

3 Modelling choices for the development of standard mortality tables

Jessica Mosher

This chapter describes the various modelling decisions required for the development of mortality tables for pensioner and annuitant populations. It explains why one approach may be selected over another, along with any potential advantages or disadvantages. As such, it should help regulators and supervisors to have a better sense of the implications and appropriateness of the modelling choices in a particular context.

The development of mortality tables requires numerous modelling decisions. Each aspect of modelling involves a certain level of judgement to determine whether one model could be more suitable than another, and therefore it is useful to have an understanding of the implications that different choices can have.

This chapter aims to describe the different options in an intuitive manner, and to explain why one approach may be selected over another, along with any potential advantages or disadvantages. As such, it should help regulators and supervisors to have a better sense of the implications and appropriateness of the modelling choices in a particular context. Its objective is not to enter into the mathematical details and the technical aspects of model implementation, though it does include some details where they might be relevant.

The chapter is organised around the main steps followed in the development of mortality tables. The first section discusses the graduation techniques used to smooth raw mortality rates. The second section looks at ways that the mortality curve can be extended to both younger and older ages. The third section considers how mortality tables account for selection effects and adjust the mortality rates to reflect the lower expected mortality for pensioner and annuitant populations. The fourth section discusses the different options available to model future mortality improvements. The final section provides some insights as to how innovations in data analysis are starting to be used to inform the development of mortality assumptions.

3.1. Graduating central mortality rates

The graduation of mortality rates involves fitting a function to the observed raw mortality rates to result in a smooth pattern of mortality across ages. Due to normal variability, raw mortality rates do not necessarily follow a smooth pattern across ages, particularly for smaller pensioner and annuitant populations. As such, graduation techniques are usually employed to smooth the raw rates and obtain a mortality curve that reflects the expected biological pattern of mortality, that is generally increasing monotonically with age.

For pensioner and annuitant populations, graduation is more commonly performed over a range of central ages (e.g. 50-95) where observations are sufficient to calculate a robust estimate of mortality. For larger populations, such as the general population of a country, the graduation can potentially be done over a larger range of ages.

Mortality rates are calculated taking the ratio of observed deaths to the number of individuals alive in the year.¹ Box 3.1 provides a more precise explanation. This ratio can also be calculated based on pension amounts rather than individual lives. The latter approach is intended to capture the economic gradient of mortality, where lower mortality rates are observed for those with higher pensions. It is therefore often the preferred approach for mortality tables used to value liabilities. Nevertheless, to limit the cross-subsidisation across members, different mortality assumptions are also commonly generated for different segments (e.g. high and low income, as in the United Kingdom, or White and Blue collar, as in the United States).

Box 3.1. Mortality rates

'Mortality rate' is a general term, and it is not always clear to which mathematical concept it refers. There are a few key concepts to understand when discussing mortality rates.

The probability of death, often written as q_x , refers to the probability that someone aged x at the beginning of the year will die during the year. This is therefore calculated as the ratio of the number of deaths at age x during the year over the number of individuals at age x alive at the beginning of the year. This is the probability that standard mortality tables normally provide.

In contrast, the force of mortality, μ_x , is the instantaneous rate of mortality for someone at an exact age x . This concept is related to the central mortality rate, m_x , which is calculated as the number of individuals aged x that died during the year divided by the average number of individuals aged x during the year. This is the quantity often calculated as the 'raw mortality rate' because data for central exposures tends to be more commonly available, particularly for the general population.

Under the assumption that the force of mortality is relatively constant over the year and that deaths follow a Poisson distribution, since deaths are random and independent, the probability of death can be approximated with the central mortality rate as follows, though this approximation is less accurate at high ages:

$$q_x = 1 - \exp(-m_x)$$

The force of mortality and central mortality rate are often used interchangeably.

This chapter tends to use the generic term of mortality rate to avoid specifying which quantity is being modelled. Graduations and projections of mortality rates can reference either term, q_x or m_x , depending on the particular model chosen.

The most well-known model to capture the pattern of human mortality across ages is the Gompertz model. This model essentially assumes that mortality increases exponentially with age, or equivalently that the log of the mortality rates is linear. It provides a good fit for the ages that are most relevant for pensioners and annuitants, that is, older ages above around age 65.

There are several variations of the Gompertz model which aim to take into account varying patterns for different age groups. Makeham extended the Gompertz model to include a positive constant that better reflects the pattern of mortality for the middle ages, as it captures the excess mortality due to drivers such as accidents and infections that affect all ages (Ramonat and Kaufhold, 2018^[1]; CMI, 2015^[2]). Heligman-Pollard further extended the model to better capture the pattern of mortality at the youngest ages, in particular the declines in infant mortality and the "accident hump" reflecting the higher mortality of those in their late teens and twenties (CMI, 2015^[2]). Beard adapted the model to improve the fit at the oldest ages, in line with the argument that the most frail of the population die sooner, and therefore the mortality at the oldest ages does not increase as quickly because the individuals surviving to these ages tend to be the strongest and healthiest (Ramonat and Kaufhold, 2018^[1]).

The appropriate variation of the Gompertz model to use will depend in part on the ages for which experience is available. Often, for pensioner and annuitant populations mortality experience is only sufficient at the middle to older ages, in which case the Gompertz or Gompertz-Makeham model are likely to provide an adequate fit. If smoothing across a larger range of ages, alternative models may be more appropriate.

Another type of model that is commonly used to graduate mortality for pensioners and annuitants is the Whittaker-Henderson model. The Whittaker-Henderson model is a special case of p-spline model, which splices together cubic polynomials at specified intervals (b-splines) and applies a penalty function to increase the level of smoothness and avoid over-fitting. For the Whittaker-Henderson model, the intervals are specified over single years (Ramonat and Kaufhold, 2018^[1]).

