Population Studies
A Journal of Demography

ISSN: 0032-4728 (Print) 1477-4747 (Online) Journal homepage: http://www.tandfonline.com/loi/rpst20

A new method of projecting populations based on trends in life expectancy and survival

Les Mayhew & David Smith

To cite this article: Les Mayhew & David Smith (2013) A new method of projecting populations based on trends in life expectancy and survival, Population Studies, 67:2, 157-170, DOI: 10.1080/00324728.2012.740500

To link to this article: http://dx.doi.org/10.1080/00324728.2012.740500

Published online: 03 Dec 2012.

Population Studies, 2013
Vol. 67, No. 2, 157–170, http://dx.doi.org/10.1080/00324728.2012.740500

Routledge
Taylor & Francis Group

A new method of projecting populations based on trends in life expectancy and survival

Les Mayhew and David Smith
Cass Business School

There is increasing concern about the lack of accuracy in population projections at national levels. A common problem has been the systematic underestimation of improvements in mortality, especially at older ages, resulting in projections that are too low. In this paper, we present a method that is based on projecting survivorship rather than mortality, which uses the same data but differs technically. In particular, rather than extrapolating trends in mortality, we use trends in life expectancy to establish a robust statistical relation between changes in life expectancy and survivorship using period life tables. We test the approach on data for England and Wales for the population aged 50 and over, and show that it gives more accurate projections than official projections using the same base data. Using the model to project the population aged 50 and over to 2020, our method suggests nearly 0.6 million more people in this age group than official projections.

Keywords: population projections; life expectancy; survivorship; evaluation

[Submitted November 2010; Final version accepted February 2012]

1. Introduction

There is increasing concern about the poor accuracy of population projections at national levels. Strictly speaking, a population projection is simply the outcome of a given set of assumptions and cannot be wrong on that basis, assuming the arithmetic is correct. In reality, population projections are essentially predictions or forecasts and are treated as such for government planning and expenditure purposes. Accuracy is crucial, especially where a projection is used to control the total for each area of expenditure; otherwise there is a danger of error becoming endemic in all areas of government policy.

Recent research has investigated why population forecasts have been so inaccurate. Keilman has been particularly vocal in raising concerns after evaluating past projections against observed outcomes (1997, 2001). He reached the damning conclusion that demographic forecasts published by statistical agencies in 14 European countries had not become more accurate over the past 25 years (Keilman 2008). Shaw (2007) and Keilman (2007) reached similar conclusions with specific reference to the UK, but concerns about population projections go back much further (Brass 1974; Preston 1974; Keyfitz 1981).

Booth (2006) undertook a comprehensive review of demographic forecasting over the previous 25 years and called for more retrospective analysis of forecasting accuracy. She portrays the field as one in which there have been many technical advances in methods and a borrowing of ideas from other disciplines but not necessarily comparable improvements in predictive accuracy. The methods used fall broadly into two types: 'extrapolative' methods and structured 'causal' projection methods. The former, which are by far the more typical and where most technical advances have been made, focus on stable patterns and trends in the data. Such projections may be subject to adjustments according to the views of experts in the field. The other type, structured 'causal' projection methods, seeks to explain demographic rates in terms of socio-economic or proximate determinants. For example, it may be possible to draw a direct link between a new cure for cancer or a decline in smoking behaviour and subsequent changes in mortality rates. Often regarded as the 'ideal', structural methods have not, on the whole, produced any greater accuracy according to Booth and others. However, we think these methods, which are not further discussed here, should be seen as 'work in progress' and that it is too early to make a definitive judgment on their efficacy.

By their very nature, population projections must take into account trends in fertility, migration, and mortality, errors in any of which have the potential to affect accuracy. However, an important problem highlighted by Shaw (2007) in his review of the UK's projections record of the last 50 years has been the systematic underestimation of improvements in mortality. This underestimation has resulted in projected populations at the older ages being too low, a problem not confined to the UK and in fact fairly common (Bengtsson and Keilman 2003). One response to the problem has been to develop stochastic projections that show a range of uncertainty in future mortality (Keilman 2002).

The best known and widely used method in this category is the Lee–Carter model, which is based on a combination of statistical time series methods to project mortality and a simple method of estimating the age distribution of mortality (Lee and Carter 1992). Lee and Miller (2001) find that the model gives good central projections of life expectancy, but over long time horizons it tends to be too pessimistic. However, it is extremely difficult to generalize about its performance, because there have been many developments and model variants published since the model was introduced (for good examples see Renshaw and Haberman 2003, 2006; Haberman and Renshaw 2008, 2009).

A practical problem is that 'probabilistic forecasts' are sometimes seen as too complex when what users really want is a single or 'best forecast' (Booth 2006). We too use extrapolative methods, but start from a different position to Lee and Carter, who use mortality data. As Booth has pointed out, the simplest method is to extrapolate life expectancy or some other life table measure and to use empirically based model life tables to obtain the age pattern. We proceed in a similar way but we concentrate on survivorship rather than mortality, which usefully reduces the amount of unexplained variation that accompanies mortality data.

Our method exploits the empirical relationship between life expectancy and the probability of survival to a given age using period life tables. We use this relationship to estimate the age-specific probability of survivorship in some future year. For reasons explained later, we fit a Gompertz–Makeham function (see Olshansky and Carnes 1977 or Forfar and Smith 1987) to each of the resultant distributions, from which we derive full single-year life tables. Using simple regression to project survivorship, each future life table then becomes the basis for deriving age-specific mortality rates that are applied to a base population in the conventional way.

Our use of the shape of the survival curve to project forward the probability of survivorship based on expectation of life can be compared with Brass's method (1974). His premise was similar to ours, the idea being that the shape of previous life tables based on observed data can be used to create future life tables through a simple transformation of the data. Whereas we work with percentiles and the Gompertz–Makeham survival function, he used a logit function to transform the data.

Obviously, these methods rely heavily on the ability to make accurate projections of life expectancy, and this has been changing rapidly. In England and Wales, for example, the life expectancy in 1960 of men at age 50 was 22.9 years, having increased by only 0.4 years since 1950, but from 1990 to 2000 it increased by 2.3 years to 25.9 years. For women at age 50, the improvement from 1950 to 1960 was 1.4 years (to 27.7 years), and from 1990 to 2000 it was 1.45 years. We show that the accelerating trend over the last 50 years for men and the more slowly increasing trend for women are very accurately modelled by a second-order polynomial function.

A key problem in using a simple extrapolative approach of this kind is that little is known about the biological limits to human longevity. Given the rapid pace of change over recent decades, it is hard to know whether the upward trend will be maintained, slow down, or plateau. To address this question, Oeppen and Vaupel (2006) analysed data from many countries from 1840 onwards and found no empirical evidence to suggest that life expectancy is reaching a limit. However, as we are more concerned with improvements in the accuracy of projections in the short to medium term, this issue is arguably of lesser concern as long as our projection horizon is relatively short and there is evidence that the indicated level of future life expectancy is achievable (e.g., by comparing with levels of life expectancy in other countries).

