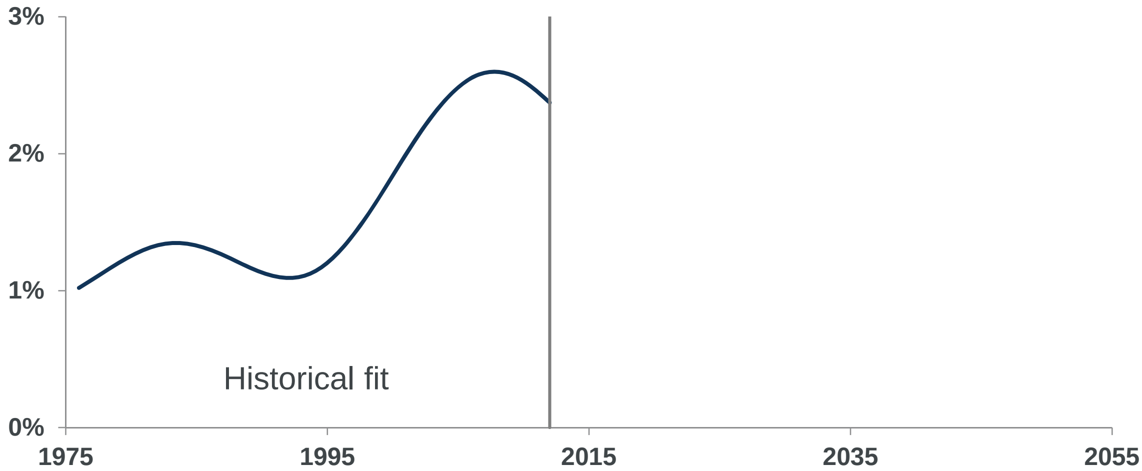

Source: Barnett Waddingham calculations based on CMI_2015

CMI Model: age-period terms

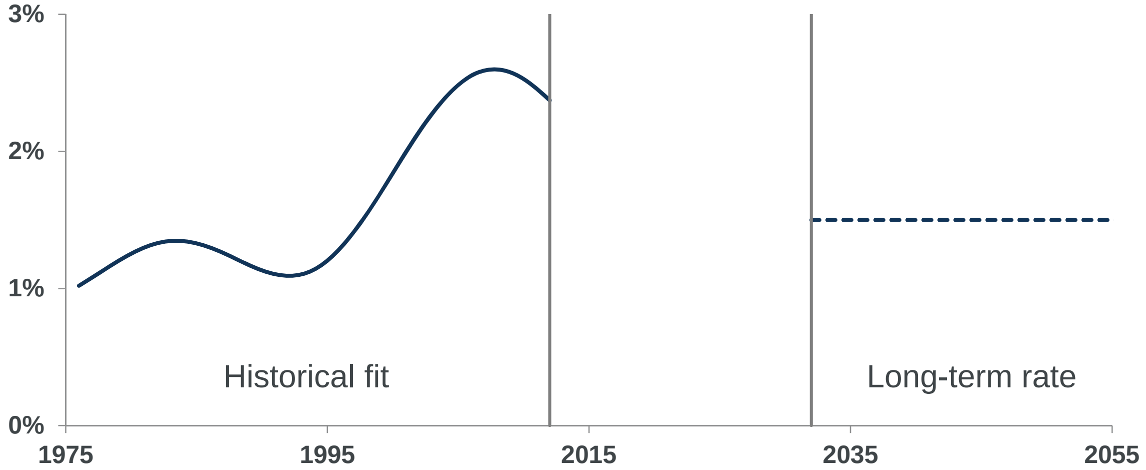

Source: Barnett Waddingham calculations based on CMI_2015

06 November 2015 7

CMI Model: age-period terms

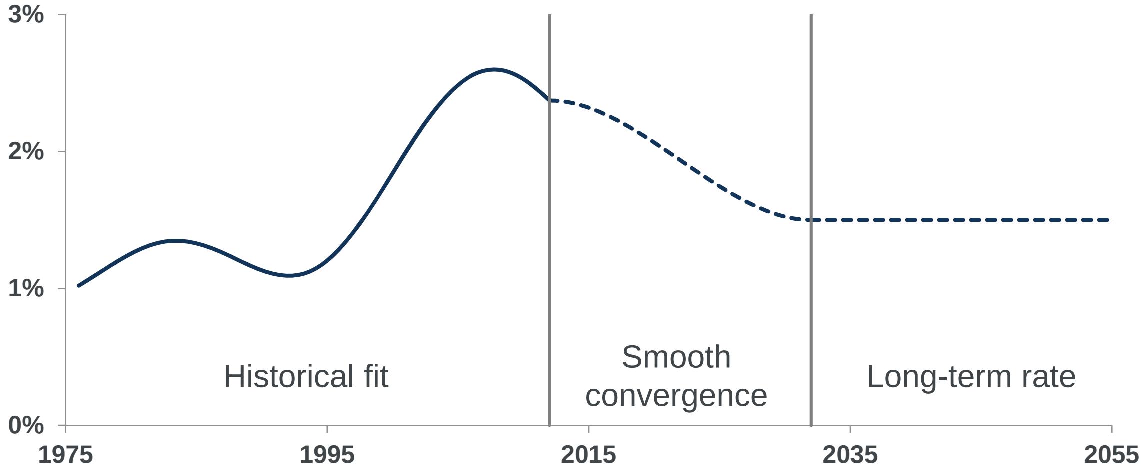

Source: Barnett Waddingham calculations based on CMI_2015

06 November 2015 8

CMI Model: Age-period and cohort components

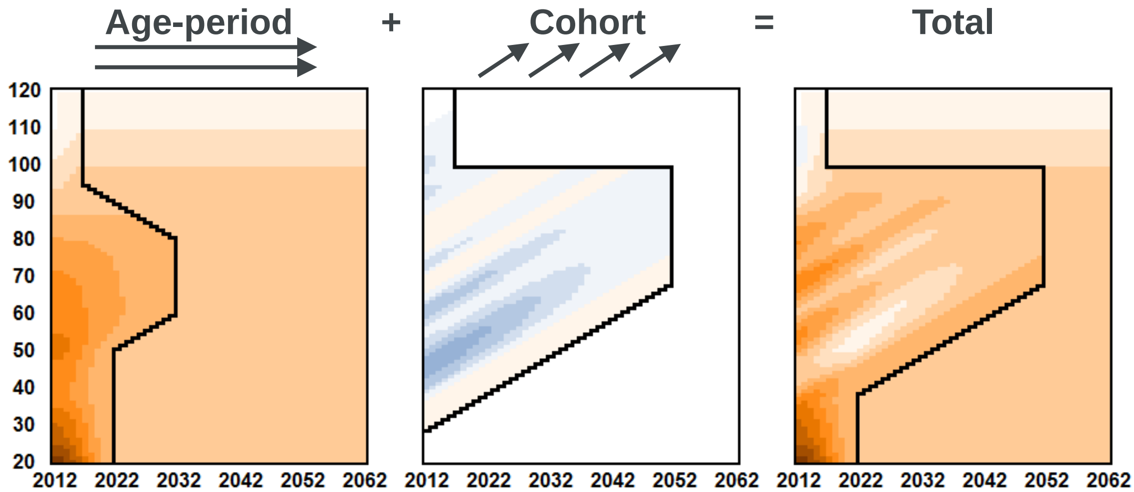

Source: Barnett Waddingham calculations based on CMI_2015

Cause of death

- Deterministic (typically)
- By cause
- Extrapolation and/or opinion, perhaps based on underlying drivers
- Basic idea – split mortality into cause of death groups, analyse historic patterns of improvements, and project them, perhaps allowing for changing conditions
- Can be used to inform the long-term rate in the CMI Model.

Criticisms of causal models

- Data quality and changes in recording
- Correlation between causes
- Importance of general ageing rather than discrete diseases
- Shifting medical resources over time
- Risk of missing new drivers
- Complexity of multi-factor models.

Cause of death data

Smoothing the raw data can give volatile improvements and undulating mortality rates.

