WP 10
23 April 2010

UNITED NATIONS STATISTICAL COMMISSION and ECONOMIC COMMISSION FOR EUROPE

STATISTICAL OFFICE OF THE EUROPEAN UNION (EUROSTAT)

CONFERENCE OF EUROPEAN STATISTICIANS

Joint Eurostat/UNECE Work Session on Demographic Projections
(28-30 April 2010, Lisbon, Portugal)

Item 4 – Constructing assumptions for mortality: data, methods and analysis

Estimating life expectancy in small population areas[1]

Jorge Miguel Bravo[2], University of Évora, Portugal,
and Joana Malta[3], Statistics Portugal

ABSTRACT

In recent years we have seen an increasing demand for indicators of mortality for smaller (sub-national, sub-regional) areas, either to examine geographic inequalities in mortality, to monitor the effects of Public Health policies, to inform local strategies or to prepare long-term sub-national population projections. The usual way to obtain life expectancy indicators involves the construction of complete or abridged life tables. Attempts to calculate mortality rates directly from small numbers of counts and deaths often results in highly erratic schedules that are very difficult to interpret. In this paper we give an in-depth overview of the method adopted by Statistics Portugal for estimating life expectancy in small population (sub-national) areas (NUTS II and NUTS III). The method uses parametric graduation techniques to smooth crude age-specific mortality rates in order to construct a survival model presented in the form of a life table. We give an overview of parametric and non-parametric graduation methods and revisit the graduation methodology developed by the Continuous Mortality Investigation Bureau (CMIB) and its extension to generalized linear models, recently adopted by Statistics Portugal. The method uses a family of parametric (generalised Gompertz-Makeham) functions estimated by means of generalized linear models in order to graduate crude mortality estimates. We discuss the statistical tests and procedures used to evaluate the goodness-of-fit of the models. The methodology is empirically tested using data for the Portuguese sub-national region of Lisbon and for the period 2006-2008. We conclude that the Gompertz-Makeham functions estimated by means of generalized linear models offer a good alternative for estimating life expectancy in small population areas. The method is flexible and applicable to mortality data for a wide range of ages from any geographical conditions.

1. INTRODUCTION

Life expectancy at birth and at adult ages has long been used as an indicator of the health status and of the level of mortality experienced by a population. It is well known that its main advantage over other methods of measuring mortality is that it does not reflect the effects of the age distribution of an actual population and does not require the adoption of a standard population for comparing levels of mortality among different populations. Life expectancy is a summary measure of mortality at every age that allows us to compare mortality/longevity between geographical areas (and time periods) that may exhibit very diverse population structures. Although there are alternative methods to derive life expectancy, the usual way to obtain it involves the construction of a life table, a process which entails rigorous and exhaustive data requirements and often requires a great deal of both personal and computational time.

In recent years we have seen an increasing demand for indicators of mortality for smaller (sub-national, sub-regional) areas, either to examine geographic inequalities in mortality, to monitor the effects of Public Health policies, to inform local strategies or to prepare long-term sub-national population projections. Given the above data and time requirements, and the particularities of small population data (ages or age bands with zero deaths, reduced population

[1] The analyses, opinions and findings contained in this paper represent the views of the authors and are not necessarily those of Statistics Portugal.
[2] Center for Advanced Studies in Management and Economics of the University of Évora, Largo dos Colegiais, 2, 7000-803 Évora, Portugal, jbravo@uevora.pt.
[3] Statistics Portugal, Department of Demographic and Social Statistics, Av. António José de Almeida, 1000-043 Lisbon, Portugal, joana.malta@ine.pt.

counts, increasing volatility of death rates, oldest-old mortality rates,...) calculating life expectancies for small geographical areas is not often possible or entails more complex methodological challenges.

