

Checking the Pulse of Mortality

E. General Actuarial Topics

Vishwajeet Sanyal
Nivedita Hanra
Priya Sibal

With growing competition, companies do realize that products need to be priced competitively and scientifically. This has led to an increased interest in the estimation and modeling of the mortality assumption. Different schools of thought exist for the interpretation of mortality behavior:

1. Some believe that mortality will continue to decrease in a way that it has decreased historically
2. Another school of thought believes that the trend might not be in line with past experiences as medical advances are a one time event and hence should not be included in the forecasts. They believe that mortality decline is a one time event and will not occur in future.
3. In contrast, it is believed that future medical advancements and technologies would delay ageing and increase lifespan of the population
4. Others debate over:
 - a. The existence over a biological limit of life
 - b. and whether we are approaching this limit

Whatever the belief, bottom line is that mortality needs to be modeled and estimated as close to actual experience as possible so that offices have a logical basis for pricing their products.

The risks involved in estimating mortality are as follows:

1. Mortality trend risk– This arises due to the uncertainty of the pattern that mortality rates would follow in future. Historical experience is a good yardstick to estimate the assumption, but it needs to be augmented with future expectations to arrive at the best estimate.
2. Mortality sampling risk– Sampling risks are generally associated with small samples, for which the estimation error could be high. For a large population, due to averaging, the estimated volatility is quite close to the mortality experienced. Thus, for a small population, one must bare in mind that:
 - a. A large estimation error is associated with any estimate of mortality based on actual experience
 - b. The successive year's experience can be far from expectation even if the true mortality rate is known with certainty.
3. Mortality jump risk– This refers to sharp movements in mortality rates from one year to the next. Actual mortality rates tend to jump to higher levels (due to epidemics, natural disasters, terrorism etc) rather than lower levels (due to medical advances, healthier lifestyles etc).
4. Mortality volatility risk– This risk is due to the annual fluctuation in mortality rates that arise due to the underlying process of mortality and noise in data collection. The impact of volatility on exposures is lower than what might be expected from crude mortality data due to smoothing across ages and smoothing over time.
5. Errors in estimating mortality rates at higher ages- The estimation of population and mortality data at high ages has always been subject to significantly increased error as:

- Birth records may be unavailable and individuals may not know their true age for accurate census reporting
- An individual may also knowingly inflate his age.
- For death data, it is frequently the case that age at death is unknown, or is inaccurately estimated by family members
- At high ages, despite increased mortality rates, the lower population size means that the absolute number of deaths in a year is typically small.

To some extent, this can be coined into a question of whether a forecasting model should rely on statistical analysis alone or expert opinion based in biology.

Forecasting methods

We summarize some of the most commonly used methods in order of increasing reliance on expert opinion.

Extrapolation: deterministic & stochastic

An extrapolative model calculates

- Estimates of future mortality using the current level of mortality
- An estimate of the rate of change in future mortality

The rate of change is based upon the changes observed in the recent to medium-term past and is assumed to continue over the forecast horizon.

Extrapolative models could be deterministic or stochastic. The Lee-Carter (1992) mortality projection model is example of a stochastic extrapolative model. Though this approach is mainly objective, it relies on a subjective judgments also.

Cause-specific extrapolation

Cause-specific models disaggregate total mortality by cause of death and forecast mortality rates for each cause separately.

Projection using relational models

Here the mortality profile of the target population is believed to be the mortality profile that the forecasted population will achieve over some future time horizon. Generally, this is applied under the assumption that mortality forecasts are considered likely to trend toward the mortality already observed in a longer-lived population. This approach, therefore, is more applicable to developing rather than developed countries.

Extrapolation with expert opinion

Using any mathematical model blindly without judgment and without an understanding of the context can be dangerous so Expert opinion by the forecaster in respect of future mortality is inevitable. This explicit expert opinion could be simply about the ultimate level of mortality, or that mortality improvements follow a specified path through time. In both the aggregate mortality rates and specific cause of death this approach is applicable

Cause-delay

These models are similar to relational models except that it targets a different sector of the mortality for the same population. It assumes improving mortality for a specific cause means delaying some fraction of the deaths from that cause.

For example, an eight-year delay in mortality by heart disease could be estimated by shifting from the observed mortality rates by heart disease to those for an age eight years younger; so for 58 year olds, the relevant mortality rates would be the pertinent estimates for 50 year olds.

Epidemiological

The approach is similar to cause-delay but here the focus is on the impact of specified risk factors such as smoking, obesity, socio-economic status and specific diseases rather than causes of death.

