B.A.J. **15**, Supplement, 267-283 (2009)

SELECTED ISSUES IN MODELLING MORTALITY BY CAUSE AND IN SMALL POPULATIONS

By S. J. Richards

ABSTRACT

Actuarial practice as regards mortality analysis and projection is changing rapidly. This paper provides a short introduction to some of the limitations and risks in using trends in cause of death as a means for projecting future mortality rates. It also covers recent developments in analysing the mortality of smaller populations, including survival models and "piggyback" models.

KEYWORDS

Cause of Death; Forecasting; Survival Models; Mortality Laws

CONTACT ADDRESS

Stephen Richards, 4 Caledonian Place, Edinburgh, EH11 2AS, U.K.
Telephone: +44 (0)131 315 4470; E-mail: stephen@richardsconsulting.co.uk;
Web: www.richardsconsulting.co.uk

1. Introduction

1.1 This paper looks at selected issues in actuarial practice regarding mortality and longevity, and illustrates some of the points raised in the workstream "Mortality: drivers for change" during the inter-disciplinary conference which took place in Edinburgh in October 2009. The paper introduces the main limitations and risks for actuaries in attempting to create mortality projections based on extrapolating trends by cause of death. The paper also looks at the limitations of using data which is collected and managed for a purpose besides mortality analysis.

1.2 Actuaries in the private sector tend to work with clearly defined sub-groups of the wider population. This paper introduces some of the background to the increasing use of survival models for stand-alone portfolio analysis, and also the "piggyback" models used when relating the mortality of a smaller group to a much larger one.

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

2. Population Trends by Cause of Death

2.1 Examining past trends in cause of death can be very instructive. As a result, some actuaries have attempted to extrapolate trends in causes of death to create a forecast of future mortality rates. This has an initial appeal: using a more detailed breakdown of mortality data feels like it ought to result in a better quality forecast. However, separate projections by cause of death need to be recombined to produce all-cause rates. This is particularly tricky, since we are projecting correlated time series. For example, smoking increases mortality due to both heart disease and numerous cancers. A further complication is that the same projection model might not be suitable for all causes of death:

> "*Causes of deaths characterized by well-defined cohort effects, such as lung cancer need to be forecast using model[s] which incorporate cohort factors.*" Di Cesare & Murphy (2009)

2.2 Trends by cause of death are not independent, and some causes of death may be linked in poorly understood ways. For example, Azambuja & Levins (2007) postulated links between influenza and coronary heart disease, as suggested by the apparent correlation in Figure 1.

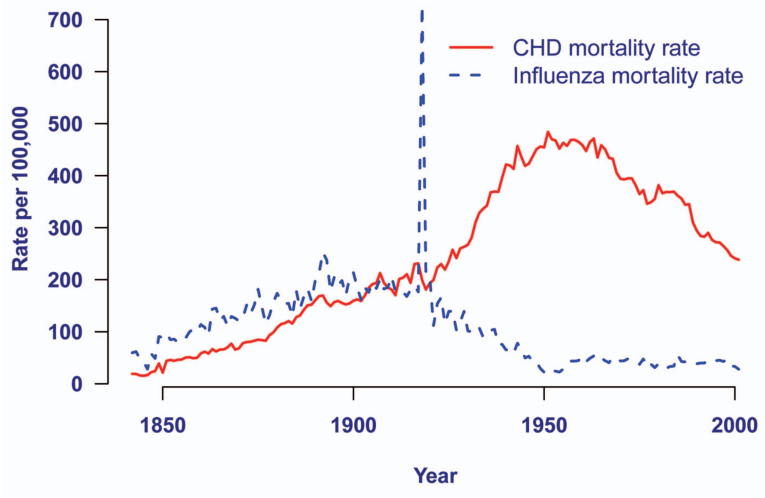


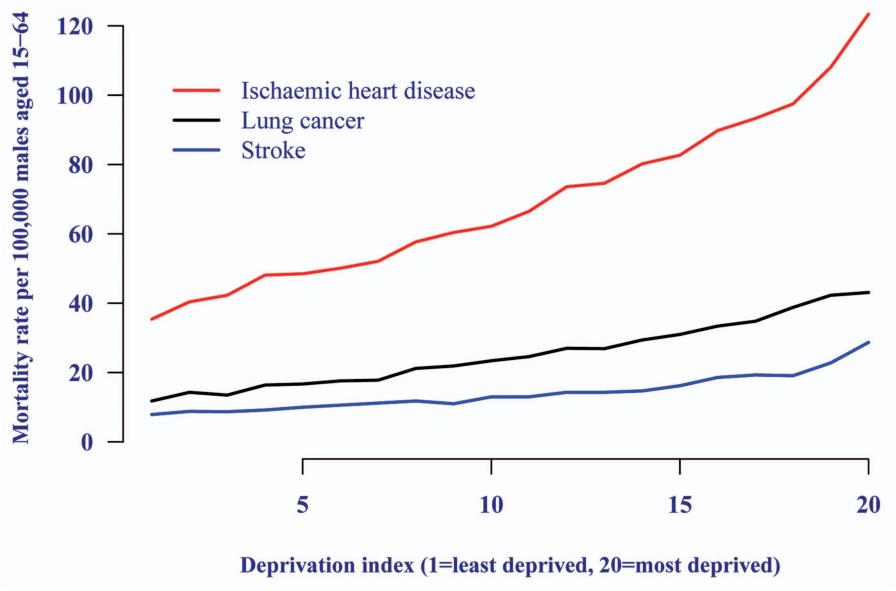
Figure 1. Mortality rates from influenza and coronary heart disease in Massachusetts. Styled after Azambuja & Levins (2007). Data from Massachusetts Department of Public Health "Registry of Vital Records and Statistics"