In theory, there are several methodologies that could be used to estimate life expectancy for small population areas. They include methods based on stable population concepts (see, *e.g.*, Coale and Demeny, 1966), methods based on biological theories of aging (see, *e.g.*, Siler, 1979), methods based on the estimation of population by age (see, *e.g.*, Irwin, 1980), regression equation methods that exploit the relationship between life expectancy and other demographic indices (see, *e.g.*, Mazur, 1969), methods based on the construction of abridged life tables (see, *e.g.*, Chiang, 1984), Brass-type relational methods (Brass, 1971) or solutions that combine traditional complete life table construction techniques with smoothing or graduation methods. The method currently used by Statistics Portugal lies in this later category.

The methodology based on stable population concepts in clearly inappropriate for small areas because it relies on the assumption of the stability of both mortality and fertility rates and on the absence of migration, since we know that migration is probably the most influential source of variation in population change in small areas. Methods based on biological theories of aging have extensive data requirements that cannot be met with the data usually available for small geographical areas. Methods based on the estimation of population by age rely on census data available once every 10 years or on postcensal estimates that suffer from high levels of inaccuracy for certain ages groups and that are available some time after the estimate. The Brass two-parameter logit model system is based on a linear relation between the logit of the survival function in the current population and the logit of the survival function in a standard or reference population. The model provides a trajectory well suited to extrapolation in a space of only two parameters but strongly relies on the stationarity of model parameters. An alternative approach is to use a model life table chosen on the basis of observed mortality, an approach that is free of the stochastic variability due to small numbers but may seriously mis-estimate life expectancies either because the stability assumption is violated or because the shape of the model table significantly departs from the mortality underlying the observed mortality rates.

Attempts to calculate mortality rates directly from small numbers of counts and deaths often results in highly erratic schedules that are very difficult to interpret. Statistical significance may be improved by aggregating data over time or age groups if there is adequate vital registration, but this is not an option when records are poor or lacking altogether.

The study of mortality and methods to estimate life expectancy in small population areas is a key element for producing high quality demographic projections. In this paper, we intend to give an in-depth overview of the method adopted by Statistics Portugal for estimating life expectancy in small population (sub-national) areas (NUTS II and NUTS III). The method uses parametric graduation techniques to smooth crude age-specific mortality rates in order to construct a survival model presented in the form of a life table. This approach can be framed under the most recent developments undertaken in mortality studies conducted in actuarial studies, where the knowledge of the dynamics of age-specific mortality indicators and of life expectancy changes is of the utmost importance.

The paper is organized as follows: in Section 2 we briefly review the main parametric and non-parametric graduation methods suggested in the literature; in Section 3 we give an in-depth review of the graduation methodology developed by the Continuous Mortality Investigation Bureau (CMIB) and its extension to generalized linear models recently adopted by Statistics Portugal. The method uses a family of parametric (generalised Gompertz-Makeham) functions estimated by means of generalized linear models in order to graduate crude mortality estimates. The criteria used to select the model order and to evaluate the goodness-of-fit are succinctly described. We present the model used to extrapolate the probability of death at very old ages and to generate complete life tables for small population areas; in Section 4 we obtain crude estimations of the probability of death in the Lisbon Region for the period of 2006-2008 and apply the CMIB graduation methodology to these estimations, commenting on their advantages and disadvantages, as well as on their suitability for the mortality analysis in question; in Section 5, the most relevant conclusions are presented in light of our main objective: to develop an easily applicable and reliable method to use in areas with small population numbers.