Smoothed mortality improvements (LH chart) and mortality rates (RH chart) for cardiovascular disease, males aged 75-79

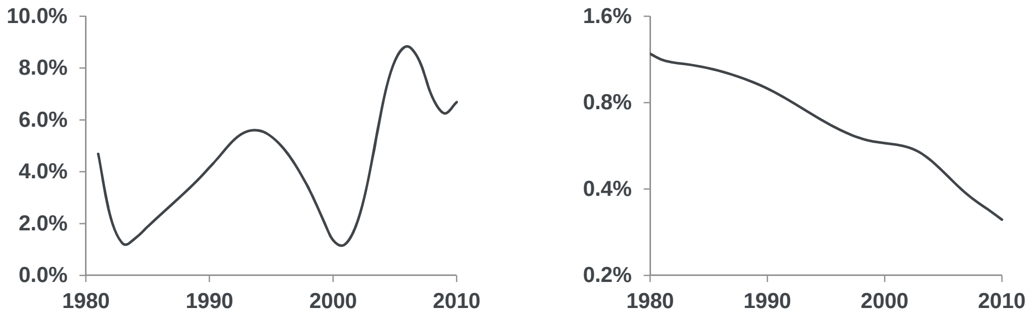

Source: Barnett Waddingham calculations based on data published by the Office for National Statistics (ONS)

Cause of death data

- Volatility arises from changes to the International Classification of Diseases (ICD) rules for coding deaths
- If we allow for coding changes, we can:
  - fit smoother lines to each period
  - produce a single consistent trend based on the current ICD definitions.

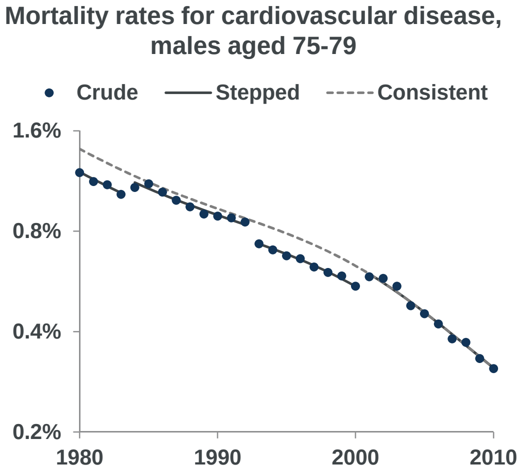

Source: Barnett Waddingham calculations based on ONS data

Implications of cause of death models

- Circulatory disease has been the major driver of improvements in recent years
- As it becomes less dominant, it will have less impact on overall improvements
- Circulatory disease has been relatively easy to solve (compared to cancer, dementia).

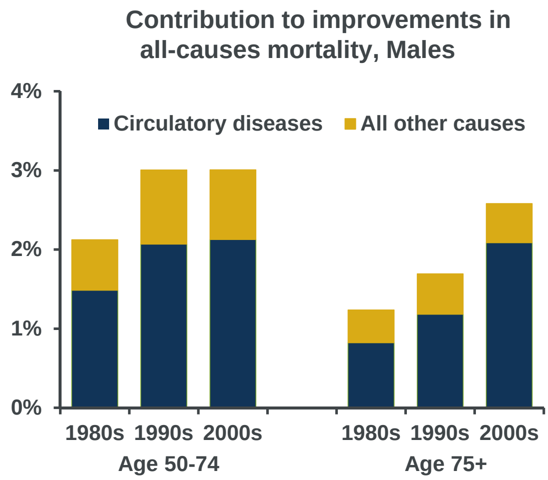

Source: Barnett Waddingham calculations based on ONS data

Impact of reduced smoking

- Mortality improvements for smokers and "never-smoked" are both lower than the combined population
- Over 0.5% p.a. of mortality improvements comes from reductions in smoking prevalence
- This cannot be repeated.

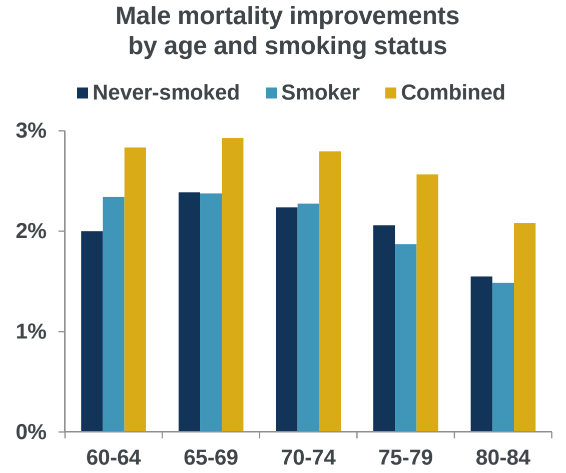

Source: Barnett Waddingham calculations based on ONS data

06 November 2015 15

Crude mortality rates by cause

- Annual mortality improvements, 1993-2012, males ages 70-74
  - circulatory 6.3% p.a.
  - cancer 1.9% p.a.
  - other 2.3% p.a.
- Cancer is now the leading cause-group.