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

2.3 Even within a cause-of-death group, Azambuja (2009) points out that "the term CHD encompasses more than one disease" and as a result that "CHD mortality trends represent a varying combination of types of CHD over time". If the underlying composition of an important cause of death is changing, this increases the uncertainty over the reliability of any projections of that cause of death.

2.4 CMI (2004) gives a good overview of some of the technical problems with cause-of-death forecasts, concluding that it is preferable to use all-cause mortality data. In a review report the Government Actuary's Department (2001) provides a more detailed review, leading to the same conclusion.

2.5 Another stumbling block is that mortality rates by cause of death are strongly linked to socio-economic group, which is an important risk factor in actuarial work. For example, Romeri *et al.* (2006) show the link between cause-of-death rates and deprivation index, and we reproduce some of their data in Figure 2.



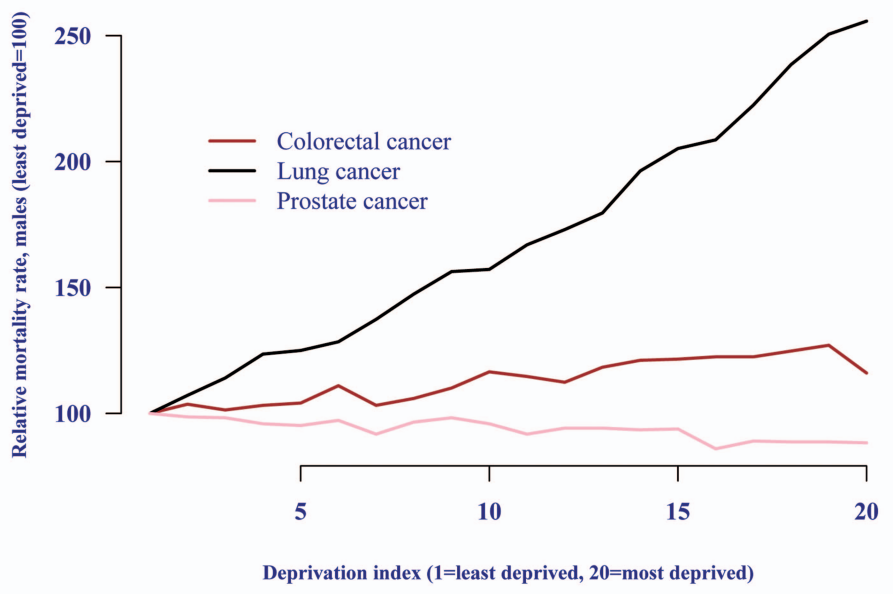
Source: ONS data for males aged 15-64 in England and Wales

Figure 2. Mortality rates by deprivation index for selected major causes of death. The lives who account for the largest proportion of actuarial liabilities have a different cause-of-death mix than the wider population

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

2.6 The implication of Figure 2 is that cause-of-death statistics need to be broken down by socio-economic group or deprivation index before being used for forecasting purposes for actuarial work. Failure to do this could give misleading results, since the various sections of society appear in cause-of-death statistics in different proportions to their presence in the population.

2.7 Figure 2 does at least raise the hope that differences by deprivation index might be similar across the various causes of death, i.e. that the rates might be in some kind of constant proportion. However, Figure 3 shows that this is not the case for the three most important cancers for males in England and Wales. Figure 2 shows that separate handling of cause-of-death data by deprivation index or socio-economic group is necessary, but Figure 3 shows that the relationship is not simple.



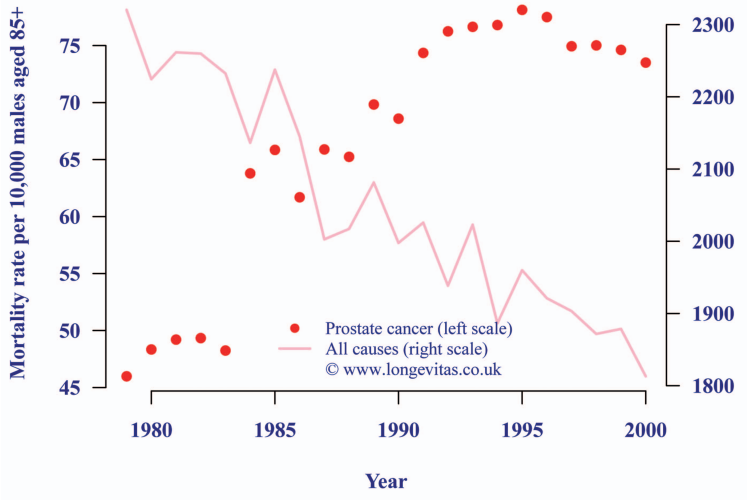
Source: ONS data for males of all ages in England and Wales

Figure 3. Relative mortality rates by deprivation index for selected major cancers. There is no simple relationship between cause incidence and deprivation index: lung-cancer rates rise sharply with deprivation, whereas prostate-cancer rates fall slightly. Projections based on extrapolating rates by cause of death in the population may be misleading for annuity liabilities, which are concentrated in the least-deprived sub-group

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

2.8 There is a corollary to this and it is linked to concentration risk, i.e. the tendency for the least deprived members of society to have the biggest pensions and sums assured. Richards (2008) notes that around half of all pensions in payment are paid to around 10% of the pensioner population, while Figures 2 and 3 show that the least deprived members of society have both different mortality rates and a different cause-of-death mix from the most deprived. Simply put, people with the largest liabilities in a pensioner portfolio have a different cause-of-death mix than the wider population.