2. AN OVERVIEW OF MORTALITY GRADUATION METHODS

A significant body of literature on the estimation of life expectancy for small populations is dedicated to the construction of life tables, i.e., to the estimation of probabilities of death q_x for ages $x \in \{x_{\min},...,\omega\}$. Beginning with a crude estimation of q_x, $\hat{Q} = \{\dot{q}_x : x_{\min},...,\omega\}$, we wish to produce smoother estimates, $\hat{q}_x$, of the true but unknown mortality probabilities q_x from the set of crude mortality rates, $\dot{q}_x$, for each age x. The crude rate at age x is usually based on the corresponding number of deaths recorded, d_x, relative to initial exposed to risk, E_x. Attempts to calculate crude rates directly from small numbers of counts and deaths often result in highly erratic schedules, which do not correspond to the reasonable hypothesis that the probabilities of death for two consecutive ages should be very close. This has motivated actuaries and demographers to search for methods that smooth these quantities in order to

better express the mortality characteristics of a given population, which we assume to be relatively regular. These procedures are known in actuarial literature as graduation methods.

Graduation is defined in the actuarial literature as the set of principles and methods by which the observed (or crude) probabilities are fitted to provide a smooth basis for making practical inferences and calculations of premiums and reserves (Haberman and Renshaw, 1996). One of the principal applications of graduation, particularly important in this context, is the construction of a survival model, normally presented in the form of a life table.

Classical (non-Bayesian) methods for graduation are usually framed under what is known as parametric and non parametric approach, depending on whether they adjust the data to a mathematical function or simply replace crude estimates by a set of smoothed probabilities. Since De Moivre, in 1725, that the representation of mortality data by means of parametric models attracted the attention of actuaries, demographers and statisticians. The principle is very simple: the probabilities of death (or mortality rates) are expressed as a mathematical function of age and limited set of parameters, estimated on the basis of mortality statistics. More formally, parametric methods are based on the hypothesis that the chosen measurement of mortality, *e.g.*, $q_x = f(x,\theta)$, is a function of age x and a set of p unknown parameters $\theta = \{\theta_1, \theta_2, ..., \theta_p\}$ to be estimated using standard regression techniques. The goal is to obtain the best possible fitting with the minimum number of parameters, i.e., by balancing the number of parameters and the goodness-of-fit.

Let $\hat{\theta}$ denote the estimator of θ. The parametric adjustment of $\hat{Q}$ consists of replacing the crude estimates with the series $\hat{Q}^{fit} = \{\dot{q}_x^{fit} = f(x,\theta), x = x_{\min}, ..., \omega\}$. A number of alternative parametric functions have been proposed in the literature, including the classical Gompertz (1825) and Makeham (1860) models, the (Heligman and Pollard, 1980) mortality laws or the Perks (1932) logistic model[4]. The Gompertz-Makeham function described in Forfar *et al*. (1988) generalizes the original models proposed by Gompertz and Makeham and constitutes the basis of the CMIB methodology described below.

Parametric graduation methods provide a particularly efficient method to smooth mortality data when we have preliminary information about the behaviour of the underlying variables. However, by summing up the behaviour of the underlying phenomena in a set of parameters, parametric methods introduce a new source of risk, namely the risk that the analytical model is inappropriately specified. For this reason, a number of alternative non-parametric methods have been proposed to smooth mortality data. These methods consist of replacing the crude estimate $\hat{Q}$ with a smoothed series, noted $\hat{Q}^{smo} = \{\dot{q}_x^{smo}, x = x_{\min}, ..., \omega\}$. Instead of resuming mortality rates by a mathematical formula of a limited set of parameters, non-parametric graduation describe mortality by considering the original $\omega - x_{\min} + 1$ probabilities, smoothed using a specific procedure. The information provided by the original data is preserved and any influence that does not come from the predictor variable is eliminated.

The representation of mortality data by means of nonparametric models attracted significant attention of actuaries, demographers and statisticians. The basic method was "bin smoothers" and moving weighted averages. More recently, the power of modern computers made more complex and powerful smoothing procedures considerably more accessible and easier to apply in graduation. These include the use of splines (cubic, B-Splines, P-Splines), Whittaker-Henderson method locally-weighted regression (LOESS), kernel smoothing and Generalised Additive Models (GAM)[5].

