“Achieving closure”: improving the estimation of life expectancy for small populations

Michael P. Grayer

31st December 2009

Department of Geography, Queen Mary University of London, Mile End Road, London E1 4NS. Email: m.p.grayer@qmul.ac.uk

Abstract

Life expectancy is a useful indicator of mortality for policy-makers, though its calculation for small populations is problematic. In the UK, two technical reports, one from the Office of National Statistics (ONS), and the other from the South East Public Health Observatory (SEPHO), were commissioned to investigate life table methods applied to small populations. As a result, the ONS published “experimental” life expectancy indicators following their recommendations, and continue to do so. However, some of the life expectancy estimates produced by this methodology are implausibly high.

This paper attempts to address a key reason why the current methodology is prone to error by examining a specific area of the life table: the method of closing the table. The applicability of the current method to small populations is critically assessed, with particular attention paid to how small numbers of deaths occurring in the final age band lead to implausibly long survival times being estimated.

Two alternative methods of closing the life table are suggested. The first comes from the literature (Silcocks, 2004). The other is proposed here. This method uses a Brass relational model to extrapolate survivorship beyond the start of the final age interval, allowing for a smoother transition from the rest of the life table until its eventual culmination. The three methods are compared using empirical data: deaths and population estimates for 625 small areas (electoral wards) in London, for the years 2001-05. The Brass extrapolation method reduces the skew in life expectancy estimates seen when using the standard method, though the method appears to work better for males than for females. Reasons why this may be the case, and possible improvements, are suggested.

1 Introduction

1.1 Context of the study

Life expectancy, which is defined as the average length of time a new-born baby in a given population could expect to live if it experienced the same age-specific mortality rates as prevailed at the time of their birth, is arguably the most readily recognised measure of population mortality among policy-makers and the general public (Gardner and Donnan, 1977; Raleigh and Kiri, 1997; Veugelers and Hornibrook, 2002; Williams et al., 2005). Though in widespread use at the national level, its use at the sub-national level to make comparisons of mortality between areas is less common. Measures such as the age-standardised mortality rate (ASR) or standardised mortality ratio (SMR) are more commonly used. However, it is argued that life expectancy is a more useful indicator of mortality for policy-makers because it is measured in tangible, concrete, quantifiable units (Silcocks et al., 2001; Veugelers and Hornibrook, 2002). Indeed, as part of the UK government's policy on reducing social health inequalities, in 2001 a target was set with the aim "by 2010 to reduce by at least 10 per cent the gap between the fifth of areas with the lowest life expectancy at birth and the population as a whole" (Department of Health, 2001). However, a drawback of life expectancy compared with standardised rates is that it is fairly complex to calculate, requiring the use of life table methods.

The Office for National Statistics (ONS) established the Neighbourhood Statistics Programme in 2000 with the goal of increasing the availability of neighbourhood level statistics, thus addressing "significant gaps in the information required for evidence-based policy making" (Office for National Statistics 2006). The development of life expectancy indicators for electoral wards during the course of this programme was regarded as "high priority" (Williams et al., 2005). This, however, represented a methodological challenge as the populations of electoral wards are far smaller than the recommended minimum size required for calculating life expectancy: in the United States, for example, the National Center for Health Statistics requires a minimum of 700 deaths for an estimate to be considered "stable". This implies a minimum population of 70,000 persons per administrative unit assuming crude death rates of 1% (National Center for Health Statistics, 1987, 1998; Cai, 2005). By contrast, typical electoral wards have populations of between 5,000 and 10,000 persons, and in remote parts of the UK, the population of these geographic units was counted at approximately 1,000 persons (at the time of the 2001 census). Even after the culmination of the Neighbourhood Statistics Programme in 2006, during which two technical reports on the methodology for calculating life expectancy for small populations were published, one by the ONS (Toson and Baker, 2003) and one by the South-East Public Health Observatory (SEPHO) (Williams et al., 2005), small-area life expectancy methodology is still considered "experimental".

This paper investigates the current "experimental" methodology, in particular identifying a component of the methodology which was largely overlooked by both reports—how the life table should be closed. An alternative method of closing the life table is proposed, and compared with the standard method and another method proposed specifically for small populations identified in the literature (Silcocks, 2004). Empirical data from electoral wards in London for the years 2001-06 are used to demonstrate the validity of these methods in practical use, rather than with simulated data.

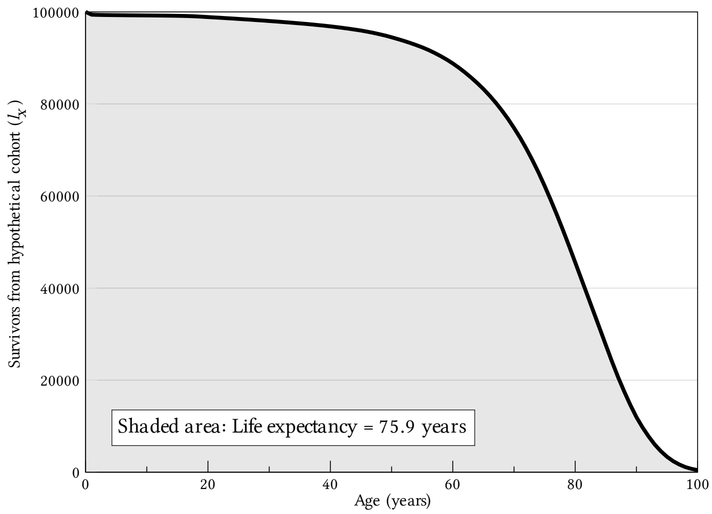

Figure 1: Survivorship curve of England males, 2000-02

1.2 Abridged life tables and life table closure

Before exploring the literature on the application of life table methods to small populations, the abridged life table method, as applied in general, is briefly described.

Calculating a life expectancy for a fixed point in time (as opposed to a cohort of births) requires the analyst to construct a *hypothetical* cohort of an arbitrary number of persons, and apply to it age-specific mortality rates from the population of interest. In this way, it is possible to generate a curve, the *survivorship* function, depicting the theoretical proportion of the hypothetical cohort surviving to different ages (in demographic notation, this is denoted by l_x. An example of a survivorship function is illustrated in Figure 1; this example shows survivorship of a hypothetical cohort of 100,000 persons based on age-specific mortality rates for males in England during 2000-02.

The total number of person-years lived by the hypothetical cohort (demographic notation T_x) is represented by the area under the curve from age x onwards. Life expectancy (demographic notation e_x) is derived by dividing this total by the number of survivors to age x; it is the average survival time past that age. Since the number alive at birth in our hypothetical cohort is arbitrary and fixed, life expectancy at birth, our outcome of interest, can be thought of as represented by the area under the whole curve.

In practice, however, it is not possible to get reliable mortality data for single years of age for small areas because the numbers of people at each age are so small. Not only does this raise questions of instability, there are issues with respect to non-confidentiality; data are unavailable due to the possibility that individuals may be identifiable. In this case, mortality rates for five-year (or in some cases ten-year) age-groups are used instead, and an abridged life table constructed. For five-year age-groups, the survivorship function is derived for ages 1, 5, 10, 15, and so on up to an arbitrarily chosen old age and the remainder of the function is derived by some method of interpolation, so as to determine the contribution to the total number of years lived by each five-year "slice". The Chiang method (Chiang, 1968), which determines each age-group's contribution using a term to represent the average length of time lived during the interval by those who die (denoted by a_x), is the most common method for doing this, and is the one recommended by both Toson and Baker (2003) and Williams et al. (2005). Conventionally, deaths are assumed to occur evenly throughout the time period ($a_x = 0.5$), with the exception of deaths under one year, where the majority of deaths occur within the neonatal period, within 28 days of birth. Here, $a_x = 0.1$ is conventionally used for developed-world populations (Newell, 1988).