Deterministic and stochastic mortality models

Many research methods have evolved as a result of the need to analyze and forecast mortality and longevity as accurately as possible. Two types of models can be used:

1. Deterministic models and
2. Stochastic models

Deterministic models are simpler than stochastic models wherein mortality rates are based on deterministic scenarios and analysis by cause of death. However, these methods tend to underestimate mortality improvements thus giving inflated rates

Stochastic models are much more popular as they involve time series analysis with parameters estimated from historical mortality rates.

Stochastic mortality models

Stochastic models can be evaluated on the basis of their general characteristics and their capability to explain historical patterns of mortality.

Before understanding the underlying principles of each model, we give below the notation to be used.

Notation

1. $\beta_x^{(i)}$ function reflect age – related effects where x denotes age
2. $\kappa_t^{(i)}$ function reflect period – related effects where t denotes period or time
3. $\gamma_c^{(i)}$ function reflect cohort – related effects where $c=t-x$ denotes birth cohort

Lee-Carter

This is an extrapolative model formulated by Ronald Lee and Lawrence Carter in 1992 where the model parameters were calibrated to historical mortality experience. It is a one-factor model of the mortality

surface (age and time), where the general level of mortality is assumed to be a stochastic process in time, and where the age distribution of mortality is assumed to be deterministic and is calibrated from historical data. The Lee-Carter model for mortality can be expressed as:

$$\mathbf{Log\ } m(t,x) = \beta_x^{(1)} + \beta_x^{(2)} \cdot K_t^{(2)}$$

Where the $\beta_x^{(1)}$ coefficients describe the age-specific pattern of mortality (I.e. the current mortality table), $K_t^{(2)}$ coefficients describe the changes through time and the $\beta_x^{(2)}$ coefficients describe the sensitivity of $\beta_x^{(1)}$ to $K_t^{(2)}$

In terms of stochastic forecasting models, Lee-Carter (LC) is the simplest implementation of a range of models that incorporate varying degrees of complexity in order to model perceived age, period and cohort effects in the data. The disadvantage of more complicated models is that they can require careful calibration and the larger number of parameters can make it more difficult to gain an intuitive sense of how the model reacts to changes in parameters.

Two aspects of it are

- It is based on the opinion that future mortality will continue to improve at the same rate as in the past.
- The forecaster must choose the historic period over which mortality improvement rates should be analyzed to serve as the basis for forecasting.

Renshaw-Haberman

The Renshaw-Haberman model is a generalized version of the Lee-Carter model which includes a cohort effect (Renshaw & Haberman, 2006). It can be expressed in the form:

$$\mathbf{Log\ } m(t,x) = \beta_x^{(1)} + \beta_x^{(2)} \cdot K_t^{(2)} + \beta_x^{(3)} \cdot Y_{t-x}^{(3)}$$

Where the addition of the final term relative to LC is used to model a cohort effect in terms of a dynamic which is dependent on year of birth, $Y_{t-x}^{(3)}$. The impact of this cohort effect can be varied by age through $\beta_x^{(3)}$.

Currie Age-Period-Cohort Model

Currie (2006) proposes what amounts to a simplified version of the Renshaw-Haberman model, where the age, period and cohort effects influence mortality rates independently. The model can be expressed in the form:

$$\mathbf{Log\ } m(t,x) = \beta_x^{(1)} + K_t^{(2)} + Y_{t-x}^{(3)}$$

P-Splines

In the UK, the P-Spline model has been particularly high profile in the last year due to the release of a free software implementation by the CMIB. The P-Spline approach is essentially a penalized fitting process using basis splines (Currie et al., 2004). Basis splines are a set of basis functions constructed from cubic splines. They are fitted to the underlying data through the optimization of a penalized likelihood or regression function and have the form:

$$\text{Log } m(t, x) = \sum_{i,j} \theta_{ij} \cdot B_{ij}^{ay}(X, t)$$

Where $B_{ij}^{ay}(X, t)$ are B-spline basis functions and θ_{ij} are weights

Associated with each function, the fitted set of splines can be extrapolated forward in time enabling forecasting. By measuring volatility in the estimation process, confidence levels around the central mortality estimate can also be produced. It is important to note, however, that the methodology does not enable direct stochastic simulation of future mortality rates, although a common misperception is that this is possible. Simulation of the actual number of deaths (based on the Poisson process applied to the central mortality estimate) is of course possible.

A cubic B-spline consists of cubic polynomial pieces bolted together on the plots at points known as knots. The B-spline pieces are continuous, and have continuous first and second derivatives at the join points on the plots.