To illustrate the merits of our method, we compare below the results of using it to produce a set of alternatives to 21-year projections produced in 1981 for England and Wales by the UK Government Actuary's Department (GAD); this was the agency responsible for national projections at the time. We used exactly the same assumptions and base population as GAD. It will be seen that our projections are closer to the actual populations than those projected by GAD. We will show that our method can be adapted to produce future period life tables that appear to be more accurate than those currently used to make official population projections.

A limitation of our method is that we have applied it only to populations aged 50 and over, since our purpose was to focus on the implications of increasing longevity at older ages. To cover the whole age range, the method would need to be integrated with ways of projecting other components of population change, namely, fertility and migration, the projection of which raise completely different analytical and technical issues.

The paper is structured as follows. Section 2 describes the theoretical and empirical basis for the method adopted. Section 3 describes how the framework is adapted to provide population projections. Section 4 compares the results with those of past official projections. Section 5 compares projections based on our methods with official projections to 2020. Finally, Section 6 discusses how the techniques could be adapted and improved in future research.

2. Method

Our method uses an assumed mathematical relationship between the probability of survival to a given age and life expectancy. The method belongs to a branch of mathematical theory known as queuing theory, in which the time predicted to be spent in a system is related to the cumulative probability of either being 'processed' within a given time or continuing to be held in a queue (Mayhew 1987; Mayhew and Smith 2007).

The analogy with survivorship is that 'time spent in the system' corresponds to the period in which a person is 'alive', but that once they have been 'processed' (i.e., have left the system) they are counted as 'dead'. As with queuing systems, populations demonstrate a strong correlation between expectation of life and the probability of survivorship (Mayhew 2001), such that if it were possible to project future life expectancy, it would be possible to predict corresponding survivorship probabilities to different ages or 'time spent in the system'.

For the demographer, survivorship, in an ideal world, would be defined by a basic mathematical function such that, if we knew the expectation of life and other required parameters, we could determine the ages by which certain percentages of the population would have died (or survived). The simplest form of such a survival function, $S(x)$, would involve it having only one parameter.

To take a simple case, suppose that $S(x)$ is represented by a negative exponential function with the special property that future life expectancy is the same at any age. We have

$$S(x) = \exp(-\lambda x)$$

where λ is a parameter and where the mean value of the negative exponential distribution is given by $\bar{x} = 1/\lambda$.

Substituting this into the previous equation and rearranging we obtain $x = -\bar{x}\ln(S(x))$.

Hence, the age to which a given proportion survive, x, is a linear function of life expectancy $\bar{x}$ and ln $S(x)$. By selecting different percentiles of the proportions of those surviving, we can derive a family of straight lines relating life expectancy to the percentage surviving to any age. For example, suppose that life expectancy at birth is 75 years. The age to which 90 per cent of the population survive (or by which 10 per cent die) would be 7.9 years ($-75 \times \ln(0.9) = 7.9$); for 50 per cent, it would be 52.0 years ($-75 \times \ln(0.5) = 52.0$) and so on.

We could also define a different distribution, giving the future expected life at some age other than zero. Let us assume a starting age of 50, with the probability of death modelled by a normal distribution with a standard deviation defined as $\sigma = \mu/3$, where μ is life expectancy at age 50.

Now suppose that the mean expected future life at age 50 is 24 years, so that the standard deviation is 8 years from the relationship above. Using the statistical reference table for the areas under the standard normal distribution, we find that 80 per cent of the population will have died, or 20 per cent survived, when $z = 0.8416$. This occurs after $24 + 8 \times (0.8416) = 30.7330$ years from age 50, that is, at age 80.7 years.

Simple relationships of the kind shown in these examples are not found in practice, and distributions with more parameters are needed to produce sufficiently accurate fits to actual survivorship data. In our method, we use the three-parameter Gompertz–Makeham function for this relationship, as described below.

Life table data for different calendar years are used to estimate the parameters for this model. Since the focus is on ages 50 and over, we set $l_{50} = 100{,}000$ to obtain future life expectancy at age 50, using the following standard expression:

$$e_{50} = \frac{1}{l_{50}} \sum_{y=50+1}^{\Omega} l_y + 0.5$$

where Ω is the age of the oldest person to die, assumed to be 110 years in our calculations, and l_y is the number of survivors to age y. It is recommended that a suitable period be chosen for investigating survivorship patterns, preferably one in which life expectancy has changed by a reasonable amount, to give a robust representation of any systematic change (as occurred in England and Wales over the last half century).

For each calendar year and every percentile proportion, the life expectancy and the percentile proportions surviving to each age are tabulated in successive columns of a spreadsheet. Linear regressions are then fitted to each vector of points corresponding to a column of data, using an equation of the following form in which x, the age at death, is the dependent variable and life expectancy e_{50} the independent variable:

$$x_p = \alpha_p + \beta_p e_{50} + u_p$$

where

x_p is the age of death of the pth percentile

α_p, β_p are the regression parameters for the pth percentile

e_{50} is the expected life expectancy at age 50

u_p is a normally distributed random error term.

The data used in the regressions were extracted from period life tables for England and Wales for 1952–2003 in the Human Mortality Database (HMD). This set of life tables was preferred to those of GAD, because the use of a common method of constructing the life tables allows comparisons of results with those from other countries. In practice, we found that the differences in results between HMD and GAD data sources were small.

3. Results

We begin by describing the results of applying our method and then explain their implications. For projection purposes (see later), we made use of survival data for all the percentiles, but for brevity only report below on selected percentile values. Using data for each calendar year, we calculated the proportion of people surviving to a given age in 1 per cent steps. We then regressed the age to which each percentile survived on life expectancy at age 50 from 1952 to 2003.

Table 1 presents the fitted regression parameters for both men and women. For example, for the 70th percentile for men who survive, the value $\alpha_{70} = 42.0002$ and $\beta_{70} = 1.1089$. If life expectancy at 50 is assumed to be 25 years, the predicted age to which 70 per cent of the population of men survive is 69.72 years (i.e., $42.0002 + 1.1089 \times 25$). The equivalent value for women is $45.0262 + 1.0113 \times 25$ or 70.31 years.

Table 1 shows that, for men, the goodness-of-fit statistic, or coefficient of determination, R^2, for the regression is always larger than 0.99 between the 10th and 90th percentiles. For percentiles outside this range, it is always larger than 0.96 for men, suggesting a high degree of precision is possible, even among percentiles at either end of the distribution. The goodness-of-fit for women is also very good, though not quite as good as that for men.

Figure 1 shows the survival percentiles and fitted regression lines for men for the 99th, 95th, 90th, 80th . . . 10th, 5th, and 1st percentiles. In 1952, the first year of the period examined, life expectancy at age 50 was 22.6 years (denoted by hatched line A); in

Table 1 Estimated life table proportions surviving to given ages from regressions fitted to observations in Figure 1, England and Wales 1952–2003

	Men			Women		
Percentile	α_p	β_p	R^2	α_p	β_p	R^2
99th	46.5921	0.2068	0.9718	43.7026	0.3018	0.9333
95th	40.4351	0.6500	0.9883	36.7884	0.7660	0.9453
90th	38.6480	0.8800	0.9945	37.7765	0.9041	0.9524
80th	39.5360	1.0553	0.9981	41.8321	0.9784	0.9795
70th	42.0002	1.1089	0.9985	45.0262	1.0113	0.9939
60th	44.9120	1.1173	0.9985	46.7969	1.0608	0.9985
50th	48.0358	1.1064	0.9987	48.9234	1.0829	0.9988
40th	51.0673	1.0925	0.9988	51.2654	1.0896	0.9976
30th	54.4762	1.0655	0.9991	54.0927	1.0785	0.9958
20th	58.8461	1.0104	0.9989	57.0801	1.0684	0.9915
10th	64.0007	0.9562	0.9964	61.8253	1.0232	0.9840
5th	67.9135	0.9145	0.9882	65.9977	0.9705	0.9773
1st	74.5729	0.8449	0.9649	72.7077	0.8959	0.9598

Source: HMD database.