Vallin and Caselli (2006) note two distinct problems faced by demographers in determining rates of mortality at the oldest ages. The first problem, they note, relates to the quality of empirically collected mortality data at old ages; the ages of the elderly are in many cases inaccurately reported, either exaggerating very old ages (Mazess and Forman, 1979; Retherford and Mirza, 1982; Rosenwaike and Stone, 2003) or "heaping" of reported ages on numbers ending in zero or five (Myers, 1940). The second problem relates to the fact that, as cohort members die out, mortality rates of the very elderly become inevitably based on observations of a very small number of people, making probability calculations derived from them less justifiable based on the law of large numbers. Vallin and Caselli note further that "this phenomenon naturally occurs earlier among men than among women, owing to excess male mortality" (Vallin and Caselli, 2006, p.121).

As a result of these problems, at some point the life table must be interrupted and some method for closing the life table be employed. The standard method is that proposed by Chiang (1968), where the number of survivors are divided by the mean mortality rate to determine the total survival time after the point at which the life table was interrupted. This method leads to the usage of terms such as "open-ended age group" to denote the end portion of the life table after the point at which it was interrupted.

1.3 Previous work

A number of studies have been conducted to propose solutions to the difficulties encountered in estimating life expectancy for use in comparisons of mortality between small populations. Silcocks et al. (2001) suggested that the sampling distribution of life expectancy was approximately normal, and that as a result life expectancy could be used for mortality comparisons, and was better for illustrative purposes than SMR. The authors also noted, however, that the formula for variance of life expectancy (and thus standard deviation and confidence intervals) were more complex. They also observed that the age at which the final age-group begins affected the calculation of life expectancy. They concluded that its lower bound should be as high as possible to avoid conflating differences in age structure among the oldest-old. However, the assumptions made in constructing their abridged life tables differed from the standard methodology as described by Chiang (1968).

Two technical reports evaluating methodological options for making routine estimates of life expectancy at the electoral ward level have been published by UK organisations. The first (Toson and Baker, 2003) was commissioned by the Office for National Statistics (SEPHO); the second (Williams et al., 2005) was commissioned by the South East Public Health Observatory (SEPHO).

The main body of both reports was a comparison between the abridged life-table methods proposed by Chiang (1968) and Silcocks et al. (2001). Both reports came up with similar conclusions regarding the most appropriate method for making point estimates of life expectancy: the Chiang method was considered the most accurate in both cases. Both papers also concluded that the minimum population threshold for reasonably reliable estimates of life expectancy to be made should be 5,000 persons, and prescribed a set of workarounds to account for instances when zero counts in either the deaths or population estimates caused the life expectancy calculation to fail. However, the reports made different conclusions as to the best method for calculating variance of life expectancy, with the difference being manifest in the method for estimating the contribution to overall variance made by the final age interval.

Life expectancy estimates at the electoral ward level for all wards in England and Wales for the period 1999-2003 were published as "experimental" data (Office for National Statistics, 2006), following the methodology suggested by Toson and Baker (2003). Five years of data were aggregated to ensure that the minimum threshold of 5,000 person-years at risk was met in all wards. Life expectancies for both sexes combined were published for all wards; life expectancy disaggregated by sex could not be calculated for wards with very small populations where gender-specific populations did not meet the minimum threshold so were only estimated for a subset of wards.

Despite being published with the caveat that the data were estimated using an experimental methodology, and being accompanied by a report detailing the limitations of the data, a number of policy papers, keen to present life expectancy data for small areas, cite the ONS small-area life expectancy estimates with little or no mention of their being "experimental" (Mindell et al., 2004; Bailey and Walrond, 2005), and a number of papers have uncritically used the experimental methodology to produce life expectancy estimates for neighbourhoods in other countries (New Zealand Ministry of Health, 2005; Hogstedt et al., 2006; Huong et al., 2006).

Before moving on to other studies, several other findings from the Williams et al. (2005) report are of relevance to this paper; these findings were published as papers in their own right in the *Journal of Epidemiology and Community Health* (Eayres and Williams, 2004; Williams et al., 2004). The first concerns an investigation on the effect of overall population size and lower-bound for the final age-interval on life expectancy and its standard error Eayres and Williams (2004). The authors noted that as lower-bound was raised, both life expectancy and standard error became more skewed, and that as population size decreased, life expectancy retained its accuracy but standard error became more skewed. They noted that the distribution of estimated survival time within the final age-interval was skewed, observing:

> "The right hand tail becomes stretched, and as population size decreases, the skewing of the transformed survival time distribution increases, shifting the distribution mean further and further to the right. This results in an overestimate of the underlying survival time, and consequently the LE [life expectancy]."

However, the authors failed to note that this was distinct from other age-intervals, suggesting:

> "Similar effects can be expected in the other, finite, age intervals where the years of life lived during the interval is related to the probability of dying (Chiang) or probability of survival (Silcocks), both of which are transformations of the mortality rate M_i. "

Silcocks (2004) revisited this problem, suggesting an alternative formula for estimating mean survival time in the final age-group. The proposed method works by multiplying the Chiang estimate for mean survival time by a "shrinkage factor". This is a number between 0 and 1, and moves further from 1 as the number of deaths upon which the original estimate was based decreases. This method, as well as the original Chiang method, is described in more detail in Section 3.2.

The other relevant observation arising from the report is one concerning the potential for local geographic factors to confound life expectancy within small, contiguous geographic units (Williams et al., 2004). The factor considered by the authors was the location of nursing homes, based on findings from a previous study of geographical differences of mortality within the borough of Croydon (Williams et al., 1995). Here it was observed that the mortality rates of nursing home residents were higher than their same-age counterparts who remained in private accommodation, and that the effect of this was to negatively bias life expectancy in wards containing a high number of nursing home residents. The authors suggested discounting deaths of nursing home residents in order to mitigate against the effects of this bias. However, there are two drawbacks of this. The first is that the bias is only considered to act in one direction, and doesn't adjust life expectancy in wards with no nursing homes where elderly residents automatically migrate out of a ward to move to a nursing home. The second is that no distinction is drawn between nursing home residents who previously lived in the same ward and those who moved from another ward to get there.

Other studies have investigated this phenomenon. Manuel et al. (1998) suggested that this effect was minimal based on a correlation of mortality in young ages with old ages. However, the strength of the correlation was not particularly strong (0.554 Pearson correlation coefficient; no p-value given) and even so, the choice of groups to act as the two variables in the correlation was rather unusual, with no justification given as to why they should accurately represent young-age and old-age mortality. A more satisfactory analysis is offered by Veugelers and Hornibrook (2002), who recoded nursing home deaths by place of previous residence, comparing the life expectancy estimates produced with those without any recoding. Differences of up to a year of life expectancy were found. However, this study was conducted in Canada, where data on previous residence are available. These data are sadly not collected universally, for example, data on previous residence of nursing home residents are not available in the UK (Williams et al., 2004).

A further contribution to the development of small-area life expectancy estimation was made by Scherbov and Ediev (2009). Here, it was observed that the reliability of life expectancy estimates not only depended on population size but also on the age-structure of the population. Thus a weakness of the previous papers emerges: simulations were conducted based solely on a single population age structure. Thus while mortality levels vary from area to area and we are principally concerned with these, we cannot ignore the fact that local factors affect the age-composition of a small area such as an electoral ward or census district as well. The "nursing homes problem" described above is an example of this. This question is also salient when considering life expectancy disaggregated by sex, as the population structure of the female population differs from that of the male population, particularly at old ages.

Thus, at present, a number of problems remain unsolved. Though "5,000 person-years" is now commonly accepted as a threshold for reliable life expectancy estimates by policy-makers, we don't yet know how local differences in population age-structure affect the validity of this generic rule. The problem of skewness in the final, open-ended, age interval still remains, and is an important hurdle to overcome if it is hoped that reliable life expectancies for even smaller populations can be obtained. Finally, previous studies have used simulated data, in most cases only conforming to a single age-structure; Scherbov and Ediev (2009) is an exception. The present study uses empirically observed data.