This model can be fitted with standard software since the Poisson distribution together with the linear structure for $\text{Log } m(t, x)$ defines a generalized linear model and the regression coefficients θ_{ij} are chosen by maximum likelihood.

Cairns, Blake & Dowd

Cairns, Blake and Dowd fit a model direct to initial mortality rates instead of central mortality rates (Cairns et al., 2006). The simplest version, of the model, has the form:

$$\text{Logit } q(t, x) = \kappa_t^{(1)} + \kappa_t^{(2)} (x - \mu)$$

Where, μ is the mean age over the range of ages used in the analysis.

There are a number of extended versions of the model, with varying additional terms to cover a cohort effect with a linear, quadratic or time-dependent variability. These models differ from the previous stochastic models (Lee-Carter, Renshaw-Haberman, and Currie) in that they assume a functional relationship between mortality rates across ages. They mainly are :

$$\text{Log } m(t, x) = \beta_x^{(1)} + n_a^{-1} \kappa_t^{(2)} + n_a^{-1} \gamma_c^{(3)}$$

$$\text{Logit } q(t, x) = \kappa_t^{(1)} + \kappa_t^{(2)} (x - \mu) + \kappa_t^{(2)} \{(x - \mu)^2 - \sigma_x^2\} + \gamma_c^{(4)}$$

$$\text{Logit } q(t, x) = \kappa_t^{(1)} + \kappa_t^{(2)} (x - \mu) + \gamma_c^{(3)} (x_c - x)$$

The latter two models developed by Cairns aims the time and the cohort effects to be taped more effectively.

Generalized Linear Model

A GLM can model Mortality as a function of age, amount of benefit etc. Linear model is given as

$$Y_i = \mu_i + \text{error}$$

μ_i based on linear combination of measured factors - which factors, and how best combined, is to be derived as

$$\mu_i = \alpha + \beta \cdot \text{age}_i + \gamma \cdot \text{age}_i^2 + \delta \cdot \text{height}_i \cdot \text{age}_i$$

The mathematics practically means:

Probability of death in year =

Base level for observed population ×
Factor 1 (based on age) ×
Factor 2 (based on sex) ×
Factor 3 (based on amount) ...

Here each factor is a series of multiplicative coefficients and all factors considered simultaneously, allowing for correlations in the data automatically. This allows for nature of the random process involved and provides information about certainty of result

So the model results would be :

- One number (the base level – everything else will be relative to this – eg the mortality of a 65 yr male with pension £1000)
- A series of multiplicative coefficients for Factor 1 (age)
- A series of multiplicative coefficients for Factor 2 (sex)
- A series of multiplicative coefficients for Factor 3 (amount) ...

It is preferred to look at these multiplicative coefficients via graphs! As the factor results are only valid in their totality – we cannot take results for one factor and use those in isolation. Mortality rates (etc) presented as 'absolute' numbers (qx etc) GLMs, results are shown as multiplicative relativities and allows us to have year of birth as a factor at the same time as age. Often we find no year of birth effect but still get much information from the calendar year factor. We can also investigate the calendar year factor trends for broad age/sex groups – often find substantially different results

Forecasting with stochastic mortality models

We start with analyzing the population size and demographics. How the exposure evolves as per the mortality rates. We compile the data, understanding the shortcomings in it. We analyse the components of the mortality risk and Graduate the data for forecasting mortality.

We take the stochastic mortality models which, on the basis of fitting to historical data, appear to be suitable candidates for forecasting future mortality for the age group under consideration, and prepare them for forecasting. To do this, we need to specify the stochastic processes that drive the age, period and (if present) cohort effects in each model.

Random-walk processes have been widely used to drive the dynamics of the period effect ever since the introduction of the original Lee-Carter (1992) model. The method used to estimate the model has been refined by subsequent authors in order to improve the fit and place the model on more secure statistical foundations.

The principal challenge we face in building a stochastic mortality model that can be used for forecasting lies in specifying the dynamic process driving the cohort effect. A simple random-walk process is unlikely to be appropriate and various alternative stochastic processes that might be suitable for the different models need to be analysed. As with previous studies (e.g., Renshaw and Haberman,

2006, and CMI, 2007), it is assumed that the cohort effect, has dynamics that are independent of the period effect.

The age effects are either non-parametric or estimated from historical data as done in model number 1, 2 and 3 or assume some particular functional form as assumed by Cairns. Further, we focus on forecasts of mortality within the same range of ages used to estimate the underlying models, so it is not necessary to simulate or extrapolate the age effects.

Following Cairns, Blake and Dowd (2006b), we use a multivariate random walk with drift to drive the dynamics of the period effect. This model appears to be consistent with the data for England and Wales. However, more general ARIMA models might provide a better fit statistically to some datasets.