3. GRADUATION OF SUB-NATIONAL MORTALITY DATA IN PORTUGAL

The method developed by the Continuous Mortality Investigation Bureau (CMIB) was adopted by Statistics Portugal in 2007, and is now the main procedure to calculate graduated mortality rates for sub-national levels (regions NUTS II and NUTS III). This method can be framed under the parametric graduation procedures, and is an extension of the Gompertz and Makeham models. In following, we present the methodology developed by the CMIB and its extension to generalized linear models.

3.1. The CMI Bureau methodology

The Continuous Mortality Investigation (CMI) Bureau of the Institute and Faculty of Actuaries of London was created in 1924, when the continuous collection of mortality data began. It is responsible for constructing standard life tables for use in Great Britain's insurance industry. The methodology that is normally used by the CMIB to produce such tables is described in detail in Forfar *et al.* (1988). This methodology is, in essence, a generalization of the Gompertz

[4] See Benjamin and Pollard (1993) and Pitacco (2004) among other for a detailed review of mortality laws.
[5] For a review of smoothing life tables see, *e.g.*, Wang (2005).

(1825) and Makeham (1860) models. It was applied to Portuguese insurance data by Bravo (2007) and adopted by Statistics Portugal in 2007.

Consider a group of consecutive ages x and the series of independent deaths $\{d_x\}$ and corresponding exposure to risk $\{E_x\}$. The graduation procedures uses a family of parametric functions know as Gompertz-Makeham of the type (r,s). They are functions with $r+s$ parameters of the form

$$GM_{\theta}^{r,s}(x) = \sum_{i=0}^{r-1} \alpha_i x^i + \exp\left(\sum_{j=0}^{s-1} \beta_j x^j \right) \quad (1)$$

where parameters r and s give the order of the model, and can only assume positive integer values, excluding the possibility of both being zero; θ is a vector containing the estimated mortality curve parameters $\theta = (\alpha_1, ..., \alpha_{r-1}; \beta_1, ..., \beta_{s-1})$. The model presented in (1) contains both a polynomial and an exponential component. We assume that when $r = 0$, (1) includes only its exponential term, whereas if $s = 0$, equation (1) is reduced to its polynomial component. In some applications it is useful to establish the following Logit Gompertz-Makeham functions of the type (r,s), defined as

$$LGM_{\theta}^{r,s}(x) = \frac{GM_{\theta}^{r,s}(x)}{1 + GM_{\theta}^{r,s}(x)} \quad (2)$$

Given a vector of parameters θ, the methodology developed by CMIB states that the expression

$$q_x = LGM_{\theta}^{r,s}(x) \quad (3)$$

results in an adequate adjustment. Note that, since $q_x \in [0,1]$, the adequate function to consider for graduation of gross death probabilities is the Logit-Gompertz-Makeham (LGM), as it guarantees values within that interval.

This method recommends the use of an orthogonal base for the polynomial component of the $GM_{\theta}^{r,s}(x)$ equation. Special emphasis is given to the Chebycheff and Legendre polynomials (Bravo, 2007). The use of orthogonal polynomials requires some form of age scaling so that range of ages lies, *e.g.*, in the $[-1,1]$ interval. Therefore, the transformations suggested are as follows:

$$x \to x' = \frac{x-u}{v}, \text{ with } u = \frac{x^{\max} + x^{\min}}{2} \text{ and } v = \frac{x^{\max} - x^{\min}}{2},$$

from which we have $x \in \left[x^{\max}, x^{\min} \right] \to x' \in [-1,1]$.

Gathering all referred aspects, the equation used for estimation is given by

$$\underset{u,v}{\overset{\{p\}}{}} GM_{\theta}^{r,s}(x) = \sum_{i=0}^{r-1} \alpha_i p_i \left(\frac{x-u}{v} \right) + \exp\left(\sum_{j=0}^{s-1} \beta_j p_j \left(\frac{x-u}{v} \right) \right) \quad (4)$$

where $\{ p_{i,j}(x) : i, j = 0,1,2,... \}$ defines the adopted polynomial base, and all other parameters have the same meaning as before.