With the exception of Silcocks (2004), no studies have been conducted investigating the effect of different methods of life table closure on life expectancy estimates for small populations, despite there being many proposed options for larger populations (Vallin and Caselli, 2006). Here, an alternative method is proposed, based on the relational life table method proposed by Brass (1971). Rather than expressing survival after the point of interruption as a single fraction, it is proposed to extrapolate survival for the oldest-old based on survival rates at preceding ages and known survivorship patterns among the oldest-old taken from a larger population.

2 Aims

The aims of this study are as follows.

1. To improve the reliability of life expectancy estimates for small populations by considering alternative options for closing the life table. Two alternatives are considered: the "shrinkage factor" method proposed by Silcocks (2004) and extrapolation of survival using a Brass relational model as outlined above.

2. To produce estimates of life expectancy and other indicators of mortality for 625 geographical units (hereafter referred to simply as "wards") in London: the 624 electoral wards that comprise the 32 boroughs of Greater London, along with the City of London taken as a whole. The wards range in size from approximately 5,000 to 18,000 persons.

3. To consider the effect of estimating life expectancy separately for males and females, and whether the methodology should differ by sex, considering the differences in age-structure.

3 Methods

3.1 Setting and data

Mid-year population estimates and annual death counts were obtained from ONS for all 625 wards London for 2001-2006. Population estimates and death counts were disaggregated by sex and also by age-group, of which there were nineteen: under 1 year, 1-4 years, 5-9 years, ..., 80-84 years, 85 years and over.

In this study, two different time scales were considered. Firstly, data were aggregated over the five years 2001 to 2005, consistent with the five-year aggregation currently employed in the ONS experimental life expectancy estimates. Additionally, three separate two-year periods (2001-02, 2003-04 and 2005-06) were considered. The purpose of this was to investigate population sizes much closer to the 5,000 person-years minimum threshold currently in common use (Toson and Baker, 2003; Williams et al., 2005).

For the purposes of implementing the Brass relational model, the English Life Table 2001 was obtained from the Government Actuaries Database. Survivorship (l_x) values for males and females separately at the start of each age-group mentioned above and additionally at ages 90, 95 and 100 were used as the standard life table.

3.2 Life table methods

Age-specific mortality rates (M_x) were estimated simply by dividing the aggregated number of deaths by the aggregated mid-year population estimates for the years concerned. The Chiang life table method was then used to calculate the survivorship (l_x) column and hence the number of person-years lived throughout the age-group (L_x) up until age 85, for males and females separately, for each of the 625 wards. From age 85, however, three separate methods were used to close the life table. Each is briefly described below.

3.2.1 The "standard" method (Chiang)

This method is widely used as the standard method for closing the life table (Newell, 1988). It is the method employed in the studies cited in Section 1.3. Here, we assume that survivors to the start of the final age-group (i.e. age 85) die out at an average rate of M_{85+} until there are no further survivors. Hence:

$$L_{85+} = \frac{l_{85+}}{M_{85+}}$$

This can equivalently be expressed as treating the final age-group as if it were a normal age-group, with width equal to $\frac{2}{M_{85+}}$ (Silcocks, 2004).

3.2.2 The "shrinkage factor" method (Silcocks)

This method can also be expressed as treating the final age-group as if it were a normal, closed, age-group, but with a width shrunk by a factor related to the number of deaths upon which the estimate of mortality was based. Silcocks' formula for estimating the width of the final age interval (z_{85+}) is:

$$z_{85+} = \frac{2}{M_{85+}}\left[1-\left(\frac{1-M_{85+}}{D_{85+}}\right)\right]$$

where D_{85+} is the number of deaths occurring throughout the final age group. Hence the number of person-years, denoted here by L^*_{85+}, is calculated by:

$$L^*_{85+} = \frac{l_{85+}}{M_{85+}}\left[1-\left(\frac{1-M_{85+}}{D_{85+}}\right)\right]$$

3.2.3 Extrapolating survival using the Brass relational method

The Brass relational model is a system for estimating a survival curve as a transformation of another, known as the standard. More correctly, the logit of the fitted survival curve is expressed as a linear translation of the logit of the standard:

$$0.5\ln\left(\frac{1-\hat{l}_x}{\hat{l}_x}\right) = \alpha + \beta \cdot 0.5\ln\left(\frac{1-l^s_x}{l^s_x}\right)$$

where $\hat{l}_x$ denotes the fitted survivorship curve and l^s_x the standard survivorship curve at age x, and α and β are parameters obtained by linear regression of observed values of l_x against corresponding values of l^s_x. A question that arises here is at what age the regression should be started. The first age should be early enough to include enough data points, but if the regression begins too early then old-age survival will be unduly influenced by mortality at old age. 50 years was chosen as a reasonable compromise start point, allowing eight data points. Hence, for this study, the parameters α and β are derived either by regressing l_{50}, l_{55}, ..., l_{85} against the corresponding $\hat{l}_x$ values in the standard population using statistical software, or can be calculated directly using the following formulae:

$$\beta = \frac{\sum(z_x-\bar{z})(y_x-\bar{y})}{\sum(z_x-\bar{z})^2} \qquad (1)$$

$$\alpha = \bar{y}-\beta\bar{z} \qquad (2)$$

where

$$z_x = 0.5\ln\left(\frac{1-l^s_x}{l^s_x}\right)$$

$$y_x = 0.5\ln\left(\frac{1-l_x}{l_x}\right)$$

$\bar{z}$ and $\bar{y}$ are the mean values of z_x and y_x respectively and the sums in equations 1 and 2 are across the eight values of x upon which the regression is based (from ages 50 to 85).

From here, fitted values of $\hat{l}_x$ are obtained for ages 50 through to 100 by taking the anti-logits:

$$\hat{l}_x = \frac{1}{1+\exp(2(\alpha+\beta z_x))}$$

The final problem is to then piece the two survivorship curves, l_x and $\hat{l}_x$ together. Here, to protect against the possibility that $l_{85} < \hat{l}_{90}$ (i.e. members of the hypothetical cohort come back to life!), survivorship at age 85 is taken to be the mean of the observed and fitted values. Hence the final survivorship curve is defined by:

$$\begin{array}{rll} l_x & \text{for} & 0 \leqslant x \leqslant 80 \\ \frac{1}{2}\left(l_x + \hat{l}_x\right) & \text{for} & x = 85 \\ \hat{l}_x & \text{for} & 85 \leqslant x \leqslant 100 \end{array}$$

and the remainder of the life table is calculated using the standard Chiang methodology up to age 100, after which point the life table is closed by assuming mean survival time beyond age 100 to be the same as in the standard population.

3.3 Data analysis

Life expectancies were estimated for males and females separately for all 625 wards for each of the time periods described in Section 3.1, using each of the three life table closure methods described above. Additionally, directly age-standardised mortality rates are calculated for each sex in each of the 625 wards. The standard population in this case is the European Standard, following precedents set in other studies (e.g. Wheller et al., 2007).

An important criterion for the validity of using life expectancy as a comparitive measure of mortality is that the sampling distribution of life expectancy is approximately normal (Silcocks et al., 2001). Tests of skewness and kurtosis were carried out on life expectancy estimates produced by each of the three closure methods along with age-standardised mortality rates as a control.

Pearson and Spearman rank correlation tests were conducted on the life expectancy estimates to provide an indication of the consistency of life expectancy estimates produced by each method. The strength of correlation between males and females, and between each method of life expectancy and age-standardised rates during the same period was tested. Additionally, for the two-year data, the strength of correlation of life expectancy between individual two-year periods was tested.

Finally, the level of spatial auto-correlation between life expectancies of neighbouring wards was tested using Moran's I co-efficient. For this analysis a 625×625 neighbourhood matrix was constructed using ArcGIS software, where the ijth element was defined as 1 if wards i and j shared a boundary, and 0 otherwise.