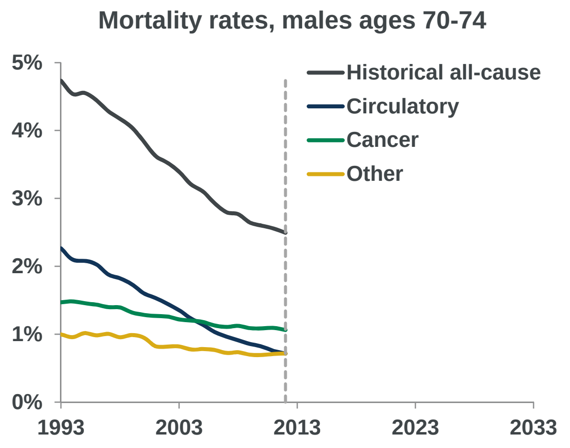

Source: Barnett Waddingham calculations based on ONS data

06 November 2015 16

A simple extrapolation of mortality by cause

- The sum of the projections by cause results in lower improvements than projecting all-causes mortality
- Over time the faster-improving causes (e.g. circulatory) have less impact as they cause fewer deaths.

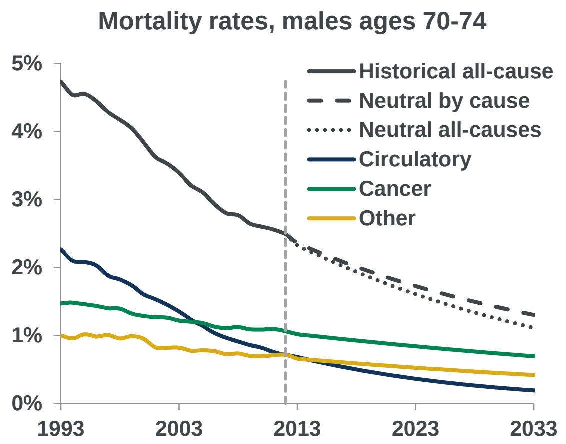

Source: Barnett Waddingham calculations based on ONS data

Data quality

- Death counts are fairly reliable – deaths have to be registered
- Population and exposure estimates are less reliable
  - older ages
  - birth patterns and specific cohorts
- Mismatches between deaths and exposures
- These concerns apply to all-cause and by-cause data.

ONS population data – older ages

- The ONS population estimates were re-stated after the 2011 Census
- The impact was greatest at old ages
- We have concerns over the data, even after restatement.

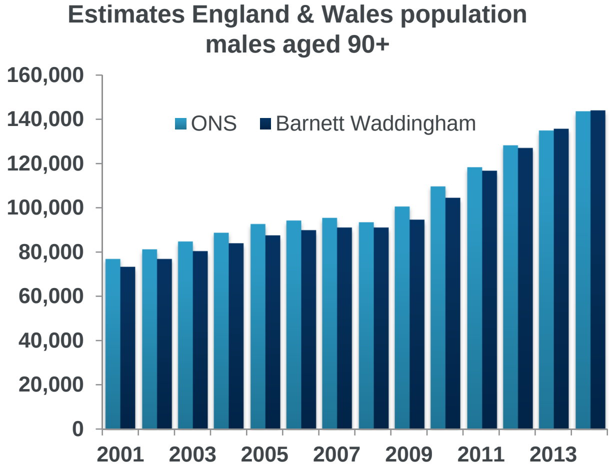

Source: ONS and Barnett Waddingham calculations based on ONS data

ONS population data – 1919/1920 cohort

- Commonly assumed that annual exposure = mid-year population
- This fails if birth patterns are unusual e.g. 1919/1920 after World War I
- Leads to artefacts in crude mortality rates
- Adjustments made in CMI_2014 and CMI_2015.