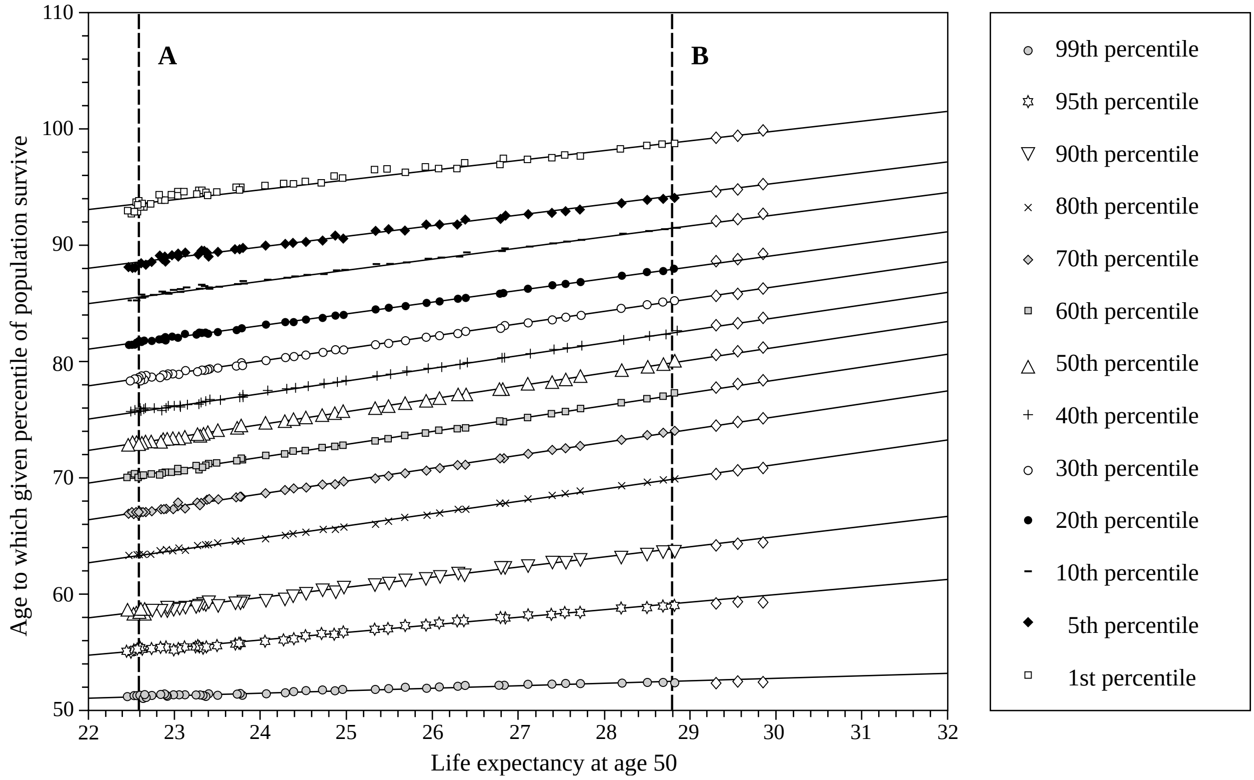

Figure 1 Period life table proportions of men surviving to different ages after age 50, shown as percentiles, England and Wales 1952–2003

Note: Vertical line A: life expectancy at age 50 in 1952. Vertical line B: life expectancy at age 50 in 2003. Diamond-shaped symbols to right of line B are based on life tables for the years 2004–06 and are not included in the regression lines.
Source: As for Table 1.

2003 it was 28.8 years (hatched line B), an increase of 6.2 years over the 51-year period. If we follow each hatched line from the bottom of the chart upwards until it crosses the 50th percentile (i.e., median life expectancy), the prediction is that 50 per cent of those reaching age 50 in 2003 will live to age 80, compared with age 73 for those reaching age 50 in 1952.

In projecting the proportions surviving to a given future year, we need to extrapolate each regression line by the amount that life expectancy is expected to increase by that year. To show that this is a reasonable step to take, three data points, for the years 2004, 2005, and 2006, which were *not* included in deriving the regression estimates, are shown as diamond-shaped symbols to the right of hatched line B. The predicted trend lines pass through, or close to, each data point, regardless of percentile, suggesting that a high level of precision is attainable if life expectancy can be accurately predicted.

There are other features of Figure 1 that should be noted. First, a useful property is that if the number of persons alive at age 50 (or at any alternative age) is 100,000, the vertical distance on the chart between deciles (e.g., between 50th and 60th percentiles) represents 10,000 deaths (or 10,000 fewer survivors), and between individual percentiles, 1,000 deaths (or 1,000 fewer survivors), and so forth (noting that the distance between these percentiles will depend on the concentration of deaths at each age in the survival distribution).

The slope of each regression line shows how quickly survivorship is increasing at each age with increasing life expectancy at age 50. If the slope for a particular percentile, that is, the β_p parameter in Table 1, is greater than one, it means that this percentile proportion surviving gains more than 1 year of life for each additional year of life expectancy. Similarly, when the β_p parameter is less than one, the percentile proportion surviving gains less than 1 year of life for each additional year of life expectancy.

For men, it can be seen that those gaining most years fall for percentiles between the 20th and 80th, all of which have values greater than one. For the oldest survivors, that is, those in the 10th to the 1st percentiles, the increase in life expectancy is not fully reflected in their additional years of life. For women, those gaining most years fall between the 10th and 70th percentiles. Because for women the slopes for the 5th and 1st percentiles are closer to a value of 1 than are those for men, the oldest survivors gain more years than their male counterparts.

Theory predicts that convergence of the percentile regression lines over the long term would lead to a more equitable distribution of life expectancy across the ages, resulting in a more rectangular-shaped survival distribution (Kannisto 2000). Specifically, for a perfect rectangular shape to occur, the lines would eventually need to converge to a point (Mayhew and Smith 2011). There is no indication of such a trend in our data, although it could occur in other countries, where the pattern of development may differ from that in England and Wales.

An important implication of these results is that those dying before say, age 65, have not benefited as much as others from the overall increase in life expectancy. There are a number of possible reasons for earlier deaths in this age range: the effect of incurable or hard-to-treat diseases where prolongation of life is harder to achieve unless effective medical cures are found; for men, the result of working in hazardous industries for part of their lives; and the effect of an unhealthy lifestyle (heavy smokers, the obese, etc.). If this pattern of relatively smaller improvements in survival persists, there will be no progress towards a more rectangular survival distribution.