Calculation of life expectancy and age-standardised rates, as well as all the statistical analysis, was carried out using Stata 10 for Unix.

4 Results

4.1 Descriptive analysis

4.1.1 Five-year data (2001-05)

Table 1 displays descriptive data on population and deaths for the 625 wards over the course of the five year period 2001-05. For both sexes, the number of population years lived during the time period ranged from just over 11,500 to over 40,000. The minimum ward size is therefore twice the size of the minimum population threshold recommended in the current ONS methodological literature of 5,000 person-years. The crude death rate during this period was slightly higher for females than for males, though this can be accounted for by the differences in age structure between the male and female populations; over twice as many deaths occurred to women aged 85 years or over than to men of the same age. The number of deaths occurring in that age-interval was also very varied, with a standard deviation around two-thirds of the mean. Some wards had particularly low numbers of deaths during that age-group, particularly among males, where the minimum number of deaths in a ward was five. The heterogeneity of the population age structures, as well as stochastic variation in numbers of deaths, therefore lead to very small numbers in the age-intervals, despite population sizes of the wards being well above the minimum threshold currently suggested.

Table 3 displays summary statistics of central tendency and variation of life expectancy calculated using the three different methods of life table closure. Weighted mean life expectancies were calculated using the following formula:

$$\bar{LE}_{weighted} = \frac{\sum_{i=1}^{625} n_i LE_i}{N} \tag{3}$$

where n_i and LE_i are the total population and life expectancy in ward i respectively, and N is the total population for London.

The weighted mean life expectancy of the 625 wards was similar to the overall life expectancy for London as a whole, though the estimates produced by the Brass extrapolation method were lower than with the other two methods. There was very little difference between the ONS and Silcocks methodologies. However, with no established method of obtaining a standard error, it is not possible to say whether or not these differences were significant.

4.1.2 Two-year data (2001-02, 2003-04, 2005-06)

Table 2 presents summary descriptive statistics of the population sizes and death counts of all wards for each two-year sub-period. The minimum population time at risk fell slightly below the 5,000 person-year threshold suggested in the ONS methodology. However, this only occurred in one ward (Darwin, Bromley) so it was not considered that this would significantly affect the validity of the results. The mean population-time at risk during each two-year sub-period was more than double the 5,000 threshold.

Table 1: Descriptive statistics of population and deaths for London wards, 2001-2005 ($n = 625$)

	Central tendency			Range	
	Mean	Median	Std. dev.	Min.	Max.
Females					
Total population-years	29,905.75	29,433	4,843.99	11,646	46,029
Total deaths	231.80	211	91.54	56	690
Deaths (85+)	93.75	79	58.01	10	411
Males					
Total population-years	29,123.36	29,038	4,772.36	11,668	43,061
Total deaths	218.21	214	58.58	70	437
Deaths (85+)	43.91	39	22.40	5	158

Crude death rate (per 1,000 persons):

Females 7.75

Males 7.49

The number of deaths occurring in each ward during each two-year sub-period also varied considerably between wards. This is a product not only of the size of the ward's population, but also of the overall level of mortality, and its age-structure. Of particular note are the number of deaths in the 85+ age group, which vary from as little as one death to nearly 200 deaths (for females) during a two-year period. The mean number of deaths in females in this age group is twice that of males. Thus small numbers in the final age interval are likely to be a greater problem among males than for females.

The life expectancy estimates produced by the ONS method, the Silcocks bias-correction method and the Brass extrapolation methods are summarised in Table 4. For males, the Brass extrapolation method produces mean and minimum values which are similar to those obtained by the other two methods. However, the maximum values are reduced by up to seven years. This enables male life expectancy estimates to be more consistent with female life expectancy: for 2003-04 and for 2005-06, the maximum life expectancy for males was estimated to be higher than for females, contrary to usual patterns. Using the Brass extrapolation method proposed here, however, the highest female life expectancy is higher than the highest male life expectancy for all three periods. Surprisingly, the Silcocks bias-correction method has very little effect on the extreme high life expectancy estimates.

However for females, the effect of using the Brass extrapolation method is more varied. In 2001-02 and in 2003-04, the minimum life expectancy calculated by the ONS method was the same, yet the Brass extrapolation method estimates over a year's difference in the minimum life expectancy between the two periods.

Table 2: Descriptive statistics of population and deaths for London wards, two-year data ($n = 1875$)

	Central tendency			Range	
	Mean	Median	Std. dev.	Min.	Max.
Females					
Total population-years	11,979.74	11,769	1,945.16	4,614	18,578
Total deaths	91.11	83	37.26	21	298
Deaths (85+)	36.99	31	23.55	2	183
Males					
Total population-years	11,673.21	11,632	1,922.62	4,660	17,536
Total deaths	86.21	84	24.70	19	203
Deaths (85+)	17.52	15	9.56	1	68

Table 3: Life expectancy estimates produced using three different life table closure methods, for 625 wards in London, 2001-2005

Closure method	Mean[a]	Min.	Max.
Females			
Life expectancy for London	81.18	—	—
Chiang (ONS method)	81.45	75.93	90.01
Silcocks bias-correction	81.44	75.93	89.99
Brass extrapolation	80.78	76.05	87.68
Males			
Life expectancy for London	76.49	—	—
Chiang (ONS method)	76.51	70.36	87.12
Silcocks bias-correction	76.50	70.36	87.10
Brass extrapolation	76.16	70.30	84.60

[a]Weighted mean, based on population size.

Table 4: Life expectancy estimates for two-year data calculated using three different methods, 625 London wards, 2001-06

	Mean[a]			Min.			Max.		
Closure method	2001-02	2003-04	2005-06	2001-02	2003-04	2005-06	2001-02	2003-04	2005-06
Females									
Chiang (ONS method)	81.32	81.56	82.58	74.92	74.92	76.21	92.55	91.10	93.36
Silcocks bias-correction	81.31	81.55	82.56	74.91	74.91	76.20	92.49	91.07	93.31
Brass extrapolation	80.51	80.92	81.62	74.09	75.13	76.24	87.87	89.42	88.61
Males									
ONS method	76.11	76.87	77.82	68.19	69.12	70.07	89.58	91.37	94.62
Silcocks bias-correction	76.09	76.84	77.79	68.18	69.11	70.06	89.41	91.18	94.56
Brass extrapolation	75.76	76.41	77.14	68.13	68.92	69.82	86.27	86.12	87.60

[a] Weighted mean, based on population size.

4.2 Discrepancies between the three methods

This section illustrates with practical examples how the three methods produce different estimates of life expectancy. Individual wards have been chosen as best illustrative examples where particularly large discrepancies occurred, and are not intended as a representative sample of all results. For the purposes of illustration, only two-year data are considered.

Differences between the standard method and the Silcocks "shrinkage factor" method were minimal in all areas. The maximum differential in any ward for any two-year period was 0.13 years for females and 0.28 years for males. However, differences between the standard method and the Brass extrapolation method were much larger. The largest reductions in life expectancy incurred by applying the Brass extrapolation method were 7.6 years for females and 15.4 years for males. The top two survivorship graphs in Figure 2 illustrate how the life expectancy has been reduced. In both cases, the "tail" of the survivorship curve (i.e. the portion beyond 85 years of age) approaches horizontal, extending the effective width of the final age interval beyond what is plausible. In the case of Green Street East, the survivorship curve doesn't reach zero until the age of 205 years. Unsurprisingly, it is this ward where the discrepancy of 15.4 years occurred.

Unlike the Silcocks method, however, the Brass extrapolation method can adjust life expectancy upwards as well as downwards. The greatest increase in life expectancy was 2.96 years for females and 2.34 years for males. Examples of this effect are shown in the lower half of Figure 2. Here, mortality in the final age group is so high that survivors to age 85 are predicted to all die out shortly afterwards, if the standard or Silcocks method of closure is followed. The Brass extrapolation method assumes a smoother end to the survivorship curve.