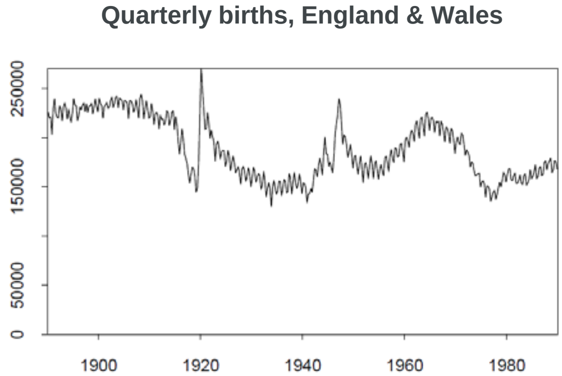

Source: "Phantoms never die", Cairns, Blake, Dowd and Kessler

Basis risk

- Projections are typically calibrated to the general population
- How to model your own book
  - annuitants vs. general population
  - enhanced annuitants
- Measuring improvements directly requires huge data volumes.

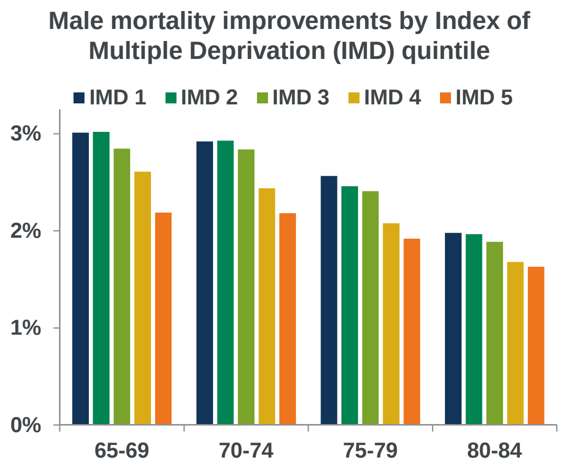

Source: Barnett Waddingham calculations based on ONS data

Common stochastic models

- Stochastic
- All-cause (typically)
- Extrapolation (typically)
- Basic idea: model historic improvements as a combination of age, period and cohort terms, and project period terms using their historical trends and uncertainty.

The Lee-Carter model

- Historical improvements have complex patterns by age, period, cohort
- Lee and Carter’s key idea:
  – reduce the problem to one factor – effectively create a ‘mortality index’

$$m(x,t) = A(x).B(x)^{K(t)}$$

- $m(x,t)$: mortality rate
- $A(x)$: ‘base’ mortality
- $B(x)$: expected reduction factor
- $K(t)$: uncertain speed of improvements

The Lee-Carter model

Because the Lee-Carter has a single source of uncertainty, mortality shapes are highly constrained.