2.9 When projecting all-cause mortality rates it is common for people to ask what sort of changes in causes of death might be required to achieve a particular scenario. Often one is asked to posit what causes of death have to be "eliminated", and the results from such calculations might lead to the mistaken conclusion that a particular projection is unlikely and therefore too prudent. However, the linkages between causes are complex; indeed, reducing mortality due to one cause of death can lead to an increase in mortality due to another. One example is prostate cancer, which was historically considered rare but by 2008 had become the fifth most common cause of death amongst males above age 85 in England and Wales (and the most common cancer cause of death, exceeding even lung cancer). Mortality rates for this age group are shown in Figure 4.



Source: ONS data in "20th Century Mortality", a CD containing cause-of-death data for England and Wales

Figure 4. Mortality rates per 10,000 males over age 85 in England and Wales

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

2.10 Figure 4 shows that although all-cause mortality rates in this age group have fallen steadily, the mortality rate due to prostate cancer has clearly risen. Prostate cancer is slow-growing, and it is often described as something a male is more likely to die *with* than *from*. As other causes of death reduced, however, the likelihood of dying from this cause actually increased. Unless this sort of linkage is allowed for, using cause-of-death "elimination" calculations make scenarios appear more prudent than they are. Such calculations are sometimes used to communicate mortality projections, but this example shows that this could lead to a false level of comfort.

2.11 A final observation on Figure 4 is that there is a clear discontinuity in prostate-cancer mortality rates between 1983 and 1984, which is suggestive of a change in diagnosis and/or classification. This serves to underline the uncertainty of the data on which cause-of-death projections are based. The impact of shifting rules for cause-of-death classification are described by Aylin (1997) as follows:

> "*Trend analysis spanning the years either side of 1984 and 1993, must take into account some important coding changes. There is a large increase in mortality from chronic diseases […] between 1984 and 1993. This is an artefact due to changes in the way ICD-9 rules for selecting the underlying cause of death were interpreted in England and Wales. […] As a result, some deaths for which bronchopneumonia in Part I of the certificate would previously have been coded as the underlying cause of death were coded to a condition mentioned elsewhere in Part I or Part II.*"

2.12 In other words, some of the "trends" by cause of death are a result of changes in classification methodology. When doing projections by cause of death we must not only take great care with socio-economic differentials, and also worry about projecting correlated time series, but we must also take note of uncertainty surrounding the classification of cause of death itself.

3. Alternative Data Sources for Trends by Cause of Death

3.1 Besides population data, there are other data sources which actuaries might seek to use, such as clinical records recorded by doctors. An example in the UK is the General Practice Research Database (GPRD). However, actuaries should always bear in mind the so-called First Law Of Informatics espoused by van der Lei (1991):

> "*Data can only be used for the purpose for which it is collected.*"

which was restated in a milder form by van Ginneken *et al.* (1993) as follows:

> "*Using data for different purposes for which they were recorded carries the risk of erroneous interpretation of these data, unless those data permit unambiguous interpretation.*"

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

3.2 Van der Lei's injunction may seem a bit harsh, but he was writing specifically about computerised medical records, which historically were not as unambiguous as one might think. For example, Jordan *et al.* (2004) conducted a systematic review of morbidity coding in computerised general practice records and found that:

> "*[…] quality of recording varied between morbidities. One reason for this may be in distinctiveness of diagnosis, e.g. coding of diabetes tended to be of higher quality than coding of asthma.*"

3.3 There are good reasons why clinical databases may not be suitable for mortality research or forecasting, and these lie in the primary purposes to which the databases are put and doctors' motivation to enter data:

> "*To make a written diagnosis readable for a database, it must be coded in some way, and this was often not done. Whilst this might seem a bit sloppy, it is best to view it from the perspective of the purpose of the record. This data was recorded as an aid to patient care and fee claims, and not as a research database.*" Martin (2010)

3.4 As a result, computerised medical records can contain all sorts of recording bias, a prominent example of which was given by Martin (2010):

> "*[A] GP is more likely to record a history of smoking than non-smoking, not least because non-smokers are less likely to be sick. If a well person never sees the doctor, their smoking status may never be recorded. Consequently, estimates of smoking rates from GP databases are likely to be skewed upwards if the denominator is taken as the number of smoking entries, or downwards if the denominator is taken as the number patient records. Statistical corrections and imputations can be made, but this changes the status of the data from observation to approximation.*