To estimate the model parameters, two optimization criteria are considered: maximum likelihood and the minimum Chi-Squared. Empirical studies show that, in practice, the two criteria produce similar graduations (see, *e.g.*, Forfar *et al.*, 1988, and Bravo, 2007). The maximum likelihood method is based on the maximization of the appropriate (Binomial, Poisson, etc.) likelihood function. The minimum Chi-Squared criteria correspond to the usual χ^2 statistic, namely the sum of squared standardized residuals.

The methodology described above can be reformulated and extended in the framework of generalized linear and non-linear models. Generalized linear models (GLM) are an extension of linear models for non normal distributions and non

linear transformations of the interest variable. Their use in the graduation of either the probability of death at age x, q_x, or the force of mortality at age x, μ_x, is justified because both response variables are not normal. As an alternative to the classical regression linear models, the GLM allows, using a link function, estimation of a mean function of the response variable, written as a linear combination of all independent variables. The use of GLM in the context of graduation is explored in Renshaw (1991), Renshaw and Hatzopoulos (1996), Haberman and Renshaw (1996), Renshaw *et al.* (1997), Verrall (1996), Delwarde *et al.* (2004) and Bravo (2007). We now give some details about modelling and probability distribution assumptions for the death probability within the GLM approach.

Assume that E_x individuals initially exposed to the risk of death come under observation at age x and continue under observation until they survive to $x+1$ or, otherwise, die before that. In addition, suppose that the probability of death during the year for each one of them is q_x, and that the death or survival of one is independent of the death or survival of the others. Let D_x denote the random variable which represents the number of deaths observed in the year. In this case, we can write $D_x \square Bin(E_x, q_x)$. The general expression of the maximum likelihood function is given by

$$\begin{aligned} L(q) &= \prod_{x=x_{\min}}^{x_{\max}} q_x^{d_x} (1-q_x)^{E_x-d_x} \\ &= \exp\left\{ \sum_{x=x_{\min}}^{x_{\max}} [d_x \log(q_x) + (E_x - d_x)\log(1-q_x)] \right\} \\ &= \exp\left\{ \sum_{x=x_{\min}}^{x_{\max}} \left[d_x \log\left(\frac{q_x}{1-q_x} \right) + E_x \log(1-q_x) \right] \right\}. \end{aligned} \tag{5}$$

The graduation of q_x is performed using the function

$$q_x = LGM_\theta(r,s) = \frac{GM_\theta(r,s)}{1+GM_\theta(r,s)} \tag{6}$$

using the *logit* transformation as the canonical link of the binomial family, i.e.,

$$\eta_x = \log\left(\frac{q_x}{1-q_x} \right) \tag{7}$$

with inverted function

$$q_x = \frac{\exp(\eta_x)}{1+\exp(\eta_x)} \tag{8}$$

To evaluate the goodness-of-fit of the model, an appropriate measure is the so-called *scaled deviance*, defined as

$$D(y;\hat{m}) = 2\sum_{x=x_{\min}}^{x_{\max}} \left[d_x \log\left(\frac{d_x}{\hat{m}_x} \right) + \left(E_x - d_x \right) \log\left(\frac{E_x - d_x}{E_x - \hat{m}_x} \right) \right] \tag{9}$$

where $\hat{m}_x = g^{-1}(\hat{\eta}_x)$ represents the adjusted values generated by the estimated model.