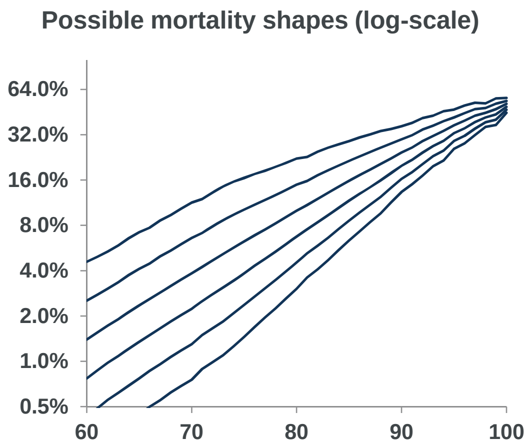

Source: Barnett Waddingham

The Lee-Carter model

1. Calibrate to historical data

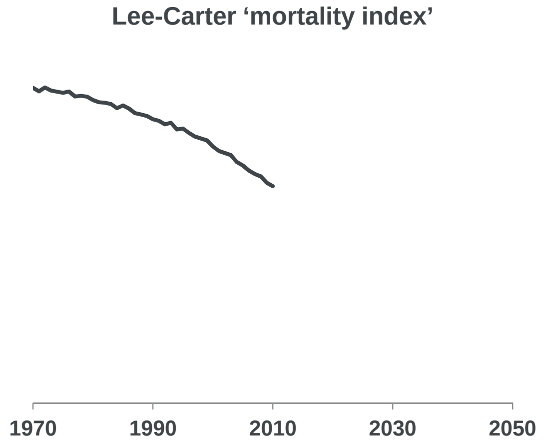

Source: Barnett Waddingham

The Lee-Carter model

1. Calibrate to historical data
2. Central projection (e.g. extrapolate recent trend)

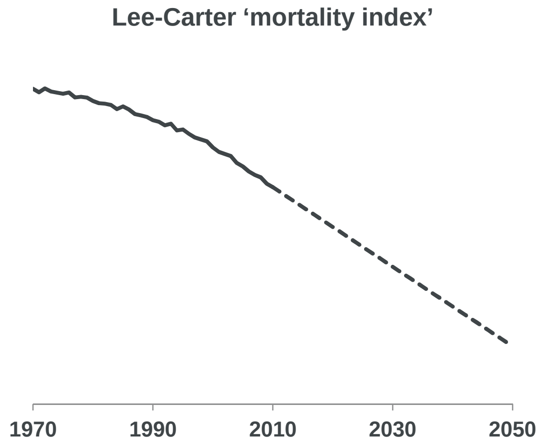

Source: Barnett Waddingham

The Lee-Carter model

1. Calibrate to historical data
2. Central projection (e.g. extrapolate recent trend)
3. Allow for uncertainty

Lee-Carter "mortality index"

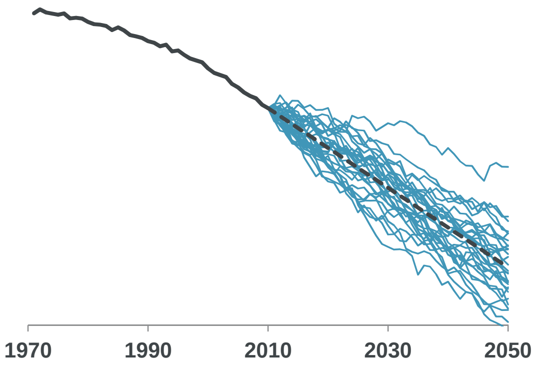

Source: Barnett Waddingham

06 November 2015 27

The Lee-Carter model

1. Calibrate to historical data
2. Central projection (e.g. extrapolate recent trend)
3. Allow for uncertainty
4. Use mortality index to calculate mortality rates, cashflows, liabilities

Lee-Carter "mortality index"

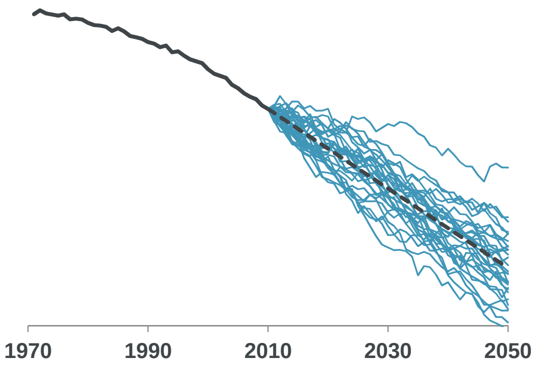

Source: Barnett Waddingham calculations

06 November 2015 28

Decisions in stochastic modelling (1)

- Model structure
  - how to assigning past mortality to components: age, period, cohort terms
  - how many sources of uncertainty
  - Lee-Carter, Age-Period-Cohort (APC), Cairns-Blake-Dowd (CBD) and many more…

$$\log m(t,x) = \beta_x^{(1)} + \beta_x^{(2)}\kappa_t^{(2)}$$
$$\log m(t,x) = \beta_x^{(1)} + \beta_x^{(2)}\kappa_t^{(2)} + \beta_x^{(3)}\gamma_{t-x}^{(3)}$$
$$\log m(t,x) = \beta_x^{(1)} + n_a^{-1}\kappa_t^{(2)} + n_a^{-1}\gamma_{t-x}^{(3)}$$
$$\log m(t,x) = \Sigma_{i,j}\theta_{ij}B_{ij}^{ay}(x,t)$$
$$\text{logit } q(t,x) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x})$$
$$\text{logit } q(t,x) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \gamma_{t-x}^{(3)}$$
$$\text{logit } q(t,x) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \kappa_t^{(3)}((x-\bar{x})^2 - \hat{\sigma}_x^2) + \gamma_{t-x}^{(4)}$$
$$\text{logit } q(t,x) = \kappa_t^{(1)} + \kappa_t^{(2)}(x-\bar{x}) + \gamma_{t-x}^{(3)}(x_c - x)$$

Source: A Quantitative Comparison of Stochastic Mortality Models Using Data From England and Wales and the United States, Cairns et al

Model structure – managing complexity

- Lee-Carter – one-factor model
  - ✓ simple and interpretable
  - × too-strong correlation between ages
  - × uncertainty is proportional to mean mortality improvement
- Cairns-Blake-Dowd M7 – three-factor model
  - ✓ richer uncertainty structure
  - × requires covariance assumptions between time series components
  - – a better fit need not give a better projection.