4. Practical Difficulties with Cause-of-Death Projections

4.1 The challenges for cause-of-death projections described in the preceding sections are compounded by an issue of particular concern to actuaries: bias. Bias is generally undesirable in any forecasting method when applied to financial calculations. However, it is particularly undesirable when projecting future mortality rates for reserving for pension liabilities, especially if there is a bias towards over-stating mortality rates which would lead to under-reserving. Unfortunately, this kind of bias appears to be a specific feature of cause-of-death projections:

> "*Mortality projections disaggregated by cause of death have been found in practice to be more pessimistic than those without disaggregation […]. The reason is straightforward: over time the overall trend becomes dominated by the trend for those causes with the slowest decline.*" Wong-Fupuy & Haberman (2004)

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

4.2 Cause-of-death methods are sometimes described as having the ability to incorporate expert medical opinions. However, this is often not the advantage it appears to be:

> "*The advantage of expert opinion is the incorporation of demographic, epidemiological and other relevant knowledge, at least in a qualitative way. The disadvantage is its subjectivity and potential for bias. The conservativeness of expert opinion with respect to mortality decline is widespread, in that experts have generally been unwilling to envisage the long-term continuation of trends, often based on beliefs about limits to life expectancy.*"
>
> Booth & Tickle (2008)

4.3 Booth & Tickle (2008) list several examples where expert opinion has under-estimated rates of mortality improvement.

4.4 In this section and the preceding ones we have seen a number of major technical, practical and fundamental difficulties with projecting mortality rates by cause of death. It is for these reasons that some actuaries prefer stochastic projection methods. The reader is directed to Booth & Tickle (2008) for a comprehensive overview of various projection methods.

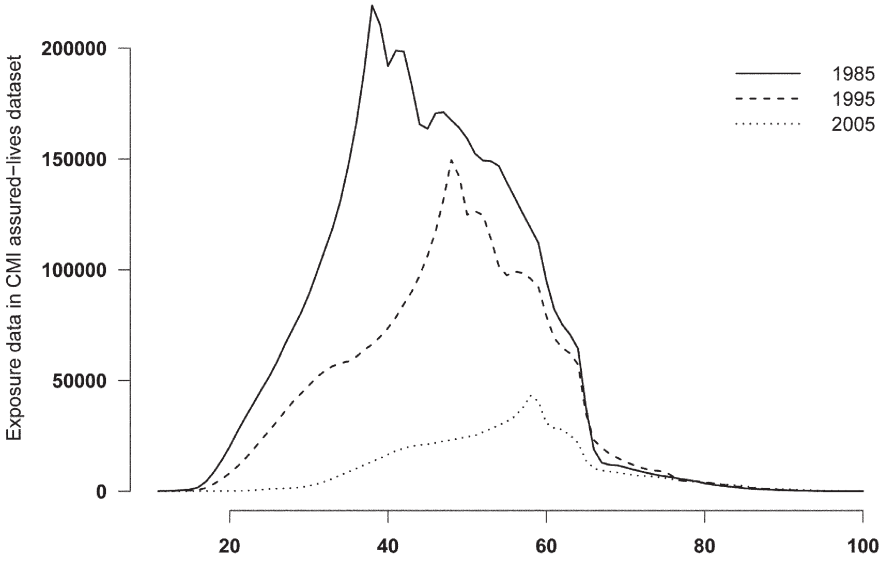
5. Working with Small Populations

5.1 A particular feature of actuarial work lies in dealing with financial liabilities to a select subset of the wider population, often the holders of private insurance policies or private pensions. An illustration of this is given by Fletcher (2009). These sub-populations can be very different from the broader population, often tending to be of higher socio-economic status and longer-lived as a result. This gives rise to two specific challenges for the actuary wishing to provide portfolio-specific mortality projections. The first problem is the scale of data available: the portfolio in question typically has a membership several orders of magnitude smaller than the population from which it is drawn. For example, in 2007 there were around half a million people in England and Wales aged around 65 alone (ONS data), which is larger than most annuity portfolios across all ages.

5.2 The second problem is of the length of historic data: good-quality data for deaths and populations at each age are available in England and Wales back to 1961. In contrast, annuity administration systems often do not have archives of historic mortality data for much more than around ten years. The availability of data is often linked to continuity of the computerised administration system — where a major system change or migration occurred in the past, mortality-experience data was often not regarded as valuable and therefore lost. The situation for occupational pension schemes is similar: in this author's experience, a change in scheme administrator often means the loss of historic mortality data.

5.3 Faced with a lack of both scale and length of time series, actuaries

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press



Reproduced from Richards & Currie (2009)

Figure 5. Exposures in CMI assured-lives data set by age. The rapid reduction in data volumes of recent years carries the risk that the composition of the data has materially changed. The distribution by age also militates against relying on this data set for applications to post-retirement mortality: there is relatively little data above age 65

will often fall back on a related data set which does have appropriate scale and length. One option in the UK is the "assured-lives" data from the Continuous Mortality Bureau (CMI). However, Figure 5 shows that this data set has shrunk considerably in recent decades, raising the question of whether the data volumes have shrunk for reasons other than death, surrender or maturity. If so, it is an open question as to whether changing composition of contributing life offices, or the composition of their policyholder base, might skew any trends observed in the recent decades.