Absolute differences between the standard method and the Brass extrapolation are on average, fairly large. The mean absolute difference (that is to say, the mean change ignoring the direction of the change) was 1.09 years for females and 0.71 years for males. The mean difference between the Silcocks method and the standard was only 0.019 years for females and 0.024 years for males. In the following sections, although the tables include analysis of the Silcocks method, the text will only describe differences between the standard method and the Brass extrapolation.

4.3 Normality

Table 5 shows the results from testing the life expectancy estimates for each of the three methods along with age standardised rates for approximation of a normal distribution, using five-year data. For both males and females, the Brass extrapolation method reduces both skewness and kurtosis in life expectancy, bringing the distribution of life expectancy estimates for wards in London closer to the normal distribution. However, the p-values suggest that the data still can't be assumed to have been sampled from a normal distribution. Figure 3 illustrates how the right-hand tail of the distribution has been curtailed when the Brass extrapolation method is used to close the life table.

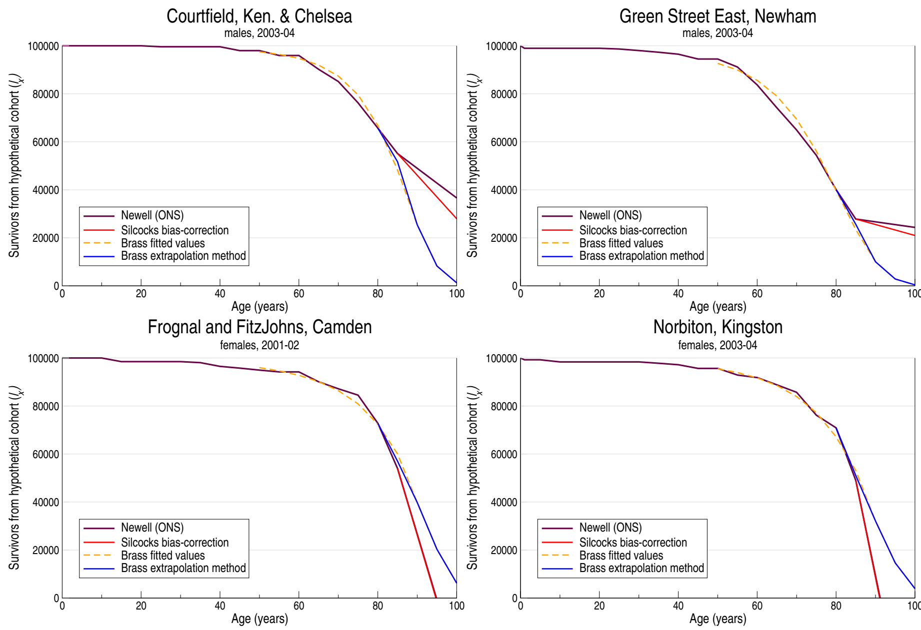

Figure 2: Illustrative examples of how the three methods of life table closure yield different estimates of life expectancy, selected London wards, two-year data

4.4 Comparing male to female life expectancy

Tables 6 and 7 show Pearson and Spearman rank correlation co-efficients for the association between male and female life expectancies as calculated using each of the three methods of closing the life table. Though the strength of correlations are slightly reduced when the Spearman rank test is employed, the correlations are still highly statistically significant (P<0.0005 in all cases). Correlations were stronger among the five-year data than for two-year periods.

For all periods, the Brass extrapolation method increases the strength of the association between male and female life expectancy to a similar level to the strength of the gender correlation of age-standardised mortality rates. The two-year data show a further interesting trend: in 2003-04, the strength of the life expectancy gender correlation dropped substantially below that of both 2001-02 and 2005-06 when measured by the standard method; in comparison, gender correlations of life expectancy when the Brass extrapolation method was employed remained constant throughout. One explanation for this may be the erroneous life expectancy estimated by the standard method for males in Green Street East (see Section 4.2).

Figure 4 shows a scatter plot of life expectancy for males against life expectancy for females for the standard method as compared with the Brass extrapolation method, for the five-year data 2001-05. This illustrates how the extreme outliers have been brought closer to the line of best fit, yielding a modest improvement in the correlation between male and female life expectancies.

Table 5: Skewness and kurtosis tests, and test for normality of life expectancy and age-standardised rates, London wards, 2001-05 (p-values in parentheses)

Measure	Skewness		Kurtosis		Test for normality[a]
Females					
Age-standardised rates	.275	(0.005)	3.26	(0.175)	(0.011)
Life expectancy:					
Chiang method (ONS)	.517	(<0.001)	3.66	(0.006)	(<0.001)
Silcocks bias-correction method	.516	(<0.001)	3.65	(0.006)	(<0.001)
Brass extrapolation method	.371	(<0.001)	3.11	(0.489)	(0.002)
Males					
Age-standardised rates	.244	(0.013)	3.02	(0.815)	(0.044)
Life expectancy:					
Chiang method (ONS)	.402	(<0.001)	3.60	(0.010)	(<0.001)
Silcocks bias-correction method	.399	(<0.001)	3.59	(0.011)	(<0.001)
Brass extrapolation method	.197	(0.043)	2.84	(0.431)	(0.096)

[a]The Stata skewness-kurtosis test is based on D'Agostino et al. (1990)

Table 6: Pearson correlations between males and females for life expectancy and age-standardised rates, London wards, five-year and two-year data.

	Five-year	Two-year		
Measure	2001-05	2001-02	2003-04	2005-06
Age-standardised rates	0.7624	0.6239	0.6036	0.6505
Life expectancy:				
standard	0.7432	0.6013	0.5396	0.5998
Silcocks	0.7434	0.6016	0.5409	0.6006
Brass extrapolation	0.7804	0.6182	0.6184	0.6120

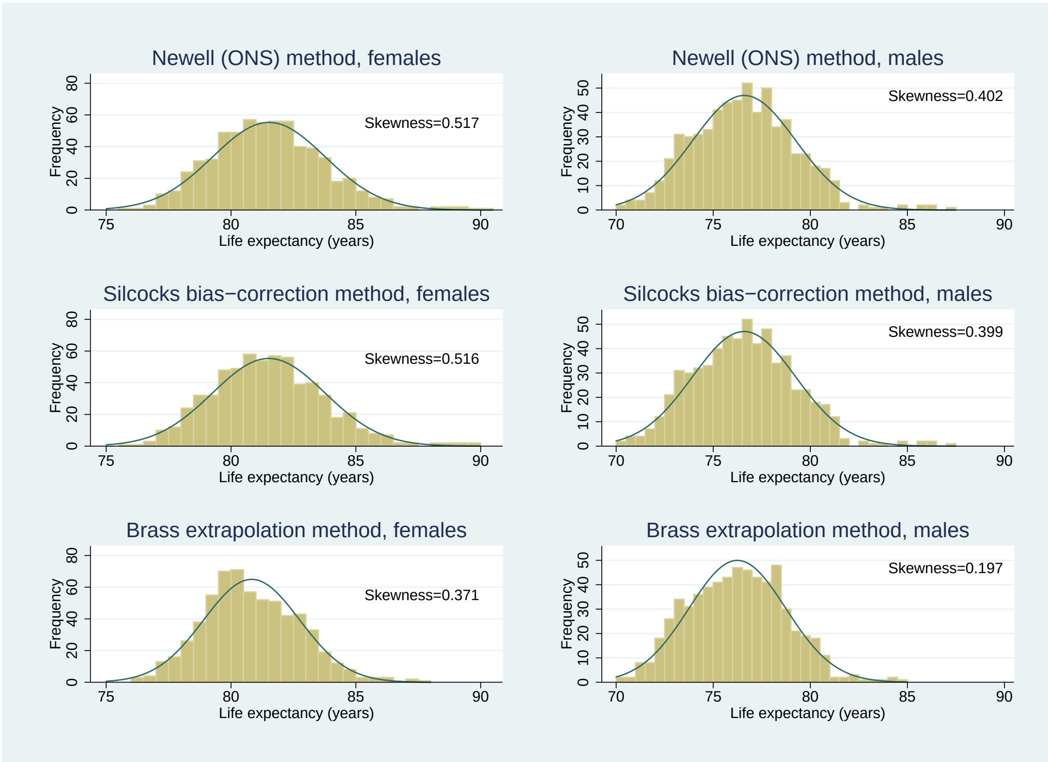

Figure 3: Distribution of life expectancy estimates, electoral wards, London, 2001-05 ($n = 625$)

Table 7: Spearman rank correlations between males and females for life expectancy and age-standardised rates, London wards, five-year and two-year data.