For example, CMI (2007) uses an ARIMA(1,1,0) process for the period effect in the Lee-Carter model (M1) and an ARIMA(2,1,0) process for the period effect in the Renshaw and Haberman model (M2).

In order to tap the dynamics of the cohort effect we use different ARIMA processes varying basically the differencing. Corresponding to each differencing being used we get the BIC. The ARIMA model giving the highest BIC is the one used to model the cohort effect under a particular model. Here BIC stands for Bayes Information Criterion which is calculated for the kth model as

$$BIC_k = l_k - 1/2 * n_k * \log N$$

Here l_k is the model maximum likelihood of the kth model
 n_k is the number of parameters and N is the number of observations.

The models have been evaluated quantitatively, but this is not sufficient and other qualitative criteria which need to be kept in mind are as follows:

- quality of fit, as measured by the Bayes Information Criterion (BIC);
- ease of implementation;
- parsimony;
- transparency;
- incorporation of cohort effects;
- ability to produce a non-trivial correlation structure between ages;
- robustness of parameter estimates relative to the period of data employed.

Some models fared better under some given criteria as compared to others and no single model can claim superiority under all the criteria considered. Thus

- There remain a large number of potentially valid stochastic mortality models, despite significant conceptual differences between them.
- The model choice depends on what priority the model user attaches to each of the assessment criteria.

Model selection

In general, no mortality projection basis can ever be considered "correct" as many of the projection methods discussed here suffer from potential drawbacks. So we select the right approach and model to project mortality rates based on:

- Data and their reliability- we have to bear in mind that regardless of the method of projection used, when projecting mortality at very old ages several problems arise, in particular, because of inaccuracies in the data available and variability due to small exposures to risk.
- Resources available for the project- univariate extrapolation of the parameters of a mortality model can be misleading, and while a multivariate time series model for the parameters is possible, it can lead to computational intractability. While the methodology suggested by Lee & Carter (1992) avoids these problems, it implicitly assumes - by using the ordinary least square to estimate the parameters - that the errors are homoskedastic, which is unrealistic in that the logarithm of the observed mortality rate is much more variable at older ages.
- Purpose for which the projection is required- whether the model should capture long-term trends or "tail-event" shocks. For example, if the forecaster is concerned with the impact of a possible pandemic, a model specifically designed to incorporate catastrophic "shock" mortality would be appropriate. As there is limited historical data on shocks, expert opinion should generally be incorporated (Van Broekhoven, 2002). Additionally, modelers should review the qualitative tradeoff between simplicity and accuracy. Cairns et al. (2007) provide a list of qualitative criteria that should be considered in the model selection process.
- Ease of scenario testing- A significant uncertainty surrounds mortality forecasts, hence the models should include "high-medium-low" scenarios, in producing stochastic forecasts

Simulation of mortality risk

Stochastic models for simulating mortality risk vary depending on the nature of the mortality exposure and its time horizon over which the risk is being measured. For example, for simulations over short time horizons it is generally not necessary to model trend risk in a sophisticated manner, as its short-term impact is small. In a similar fashion, when simulating mortality for large populations it is not necessary to model the sampling risk, but this must be incorporated in simulations for small populations.

Model risk and parameter risk

When simulating mortality, another risk arises that we have not discussed above, namely, model risk – the risk that we have incorrectly chosen the model. The different types of risks associated with this are:

1. The risk of choosing the wrong model
2. Parameter estimation risk – This refers to the uncertainty associated with the calibration of the model, in which the model's parameters are fitted to appropriate data
3. Trend Risk – This is the risk associated with the long term pattern associated with the mortality rates and the parameter estimation risk.

Mortality catastrophe risk v/s longevity risk

The nature of the simulation may also change depending on the exposure. For example, the risk associated with mortality catastrophe (or "cat") bonds (e.g., Swiss Re's Vita bonds) requires very different simulation models from the longevity risk associated with a pension plan or annuity. Mortality cat bonds are relatively short-term investments that are subject to losses if mortality rates jump significantly higher than current levels. For example, the original Vita bond issued in 2003 was a three-year investment in which investors started suffering losses if a customized mortality index exceeded 130% of the 2002 baseline level. To capture this jump risk requires a so-called "jump-diffusion" model of mortality rates which captures the probability of large jumps in

mortality due to pandemics, natural disasters, etc. along with the mortality volatility (see, for example, Lin and Cox, 2006). By contrast simulation models for the longevity risk of pension plans and annuity books usually do not need to incorporate jump risk, as it has limited impact on the downside value of the liabilities. However, simulation models for these exposures must incorporate trend risk, as well as the parameter estimation risk associated with other parameters in the model.