5.4 Currie (2009) and Cairns (2009) present different models which enable stochastic projections for a smaller data set by reference to a larger one. Currie (2009) refers to his models as "piggyback" models as the smaller data set is in a sense riding on the larger one. Models linking a small population to a larger one necessarily involve making a number of strong assumptions, such as the gap between mortality rates being constant in time and linear in age. These assumptions may seem simplistic, but, as Currie (2009) notes, "doing nothing is also an assumption". One merit of such

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

models is that these assumptions are explicit, and are therefore preferable to the implicit assumptions sometimes made in actuarial work.

6. Models of Mortality

6.1 When fitting a model to a body of mortality data, actuaries and gerontologists are superficially looking to achieve the same thing: the best-fitting model which explains the data. However, actuaries and gerontologists often have different aims: actuaries often want the best-fitting curve to smooth or explain the data, whereas gerontologists want the model to reveal deeper insights into the mortality process. This can often lead to different parameterisations of the same model. For example, actuaries may want their model parameterisations to behave consistently with respect to a risk factor: Richards (2010) gives parameterisations of sixteen mortality models designed in such a way that an increase in risk is always represented by a positive increase in a parameter value. In contrast, gerontologists are concerned with models being structured so that each parameter has a real-world interpretation. As an illustration, Richards (2010) shows that the three-parameter logistic model used by the gerontologists Vanfleteren *et al.* (1998) is actually the same as the model proposed by the actuary Beard (1959). Gerontologists will prefer the parameterisation by Vanfleteren *et al.* (1998), whereas actuaries will prefer the parameterisation from Beard (1959). The underlying model is identical, but the preferred parameterisation depends on the model's purpose.

6.2 While different parameterisations of the same underlying model are justifiable, it is less helpful that different terminology is used in different academic communities. For example, actuaries define the force of mortality, μ_x, as:

$$\mu_x = \lim_{h \to 0^+} \frac{{}_hq_x}{h} \tag{1}$$

where ${}_hq_x$ is the probability of a life aged x dying in a small interval of time of width h. To statisticians, μ_x is known as the continuous-time hazard function (or hazard rate), while engineers call the same thing the failure rate. Similarly, the total time exposed to such a risk is called the central exposed-to-risk by actuaries, whereas statisticians call it the waiting time.

6.3 An important aspect of mortality models lies in avoiding the ecological fallacy (Robinson, 1950), whereby a model based on aggregate statistics for a group can lead to erroneous inferences about the behaviour of the individuals. For example, consider a population of lives where each follows a Gompertz mortality law (Gompertz, 1825):

$$\mu_x = e^{\alpha+\beta x}. \tag{2}$$

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

If the individuals are heterogeneous with respect to the value of α, then the mortality law apparently followed by the group will not be Gompertzian. The term e^{α} is presumed to be fixed at birth and is different for each individual, and this term is often known as the frailty. Specifically, if e^{α} has a gamma distribution, then the group exhibits mortality according to the Beard law (Beard, 1959):

$$\mu_x = \frac{e^{\alpha' + \beta x}}{1 + e^{\alpha' + \rho' + \beta x}} \tag{3}$$

where α' and ρ' are functions of β and the parameters of the gamma distribution. Similarly, an individual life may follow a Makeham law (Makeham, 1859):

$$\mu_x = e^{\epsilon} + e^{\alpha + \beta x}. \tag{4}$$

However, if the individuals have values of e^{α} drawn from a gamma distribution then the mortality of the group as a whole will appear to follow the Makeham–Beard law:

$$\mu_x = \frac{e^{\epsilon} + e^{\alpha + \beta x}}{1 + e^{\alpha + \rho + \beta x}}. \tag{5}$$

6.4 This basic result is discussed in Beard (1971) and Horiuchi & Coale (1990), while Richards (2008) gives worked equivalences of both these examples. The effects of these so-called frailty models are well-known to actuaries: in a population of mixed health, those in the poorest health tend to die first, leading to a healthier population with changing dynamics as the group ages. Vaupel & Yashin (1985) illustrate a number of ways in which heterogeneity amongst individuals can result in aggregate hazard rates which look very different from the actual hazard rates of the individuals.

6.5 The ecological fallacy lies in making erroneous assumptions about individuals based on a model for groups. The reverse is also possible: the atomistic fallacy — sometimes known as the individualistic fallacy (Diez-Roux, 2003) — is the name given to false assumptions about a group based on data or models for individuals. The same example for the ecological fallacy can be used in reverse to illustrate the atomistic fallacy: a population of heterogeneous individuals each following Gompertz mortality with a gamma-distributed frailty would produce population-level mortality according to the Beard law. In general it is preferable to work with data and models based on individuals, since it is easier to avoid these kinds of fallacies than if working with grouped data.

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

7. Modelling Mortality Differentials

7.1 Actuaries are keenly interested in mortality differentials for accurate risk-based pricing and reserving. Milne (2009) describes three different types of mortality differential which could occur, as shown in Figure 6 in stylised form.

7.2 Type A differentials are perhaps the least likely to crop up in actuarial work, i.e. where mortality is largely unchanged at younger ages, but differences are greatest at the oldest ages. Type B differentials are the most common in actuarial work on contemporaneous populations: different sub-groups start with different mortality rates at a "young" age such as 60, but these differentials narrow with increasing age to the point where they largely vanish by age 95. This mortality convergence (Gavrilov & Gavrilova, 2001) is exhibited by many of the risk factors used by actuaries in rating pensioner and annuitant longevity (Richards, 2008). An example of this is shown in Figure 7.