	Five-year	Two-year		
Measure	2001-05	2001-02	2003-04	2005-06
Age-standardised rates	0.7467	0.6172	0.5921	0.6194
Life expectancy:				
standard	0.7131	0.5695	0.5434	0.5698
Silcocks	0.7134	0.5698	0.5443	0.5707
Brass extrapolation	0.7732	0.5949	0.6240	0.6021

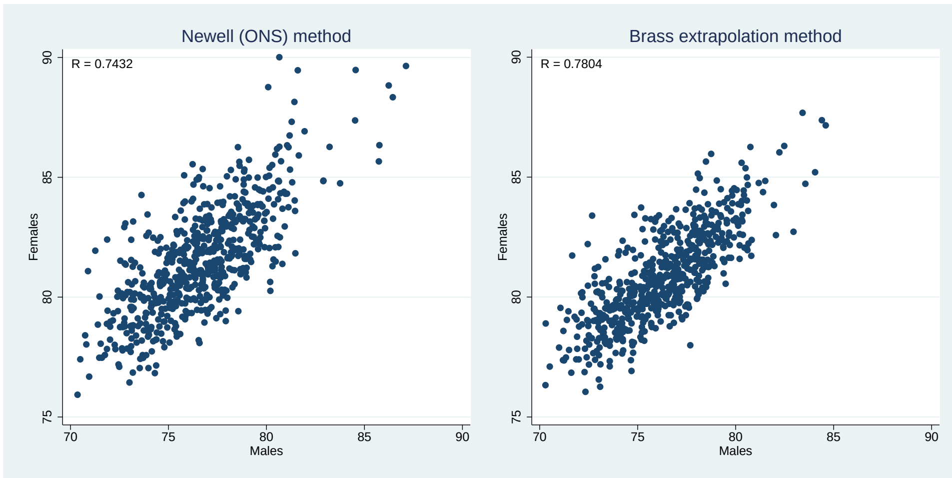

Figure 4: Scatterplots of female against male life expectancy, ONS method and Brass extrapolation method, electoral wards, London, 2001-05 ($n = 625$)

4.5 Correlations between time-periods

Table 8 shows the correlations between consecutive periods of different methods of estimating life expectancy, as well as age-standardised rates. Whereas the correlation between consecutive years improved when the Brass extrapolation method was applied to estimate male life expectancy, the equivalent correlation for females was weakened.

4.6 Spatial auto-correlation

Table 9 shows the results of testing for spatial auto-correlation of life expectancy between neighbouring wards, using each of the three life table closure methods. Life expectancy for both males and females exhibited strong evidence of auto-correlation regardless of what life table closure method was used (P<0.0005 in all cases). Spatial auto-correlation of life expectancy was stronger for males than females. Age-standardised mortality rates were less strongly auto-correlated than life expectancy for both males and females. Using the Brass extrapolation method to close the life table resulted in the spatial auto-correlation of life expectancy appearing to be stronger.

Figure 6 shows a cartogram of life expectancy estimates for electoral wards across London, 2001-05.

This spatial auto-correlation is illustrated using a cartogram (Figure 6). This is not a map of London based on geographic location, however, as this does not give an accurate portrayal of the number of people living in areas of high or low life expectancy. Indeed, as many inner London wards are densely populated, cover a small geographic area and have a low life expectancy (apart from in two boroughs

Table 8: Spearman rank correlations between consecutive two-year periods of life expectancy, 625 London wards

	Females		Males	
	2001-02 with 2003-04	2003-04 with 2005-06	2001-02 with 2003-04	2003-04 with 2005-06
Age-standardised rates	0.6297	0.6802	0.6891	0.6889
Life expectancy:				
standard	0.6021	0.6468	0.6883	0.6691
Silcocks	0.6024	0.6471	0.6893	0.6703
Brass extrapolation	0.5698	0.6242	0.7212	0.7329

Table 9: Moran's I statistic for spatial auto-correlation for life expectancies in London wards

	Five-year	Two-year		
	2001-05	2001-02	2003-04	2005-06
Females				
Age-standardised mortality rate	0.347	0.273	0.286	0.295
Life expectancy:				
standard	0.393	0.290	0.338	0.297
Silcocks	0.394	0.291	0.339	0.298
Brass extrapolation	0.463	0.334	0.365	0.333
Males				
Age-standardised mortality rate	0.414	0.353	0.349	0.351
Life expectancy:				
standard	0.447	0.379	0.345	0.374
Silcocks	0.448	0.380	0.347	0.376
Brass extrapolation	0.499	0.413	0.426	0.424

(P<0.0005 for all correlations)

Figure 5: PCT districts within Greater London, resized by ward population size, 2001-05

as described below) this visualisation tends to under-represent the number of people living in areas of low life expectancy. A cartogram was thus constructed representing each ward with an area equal to its estimated population size during the period 2001-05. This was done using an implementation of the "Rubber Sheet" algorithm for cartogram construction (Wolf, 2005) in the ArcGIS software package. The thick lines represent the boundaries of the regions of responsibility of Primary Care Trusts (PCT), the local bodies responsible for health planning in the local community. These areas are contiguous with the London boroughs, with two exceptions: Sutton and Merton are merged into one PCT, and the City of London is combined with the borough of Hackney. The names of PCT areas are illustrated in Figure 5.

The overall geographic trend for both sexes is that wards with low life expectancies are concentrated in East, North and South inner London (including Islington, Hackney, Tower Hamlets, Newham, Greenwich, Lewisham, Southwark and Lambeth). For males, Camden can also be included in that group, though for females, life expectancy in this borough is fairly mixed. Areas of high life expectancy are broadly Westminster and Kensington & Chelsea (except the most northern wards in those boroughs) as well as the more suburban outer London boroughs.