Although for most of the period under investigation (1952–2003), there were almost continuous annual gains in life expectancy, it is possible to postulate that a homeostatic relationship exists between life expectancy and individual percentiles through time. In other words, if life expectancy were to fall instead of increase, the trend would reverse, and small annual perturbations in life expectancy suggest that this is the case. However, this hypothesis would need to be tested more thoroughly, using data sets for other countries for periods of time in which annual fluctuations in life expectancy have been more pronounced.

The above results can be presented diagrammatically in an alternative form to allow other insights. For example, Figure 2 shows life expectancy at age 50 as before, but it is now plotted on the vertical axis. The variable on the horizontal axis is replaced by the percentage of the population still alive. Each curve represents a given age, as indicated on the top horizontal line. To maintain consistency with Figure 1, the horizontal hatched lines indicate life expectancies in 1952 and 2003. The upper limit of the life expectancy at age 50 on the vertical axis is set at 32 years (based on current trends, this is not expected to be reached until 2013).

The bends in the curves in Figure 2 are consistent with the observation that, based on the most recent

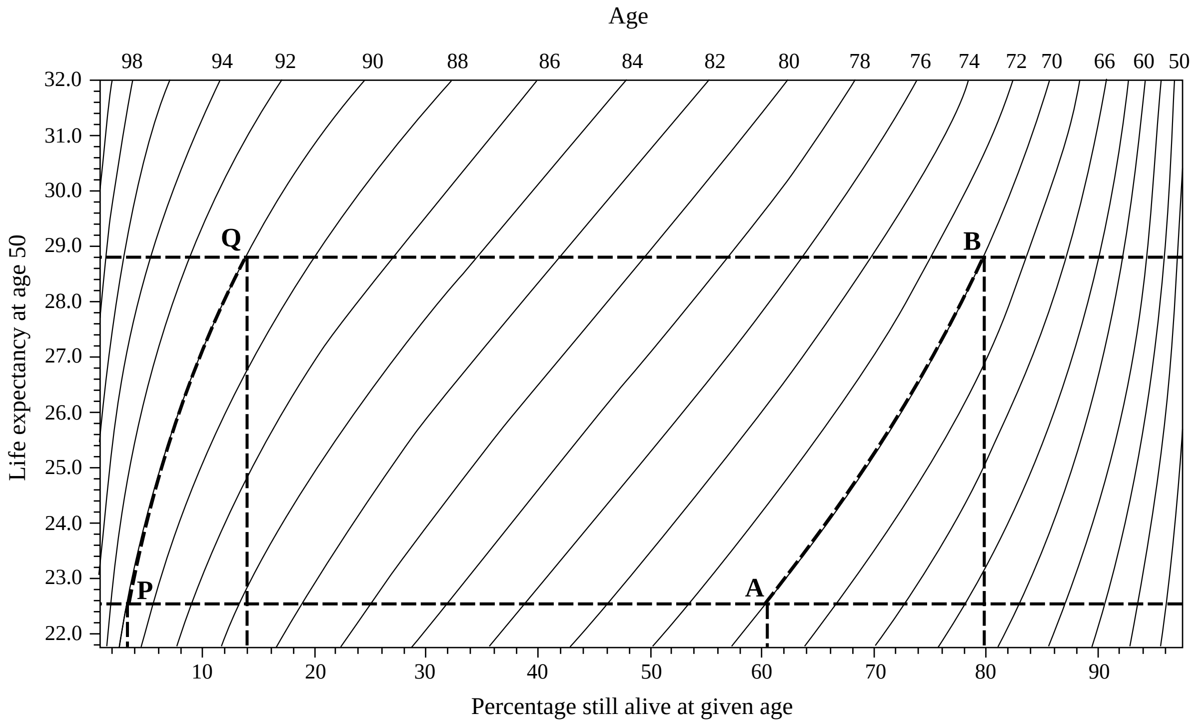

Figure 2 Life expectancy of men at age 50 and percentage alive at each age on top horizontal axis, using data from England and Wales 1952–2003 (see text for details of annotation)
Source: As for Table 1.

data, the largest gains in survival from the overall increase in life expectancy have been occurring from about age 70. For example, point P corresponds to a life expectancy of 22.6 years in 1952, when fewer than 3 per cent of men could expect to live until age 90; by 2003, point Q, this proportion had increased to around 14 per cent when life expectancy was 28.8 years.

To take a second example, in 1952, 60 per cent of men could expect to live until age 70 (point A); by 2003, this proportion had increased to around 80 per cent (point B). A further noteworthy trend is the growth in the proportions surviving beyond 90 years; in 2012, for example, when life expectancy is projected to be 31.5 years, around 22 per cent of men reaching age 50 that year would live at least into their 90s.

Fitting the survival function

The above results enable us to estimate the age to which a given percentile of the population will survive (i.e., the percentage of people surviving to age x). A life table, however, requires the results to be in the form of the proportion of people surviving to exact age x, $x+1, \ldots$, etc., that is, in strict 1-year intervals, and so we need a method of converting the survival percentiles into those for integer years of age. One option would be to use interpolation techniques between individual percentiles to find the number surviving to an exact age, x, say. However, our preferred method is to fit a survival curve, from a suitably parameterized function, to the expected percentile values, so that it becomes straightforward to extract the required information.

For this purpose, we make use of the Gompertz–Makeham function, which is well known, often used, and highly flexible, for fitting survival curves (for a historical review, see Olshansky and Carnes 1977). According to this function, the death rate is the sum of an age-dependent component (named after Gompertz 1825), which increases exponentially with age, and an age-independent component (named after Makeham 1860).

Using standard notation, the function for the force of mortality, μ_x, is given by $\mu_x = A + Bc^x$ or $\mu_x = A + Be^{\gamma x}$ where $\gamma = \ln(c)$ and A, B, and c are empirically determined parameters defining the shape of the curve, and x is age. The survival function for a life aged t to survive x years is then simply

$$S(x) = \exp\left\{ -\int_t^x \mu_s \, \mathrm{d}s \right\}.$$

Using the Gompertz–Makeham function for μ_s, and an age t, which is the starting age for future life expectancy (in this example, $t = 50$), we obtain

$$S(x) = \exp\left[A(t - x) + \frac{B(e^{\gamma t} - e^{\gamma x})}{\gamma} \right].$$

For a starting population of 100,000 people aged 50, multiplying the equation above by 100,000 gives the number of people who survive to age x, that is, we have simply $l_x = 100{,}000 \times S(x)$. Conventionally, the function is fitted over all ages, except the youngest, and it is assumed that the constant term, A, caters mainly for the non-age-related deaths, such as accidents, which mainly occur at younger ages. It can be observed that when $A = 0$, the formula reduces to the Gompertz function, and if $B = 1$ we obtain a simple exponential distribution.

The parameters A, B, and c were estimated using an iterative heuristic optimization technique of our own design. Estimated values are given in Table 2 from 1953 to 2003 in 10-year steps. As may be seen, after 1963 the B parameter becomes smaller, which is consistent with the fall in mortality over the period, whereas the corresponding increase in c means an increase in mortality at older ages. The main contribution of A is to moderate the number of deaths in the 50–60 age range, which would otherwise be too high. As B decreases after 1963, A is not needed to counteract the effects of B and so it approaches a value of zero.