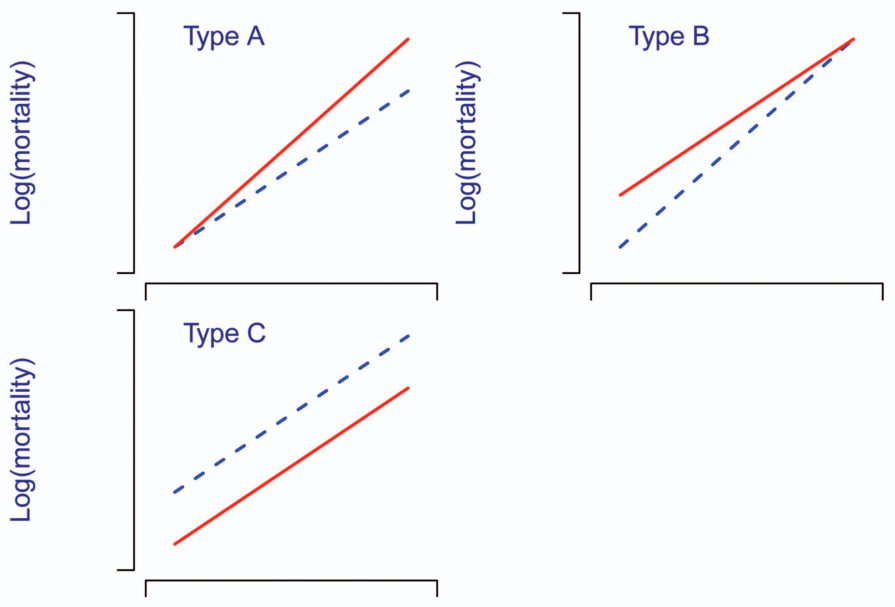


Figure 6. Three different possibilities for mortality differentials between two populations, styled after Milne (2009). Type A is a pivot of the log(mortality) curve around the rate at the youngest age, while Type B is a pivot around the oldest age. Type C is a simple shift of the log(mortality) curve up or down

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

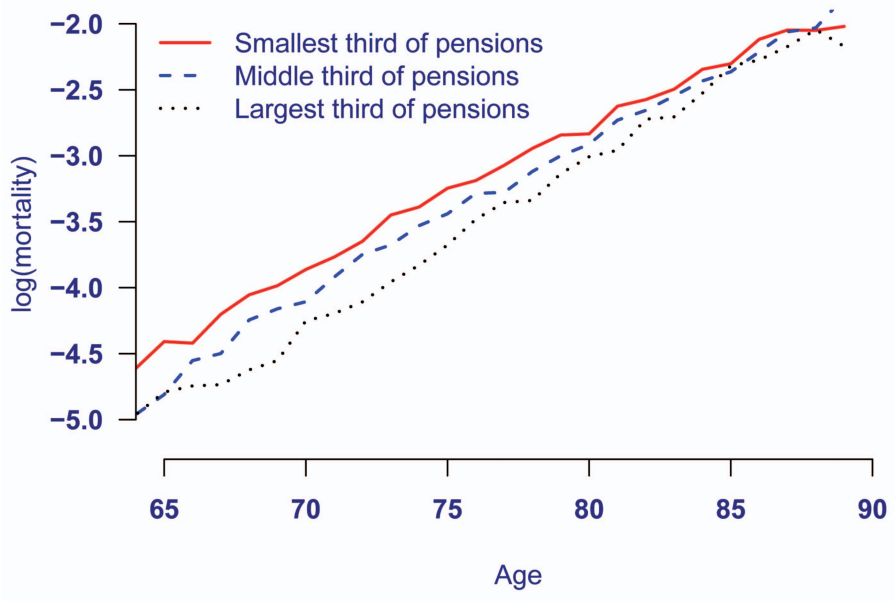


Figure 7. Example of a Milne Type B differential in mortality by pension size in a large annuity portfolio. Pension annuities are deduplicated using the algorithm described in Richards (2008) to create a data set of pensioners from the set of annuity records. The data set is then sorted by total pension income and split into three equal-sized groups. The mortality of those with the highest income is clearly lower than those with the lowest income, but the difference decreases steadily with increasing age, a phenomenon known as *mortality convergence* (Gavrilov & Gavrilova, 2001)

7.3 Actuarial models of mortality must accommodate Milne's Type B differentials, but not all models behave as required. For example, the Lognormal distribution for future lifetimes can yield hazard functions by age which are not easily consistent with Milne Type B changes, as shown in Figure 8. Richards (2010) reviews sixteen survival models and assesses them for their suitability in modelling post-retirement mortality patterns.

7.4 Milne's Type C differentials are simple vertical shifts in log(mortality), and are not commonly encountered by actuaries in contemporaneous populations. However, such vertical shifts can be observed when comparing populations separated widely in time, as shown in Figure 9.