The Whittaker-Henderson model has some advantages over Gompertz-type models. First, the model can fit the particular pattern of mortality observed across all age groups. Secondly, the graduations can also be done over two dimensions to obtain consistent mortality curves across both ages and years. However, it involves significantly more parameters, and requires more judgement by the modeller. At a minimum, it requires assumptions to set the desired balance between smoothness and fit and the order of the difference equation used to express smoothness.

The modeller can use statistical tests to aid in the selection of the best-fitting model. There are several tests and metrics commonly used in this regard. One is the Pearson's chi-squared test, which indicates the likelihood that the observed distribution was due to chance. The sign test shows whether the fitted curve is equally likely to fall above or below the observed curve, and the runs test extends this to show how often the sign of this difference changes to ensure that the fitted curve is capturing the observed shape. Information criteria are also commonly used to balance goodness of fit with the desire to have a parsimonious model with fewer parameters.

3.2. Extrapolating mortality to low and high ages

Where the observed mortality is only sufficient to be calculated for central ages, as is usually the case for pensioner and annuitant populations, it is necessary to make assumptions regarding the mortality rates at younger and older ages. This is normally done by extrapolating or interpolating the graduated mortality rates from the central age range to the extreme age ranges.

3.2.1. Modelling mortality for younger ages

Simple approximations are usually favoured for the extrapolation of mortality rates to younger ages below the central age range on which mortality experience was directly graduated (e.g. below around age 50 to 60). This is also because mortality rates at younger ages do not normally have a material impact on liabilities for pensions and annuities. For many pensioner populations, it may not even be necessary to have mortality assumptions for younger ages. These would only be needed in the context of a retirement income arrangement covering the active employed population or for survivor pensions where the beneficiary could be a younger spouse or dependent child.

The easiest approach is to simply assume that the mortality rates for younger ages are the same as those for the general population. The rates for younger ages could then be interpolated with the mortality rates for central ages. This approach takes the view that the selection effect observed for central ages, where the mortality of the pensioner or annuitant population is lower than that of the general population, is not material at younger ages. This could be a reasonable assumption where the drivers of socio-economic differences in mortality tend to manifest themselves more at older ages. Indeed, this is the case for causes of death such as heart disease and lung cancer linked to smoking. Nevertheless, other causes of mortality such as suicide or drug use that can have an impact on mortality rates at younger ages may demonstrate a socio-economic gradient.

A common approach to set mortality assumptions for younger ages is to base them on a ratio of the graduated mortality rates of the pensioner or annuitant population to some reference population. This could be based on the ratio observed at the youngest age included in the graduation of mortality for central ages, for example. The reference population could be either the general population, or another mortality table

developed for an insured population. This approach reflects an assumption that younger ages should demonstrate a similar selection effect as middle ages. It also maintains the shape of the mortality curve of the reference population for younger ages.

A less common approach for pensioner and annuitant mortality tables is to extend the graduation function used to smooth the central mortality rates to the youngest ages. This has the advantage of joining coherently with the central age range, however the selected model may not reasonably reflect the shape of the mortality curve observed elsewhere if it has not been fitted with the mortality for younger ages.

3.2.2. Modelling mortality for older ages

The mortality assumptions at older ages are particularly relevant for pensioners and annuitants. While the financial impact that these assumptions have for the valuation of liabilities may not be significant for a newly retired individual around age 65, the impact will increase with the age of the individual, and is relatively more material for deferred retirement income obligations beginning payments at older ages.

To establish mortality assumptions at the oldest ages above the central age range on which mortality experience was directly graduated (e.g. above ages 85-95), the chosen model needs to reflect the desired shape of the mortality curve at these ages. This requires forming a view not only regarding the extent to which mortality continues to increase with age and any maximum age to be attained, but also regarding the expected shape of the mortality curve of one population relative to another.

Significant uncertainty remains regarding the pattern of mortality at the oldest ages, even at the level of the general population. This is due first to the fact that the number of individuals who attain age 100, and more so age 110, is not sufficient to derive robust estimates of mortality, although this is gradually changing as individuals continue to live longer. Secondly, and likely more importantly, the data quality for these ages tends to be poor, with problems of delayed or misreported deaths and misrepresented ages. Given the difficulty of observing a clear pattern in mortality at these ages, opposing views regarding the pattern that it should follow at these oldest ages have emerged.

The main debate around the pattern of mortality at the oldest ages is whether mortality rates continue to increase exponentially with age, in line with the Gompertz model, or whether they decelerate at very old ages. Gavrilov and Gavrilova are the most cited proponents of the Gompertz model for old ages, beginning with their seminal work from 1991 (Gavrilov and Gavrilova, 1991^[3]). They later find that mortality increases exponentially until at least age 106 (Gavrilov and Gavrilova, 2011^[4]). In a study analysing data for supercentenarians (ages 110 and over), they conclude that the exponential model is still appropriate even for these very old ages (Gavrilov, Gavrilova and Krut'ko, 2017^[5]). The theory of an exponential pattern of mortality at old ages is also supported by the mortality experience of some other species, such as primates and rodents, for whom the Gompertz model demonstrates a good fit (Gavrilova and Gavrilov, 2014^[6]) (Bronikowski et al., 2011^[7]).

Nevertheless, other studies have disagreed that the exponential Gompertz model is appropriate for very old ages, finding evidence of a deceleration of mortality beyond around age 110 and indicating that a logistic model is more appropriate. One frequently cited study concludes that the annual probability of death actually plateaus at around 50% after age 110, implying a constant force of mortality of around 0.7 (Gampe, 2010^[8]). Similar patterns were confirmed in later studies. A study on several countries concluded that the force of mortality plateauing at around 0.8 for females and 1.2 for males, translating into a 55% and 70% annual probability of dying, respectively (Rau et al., 2017^[9]). Nevertheless, the authors note that this conclusion was stronger for females than for males, potentially due to a lower number of observations for the latter. Another study finds evidence of mortality deceleration in Canada (Ouellette and Bourbeau, 2014^[10]). Most recently, the Continuous Mortality Investigation in the United Kingdom concluded that mortality patterns in England and Wales do not follow a Gompertz pattern at very high ages, and that a

pattern of deceleration is more appropriate. They suggest to assume a force of mortality at age 120 of around 1 (CMI, 2017_[11]).