Figure 3 shows the population curve for the year 1973: the fitted and the actual proportions surviving; 1973 was chosen as it is midway through the period being examined. The fit is a good one. A similar close fit is achieved for all other years, and for both men and women.

Table 2 Parameter values for the fitted Gompertz–Makeham functions

Parameter	1953	1963	1973	1983	1993	2003
A	−0.00318	−0.00639	−0.00542	−0.00331	−0.00317	0.00062
B	0.00016	0.00024	0.00021	0.00011	0.00008	0.00001
c	1.08803	1.08320	1.08346	1.09091	1.09270	1.11208
γ	0.08437	0.07992	0.08016	0.08701	0.08865	0.10623

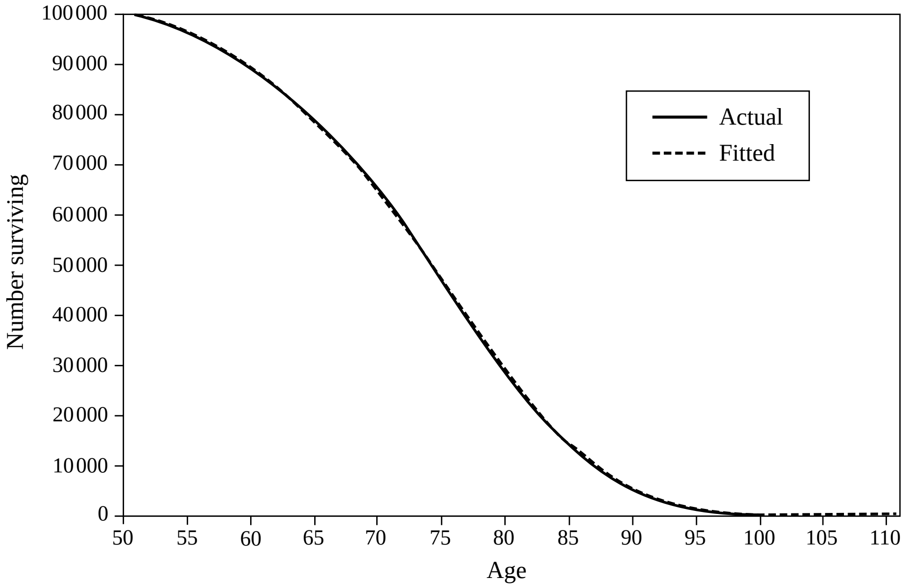

Figure 3 Actual and fitted number of survivors (men and women) at each age from age 50, England and Wales 1973
Source: As for Table 1.

4. Comparison with GAD projections from a 1981 base

For population projections, GAD makes assumptions about future rates of fertility, mortality, and migration into and out of the country. For mortality, GAD constructs life tables based on projected age-specific mortality rates. How do projections using the method described above compare with the official projections produced by GAD?

We investigated the accuracy of official population projections by using the same historical data that would have been available to statisticians and actuaries at the time, together with the same base population. We compared the 1981-based GAD projections with the actual subsequent populations for each year from 1982 to 2003, and also with our own model projections. (We also undertook similar comparisons using GAD's 1991-based projections. They were closer to the observed populations, but because the projection period was shorter, the results were not unexpected and so do not affect our general conclusions.)

Extrapolative methods were used by GAD for each of their mortality projections throughout the 1980s and 1990s, a description of which may be found in the historical projections section of the now archived GAD website. The 1981-based projection assumed that a life expectancy at birth of 74 for males and 77 for females would be achieved by 2040; it was 71 years and 77 years, respectively, in 1981 (Benjamin and Overton 1980). In these projections, scenarios that produced lower mortality rates than comparison countries were assigned lower importance than those that produced the same or higher rates.

Since our method relies on exogenous estimates of life expectancy, it was first necessary to predict life expectancy at age 50 from a 1981 base. Figure 4 shows the actual and predicted life expectancies for men in England and Wales from 1952 to 1981. After experimenting with different functions, we found that a quadratic equation consistently gave the best fit with a coefficient of determination R^2 in this particular example of 0.91. The fitted equation for the male population is given by $y = 8160 - 8.3269 + 0.0021302x^2$
where
y is predicted life expectancy at age 50 and
x is calendar year.

Figure 4 also shows 95 per cent prediction limits either side of the trend line. These indicate that the 'true' value of life expectancy should be within 0.36 years of the observed value 95 per cent of the time. Similar results, but using other time periods, enabled us to be confident that a quadratic would be equally appropriate for projecting forward from alternative base years.

However, the convex property of this particular polynomial form is a problem in that expected future life increases indefinitely into the future, which means that it must be used cautiously, especially for longer-term projections. Unlike GAD, however, we did not impose any limits on possible future

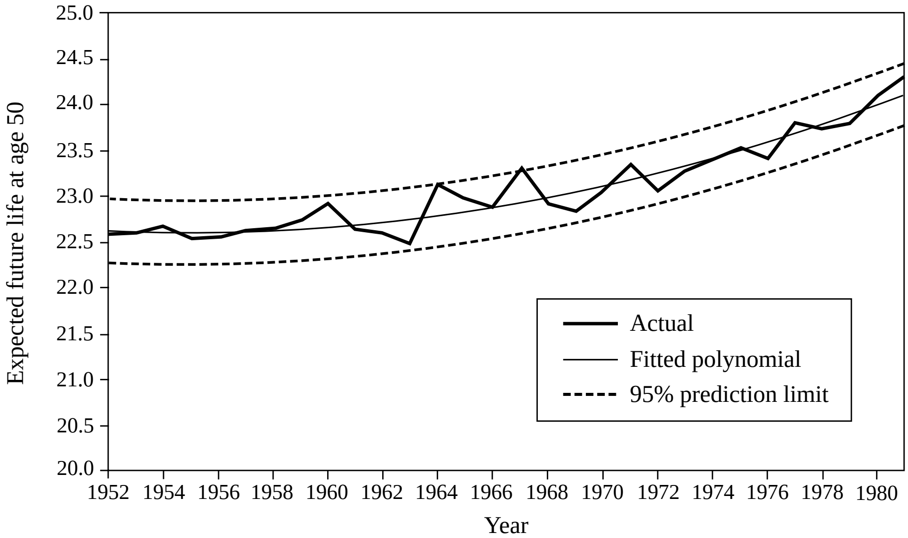

Figure 4 Fitted regression curves showing expected years of life at age 50 by calendar year for men, England and Wales 1952–81
Source: As for Table 1.

values of life expectancy, or on possible reductions in mortality rates. This seemed a reasonable decision given the observed trends in life expectancy during recent decades, but the particular polynomial fitted cannot predict turning points, so that longer-range projections are inevitably more uncertain (see conclusions for a further brief discussion on this point).

Using the life expectancies at age 50 derived above, we constructed life tables, one for each year between 1982 and 2003, using the information from projecting forward the trend in life expectancy based on the fitted polynomial equation. We then compared our projected life tables with GAD's and with published life tables for the same years (i.e., the actual values). We found that both our projection and GAD's tended to underestimate the extent of survival, but that our projections were much closer to the actual outcome than GAD's.