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

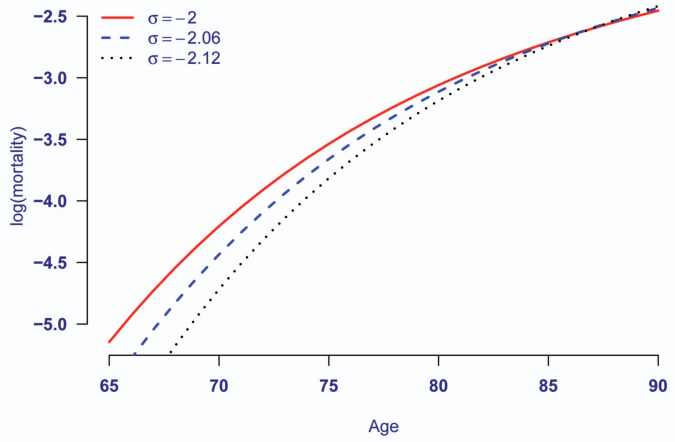


Figure 8. Log(hazard) functions for a Lognormal distribution for future lifetime, T, where log T is distributed normally with mean $-\alpha$ and standard error e^{σ}. Sample hazard curves for α, -4.46 and -4.47 with $\sigma = -2.0$, -2.06 and -2.12. The shape of the log(hazard) is different to the largely log-linear patterns in Figure 7, thus making it difficult for the Lognormal model to replicate the simple Milne Type B mortality patterns required in actuarial applications

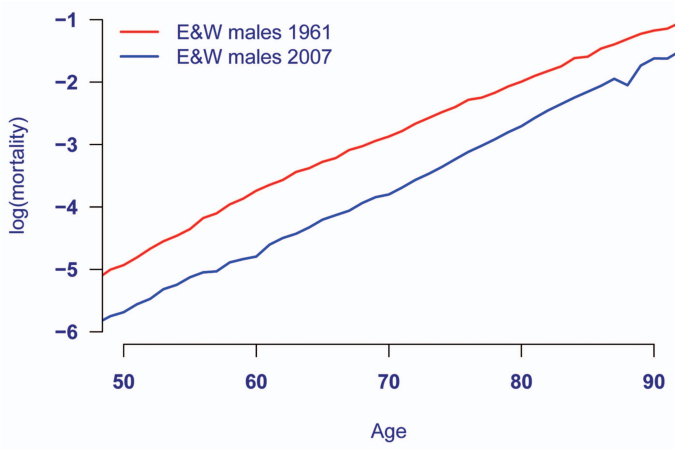


Figure 9. Example of a Milne Type C shift in male mortality in England and Wales (ONS data). Mortality rates have fallen at all ages and a first approximation of the changes is that the mortality curve on a log scale has simply been shifted downwards

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

8. Conclusions

8.1 Actuaries must take great care in using cause-of-death data as a basis for creating forecasts of future trends. Mortality by cause is strongly linked to socio-economic status, which itself is a major risk factor for mortality and longevity. Furthermore, cause-of-death data in England and Wales is subject both to changes in classification system and also to classification guidelines within a cause system. Computerised medical records might appear to be a good substitute, but actuaries must take great care in relying on results derived from them: such databases were created and run for non-research purposes and the data contained therein can be subject to serious recording bias. Finally, cause-of-death methods are vulnerable to systematic bias, leading towards over-estimation of forecast mortality rates. This is of particular concern when creating reserving bases for pension liabilities.

8.2 Actuaries often have access to rich individual data within the portfolios they manage or advise, and such data naturally lends itself to survival models. There is a wide choice of possible survival models, and actuaries must satisfy themselves that the selected model is capable of handling the various patterns of mortality shift which occur in different risk sub-groups.

8.3 While portfolio data is individually rich, it is usually limited to too short a period of time to build a meaningful model for long-term trends. Actuaries have therefore often had to rely on trends in other data sets, thus introducing basis risk relative to the portfolio being valued. Modern methods, such as piggybacking, enable actuaries to replace these implicit assumptions with explicit modelled ones to control for basis risk.

Acknowledgements

The author thanks Dr Chris Martin for his insights into the early genesis of the GPRD. Any errors or omissions remain the sole responsibility of the author. Graphs were done in R (2004).

References

Aylin, P. (1997). *Background note on twentieth century mortality files*. Office for National Statistics.

Azambuja, M.I. & Levins, R. (2007). Coronary heart disease — one or several diseases? *Perspectives in Biology and Medicine*, **50**(2), 228-242.

Azambuja, M.I. (2009). *Connections between causes of diseases: influenza and Sutherland's "epidemiology of constitutions"*. Presentation to Joining Forces on Mortality and Longevity, 22 October 2009, Edinburgh.
http://www.actuaries.org.uk/research-and-resources/documents/connections-between-causes-diseases-influenza-and-sutherlands-epide [Accessed 7 October 2010]

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

Beard, R.E. (1959). Note on some mathematical mortality models. In: G.E.W. Wolstenholme & M. O'Connor (eds.). *The Lifespan of Animals*. Little, Brown, Boston, 302-311.

Beard, R.E. (1971). Some aspects of theories of mortality, cause of death analysis, forecasting and stochastic processes. In: W. Brass (ed.). *Biological Aspects of Demography*, Taylor and Francis Ltd., London.

Booth, H. & Tickle, L. (2008). Mortality modelling and forecasting: a review of methods. *Annals of Actuarial Science*, **3**(Parts 1 and 2), 8.