Historical simulation of mortality

A particularly simple simulation method that can be applied to mortality rates is known as 'historical simulation'. Historical simulation involves first calculating the relative changes in historical mortality rates and then applying those changes to current base mortality levels to generate future scenarios (see, for example, Mina and Xiao 2001). It is widely used by financial institutions in risk simulations of F(X) rates, interest rates, equity prices, etc.

It can be applied in two different ways (Coughlan, Kolb and Emery, 2003):

- "Longitudinal historical simulation", which produces a single path for longevity over a long time horizon
- "Cross-sectional historical simulation", which produces many scenarios for longevity over just a single period (the next year)

As an example, suppose we have 41 annual historical mortality rates together with the current rate: q_{1965} , q_{1966} ... q_{2005} . Here the subscript denotes not age but year (period). The first step in historical simulation is to calculate the relative changes (or 'returns') in historical rates, which can be done in a number of different ways. A common way of calculating these relative changes is to take the natural logarithm of the ratio between consecutive mortality rates as follows:

$$r_t = \log(q_{1965+t} / q_{1965+t-1}), \text{ where } t = 1, 2, 3, \dots, 40$$

The next step is to apply these changes to the current mortality rate. The result is a series of forward-looking scenarios for the mortality rate, denoted z_t .

$$z_{2006} = \exp(r_1) \cdot q_{2005}$$

$$z_{2007} = \exp(r_2) \cdot z_{2006}$$

$$z_{2008} = \exp(r_3) \cdot z_{2007} \quad \text{etc.}$$

This procedure generates a single scenario path for this mortality rate at each future time point over the next 40 years: z_{2006} , z_{2007} ... z_{2045} . Because the scenarios define a single path for future rates we call this longitudinal historical Simulation. There is another way of generating scenarios from the relative changes in historical rates, which produces multiple scenarios but for just one future date. This is called cross sectional historical simulation and involves applying all the historical relative changes in rates to the current mortality rate.

For example:

$$z_{2006, 1} = \exp(r_1) \cdot q_{2005}$$

$$z_{2006, 2} = \exp(r_2) \cdot q_{2005}$$

$$z_{2006, 3} = \exp(r_3) \cdot q_{2005} \quad \text{etc.}$$

This procedure generates 40 scenario paths for this mortality rate at one future date in one year's time: $(z_{2006, 1})$, $(z_{2006, 2})$... $(z_{2006, 40})$

Historical simulation is a simple and transparent technique for simulating future mortality rates. It is particularly useful in short term (1-year) simulations to capture volatility in mortality rates, but it has its limitations. In particular, it cannot account for jump risk, unless such a jump has occurred in the historical data set, and it assumes that historical volatility will be repeated. Moreover, for long-term

simulations the method only produces one path (at best a small number of slightly shorter paths) and, therefore, does not permit the calculation of long-run statistical risk metrics.

Other stochastic simulation models

We gave an overview of a number of different stochastic mortality models that can be used to produce simulations of possible future paths for mortality rates over the long term. These paths enable us to make an assessment not only of the expected level of mortality rates far into the future, but also of the risk associated with those rates.

Characterizing the nature of the exposure

Life Insurance

The expected future experience of the policyholders will depend crucially on three things:

- the target market for the contract – this will be dependent on the distribution channel involved
- the underwriting controls applied (or not applied)
- the expected change in the experience since the time of the last historical investigation to the point in time at which the assumption will on average apply (typically you would be looking about ten to fifteen years into the future).

For example, suppose you want to price a new with-profits whole life contract. With-profits whole life and long-term with-profits endowment assurances may have similar mortality experience.

However, move a little further (*eg* to a pure endowment) and different underwriting procedures and target markets may make experience much more significantly different.

Move even further (*eg* to a term assurance, where the customer needs differ greatly) and experience will almost certainly be very different.

Homogeneous groups are likely to be derived categorising according to:

- age
- sex
- medically underwritten (normal) \vee medically underwritten (abnormal) \vee not medically underwritten
- smoker \vee non-smoker
- distribution channel
- some form of broad grouping by occupation and/or geographical area.

It should also be noted that the rates are more important if the premiums are guaranteed (non-reviewable).

That is, it is more important that we get the mortality assumption correct when there is no chance of correcting our mistakes later.