We compared the goodness of fit between GAD's projected survivorship, our model's projected survivorship, and the actual survivorship at each age. We found that our model gave a consistently better fit at each age, and for each year. As a visual illustration, Figure 5(a) shows the results based on our model, and Figure 5(b) on GAD's. The survival age for men is plotted on the vertical axis and calendar year on the horizontal axis. For each age and calendar year, we compared the percentage difference between the predicted survivorship and the actual survivorship, from a 1981 base. These differences are represented as contours, such that a contour value of 5 per cent indicates that the actual survivorship was 5 per cent longer than the predicted survivorship at all points along the contour etc.

Both charts correctly indicate that actual survivorship is higher than predicted survivorship. The differences become larger for longer projections, but the differences associated with our model are substantially smaller than GAD's. For example, at point P in Figure 5(a), the actual survivorship of men to age 75 in 1992 was 5 per cent higher than our model predicted in 1981. In contrast, Figure 5(b) shows that the corresponding difference between the actual survivorship and GAD's projected survivorship was just over 10 per cent; see corresponding data point P′, that is, age 75, in 1992, in Figure 5(b).

In Figure 5(a), the model error is generally around 5 per cent up to age 70. At ages over 70, the percentage error remains fairly small, but, from 1987 onwards, it gradually increases. Above age 80, the percentage error starts to rise after 1990, when it is of the order of 10 per cent, rising to 20 per cent after 1998. In Figure 5(b), the GAD case, larger errors are evident sooner and at a younger age, with an error of 10 per cent being typical as early as 1987. The error becomes larger in subsequent years; around 30 per cent for those aged 80 and over in the late 1990s.

5. Projecting forward to 2020

The results of our test persuaded us that it would be reasonable to use our method to make further projections for the population aged 50 and over. The results suggest that, whichever method is used, the error increases the longer the projection period, but whether that is important depends on the

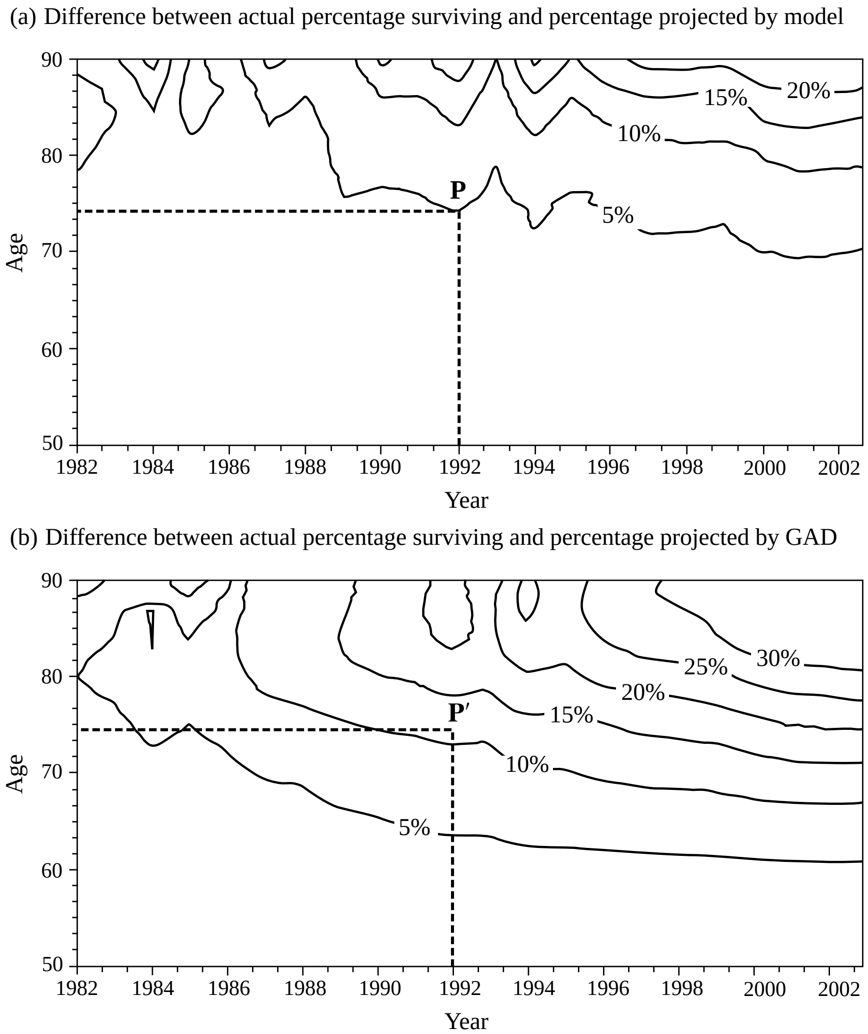

Figure 5 Differences between the actual and predicted percentage of men surviving to a given age: (a) projection by model and (b) projection by GAD, England and Wales 1982–2003
Source: As for Table 1 and Government Actuary's Department.

purpose of the projections. For example, an error of 10 per cent 5 years hence may not be as important as an error of 10 per cent 1 year hence. Thus, we need to consider how far ahead we should attempt to project survivorship, and how accurate the projection needs to be. In the UK, the Government, using the spending review process, plans its finances for the medium term, with government departments agreeing 3-year programmes. In many areas of administration, such as pensions, health, and social care, however, it is necessary to take a longer view of perhaps 15–20 years.

A comprehensive set of GAD population projections was published in 2001 (GAD 2001), and this afforded the opportunity to compare the results of our model with GAD's over a longer time period. As before, we fitted a polynomial to life expectancy at age 50 for men and women for the period 1952–2003, and used this to project life expectancy over the period 2003–20. The predicted life expectancy was then used to derive the survivorship percentiles, before fitting the survival curves. We used the same assumptions as GAD, documented on the GAD website, in particular their assumption on immigration at older ages in their principal projection of 2001 (GAD 2001). By using GAD's projections for ages below 50, we effectively ensured that the starting number of 50-year-olds in each calendar year was the same as GAD's.

Figure 6 shows the actual and fitted curves of life expectancy by year for men from 1952 to 2003. The coefficient of determination, R^2, is 0.99 and the 95 per cent confidence limits are ± 0.22 years. We assumed that life expectancy would continue to increase according to the fitted curve and that small variations would follow a similar pattern to the past. The results predict that, between 2001 and 2020, men's life expectancy will increase from 28.5 to 34.6

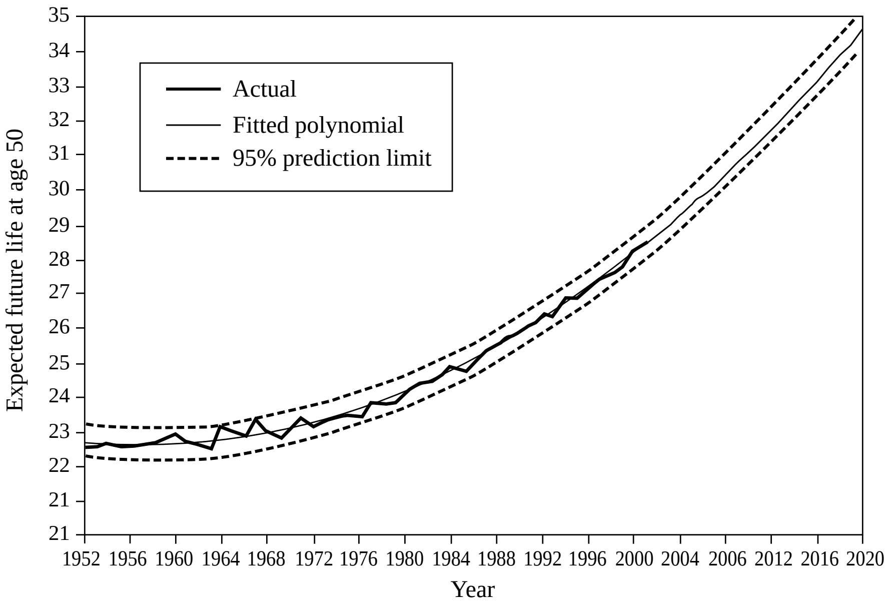

Figure 6 Fitted regression curve showing life expectancy of men at age 50, England and Wales 1952–2001
Source: As for Table 1 and Government Actuary's Department.