Model structure – spurious cohorts

- Historical UK mortality data has cohort effects i.e. mortality improvements vary by birth year
- So it's appropriate to put cohort parameters in mortality models
- But some stochastic models fit cohort parameters to data which is constructed to have no cohort effects
- So fitted cohort parameters can be "spurious" – unrelated to genuine cohort effects
- This can lead to misleading projections.

Decisions in stochastic modelling (2)

- Time series properties
  - how to project period components (and perhaps cohort components)
  - coherent modelling of multiple populations.

Short-term volatility

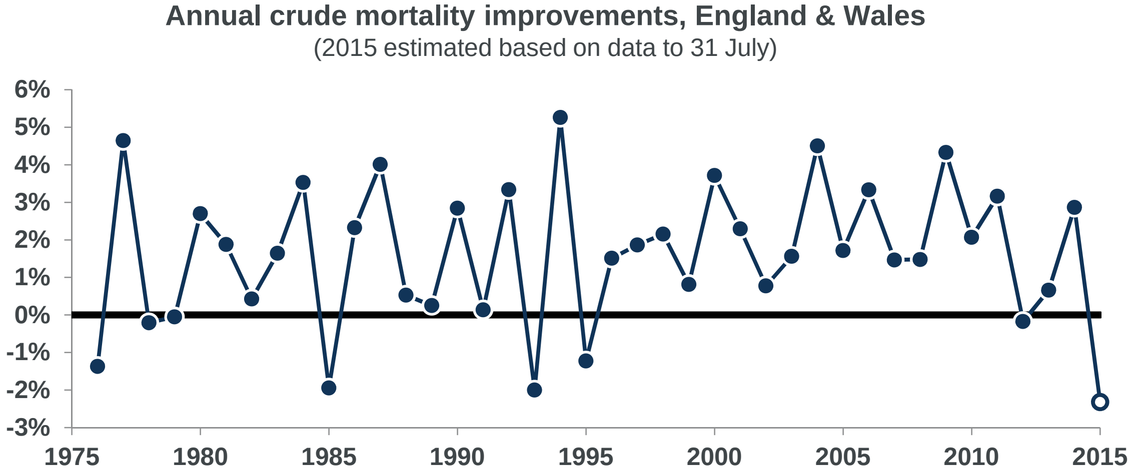

Source: CMI Working Paper 83

Time series – short/medium/long-term volatility

Age group	60-69	70-79	80-89	90-99
Mean improvement	2.5%	2.3%	1.4%	0.7%
1-year standard deviation	1.8%	2.3%	2.7%	3.5%
1-year auto-correlation	-14%	-29%	-44%	-58%
5-year standard deviation	4.9%	5.2%	5.0%	5.1%
long-term standard deviation	highest?			lowest?

Source: Barnett Waddingham calculations based on ONS data

Short- versus medium/long-term

- Short-term volatility due to temperature and infectious diseases
- Medium/long-term trends typically driven by (steadier) lifestyle, medical, and economic changes
- A simple random walk can't reflect the difference between short- and medium/long-term influences and age patterns
- Long-term volatility and model structure
  - Lee-Carter assumes it is proportional to expected
  - APC model assumes the same at all ages.

Aside: recent experience – blip or trend?

Standardised mortality ratio, England & Wales, and trend
(2015 estimated based on data to 31 July)