Gavrilov and Gavrilova argue that the observed deceleration is largely due to poor data quality at old ages (Gavrilov and Gavrilova, 2011_[4]). However, others note that they study the mortality on a cohort basis, which would tend to soften any observed deceleration due to the mortality improvements over time (CMI, 2017_[11]).

The view taken regarding the shape of mortality at the oldest ages implicitly informs the assumption around the ultimate age of the table. A mortality plateau at high ages implies that there is no maximum age of survival, but this is also a matter of much debate. One side argues that the human body is subject to biological limits that prevent it from surviving beyond a certain maximum age. This view was supported most recently in a study of various biomarkers, which concluded that beyond a certain age – somewhere around 120 to 150 – the body can no longer recover from negative shocks like illness (Pyrkov et al., 2021_[12]). But those of the opposing view point out that past estimates of maximum life expectancy have continually been disproven, and near-linear increases in the maximum life expectancy are consistently observed (Oeppen and Vaupel, 2019_[13]; Oeppen, 2002_[14]).

Another consideration when modelling old-age mortality is the expected relationship of mortality across different populations. Some evidence points towards a convergence of mortality with age (CMI, 2015_[15]). That is, pensioner and annuitant mortality will approach that of the general population, and the difference in mortality between different groups of the population, such as across socio-economic groups or genders, will reduce with age. The main argument behind a convergence of mortality at old ages is one of selection: the most frail tend to die earlier, and only the healthiest and strongest individuals survive to the oldest ages. As such, there is less heterogeneity in the population, and mortality converges.²

Regardless, the view taken regarding the expected pattern of mortality at older ages should inform the model that is used to derive the mortality assumptions at these ages. When taking the view that mortality continues to increase exponentially with age, the most common model is the Gompertz/Makeham model. When taking the view that it slows at high ages, the most common model used is the Kannisto model, which is a type of logistic model that assumes that the logit force of mortality is linear and converges to 1.

When modelling directly the mortality at the oldest ages for pensioner or annuitant populations, the model can be a different model than that used to smooth the mortality rates at central ages that better reflects the expected shape of mortality at the oldest ages. The most common option to do this is to calibrate the model to the central age range by regressing the model on the oldest ages of the central range, for example the last 10 to 15 ages, and then extrapolating mortality based on the fitted model. With this approach, statistical tests such as the Pearson's chi-squared test, the sign test, the runs test, and information criteria can aid in choosing the model that also fits the shape of the age range over which the model is regressed. The central age range and the old age range then need to be joined in some manner in order to have a smooth transition of mortality with age, which is typically done by blending or interpolating the two ranges.

Alternatively, mortality can be extrapolated directly with the chosen model from the oldest age of the central range by imposing certain constraints, such as the slope of the curve, to smooth the progression of mortality from the central to old age ranges. However, with this approach additional constraints – such as specifying the desired level of mortality at some maximum age – are usually needed to ensure that the shape of the mortality curve is also reasonable (CMI, 2015_[15]).

The argument of convergence favors the alternative approach of using population experience to set mortality assumptions at the oldest ages rather than modelling the old-age mortality of pensioners and annuitants directly. This is typically done by blending or interpolating the mortality experience from the central age range to the population mortality experience at old ages. In this way the shape of the mortality curve for pensioners and annuitants will be consistent with that of the general population. Nevertheless, the shape of the mortality curve for the general population will still have to have been established in light

of the considerations discussed above. This view will therefore inform the choice of which population mortality table the pensioner and annuitant rates should converge to.

For the sake of practicality, mortality tables include assumptions only up to some specified ultimate age regardless of the chosen model. Mortality probabilities at this age are generally fixed at 100% to have closure, even if the model selected implies a lower rate.

3.3. Accounting for selection in pensioner and annuitant populations

When there is not sufficient mortality experience for the pensioner or annuitant population, mortality assumptions must be based on an alternative population, such as the general population. However, pensioners and annuitants tend to have higher life expectancy on average compared to the general population, so the derived assumptions must be adjusted to reflect their lower mortality risk. These adjustments are referred to here as selection factors.

Lacking the pensioner or annuitant mortality experience in a particular jurisdiction, the difference between pensioner and population mortality in another country is often used as a basis for calculating selection factors. However, this involves a significant amount of judgement, as the underlying factors affecting the extent of selection can vary widely across jurisdictions. In particular, the proportion of the population that the pensioners or annuitants represent is a major factor impacting the level of selection. The smaller the population, the larger the selection effect. For example, the selection effect is generally larger for individual or voluntary arrangements compared to group or mandatory arrangements.

Selection can also vary across ages and time. Selection effects tend to be largest for the middle ages leading up to retirement age, increasing until around age 50-60, then decreasing again. For annuities, it is also expected to be higher in the initial years of the contract, as those who feel that they are in better health are also more likely to purchase an annuity. However, this effect wears off with time and mortality eventually converges to that of the general annuitant or pensioner population. At very high ages, selection may disappear altogether as the mortality of the pensioner and annuitant populations converge with that of the general population, in line with the frailty arguments that the surviving members of the population are all less frail on average.

Selection may also depend on gender, with male populations demonstrating larger selection effects than female populations. This is in line with the observation that the differences in life expectancy across subpopulations, such as different socio-economic groups, are generally larger for men than for women.