Table 3 Projection of population of men in 2020 from a base year of 2001: comparison of model projections with those of GAD, England and Wales

			Difference	
Age	GAD 2020	Model 2020	Number	Per cent
50–59	3,788,205	3,809,512	21,306	0.56
60–69	3,014,841	3,111,925	97,084	3.22
70–79	2,324,314	2,504,966	180,653	7.77
80–89	978,574	1,164,099	185,525	18.96
Total	10,105,934	10,590,502	484,568	4.79

Source: As for Figure 6.

Table 4 Projected population of women in 2020 from a base year of 2001: comparison of model projections with those of GAD, England and Wales

			Difference	
Age	GAD 2020	Model 2020	Number	Per cent
50–59	3,962,913	3,963,685	771	0.02
60–69	3,203,880	3,202,956	−924	−0.03
70–79	2,632,919	2,645,805	12,886	0.49
80–89	1,344,369	1,442,493	98,124	7.30
Total	11,144,081	11,254,939	110,858	0.99

Source: As for Figure 6.

years, that is, by 6.5 years, and for women, from 32.2 years in 2001 to 35.7 years in 2020.

Table 3 compares our projection results for men with those of GAD's. Our model predicts larger numbers in each age group, with proportionately larger numbers the older the age group. For those in their 50s and 60s, the percentage difference is relatively small as mortality rates have only a small effect. For those in their 70s and 80s, the percentage difference becomes much larger because of the larger differences between GAD's and the model's mortality rates. By far the largest difference occurs in the age group 80–89.

A similar projection was made for women (see Table 4). It was found that the difference in projected numbers between our model and GAD's was smaller, because life expectancy improvements for women have not been as large in recent years. The difference in projected numbers is largest for women in their 80s, whereas there were also significant differences for men in their 70s. Taking men and women together, our model predicts approximately 595,000 more persons than GAD's—484,000 men and 111,000 women. Had we used the Office for National Statistics' (ONS) 2006-based principal projections as our comparison set, instead of GAD's 2001-based ones, the results would have been closer.

6. Conclusions

In an era when the population is ageing rapidly, it is important that official population projections are as accurate as possible in order to inform policy and plan public finances. However, many national demographic agencies have struggled to improve the accuracy of their projections, especially for the older age groups, even over relatively short projection periods.

In this paper, we have offered an alternative way of projecting older populations based on trends in life expectancy and survivorship at ages of 50 and over. Our procedure differs from traditional methods, which base assumptions on trends in age-specific mortality rates. Using life table data for England and Wales to estimate survival, we found that survival data showed consistent and predictable trends.

Further, we tested whether our model would have produced more accurate estimates of survival than the UK Government's own projections, using the same data. We found that our method produced considerably more accurate projections than GAD's 1981-based projections. We also tested our method against the 1991-based GAD projections (results not reported here), and found that our method again gave more accurate results than GAD's, although the differences were less.

These tests of the method were applied only to the population aged 50 and over and cannot be regarded as an alternative method of making population projections across the entire age range. Two key findings, however, are that accuracy was generally improved in each age group for both men and women, and for projections further into the future.

The differences in accuracy between our model and GAD's arise from two possible sources. Using mortality data as a basis for determining life expectancy, GAD projections assumed a slower improvement than was assumed by our method. It is uncertain whether this slower improvement was due to the extrapolation methods used to project the mortality rates, or to the improvements judged likely in mortality, since the two approaches differ in their starting assumptions.

The projected populations were also underestimated in GAD's 1991-based projections, in which the assumed improvements in mortality were different. This might suggest that the assumptions used in the methods of extrapolating age-specific mortality rates at the time erred on the conservative side. However, an underlying cause might be the irregular nature of time series for age-specific mortality rates, making them difficult to extrapolate and thus susceptible to error. Indeed the benefit of our method is that it is easier to extrapolate survival rather than mortality, since the former is smoother than the latter in statistical terms.

In our projections, the assumption that there will be a continuation of the increase in improvement in life expectancy is based on the observed steady upward long-term trend. Arguably, we have no reason to assume that the growth in life expectancy will diminish in the near term. The ONS's recent (unofficial) long-term projections (personal communication) project the life expectancy for men at age 50 in 2020 to be 1.5 years shorter, at 33.1 years, than is projected by our method. For women, ONS's projection of 36.2 years is slightly longer than ours, which is 35.8 years.

On the question of a possible limit to future improvements in life expectancy, we take a pragmatic view. The fact that countries such as Japan, where life expectancy is still increasing, have already recorded more years of expected life than the UK is evidence that natural limits have not yet been encountered. Where those limits occur is an open question but one that needs to be addressed in framing long-term projections (see Wilmoth and Robine 2003; Oeppen and Vaupel 2006; Vaupel 2010).

Even over the short to medium term, the implications of our results remain considerable. By 2020, life expectancy for men at age 50 is projected to be just below that for women, who are projected to live a further 35 years. In 1960, the life expectancy of men was 22.5 years—12.5 years less than the 2020 figure. Moreover, a man reaching age 50 in 2020 is projected to have a 4.5 per cent chance of reaching age 100, while a woman is projected to have an 8.8 per cent chance (see Table A1).

These findings underline the speed at which the ageing population will grow during the next few decades. For example, our projections for the size of the population aged 50 and over in 2020 were 595,000 larger than GAD's 2001-based principal projection, which will have significant implications for pensions and other spending priorities at older ages (Blake and Mayhew 2006).

However, it is important to sound three cautionary notes. Firstly, our method has focused on creating period life tables for future years using age-specific survivorship rates at a given time with a series of linked period life tables. A different approach would involve the creation of a cohort life table, which would theoretically be more appropriate but would require further investigation to see if it were possible to overcome some of the technical difficulties of using a cohort-based approach.

Secondly, further work is needed to extend our method to younger age groups. Survivorship rates at ages up to 50 are very high, and the deaths that do occur are either concentrated in the first years of life, or in early adulthood. A start age of 30 would also not have these problems and so can be considered. It may be that a modestly adapted method would suffice for younger ages, although it is not clear that it would provide a significant improvement because the accuracy of population projections in these age groups is more dependent on the assumptions about fertility and migration, than those about mortality.