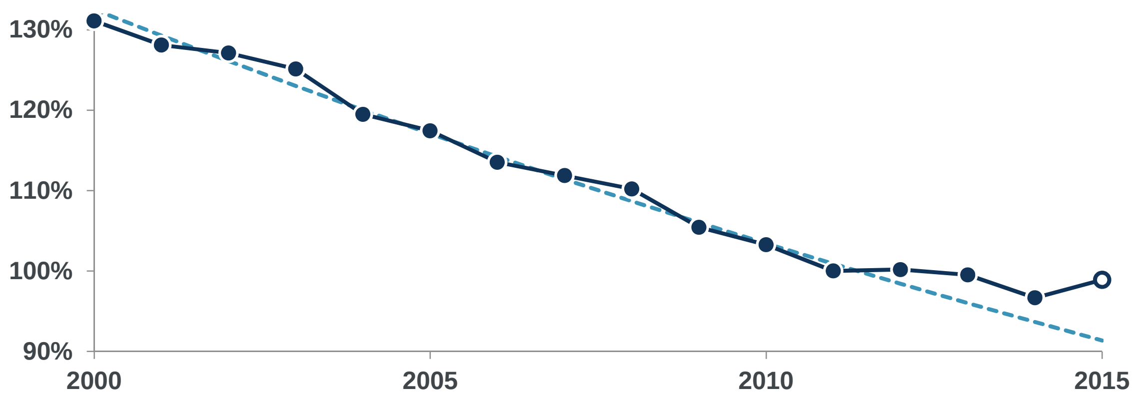

Source: CMI Working Paper 83

Solvency II one-year risk

How bad can things get over one year?

- Risks from:
  - new mortality data, so recalibrate the model
  - other information, outside the model
- Repeat N,000 times
  1. simulate one year of new data
  2. re-fit the model and calculate best-estimate
  3. take 1-in-200 value of new best-estimate.

06 November 2015 37

Solvency II one-year risk

- Repeat N,000 times
  1. **simulate one year of new data**
  2. **re-fit the model and calculate best-estimate**
  3. take 1-in-200 value of new best-estimate
- We need a good model of one-year risk rather than long-term risk
- Why use a stochastic model if we just want a best-estimate?

06 November 2015 38

Multi-population modelling

- Model related populations independently?
- Constrain the difference between them?
- What should the long-term difference between males and females be?

Female minus male life expectancy at age 65

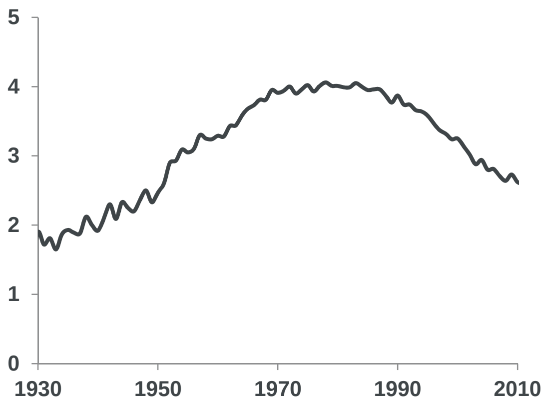

Source: Barnett Waddingham calculations based on ONS data

Multi-population modelling

- Difference in male/female life expectancy
- Different countries
  - Dutch model = “international + difference”
- Annuitant versus general population
  - higher social classes have historically had higher mortality improvements
  - will this continue, or will lower social classes catch up?
  - basis risk for index-based transactions.

Complexity budget

- Ideal case?
  - The Cairns-Blake-Dowd (CBD) M7/M9 with three time series
  - Multiple random variables per time series to allow for short/long term risk
  - Coherent modelling of males and females
- That means 12+ random variables to calibrate – too complex/spurious?
- When modelling multiple populations, we need to use simpler models for each population.

Model risk

- Each model has its own (implicit or explicit) assumptions
- We can't know if a particular model reflects the future well
- Using a single model exposes us to that model's assumptions.

Combining models

- Combine structures – e.g. CMI best estimate plus stochastic volatility
- Combine results – e.g. simulations from multiple stochastic models
- Which models can we trust?
  - Lee-Carter and M5 – no cohort term
  - Lee-Carter + cohort, APC and M6 – spurious cohorts
  - Renshaw-Haberman and M8 – difficult to fit, sensitive to data used
  - only M7, M9 and “APC-MI” seem to give robust results
- Many of these share common assumptions e.g. extrapolation.

Which model(s) to use?

- What is the model to be used for?
  - individuals or populations; broad risk assessment or assessing transactions?
  - spend the complexity budget in important areas
- Apply judgement
  - consider adjusting a simple model rather than using a complex model
- Multi-model approach
  - compare results from multiple models with different approaches, to bring different insights
- Regulator view?

Physics envy

- Physics problems:
  - stable conditions over time
  - reproducible experiments
  - ability to test/refute a prediction

"Imagine how much harder physics would be if electrons had feelings!"

Richard Feynman

- Longevity modelling isn't physics – helpful to blend judgement with statistical models