The most common way to account for selection is to apply a multiplicative factor to the mortality curve of the reference population. This is done by applying a reduction to the mortality of the reference population by multiplying it by a constant factor. Applying a single factor to all ages has the advantage of maintaining the original shape of the mortality curve, though it is less realistic and does not account for the convergence of mortality at high ages. Applying factors that vary across age may be more in line with the observed selection effects, however this approach may distort the shape of the resulting mortality curve if these effects are largest at middle ages. Similarly, applying a larger reduction to male mortality than to female mortality could result in males having a higher life expectancy than females, which is not a realistic scenario. Factors may therefore need to be adjusted to maintain coherent mortality curves that follow the expected or desired patterns and relationships.

Basing mortality assumptions on a proxy population that is expected to have a similar life expectancy to the pensioner or annuitant population is an alternative way to account for selection. This could be based, for example, on those in the population having a certain income level or occupation. Since the criteria used can be based on the actual characteristics of the pensioner or annuitant population, this approach involves less subjective judgement than using factors that are based on the experience of a different jurisdiction.

Another less common approximation is to use the reference mortality curve but assume a younger age for the pensioner or annuitant, effectively capturing the expected difference in life expectancy between the two populations. While a simplification, this approach maintains coherence between the mortality curves of the pensioners and the population.

Occasionally, selection factors are also applied to mortality improvement assumptions that are based on the general population to account for an expectation that the life expectancy of pensioners and annuitants will improve at a faster rate than that of the general population. Such assumptions, however, involve a significant level of judgement, particularly when there is not sufficient data on which to assess a mortality trend for the pensioner or annuitant population. As such, they are normally only used when it is preferable to err on the side of conservatism.

3.4. Projecting future mortality improvements

Assumptions regarding future mortality improvements are necessary to account for the continued future increases in life expectancy. This is needed to avoid underestimating life expectancy and to ensure that there will be sufficient assets to finance future pension and annuity payments. Improvement assumptions are most often expressed as an annual percentage reduction in the age-specific probability of death. In addition to varying by gender and age, they can also vary over time.

3.4.1. Data used

While base mortality assumptions can be calibrated directly to pensioner or annuitant data, there is rarely sufficient data for these populations on which to calibrate a mortality trend. A large amount of individual data and historical years of observation are required to ascertain robust trend assumptions.

In addition, populations of pensioners or annuitants may be more prone to changes in demographic composition than the general population that could make it difficult to assess the true underlying trend over time. This could be due, for example, to regulatory changes to the pension system such as expanding coverage to low-income individuals or different employment sectors, or removing any requirement to purchase an annuity at retirement. It could also be due to an economic shock that could impact the employment – and therefore pension coverage – of certain sectors or income groups. Any external shock that changes the demographic or socio-economic composition of the pensioner or annuitant population will have ramifications for the observed mortality trend of that population.

General population mortality is therefore normally used to derive mortality improvement assumptions. This means that different data sets are often used to calibrate the base assumptions and the future improvement assumptions for pensioners and annuitants. In this case, mortality improvement assumptions are developed separately and then applied to the graduated base mortality rates.

3.4.2. Mortality projection models

The types of models most often used to generate mortality improvement assumptions vary in terms of their complexity and functionality. The simplest approach is to apply a linear regression to historical mortality rates to derive the historical trend, and apply this trend going forward. Interpolative models, often using techniques more complex than linear regression to derive historical trends, have expanded on this approach to incorporate more judgement regarding expected future trends, and in particular the expected long-term rate of mortality improvement. Age period cohort (APC) models are commonly used extrapolative models that can project patterns of mortality by age and/or generation over time, and for the most part can also model stochastic projections. These have been more recently extended to accommodate stochastic projections for multiple populations simultaneously.

Simple regression models

The simplest, but also one of the most common, approaches to derive mortality improvement assumptions is to fit a linear regression to historical mortality rates. The regression is usually performed for individual ages or small groups of ages. The resulting trend can then be extrapolated forward for those ages into the future.

While this method is easy to execute and implement, it does come with a few disadvantages. First, regressing each age or age group separately means that there is no imposed relationship of improvements by age. As such, the initial pattern of mortality by age could become distorted over time, and result in unrealistic periodic mortality curves. Secondly, as with any extrapolative model, it assumes that past trends will continue indefinitely into the future. This may not be a realistic assumption, particularly as the drivers of the reductions in mortality have changed over time and impacted various age groups differently. For example, the large improvements in childhood mortality driven by increased vaccination have largely been realised in developed countries, and therefore will not likely continue at the same pace going forward. Similarly, the significant gains in life expectancy at older ages driven by medical advances and improvements in cardiovascular mortality have only emerged more recently, and would not necessarily be fully captured when regressing over a long range of historical data. Another example is the transition from a developing to a developed country, as occurred in South Korea over the past decades, when life expectancy caught up to the level observed in developed countries. Mortality improvements would be expected to slow down once life expectancy levels observed in other advanced economies has been achieved.

Interpolative models

While extrapolative models have the advantage of being objective, projecting forward historical trends indefinitely may not be realistic. Models allowing for more judgement and user input to shape future projections are therefore becoming increasingly common. In particular, mortality improvement assumptions are often assumed to converge to a lower long-term rate of mortality improvement for all ages, with the view that the high average improvements observed over the last several decades cannot be sustained indefinitely going forward.

The simplest variation on this approach is an extension of the linear regression discussed in the previous session, where the regressed trend is assumed to gradually converge – often linearly – to a long-term rate defined by the user.

More complex models smooth the historical mortality experience across ages and over time to derive initial mortality improvement assumptions. The Whittaker-Henderson model is one common approach used to smooth historical experience along two-dimensions. The smoothed improvement rates from the latest year(s) of historical data are used as the initial improvement rates, and are interpolated with a long-term rate defined by the user, often using polynomial interpolation. The length and slope of interpolation is defined by the model, and can incorporate convergence along birth cohorts in addition to along ages. The model developed by the Continuous Mortality Investigation (CMI) in the United Kingdom is the most well-known model of this type (Box 3.2).