Cairns, A. (2009). *Modelling multi-population mortality with cohort effects*, Presentation to Joining Forces on Mortality and Longevity, 22 October 2009, Edinburgh. http://www.actuaries.org.uk/research-and-resources/documents/modelling-multi-population-mortality-cohort-effects-slides [Accessed 7 October 2010]

CMI (2004). *Projecting future mortality*. A discussion paper, Continuous Mortality Investigation Mortality sub-committee Working Paper 3.

Currie, I.D. (2009). *Adjusting for bias in mortality forecasts*, Presentation to Joining Forces on Mortality and Longevity, 22 October 2009, Edinburgh. http://www.actuaries.org.uk/research-and-resources/documents/adjusting-bias-mortality-forecasts-slides [Accessed 7 October 2010]

Di Cesare, M. & Murphy, M. (2009). *Forecasting mortality, different approaches for different cause of deaths? The cases of lung cancer; influenza, pneumonia, and bronchitis; and motor vehicle accidents*, Presentation to Joining Forces on Mortality and Longevity, 22 October 2009, Edinburgh. http://www.actuaries.org.uk/research-and-resources/documents/forecasting-mortality-different-approaches-different-cause-deaths-s [Accessed 7 October 2010]

Diez-Roux, A.V. (2003). A glossary for multilevel analysis. *Epidemiological Bulletin*, **24**(3), Pan American Health Organization.

Fletcher, G. (2009). *Uncovering scheme-specific mortality*. Presentation to Joining Forces on Mortality and Longevity, 22 October 2009, Edinburgh. http://www.actuaries.org.uk/research-and-resources/documents/uncovering-scheme-specific-mortality-slides [Accessed 7 October 2010]

Gavrilov, L.A. & Gavrilova, N.S. (2001). The reliability theory of aging and longevity. *Journal of Theoretical Biology (2001)*, **213**, 527-545.

Gompertz, B. (1825). The nature of the function expressive of the law of human mortality. *Philosophical Transactions of the Royal Society*, **115**, 513-585.

Government Actuary's Department (2001). *National Population Projections: Review of Methodology for Projecting Mortality*. Government Actuary's Department, NSQR No. 8, pages 18-21.

Horiuchi, S. & Coale, A.J. (1990). Age patterns of mortality for older women: an analysis using the age-specific rate of mortality change with age. *Mathematical Population Studies*, **2**(4), 245-267.

Jordan, K., Porcheret, M. & Croft, P. (2004). Quality of morbidity coding in general practice computerized medical records: a systematic review. *Family Practice*, **21**, 396-412.

Makeham, W.M. (1859). On the law of mortality and the construction of annuity tables, *Journal of the Institute of Actuaries*, **8**, 301-310.

Martin, C. (2010). *For the record.* http://www.longevitas.co.uk/site/informationmatrix/fortherecord.html

Milne, E. (2009). *A new model of ageing and mortality*. Presentation to Joining Forces on Mortality and Longevity, 22 October 2009, Edinburgh. http://www.actuaries.org.uk/research-and-resources/documents/new-model-ageing-and-mortality-slides [Accessed 7 October 2010]

R Development Core Team (2004). *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.r-project.org

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press

RICHARDS, S.J. (2008). Applying survival models to pensioner mortality data. *British Actuarial Journal*. **14**(II), 257-326.

RICHARDS, S.J. & CURRIE, I.D. (2009). Longevity risk and annuity pricing with the Lee–Carter model. *British Actuarial Journal* (to appear).

RICHARDS, S.J. (2010). A handbook of parametric survival models for actuarial use. *Scandinavian Actuarial Journal* (to appear).

ROBINSON, W.S. (1950). Ecological correlations and the behavior of individuals. *American Sociological Review*, **15**, 351-357.

ROMERI, E., BAKER, A. & GRIFFITHS, C. (2006). *Mortality by deprivation and cause of death in England and Wales, 1999-2003*. Health Statistics Quarterly 32, National Statistics.

VAN GINNEKEN, A.M., VAN DER LEI, J. & MOORMAN, P.W. (1993). *Towards unambiguous representation of patient data*, Proceedings of the Sixteenth Annual Symposium on Computer Applications in Medical Care (SCAMC), Baltimore, MD, 1992; McGraw-Hill, New York, 1993 pp. 69-73.

VAN DER LEI, J. (1991). Use and abuse of computer-stored medical records. *Methods of Information in Medicine 1991*, **30**(2): 79-80.

VANFLETEREN, J.R., DE VREESE, A. & BRAECKMAN, B.P. (1998). Two-parameter logistic and Weibull equations provide better fits to survival data from isogenic populations of caenorhabiditis elegans in axenic culture than does the Gompertz model. *Journal of Gerontology Series A*, **53**, B393-B403.

VAUPEL, J.W. & YASHIN, A.I. (1985). Heterogeneity's ruses: some surprising effects of selection on population dynamics. *The American Statistician*, **39**(3), 176-185.

WONG-FUPUY, C. & HABERMAN, S. (2004). Projecting mortality trends: recent developments in the United Kingdom and the United States. *North American Actuarial Journal*, **8**(2), 56-83.

https://doi.org/10.1017/S1357321700005602 Published online by Cambridge University Press