Health Insurance

Separate rates of mortality, among healthy lives for Income protection, Long term care insurance and Critical Illness and among lives in claim for IP and LTCI needs to be kept. Special consideration needs to be given to the survivorship requirement for stand-alone CI. Where policy conditions do not have a significant death benefit, mortality rate should not be overestimated, else the whole contract could be

undercosted. As a high level of mortality reduces the cost of sickness claims in such cases, so it is clearly important not to overstate the mortality rate assumptions.

It is usual for the effect of mortality during the survivorship period to be allowed for by an adjustment to the claim incidence rate. The adjustments will be specific to the different disease categories (as the post-diagnosis survivorship can differ quite dramatically for the different diseases).

Mortality rates should be taken from the recent experience of a credible body of policyholders for the same contract, though mortality is less significant than morbidity which gives rise to benefit.

Reducing mortality rates by 10% will increase the population exposed to risk of critical illness by *far less than* 10%. On the other hand, claim costs will be much more sensitive to mortality improvement after claims have started (*ie* for IP and LTCI policies in payment).

There are two main influences at work:

- The difference between the total claim cost between those that do or do not die after starting benefits may be very significant, whereas someone who dies in the non-claiming state may never have claimed anyway
- The mortality rates of claiming (*ie* sick) policyholders are much higher than for non-claiming (*ie* healthy) policyholders, so the absolute impact of, say, a 10% reduction in mortality will be much greater.

So the significance of the mortality assumption does depend on whether we are considering non-claiming or claiming policyholders. This therefore has an impact on how much effort we would normally expend on determining the assumption. For stand-alone CI, IP and LTCI, the non-claimants' mortality would merit the least research, and we would probably base the rates on a standard mortality table – quite possibly for assured lives – but with adjustments to allow for full future mortality improvements.

For IP and LTCI claims in payment, the mortality assumption would merit much more effort in order to estimate the future rates with more confidence. This is because the assumption has a bigger financial impact, and because the mortality of sick lives is a lot less predictable than for the population as a whole. The subdivisions of data used would probably be the same as for claim recovery rates. Claim termination rates for given ages can be obtained by combining the appropriate recovery and mortality rates, if required.

Employee Benefits

The value of the benefits is reviewed every year. This keeps the risk associated with mortality very low. This risk further gets adjusted due to the risk being distributed over the group. Mortality risk over the group is also low because the people in the group do not come together for the purpose of insurance which solves the problem of anti selection. Anti selection is also low because nobody can opt for higher sum assured than what the scheme rules allow. The members of the group are expected to be healthy being actively at work. In Group life insurance the mortality is estimated at inception based on the type of industry, geographical location etc.

A "long" exposure means that one will benefit from an increase in the underlying variable and a "short" exposure means that one will benefit from a fall in the underlying variable. According to this terminology, DB pension plan holders (or more accurately, their sponsors) are naturally "short longevity", meaning that they benefit if longevity falls and lose if longevity increases. Alternatively, one

could also describe the DB pension plan as being “long mortality”, since they benefit if the mortality rate increases. Similarly, the annuity business of an insurer has the same short exposure to longevity as a DB pension plan. On contrary for Group term insurance portfolio has the opposite exposure and is “long longevity”, or equivalently, it is “short mortality”.

Defined benefit pension plans are post retirement scheme. Here the level of benefit paid is prescribed by the plan and paid for whole life. Although at first glance the longevity risk associated with a DB pension plan is the same as that for a simple annuity portfolio, it is in fact more complicated than that. A typical DB pension plan also has other longevity and mortality-linked features that alter the risk profile from that of an annuity portfolio. In particular, most DB plans are exposed to additional longevity and mortality risks through the following features:

- Spouse’s pension
- Dependents’ benefit
- Death benefit

These additional dependencies on longevity and mortality complicate the overall exposure of the pension plan to longevity and mortality risk and need to be taken into account. The population of members of a pension plan is often classified into three groupings:

- “Pensioners” or “retirees”: former employees who have retired and are drawing a pension
- “Actives”: employees who are still working and accruing benefits
- “Deferred” or “deferred vested”: employees who have accrued pension benefits in the past but have not yet retired

Longevity risk results in annuity payments are made for longer than expected. By contrast, interest rate risk impacts on the value of the annuity and not the cash flows. In situations where interest rates are stable and at relatively low levels, such as in Western Europe, North America, Japan, etc., longevity is the dominant risk for immediate annuities. This is because the combination of low interest rate volatility and the relatively short interest-rate duration of immediate annuities lead to a relatively modest impact on the volatility of the annuity value. However, for deferred annuities it is interest rate risk that is generally the dominant risk, because of the much longer duration of these contracts.