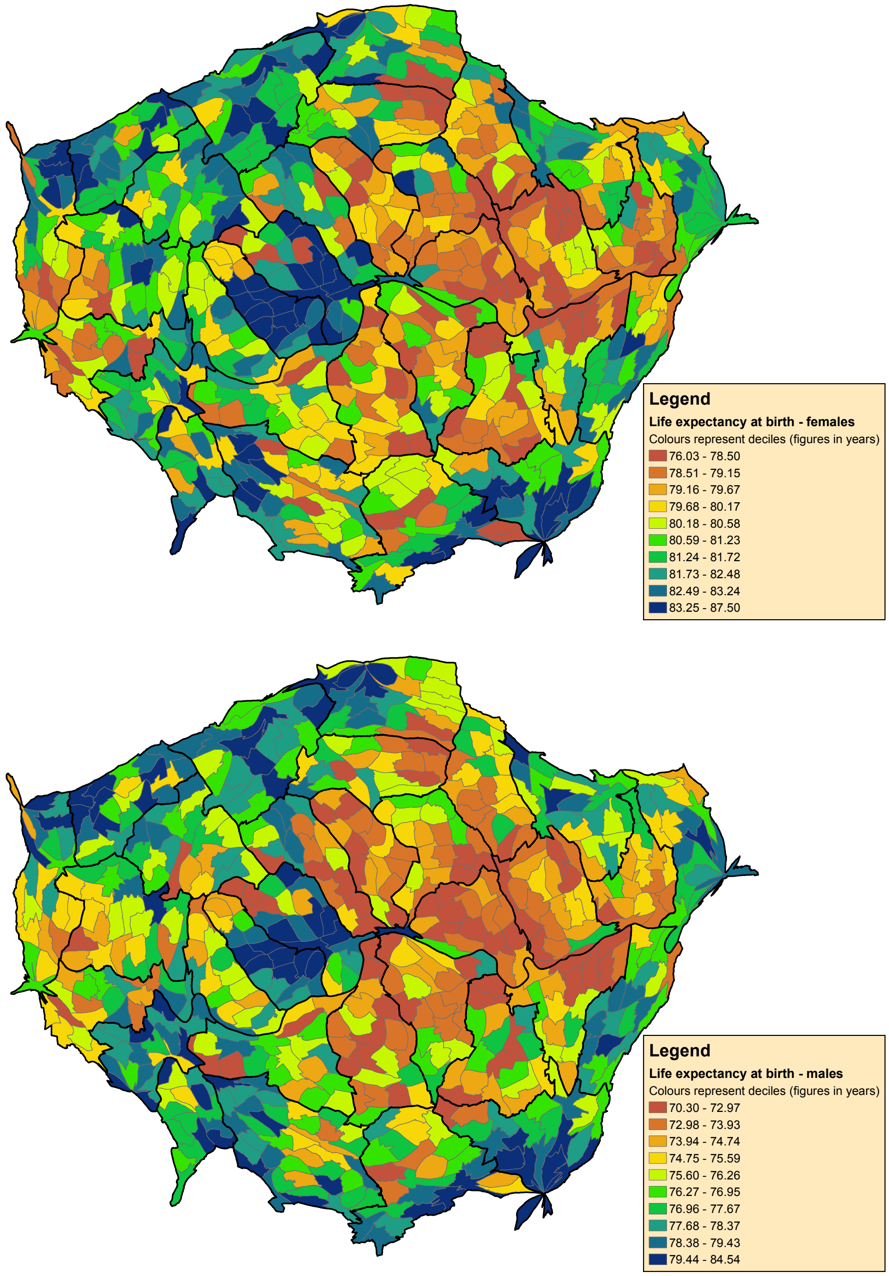

Figure 6: Cartogram of life expectancy, females (top) and males (bottom), using the Brass extrapolation method of life table closure.

5 Discussion

5.1 Conclusion

Survivorship curves for wards do fit reasonably closely to “typical” survivorship curves for larger populations up to age 85, despite the small numbers of deaths and the small populations involved. However, the trajectory of the survival curve after the point of interruption (in this case age 85 years) has a tendency to deviate from its expected path were it to follow a typical pattern through to its culmination. It is proposed here that this confounds estimates of life expectancy, and there is evidence from the empirical data used here to show that this is the case. The method proposed here, of closing the life table by extrapolating from a linear-logit transformation of a standard life table (as per Brass, 1971), as well as the “shrinkage factor” method proposed by Silcocks (2004), are intended to mitigate against this, enabling more accurate estimates of life expectancy to be made.

The choice of life table closure method does have a profound effect on life expectancy estimates for small areas. The Silcocks “shrinkage factor” method gave estimates that were mostly almost identical to those produced by the standard Chiang method, with a few exceptions in areas with few residents over the age of 85. However, there remained several wards where the trajectory of the life expectancy beyond the start of the final age interval was not consistent with what would typically be expected. This indicates that biases in life expectancy incurred through the final age group are not due to small numbers alone.

Extrapolation using a Brass relational model reduces the skewness of estimates in life expectancy, bringing the sampling distribution closer to the normal. This effect was more profound in males than in females. The trajectory of the survivorship curve beyond age 85 was more consistent with typical survivorship curves (as would be expected, since they were used as the standard) and there were no wards, even using only two years of data, with implausibly high life expectancy. The number of wards where male life expectancy was estimated to be higher than female life expectancy was also reduced.

However, as indicated by the correlations between life expectancy estimates in different two-year periods, the picture is not the same across both sexes. The Brass extrapolation method reduced the correlation in life expectancy between consecutive two-year periods for females. One explanation for this is the difference in age structure between the male and female populations, and in particular in the distribution of deaths, as shown in Figure 7. One limitation of the Brass extrapolation is that it discards deaths data beyond the age at which the life table is interrupted. While the average number of deaths discarded for males is not a large proportion of overall male deaths, more women survive beyond age 85, meaning that more deaths occur in that final age interval, and thus a greater proportion of the deaths data is discarded. This is consistent, however, with the notion that the point at which the life table is interrupted should be later for females than for males (Vallin and Caselli, 2006).

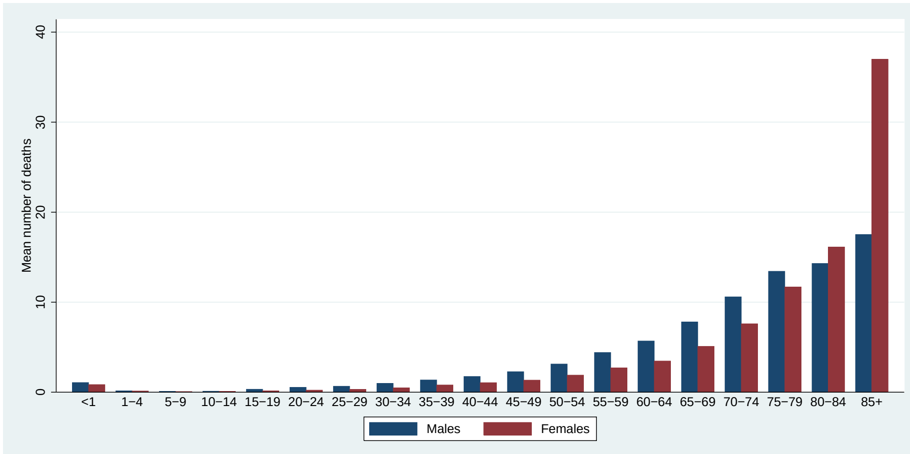

Figure 7: Mean number of deaths by age-group, two-year data (2001-02, 2003-04 and 2005-06 combined)

5.2 Limitations

While this method shows promise for the estimation of life expectancy for small populations as an indicator of mortality for use in comparative studies, there are a number of limitations. Possible avenues of investigation for overcoming these limitations are also considered here.

Firstly, most importantly, the issue of calculating variance (and hence standard error and confidence intervals) remains unsolved. This is because the contribution towards overall variance from the sections of the survivorship curve that have been extrapolated are as yet unknown. Methods of estimating standard error of a linear regression such as the ones used to predict survivorship beyond the interruption point of the life table are available, though with the formulae for calculating standard error of life expectancy already being complex and based on debated assumptions, such an approach may only serve to further muddy the waters. An alternative method would be to derive life expectancy using a Bayesian framework, thus directly estimating uncertainty (through Monte Carlo Markov Chain simulation) rather than relying on long, obfuscatory formulae.

Applying a Bayesian framework to the estimation of life expectancy for small populations would also address a second limitation: the underlying mortality rates themselves. Here, mortality rates are derived simply by dividng the number of deaths by the estimated population at risk, taking no account of age-patterns of mortality or the spatial structure of the data. A hierarchical Bayesian model allows the "intelligent smoothing" or "borrowing of strength" from the rest of the dataset, according to assumptions made based on the structure of the data (Graham, 2008). A Bayesian spatial smoothing technique was applied by Congdon (2002) to age-specific mortality rates that formed the basis of estimating life expectancy for use in informing policy on the allocation of health resources in north-east Greater London.

Another limitation in this study is the arbitrary choice of both starting age for the basis of the extrapolation and of the standard life table. Both of these were chosen on the basis that they were sensible options, though it is recognised that these choices should not be seen as definitive. Ideally, a sensitivity analysis should be carried out to determine what effect the choice of starting age and standard life table have on the resultant estimates of life expectancy, and if the effect is significant, what choice is most favourable. However, it is assumed here that these choices are not critical, and other plausible alternatives could have been chosen without drastically altering the results.