Box 3.2. CMI mortality projections model

The Continuous Mortality Investigation (CMI), supported by the UK Institute and Faculty of Actuaries (IoFA), is widely regarded as being a leader in mortality research and modelling. Their mortality projections model (hereafter the 'CMI model') has been used as a reference for the development of standard mortality improvement assumptions in numerous jurisdictions beyond the United Kingdom. The CMI model is updated annually with the latest mortality experience, and allows for significant tailoring of inputs by the user to shape projections in line with their judgement and expectations for specific populations.

The CMI model adopts the following process to project future mortality improvements (CMI, 2021^[16]):

- Adjust historical population mortality experience for each year by adjusting the raw rates at high ages and smoothing out any observed anomalies
- Fit an Age Period Cohort Improvement model
- Interpolate mortality improvements from fitted improvements to the long-term improvement assumption along both period and cohort dimensions, where the age-period and cohort components of improvements are summed to obtain the total improvement
- Convert improvements to reference q_x rather than m_x (see Box 3.1).

The CMI model allows for the following user inputs:

- Long-term rate of improvement
- Addition to initial mortality improvements if recent improvements resulting from historical mortality are judged to be too low
- Period smoothing parameter, which controls smoothing by calendar year of historical data when fitting the Age Period Cohort model, and thereby how sensitive the model is to recent experience, in determining the initial improvement rates
- Slope of interpolation to long-term rate of improvement, to speed or slow convergence
- Length of convergence periods along age-period and cohort dimensions
- Weight given to individual years of experience, to allow for example to exclude the impact of the COVID-19 pandemic.

Convergence periods are shorter for younger ages, increasing at ages around retirement and reducing again for older ages. Improvements at high ages are assumed to converge linearly to zero between ages 85 and 110.

The biggest advantage of these types of models is that they allow users to adapt the model to align with their expectations regarding future mortality. Nevertheless, determining the value of all of the different input parameters requires significant judgement by the modeler. Some objective measures can be used to set their values, however, such as looking at the historical experience over a large historical period to set the long-term rate of improvement.

One of the main disadvantages of models is the level of user input complicates comparability. Indeed, where individual users are allowed to adapt the model, as with the CMI model, this reduces comparability when assessing the liability values for different providers.

Age period cohort models

Age period cohort (APC) models deconstruct the patterns of mortality along some or all of age, period and cohort dimensions. They can thereby project mortality in a way that should better reflect the expected

dynamics of the evolution in mortality compared to assessing the historical trend by individual ages or age groups. However, they remain extrapolative models, and therefore assume that the historical trends will continue indefinitely into the future.

These models fit a structure of mortality rates across ages, and can also capture patterns linked to the evolution of mortality for specific birth cohorts where this effect is included in the model. They are typically fitted to a two-dimensional range of historical data – by age and period – beyond which the parameter determining the mortality trend over time (aka 'kappa') is projected forward following the fitted trend. Box 3.3 explains the technical details of APC models. There are key trade-offs related to the decisions regarding the different components of these models that require additional description in order to understand their use.³

Box 3.3. Age Period Cohort (APC) mortality models

Hunt and Blake (2020^[17]) provide a useful formulation to understand the different components of commonly used extrapolative mortality models. These types of models can be written with the following structure:

$$\eta_{x,t} = \alpha_x + \sum_{i=1}^N \beta_x^{(i)} \kappa_t^{(i)} + \beta_x^{(0)} \gamma_{t-x}$$

Where:

- $x = \text{age}$, $t = \text{period}$, and $t-x = \text{year of birth}$;
- $\eta_{x,t}$ is a function transforming the mortality rate (see Box 3.1) to the form used for modelling;
- α_x is an age function that defines a constant shape of mortality across ages;
- $\kappa_t^{(i)}$ is the period term driving the trend of mortality over time, with $\beta_x^{(i)}$ (the age-period term) controlling the magnitude of the period effect for each age, which can be non-parametric and fitted freely by the model, or parametric and defined as a function of other variables;
- γ_{t-x} is a cohort term determining enduring mortality effects specific to a generation, with the magnitude of the cohort effect by age defined by $\beta_x^{(0)}$.

Two of the main ‘families’ of APC models are those following from the Lee-Carter model and those following the Cairns-Blake-Dowd model.

The Lee-Carter model takes the form:

$$\ln(\mu_{x,t}) = \alpha_x + \beta_x \kappa_t$$

It models the log of the force of mortality, includes an explicit age function, and has a non-parametric age-period term that is fitted freely by the model with no pre-specified structure.

The Cairns-Blake-Dowd model takes the form:

$$\text{logit}(q_{x,t}) = \kappa_t^{(1)} + (x - \bar{x}) \kappa_t^{(2)}$$

It models the logit of the probability of death, omits an explicit age function, and includes a parametric age-period term having a structure pre-defined as a function of age.

Both types of models are easily extended to include a cohort effect.

The first aspect of an APC model to consider is which mortality rate the model will project (Box 3.1). The choice of mortality rate should reflect the format of the available data on which the model is calibrated, and therefore also which distribution the number of deaths is expected to follow. If using central mortality rates calculated based on central exposures, deaths are normally assumed to follow a Poisson distribution. If using annual probabilities of death calculated based on the initial exposures for the year, deaths are assumed to follow a binomial distribution. The family of Lee-Carter (LC) models refers to the log of the

force of mortality, whereas the family of Cairns-Blake-Dowd (CBD) models uses the logit of the probability of death.