Evaluating the impact of longevity and mortality risk

Essentially risk management is to evaluate the impact of key risk factors of the exposures. For a DB pension plan it requires evaluation of sensitivities of its Liability in cash flows and their value to changes in longevity. Similarly, for a life insurance book, evaluation of the sensitivity in cash flow and value to changes in mortality. The starting point for this is a set of forecasts for the evolution of future mortality rates which are derived from both of the following:

- A simulation of future paths for mortality using a stochastic mortality model, and
- A set of stress tests in which the best estimate path of mortality is shocked in adverse way

These paths are then applied to the exposures to determine the impact on cash flow and value, along with the risk to exposures in terms of both cash flow and value.

There are different types of risk and different ways to classify them. One simple classification pertains to the current situation as the distinction between financial and demographic risks.

1. Financial Risks - Risks based on the movements in prices of assets that are traded in the financial markets including
 - a. Interest Risk
 - b. Inflation Risk

- c. Credit Risk
 - d. Equity Risk
 - e. Other Investment Risks
2. Demographic Risks - Risks associated with an underlying population, have until recently only been traded in the insurance markets. They comprise of:
- a. Longevity Risk
 - b. Mortality Risk
 - c. Early Retirement Risk
 - d. Lump sum Election Risk
 - e. Risk of Early Surrender – Here, demographic risk involves policyholders “surrender” their policies early before they mature
 - f. Prepayment Risk – these refer to the risks for mortgage portfolios and mortgage-backed securities (MBS).
3. Population Basis Risk - Risk facing pension plans and life-contingent insurance products, this is the risk associated with the difference in longevity/mortality experience between two populations. . For example, populations with differing profiles of age, gender, socioeconomic group, lifestyle and geography will have differing profiles of longevity and mortality risk. The longevity and mortality risk for these specific populations can be decomposed into two components as:
- a. Risk associated with entire population
 - b. Basis risk of sub population relative to the entire population

Practical applications of mortality models

National and international bodies make periodic population projections. These involve estimation of future mortality rates and, here, we review their methodologies in these terms. In addition, we review the practice of the actuarial profession in both the UK and the US.

United States - Social Security Administration

Objective: US Social Security Administration (SSA) needed projection of mortality for Americans participating in the Insurance and pension programs

Method: The model for projecting mortality in the SSA's reports is an extrapolative model which includes expert opinion. The rates are based on cause-specific death rates. Specifically, the trend in mortality improvements by age, sex, and cause of death are first calculated for the past 20 years. These reductions in mortality are assumed to continue for a short period, such as two years. After that period, the reductions in mortality are projected to trend down to some predetermined ultimate improvement rate, which applies in, say, 30 years time. By applying these projected improvements to the current mortality rates, future mortality rates and hence life expectancy can be estimated. Therefore, these projections assume that the decline in mortality will not continue at its current pace, but will slow and then level off. In the SSA projects three separate scenarios of future mortality improvement to take into account the potential variability of future forecasts: an expected case, a high mortality case, and a low mortality case.

Result: As discussed in the 2005 Annual Report from the SSA's Board of Trustees (2005), the ultimate mortality improvement rates are 0.71% in the expected case, 0.33% for the high case and 1.23% for the low case.

United Kingdom - Office for National Statistics

Objective: Government Actuary's Department (GAD) produces semi annual population projections for the United Kingdom.

Method: The GAD determines the current level of mortality and then projects future rates of mortality improvements. These improvements are initially set to the recent average of mortality improvements, and then trend down to a fixed level of 1% per annum from 2029 onwards. This trend is assumed to be more rapid than a linear trend for males and less rapid than a linear trend for females. Projections are undertaken by cohort for those born before 1960 and by calendar year (i.e., period) for the remainder. The GAD projects high and low variant scenarios. In these, the final fixed improvement rates are set to 2% and 0% instead of the 1% assumed in the base scenario. In addition, the initial improvements are assumed to be 2% higher and 2% lower than the base level respectively.

Global - World Health Organization

Objective: To estimate future mortality over a period of 30 years for approximately 100 countries worldwide.

Methodology: Models based on epidemiological structures. Mortality rates were assumed related to factors specifically: average income per capita, average number of years of education, time (as a proxy for scientific and technological advancement) and tobacco consumption, using historical observation between factors and cause-specific mortality rates

Projections prepared by the actuarial profession

United Kingdom - Methods developed by the CMIB

Objective: Since the publication of the a(55) tables in 1953, the CMIB has customarily provided mortality projections for annuitants and pensioners when it has published mortality tables. For the "80" and "92" Series (two sets of mortality tables based on the 1979-82 and 1991-94 experiences respectively) an extrapolative projection method was used to introduce an explicit formula for the reduction factor to allow for mortality projection for annuitants and pensioners.