However, despite these limitations, the method of closing the life table by extrapolating the survivorship curve shows promise. It can be implemented without high-powered statistical software as the co-efficients α and β can be derived algebraically. It is shown to reduce the skewness of the distribution of life expectancy for small areas (thus increasing the validity of the assumption that life expectancy is normally distributed), and for males, to improve consistency in life expectancy both over time and with other measures of mortality. However, that this is not the case for females indicates that too much mortality data is lost when the life table is interrupted at age 85. Given that life expectancy for females is approximately five years higher than males, interrupting the female life table at age 90 should be a more appropriate point.

A logical next step in the development of this methodology is to apply it to age-specific mortality rates for small areas that have been derived using Bayesian hierarchical modelling, as these represent less crude estimates of mortality rates and allow confidence intervals to be evaluated directly without the need for long and complex formulae.

References

Bailey, K., Walrond, S., 2005. Life expectancy in Hartlepool. Tech. Rep. 18, North East Public Health Observatory.

Brass, W., 1971. On the scale of mortality. In: Biological Aspects of Demography. Taylor and Francis, London, pp. 69–110.

Cai, Y., 2005. National, provincial, prefectural and county life tables for China based on the 2000 Census. CSDE Working Paper 05-03, Centre for Studies in Demography & Ecology, University of Washington.

Chiang, C., 1968. Introduction to Stochastic Processes in Biostatistics. Wiley, New York.

Congdon, P., 2002. A life table approach to small area health need profiling. Statistical Modelling 2 (1), 63–88.
URL `http://smj.sagepub.com/cgi/content/abstract/2/1/63`

D'Agostino, R., Belanger, A., D'Agostino Jr, R., 1990. A suggestion for using powerful and informative tests of normality. American Statistician, 316–321.

Department of Health, 2001. The national health inequalities targets. Tech. rep., Crown Copyright.

Eayres, D., Williams, E. S., 2004. Evaluation of methodologies for small area life expectancy estimation. Journal of Epidemiology and Community Health 58 (3), 243–249.

Gardner, M., Donnan, S., 1977. Life expectancy: variations among regional health authorities. Population Trends 10, 10–12.
URL http://www.popline.org/docs/0427/213544.html

Graham, P., May 2008. Intelligent smoothing using hierarchical bayesian models. Epidemiology 19 (3), 493–495.
URL http://dx.doi.org/10.1097/EDE.0b013e31816b7859

Hogstedt, C., Lundgren, B., Moberg, H., Vingård, E., 2006. Medellivslängd och ohälsotal utmed spårtrafiken i Stockholm. Tech. rep., Statens Folkhälsoinstitut, Stockholm.

Huong, D. L., Minh, H. V., Vos, T., Janlert, U., Van, D. D., Byass, P., 2006. Burden of premature mortality in rural Vietnam from 1999–2003: analyses from a demographic surveillance site. Population Health Metrics 4 (1), 9.

Manuel, D. G., Goel, V., Williams, J. I., 1998. The derivation of life tables for local areas. Chronic Diseases in Canada 19 (2), 52–56.
URL http://view.ncbi.nlm.nih.gov/pubmed/9664025

Mazess, R. B., Forman, S. H., 1979. Longevity and age exaggeration in Vilcabamba, Ecuador. Journal of Gerontology 34 (1), 94–98.
URL http://geronj.oxfordjournals.org/cgi/content/abstract/34/1/94

Mindell, J., Fitzpatrick, J., Seljmani, F., 2004. Health inequalities in london: Life expectancy and mortality 1998–2002. Tech. rep., London Health Observatory.
URL http://eprints.ucl.ac.uk/902/

Myers, R. J., 1940. Errors and bias in the reporting of ages in census data. Transactions of the Actuarial Society of America 41 (Pt 2), 395–415.

National Center for Health Statistics, 1987. Methodology of the national and state life tables. u.s. decennial life tables for 1979-81. Tech. rep., NCHS.

National Center for Health Statistics, 1998. Methodology of the national and state life tables. u.s. decennial life tables for 1989-91. Tech. rep., NCHS.

New Zealand Ministry of Health, 2005. Monitoring health inequality through neighbourhood life expectancy: Public health intelligence occasional bulletin no. 28. Tech. rep., Ministry of Health, Wellington.
URL http://www.moh.govt.nz/moh.nsf/pagesmh/4142?Open

Newell, C., 1988. Chapter 6: Mortality and life tables. In: Methods and Models in Demography. Belhaven, London, p. 217.

Office for National Statistics, 2006. Life expectancy at birth for wards in England and Wales (experimental). http://www.statistics.gov.uk/statbase/Product.asp?vlnk=14466.
URL http://www.statistics.gov.uk/statbase/Product.asp?vlnk=14466

Raleigh, V. S., Kiri, V. A., 1997. Life expectancy in England: variations and trends by gender, health authority, and level of deprivation. Journal of Epidemiology and Community Health 51 (6), 649–658.
URL http://jech.bmj.com/cgi/content/abstract/51/6/649

Retherford, R. D., Mirza, G. M., Jul. 1982. Evidence of age exaggeration in demographic estimates for Pakistan. Population Studies 36 (2), 257–270.
URL http://www.jstor.org/stable/2174200

Rosenwaike, I., Stone, L. F., Nov. 2003. Verification of the ages of supercentenarians in the united states: Results of a matching study. Demography 40 (4), 727–739.
URL http://www.jstor.org/stable/1515205

Scherbov, S., Ediev, D., 2009. Life table techniques for small populations: effects of population size, of the mortality level, and of age composition. In: 26th IUSSP International Population Conference, Marrakech.

Silcocks, P. B. S., 2004. Improving estimation of the variance of expectation of life for small populations. Journal of Epidemiology and Community Health 58 (7), 611–612.
URL http://dx.doi.org/10.1136/jech.2003.011338

Silcocks, P. B. S., Jenner, D. A., Reza, R., 2001. Life expectancy as a summary of mortality in a population: Statistical considerations and suitability for use by health authorities. Journal of Epidemiology and Community Health 55 (1), 38–43.
URL http://dx.doi.org/10.1136/jech.55.1.38

Toson, B., Baker, A., 2003. Life expectancy at birth: Methodological options for small populations. Tech. Rep. 33, Office for National Statistics.

Vallin, J., Caselli, G., 2006. Chapter 11: Cohort life table. In: Demography: Analysis and Synthesis. Vol. 1. Academic Press, pp. 103–129.

Veugelers, P. J., Hornibrook, S., 2002. Small area comparisons of health: applications for policy makers and challenges for researchers. Chronic Diseases in Canada 23 (3), 100–110.
URL http://view.ncbi.nlm.nih.gov/pubmed/12443566

Wheller, L., Baker, A., Griffiths, C., Rooney, C., 2007. Trends in avoidable mortality in england and wales, 1993-2005. Health Statistics Quarterly (34), 6–25.
URL http://www.ncbi.nlm.nih.gov/pubmed/17580644

Williams, E. S., Dinsdale, H., Eayres, D., Tahzib, F., 2004. Impact of nursing home deaths on life expectancy calculations in small areas. Journal of Epidemiology and Community Health 58 (11), 958–962.
URL http://dx.doi.org/10.1136/jech.2003.017608

Williams, E. S., Eayres, D., Dinsdale, H., Tahzib, F., 2005. Calculating life expectancy in small areas. Tech. rep., South East Public Health Observatory, Oxford.

Williams, E. S., Scott, C. M., Scott, S. M., 1995. Using mortality data to describe geographic variations in health status at sub-district level. Public Health 109 (1), 67–73.
URL http://view.ncbi.nlm.nih.gov/pubmed/7871148

Wolf, E., 2005. Creating contiguous cartograms in ArcGIS 9. In: 25th Ann. ESRI Internat. Users Conf.