A second modelling choice is how to incorporate age effects and model the pattern of mortality across ages. The LC model includes an explicit age function that captures the constant features of the age structure of mortality over time. The age-period parameter that moderates the effect of period parameter across ages is a second freely set variable. In contrast, the CBD model eliminates the static age function and also replaces the free age-period parameter with one defined in advance as a linear function of age. The former approach can improve the fit of the model, because it has dedicated parameters that capture the shape of mortality by age both independent of time and over time. But the latter approach results in a more parsimonious model with fewer free parameters. It also allows for more flexibility to shape the model to fit certain expectations as to how time trends should impact mortality by age through the function for the age-period parameter. The LC model does not allow for this, as the single parameter groups the various drivers in past mortality trends that could change over time, so the model is not able to capture changing trends by age over time. This can also make it very sensitive to the historical period used for the calibration. However, the CBD approach of having a parametric age-period parameter potentially reduces the applicability of the model for certain age ranges. The CBD model itself, for example, is only appropriate for ages over around 50 where the assumption that mortality increases linearly with age holds.

Finally, there is a choice around the inclusion of a cohort effect. In general, this is optional and it should only be included if the historical data demonstrates clear patterns by cohort, that is certain cohorts demonstrating consistently higher or lower improvements. When relevant, the cohort effect should normally be secondary to the period effect, and therefore modeled more simplistically. To improve parsimony, the age-cohort parameter can simply be set to one. Nevertheless, cohort parameters can be challenging to fit, particularly for the youngest and oldest cohorts for whom there is less data. Furthermore, if the age-period parts of the model are poorly specified, the cohort term could simply capture the remaining noise and bias future projections.

APC models can normally accommodate stochastic projections, which are useful if the models will be used to assess longevity risk. The most common approach is to assume that the variables driving the periodic trend follow a random walk with a drift. Cohort parameters can also be projected stochastically. However, because the LC model only has a single period term, the changes in mortality each year are perfectly correlated across all ages, which is not a realistic outcome. In addition, the LC model results in relatively narrow confidence intervals. CBD-type models allow for more complex correlation structures and wider confidence intervals, and therefore may be better suited to assessing risk.

Multi-population models

While the models discussed up to now calibrate and project mortality improvements for a single population – usually a specific gender of the general population – multi-population models simultaneously project mortality for two or more related groups. These models are usually extensions of the stochastic APC models discussed in the previous section, though some are extensions of deterministic approaches involving regression and smoothing.

There are two main reasons for using multi-population models rather than single-population models. The first is to overcome the lack of data for a small population of interest, such as a small pensioner population, by modelling its mortality in reference to a larger, but related, population. The second is to ensure coherent mortality projections for related populations, such as for males and females, or any other sub-populations of a larger population where one group is expected to consistently have higher mortality than another.

Modeling small data sets in reference to a larger population can result in more robust estimates of future improvement assumptions. Relatively small data sets, even at the general population level, can be problematic for calibrating mortality projection models, as the higher levels of volatility in the historical data can make long-term projections highly sensitive to the choice of input data. Nevertheless, calibrating a

multi-population model still requires a sizeable data set for the target population, which is not often the case for annuitant and pensioner data sets. Villegas et al. (2017_[18]) suggest that 8-10 years of historical experience with an annual exposure of 20 000-25 000 individuals for the target population is necessary to calibrate a multi-population model.

Numerous extensions of the Lee-Carter model to multiple populations have been proposed. There are three main approaches to doing so (Villegas et al., 2017_[18]):

1. Calibrate two models separately for the reference and target populations, then assess their dependence
2. Assume a common parameter that drives the periodic trend for both populations, along with population specific parameters, or the “Joint- κ model”
3. Jointly estimate the two models using co-integration techniques

There are various considerations in choosing which approach to go with. The first approach ignores interdependence, so additional assumptions are still required regarding the relationship of the trend of the two populations to obtain coherent and integrated projections (Villegas and Haberman, 2014_[19]). The second approach is the most common and is more transparent, parsimonious, and consistent across populations. It also allows for first calibrating a model for the reference population, and subsequently calibrating a model for the target population, which is appropriate where the reference population is substantially larger. However, simplified versions have as a disadvantage that the target and reference populations will always experience the same mortality improvements, which is unrealistic. Extensions have therefore included an additional term to allow for stochastic deviations in the mortality of the reference population from that of the target population, even if the trends for the two populations tend to converge in the long run. For the third approach, joint-estimation of the parameters is difficult with short data sets, and is better adapted to larger sets of data where the two populations are of similar size, which is normally not the case.

Academics have also proposed similar extensions of the Cairns-Blake-Dowd model (Villegas et al., 2017_[18]). Additional approaches, such as the Saint model used in Denmark, aim to model the spread between the two populations directly, while limiting any long-term divergence of the mortality of the two populations (Jarner and Kryger, 2013_[20]).

Model selection should also consider how the model will be used. If the purpose is only to establish mortality assumptions with which to value liabilities, there is no need to model additional deviations between the reference and the target populations. However, if there is a need to assess longevity risk, and in particular the risk of deviations in experience between the two populations, the model needs to allow for a non-perfect correlation between them. The additional advantages and drawbacks of the Lee-Carter and Cairns-Blake-Dowd models discussed in the previous section also apply here.

A final drawback of multi-population models currently used in practice is that they most often assume no long-term divergence between the target and the reference populations. Therefore, alternative time series models would need to be considered if the two trends are expected to diverge, as could be the case with a high-income annuitant population relative to the general population. This would increase the complexity of the model, as well as the judgement involved in its calibration.

3.4.3. Old-age improvements

There is not usually sufficient historical data on which to calibrate mortality improvements for very old ages, typically over 85-90, even for the general population. As such, setting assumptions at these ages requires significant judgement regarding the magnitude of improvements and the pattern of improvements across ages.

There is some evidence of positive mortality improvements at the oldest ages. Japanese women – who have long been world leaders in terms of life expectancy – experienced accelerating mortality improvements for ages 80 to 99 since the 1960s, reaching annual improvements exceeding 3% in the early 2000s (Rau et al., 2008^[21]). Positive improvements have been observed in Japan even beyond age 100, with women aged 100 to 104 experiencing improvements exceeding 1% (Robine, Saito and Jagger, 2002^[22]). Accelerating patterns, albeit at lower magnitudes, were also observed for age 80 to 99 in East Germany and Italy (Rau et al., 2008^[21]). Combined experience in France, Japan, Switzerland and Sweden from the 1980s to the 1990s also show significant positive improvements, though declining with age, with females and males aged 95-99 experiencing a material average annual mortality decline of 1.25% and 0.9%, respectively (Vaupel, Rau and Jasilionis, 2006^[23]).