Methodology: To estimate future rates of mortality at time t they multiply the mortality rate from the base table ($qx, 0$) by its corresponding reduction factor $RF(x, t)$. The reduction factor being expressed as follows:

$$RF(x, t) = \sigma(x) + [1 - \sigma(x)] \cdot [1 - f_n(x)]^{t/n}$$

Observations: The most recent set of projections (those proposed to be used with the "00" series) were distinctive in three important aspects as compared to the corresponding sets for the "80" and "92" Series:

- They recognized the cohort effect; defined as the dependence of mortality improvement rates on a person's year of birth

- Three alternative projections were offered, instead of the traditional single projection, based on alternative assumptions on the longevity of the cohort effect
- Base projections were not directly quoted but software for their analysis was provided

United States - Methods developed by the Society of Actuaries

Objective: To recommend a projection basis to be used for statutory reserving with the GAR94 tables for annuitants.

Method: The Society of Actuaries adopted an extrapolative approach similar to that used by the CMIB for its "80" and "90" series. The 1994 mortality tables were published as a base table with associated age-dependent mortality improvement factors (Scale AA_x). The projected probabilities of death were calculated as follows:

$$q_x(y) = q_x(1994) (1 - AA_x)^{y-1994}$$

End of Paper

About the Authors

Vishwajeet Sanyal

Vishwajeet Sanyal joined the Watson Wyatt Benefits Practice in Mumbai in September 2008. Prior to his appointment, he spent eight years in LIC of India as an Analyst in Employee Benefits. Since joining Watson Wyatt, Vishwajeet has been involved in the full range of consulting projects undertaken by the practice. These cover valuation of various employee benefits schemes, financial reporting, survey Documentation for clients across west India. For his experience in Employee Benefits Vishwajeet gets involved in many technical projects and related presentations

Risk and value capabilities

- Technical inputs into converting the valuation approach from the Attained Age Method to the Projected Unit Credit Method for Gratuity, Pension and Leave Encashment schemes at LIC.
- Analysed the Group Health Insurance Market for LIC's Venturing into Group Health Insurance.
- Product development for various annuities

Project experience

- Participation on Clients Seminar for training the Clients in financial reporting of Employee Benefits
- Developed tools for expediting valuation and reporting.
- Involved in checking and correcting Benefit survey documents and survey on assumption setting.
- Has extensively taken part in product design talks for various group schemes.

Education & professional qualification

- B.Sc. (Maths) degree from the University of Mumbai, India
- Member of the Institute of Actuaries of India and London

Nivedita Hanra

Nivedita Hanra joined the Watson Wyatt Benefits Practice in Mumbai on 28 January 2008. Prior to her appointment, she worked as an actuarial associate with LIC and an actuarial analyst in another Actuarial firm. Since joining Watson Wyatt, Nivedita has been involved in a range of consulting projects undertaken by the practice. These cover actuarial valuations under various employee benefit schemes, accounting for employee benefits under AS15 (revised, 2005) and IAS19, consulting to domestic and international clients on employee benefits issues.

- Risk and value capabilities**
- Actuarial valuation of schemes like Gratuity, Leave encashment and other Long term employee benefits.
 - Consulting Clients on various employee benefits issues.
- Project experience**
- Undertaken the 2008-2009 Benefits Survey of various Defined Benefits plans most prevalent in India: Plan Provisions and Assumption Setting.
 - Identified actuarial assumption setting of Attrition rate for Capita Offshore Pvt Ltd using Statistical tools.
 - Comparative analysis for the purpose of Insurance provider selection.
- Education & professional qualification**
- Post Graduate Diploma in Actuarial Science from S.N.D.T University 2006-2007.
 - BSC in Statistics from University of Mumbai 2003.
 - Member of Institute of Actuaries India and UK.

Priya Sibal

Priya Sibal joined the Watson Wyatt Employee Benefits Practice in Mumbai, India in March 2008. Since joining Watson Wyatt, Priya has been involved in the full range of consulting projects undertaken by the practice. These cover actuarial valuations, due diligence projects and consulting to domestic and international clients on employee benefits issues.

- Risk and value capabilities**
- Actuarial valuation of all types of employee benefit schemes
 - Actively involved in internal process developments and quality assurance
- Project experience**
- Due diligence for employee benefits conducted for an M&A deal between a German and Indian company
 - Liased with the local offices for an international client for conducting their global benefits audit
- Education & professional qualification**
- BA Economics & Statistics, St. Xavier's College, Mumbai
 - MBA Actuarial Science, NMIMS University, Mumbai