Finally, there are other methods of projecting mortality in common use, and we have not compared the accuracy of theirs with that of ours. The Lee–Carter model did not exist in published form at the time of the base years, 1981 and 1991, considered in this paper. Repeating our study using the Lee–Carter model instead of GAD's might be a possibility, although it would not be easy to select an appropriate version from the many available.

New techniques for projecting mortality will undoubtedly appear in the literature, but it may take years before they can be fully evaluated. We agree with Booth (2006) that the accuracy of population projections should be regularly tested and the results published, so that evidence of improvement can be established. Currently, researchers appear to be more concerned with technical advances in methods than with the accuracy of the projections they produce, in particular the central projections used for government financial planning.

As Brass (1974) commented: '[demographers] accept responsibility for the formal processes of projection . . . but are not prepared to take the further step of specifying (however cautiously) the plausibility of the assumptions, and thus to change the projection into a prediction'. Clearly, there has been some progress in the last three decades but continued criticism in more contemporary literature shows that the problem has not gone away. The current rapid improvements in longevity are likely to bring the issue into even sharper focus as government faces up to the economic and social impacts of an ageing population.[1]

Note

1 Les Mayhew and David Smith are at the Faculty of Actuarial Science, Cass Business School, 106 Bunhill Row, Moorgate, London EC1Y 8TZ, UK. E-mail: lesmayhew@googlemail.com

References

Bengtsson, T. and N. Keilman (eds.). 2003. *Perspectives on Mortality Forecasting: Current Practice*, Social Insurance Studies No. 1. Stockholm: National Social Insurance Board.

Benjamin, B. and E. Overton. 1980. *The Prospects for Mortality Decline in England and Wales*. London: Policy Studies Institute.

Blake, D. and L. Mayhew. 2006. On the sustainability of the UK state pension system in the light of population ageing and declining fertility, *Economic Journal* 116(51): F286–F305.

Booth, H. 2006. Demographic forecasting: 1980 to 2005 in review, *International Journal of Forecasting* 22(3): 547–581.

Brass, W. 1974. Perspectives in population prediction: illustrated by the statistics of England and Wales, *Journal of the Royal Statistical Society, Series A* 137(4): 532–583.

Forfar, D. O. and D. M. Smith. 1987. The changing shape of English life tables, *Transactions of the Faculty of Actuaries* 40: 98–134.

GAD. 2001. 2001-based principal projections. Available: http://www.gad.gov.uk/Demography%20Data/Population/index.aspx?y=2001&v=Principal (accessed: 13 November 2012)

Gompertz, B. 1825. On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies, *Philosophical Transactions of the Royal Society of London* 115: 513–585.

Haberman, S. and A. Renshaw. 2008. Mortality, longevity and experiments with the Lee–Carter model, *Lifetime Data Analysis* 14(3): 286–315.

Haberman, S. and A. Renshaw. 2009. Measurement of longevity risk using bootstrapping for Lee–Carter and generalised linear Poisson models of mortality, *Methodology and Computing in Applied Probability* 11(3): 443–461.

Human Mortality Database. University of California, Berkeley (USA) and Max Planck Institute for Demographic Research (Germany). Available: www.mortality.org or www.humanmortality.de (accessed: 13 November 2012)

Kannisto, V. 2000. Measuring the compression of mortality, *Demographic Research* 3(6). doi:10.4054/DemRes.2000.3.6

Keilman, N. 1997. Ex-post errors in official population forecasts in industrialized countries, *Journal of Official Statistics* 13(3): 245–277.

Keilman, N. 2001. Data quality and accuracy of United Nations population projections, 1950–95, *Population Studies* 55(2): 149–164.

Keilman, N. 2002. Why population forecasts should be probabilistic—illustrated by the case of Norway, *Demographic Research* 6(15): 409–454.

Keilman, N. 2007. UK national population projections in perspective: how successful compared to those in other European countries?, *Population Trends* 129: 20–30.

Keilman, N. 2008. European demographic forecasts have not become more accurate during the past 25 years, *Population and Development Review* 34(1): 137–153.

Keyfitz, N. 1981. The limits of population forecasting, *Population and Development Review* 7(4): 579–593.

Lee, R. D. and Lee Carter. 1992. Modelling and forecasting U.S. mortality, *Journal of the American Statistical Association* 87(149): 659–671.

Lee, R. and T. Miller. 2001. Evaluating the performance of the Lee–Carter method for forecasting mortality, *Demography* 38(4): 537–549.

Makeham, W. M. 1860. On the law of mortality and the construction of annuity tables, *Journal of the Institute of Actuaries and Assurance Magazine* 8(6): 301–310.

Mayhew, L. 1987. Resource inputs and performance outputs in social security, *Journal of the Operational Research Society* 38(10): 913–928.

Mayhew, L. 2001. *The Japanese Longevity Revolution and the Implications for Health Care Finance and Long-Term Care*, Interim Report IR-01-010. Laxenburg, Austria: IIASA.

Mayhew, L. and D. Smith. 2007. Using queuing theory to analyse the Government's 4-hour completion time target in Accident and Emergency departments, *Health Care Management Science* 11(1): 11–21.

Mayhew, L. and D. Smith. 2011. Human survival at older ages and the implications for longevity bond pricing, *North American Actuarial Journal* 15(2): 248–265.

Oeppen, J. and J. W. Vaupel. 2006. The linear rise in the number of our days, in T. Bengtsson (ed.), *The Linear Rise in Life Expectancy: History and Prospects. Perspectives on Mortality Forecasting*, Social Insurance Studies No. 3. Stockholm: Swedish Social Insurance Agency, pp. 9–18.

Olshansky, S. J. and B. A. Carnes. 1977. Ever since Gompertz, *Demography* 34(1): 1–15.

Preston, S. H. 1974. An evaluation of postwar mortality projections in Australia, Canada, Japan, New Zealand, and the United States, *World Health Organisation Statistical Report* 27: 719–745.

Renshaw, A. and S. Haberman. 2003. Lee–Carter mortality forecasting: a parallel generalized linear modelling approach for England and Wales mortality projections, *Journal of the Royal Statistical Society, Series C* 52(1): 119–137.

Renshaw, A. and S. Haberman. 2006. A cohort-based extension to the Lee–Carter model for mortality reduction factors, *Insurance: Mathematics and Economics* 38(3): 556–570.

Shaw, C. 2007. Fifty years of United Kingdom national population projections: how accurate have they been?, *Population Trends* 128: 8–22.

Vaupel, J. W. 2010. Biodemography of human ageing, *Nature* 464(7288): 536–542.

Wilmoth, J. R. and J.-M. Robine. 2003. The world trend in maximum life span, *Population and Development Review* 29(Suppl. 3): 239–257.

Appendix: percentage projected by model to reach age 100 for those aged 50 and 80 in different years

Table A1

	Year of projection				
Age	1951	1981	2001	2010	2020
(a) Men					
50	0.014	0.129	0.535	1.152	4.483
80	0.063	0.411	1.103	1.967	6.370
(b) Women					
50	0.098	0.738	1.741	4.876	8.801
80	0.264	1.381	2.730	7.297	12.298

Source: As for Table 1.