Nevertheless, positive improvements at the oldest ages have not been observed in all countries. Women over 90 in the United States do not seem to have experienced any improvements over the 1990s, though improvements picked up slightly over this period for men (Rau et al., 2008^[21]). Mortality for centenarians in the United States seems to have plateaued since the 1950s (Gavrilov, Gavrilova and Krut'ko, 2017^[24]). Mortality has also plateaued for centenarians in Sweden and the United Kingdom (Drefahl et al., 2012^[25]) (CMI, 2015^[15]). Slightly negative improvements have been observed at the oldest ages in Canada (Adam, 2012^[26]).

There does, however, seem to be consistent evidence that mortality improvements decline with age for the oldest ages. A common approach for mortality tables is therefore to impose a pattern for this decline, often simply a linear convergence to 0% at a certain age. Assuming no mortality improvements beyond a certain age is consistent with the view that there is a limit to life expectancy and that we will not observe increases in the ultimate age of mortality. Otherwise, mortality improvements could be assumed to reduce to some positive constant value.

An alternative approach when APC models are being used would be to extrapolate the calibrated age effect to extend to older ages (Dowd, Cairns and Blake, 2019^[27]). This allows for less subjectivity in setting the assumptions for high ages, and results in future projections of mortality that remain consistent across ages.

3.5. Recent innovations applied to the development of mortality assumptions

Some recent proposals to establish mortality assumptions for pensioner or annuitant populations have sought to exploit developments in technology and data analysis to overcome some of the limitations of existing models and the challenge of the lack of data. The proposals are using advanced techniques, often employing machine learning, to improve both the calibration of base rates and the projections of future mortality.

Some approaches aim to inform the development of the base mortality rates. One example looks to overcome the lack of annuitant mortality data. The methodology used to develop a mortality table for annuitants in Cambodia relied on data science to train various models using insured lives mortality tables from the region combined with macroeconomic variables such as GDP, which is strongly correlated with life expectancy (Yeo Chee Lek, 2020^[28]). Another example applies machine learning techniques to improve the assessment of differences in mortality across socio-economic groups (Wen, 2019^[29]). The analysis groups geographic areas in England by their common socio-economic characteristics in order to model difference in mortality for these groups. The techniques aid in the selection of the most relevant variables on which to base these groupings, and therefore leads to more homogenous groupings than simply ranking the regions by decile would produce.

Innovative proposals are also being put forward to improve the estimation of future mortality. One uses machine learning to assess the adequacy of fitted models through back testing to better identify their

shortcomings and improve their fit (Deprez, Shevchenko and Wüthrich, 2017^[30]). Similarly, another proposal employs machine learning algorithms to better identify historical patterns in mortality and improve the goodness of fit of a Lee-Carter model (Levantesi and Pizzorusso, 2019^[31]).

3.6. Concluding remarks

Developing mortality tables for pensioners and annuitants involves several steps. To calculate the base mortality rates, the raw mortality rates are graduated at central ages. The rates must then be extrapolated to younger and older ages. The assumptions must also account for any selection effect and the difference in mortality between the pensioner or annuitant population and the population on which the estimated rates were based.

Assumptions for future mortality improvements are also necessary. Numerous models exist to project future mortality, including simple regression models, age period cohort models such as the Lee-Carter and Cairns Blake Dowd models, interpolative models such as the CMI model, and multi-population models. Separate assumptions are necessary for improvements at the oldest ages since historical data at these ages are limited. A common approach is to assume a linear decline to zero at a terminal age.

Developing mortality tables for pensioners and annuitants therefore requires numerous modelling choices. These decisions involve trade-offs with respect to model complexity and the level of judgement required. They also require taking a stance on expected mortality patterns, both current patterns for ages where less data may be available as well as how future mortality improvements will evolve. Choosing the appropriate model will therefore always require a certain level of expert judgement.

Understanding the trade-offs involved and what they imply for the expected mortality patterns should help regulators and supervisors assess whether the process to establish the mortality tables for pensioners and annuitants is appropriate for a given context.

References

- Adam, L. (2012), *The Canadian Pensioners Mortality Table: Historical Trends in Mortality Improvement and a Proposed Projection Model Based on CPP/QPP Data as at 31 December 2007*. [26]
- Bronikowski, A. et al. (2011), "Aging in the Natural World: Comparative Data Reveal Similar Mortality Patterns Across Primates", *Science*, Vol. 331/6022, pp. 1325-1328, <https://doi.org/10.1126/science.1201571>. [7]
- CMI (2021), *CMI_2 020 v01 methods*. [16]
- CMI (2017), *A second report on high age mortality*. [11]
- CMI (2015), *Initial report on the features of high age mortality*. [15]
- CMI (2015), *Report of the Graduation and Modelling Working Party*. [2]
- Deprez, P., P. Shevchenko and M. Wüthrich (2017), "Machine learning techniques for mortality modelling", *European Actuarial Journal*, Vol. 7/2, pp. 337-352, <https://doi.org/10.1007/s13385-017-0152-4>. [30]

- Dowd, K., A. Cairns and D. Blake (2019), "A Simple Approach to Project Extreme Old Age Mortality Rates and Value Mortality-Related Financial Instruments", *SSRN Electronic Journal*, <https://doi.org/10.2139/ssrn.3552230>. [27]
- Drefahl, S. et al. (2012), "The era of centenarians: mortality of the oldest old in Sweden", *Journal of Internal Medicine*, Vol. 272/1, pp. 100-102, <https://doi.org/10.1111/j.1365-2796.2012.02518.x>. [25]
- Gampe, J. (2010), "Human mortality beyond age 110", in *Demographic Research Monographs, Supercentenarians*, Springer Berlin Heidelberg, Berlin, Heidelberg, https://doi.org/10.1007/978-3-642-11520-2_13. [8]
- Gavrilova, N. and L. Gavrilov (2014), "Biodemography of Old-Age Mortality in Humans and Rodents", *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, Vol. 70/1, pp. 1-9, <https://doi.org/10.1093/gerona/glu009>. [6]
- Gavrilov, L. and N. Gavrilova (2011), "Mortality Measurement at Advanced Ages: A Study of the Social Security Administration Death Master File", *North American Actuarial Journal*, Vol. 15/3, pp. 432-447. [4]
- Gavrilov, L. and N. Gavrilova (1991), *The Biology of a Life Span: A Quantitative Approach*, Harwood Academic Publisher. [3]
- Gavrilov, L., N. Gavrilova and V. Krut'ko (2017), *Historical Evolution of Old-Age Mortality and New Approaches to Mortality Forecasting*, Society of Actuaries. [24]
- Gavrilov, L., N. Gavrilova and V. Krut'ko (2017), *Mortality Trajectories at Exceptionally High Ages: A Study of Supercentenarians*, Society of Actuaries. [5]
- Hunt, A. and D. Blake (2020), "On the Structure and Classification of Mortality Models", *North American Actuarial Journal*, Vol. 25/sup1, pp. S215-S234, <https://doi.org/10.1080/10920277.2019.1649156>. [17]
- Jarner, S. and E. Kryger (2013), "Modelling Adult Mortality in Small Populations: The Saint Model", *ASTIN Bulletin*, Vol. 41/2, pp. 377-418, <https://doi.org/10.2143/AST.41.2.2136982>. [20]
- Levantesi, S. and V. Pizzorusso (2019), *Application of Machine Learning to Mortality Modeling and Forecasting*. [31]
- Oeppen, J. (2002), "DEMOGRAPHY: Enhanced: Broken Limits to Life Expectancy", *Science*, Vol. 296/5570, pp. 1029-1031, <https://doi.org/10.1126/science.1069675>. [14]
- Oeppen, J. and J. Vaupel (2019), "The Linear Rise in the Number of Our Days", in *Demographic Research Monographs, Old and New Perspectives on Mortality Forecasting*, Springer International Publishing, Cham, https://doi.org/10.1007/978-3-030-05075-7_13. [13]
- Ouellette, N. and R. Bourbeau (2014), *Measurement of Mortality among Centenarians in Canada*, Society of Actuaries. [10]
- Pyrkov, T. et al. (2021), "Longitudinal analysis of blood markers reveals progressive loss of resilience and predicts human lifespan limit", *Nature Communications*, Vol. 12/1, <https://doi.org/10.1038/s41467-021-23014-1>. [12]


- Ramonat, S. and K. Kaufhold (2018), *A Practitioner's Guide to Statistical Mortality Graduation*, Society of Actuaries. [1]
- Rau, R. et al. (2017), *Where is the level of the mortality plateau?*, Society of Actuaries. [9]
- Rau, R. et al. (2008), "Continued Reductions in Mortality at Advanced Ages", *Population and Development Review*, Vol. 34/4, pp. 747-768, <https://doi.org/10.1111/j.1728-4457.2008.00249.x>. [21]
- Robine, J., Y. Saito and C. Jagger (2002), *Living and dying beyond age 100 in Japan*, Society of Actuaries. [22]
- Vaupel, J., R. Rau and D. Jasilionis (2006), *The remarkable, accelerating decline in mortality at older ages and the prospects for further improvement in life expectancy*. [23]
- Villegas, A. and S. Haberman (2014), "On the Modeling and Forecasting of Socioeconomic Mortality Differentials: An Application to Deprivation and Mortality in England", *North American Actuarial Journal*, Vol. 18/1, pp. 168-193, <https://doi.org/10.1080/10920277.2013.866034>. [19]
- Villegas, A. et al. (2017), "A comparative study of two-population models for the assessment of basis risk in longevity hedges", *ASTIN Bulletin*, Vol. 47/3, pp. 631-679, <https://doi.org/10.1017/asb.2017.18>. [18]
- Wen, J. (2019), *Factor-Based GLM Model with Socio-Economic Inputs*, https://www.actuaries.org.uk/system/files/field/document/ARC%20workshop%202019_S3_Wen.pdf. [29]
- Yeo Chee Lek, N. (2020), *Cambodian Insured Lives Mortality with Data Science*, <https://theactuarmagazine.org/cambodian-insured-lives-mortality-with-data-science/>. [28]

Notes

¹ Getting to this point involves a significant amount of work to clean the data and calculate the correct number of deaths and exposures, but this chapter does not cover the details of these preparatory steps for modelling mortality.

² This same argument supports a model of decelerating mortality at older ages, with the frailer members of society passing away earlier and leaving a more strong and homogenous group of survivors with lower mortality.

³ The discussion that follows largely draws from Hunt and Blake (2020_[17]).



From:
Mortality and the Provision of Retirement Income

Access the complete publication at:
<https://doi.org/10.1787/a10a6c09-en>

Please cite this chapter as:

Mosher, Jessica (2023), "Modelling choices for the development of standard mortality tables", in OECD, *Mortality and the Provision of Retirement Income*, OECD Publishing, Paris.

DOI: <https://doi.org/10.1787/abf3c67f-en>

This work is published under the responsibility of the Secretary-General of the OECD. The opinions expressed and arguments employed herein do not necessarily reflect the official views of OECD member countries.

This document, as well as any data and map included herein, are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area. Extracts from publications may be subject to additional disclaimers, which are set out in the complete version of the publication, available at the link provided.

The use of this work, whether digital or print, is governed by the Terms and Conditions to be found at <http://www.oecd.org/termsandconditions>.