Leaving aside the constant values and taking the logarithm from the maximum likelihood function, the $\theta = (\alpha_s, \beta_r)$ parameters of the Gompertz-Makeham function are estimated via maximum likelihood or, in other words, by resolving the following minimization problem

$$\hat{\theta} = \arg\min_{\theta} \left\{ \sum_{x=x_{\min}}^{x_{\max}} E_x \log\left(1+GM_\theta^{r,s}(x)\right) - d_x \log\left(GM_\theta^{r,s}(x)\right) \right\} \tag{10}$$

3.2. Assessing model fit

The goodness-of-fit is evaluated by means of the classical non-parametric tests. For completeness, we briefly review the use of these tests in a graduation context[6]. The first indicator of model performance involves comparing the absolute and relative deviations generated by the model. Let Dev_x denote the absolute deviance for each age x, i.e.,

$$Dev_x = d_x - d_x^{\exp} \tag{11}$$

where $d_x^{\exp}$ represents the expected number of deaths, estimated by multiplying the mortality rates by the number of individuals initially exposed to risk. The corresponding relative deviances (or Pearson residuals), z_x, is given by

$$z_x = \frac{Dev_x}{\sqrt{Var(d_x)}} \tag{12}$$

where $Var(d_x)$ stands for the variance of deaths, obtained under the distributional hypothesis established for the probability of the random variable. Therefore, the goodness-of-fit is assessed by whether or not deviances are randomly distributed when taken sequentially, and if its distribution is in accordance to the established hypothesis adopted for the model under analyses.

Other useful tests to evaluate the model adjustment in a graduation environment are the Signs Test and the Runs tests. The Signs Test considers only the deviation sign (positive or negative). Considering the binomial distribution, it is possible to assess the probability for the number of positive signs not to exceed its observed value. If this probability is too low, the number of positive deviations is unexpectedly low; if it is too high, we have an excessively high number of positive deviations. In both cases, it should be concluded that the graduated probabilities of death are too far from the observed values, resulting in a poor fit. The Runs Test takes into account the number of same sign consecutive deviations. The test assumes that deaths follow a normal distribution by specific ages, and that the deviations are independent and their signs randomly distributed. It is possible to obtain the exact probability for the number of consecutive same sign deviations to be the same or less than the observed equivalent. A small probability is associated with an excessive linearity of the estimated graduation model, resulting in a poor adjustment.

The Kolmogorov-Smirnov test can be useful to detect abnormal graduation models, *e.g.* in cases where the deviations are great or very small (Forfal *et al.*, 1988). The KS test is usually used to test or compare the distribution of the interest variable with a known distribution. In this case, the KS test is used to compare the expected and observed deaths accumulated distributions between the maximum and the minimum ages. It should be noted that, since the two series are not independent, the test has to be approached with caution. The Chi-Square test is a widely used non parametric test to assess model fit. This test can be used to test independency between observations or assess goodness-of-fit. A model can be considered to have a better adjustment if the value of the Chi-Square statistic is higher.

Finally, other tests used to asses model fit are the auto-correlation tests, based on the analyses of the correlations between the relative deviations in each age. Every relative deviation should be normally distributed, with zero mean and unitary variance. There are two auto-correlation tests usually used: the Portmanteau Test and the Ljung-Box Test. Both tests, using different statistics, test the hypothesis that all auto-correlation coefficients for different levels of deviations equal zero. Lastly, another way to evaluate the adjustment of the graduation model is simply to observe the graphical representation of both the crude and the graduated mortality curves. In practical terms, this is a very good way to immediately assess the overall model fit to the original series of data considered.

3.3. Projecting probabilities of death at older ages

The calculation of crude age specific mortality rates at advanced ages suffers from several problems. The main issue concerns the quality and the availability of data on population estimates for the oldest-old. Effectively, although data on deaths are in general of good quality, mortality rates may be contaminated by random fluctuations due to either the small number of those surviving up to very old ages, to age misreporting problems or to the lack of coherence between deaths and the number of those exposed to risk.

In order to construct complete life tables for sub-national levels, it was decided to remove fluctuations by smoothing crude estimates via a projection method. Various methodologies have been proposed for estimating mortality rates at oldest ages[7]. From these, we adopted the method proposed by Denuit and Goderniaux (2005), a method that is applied directly to crude death probabilities and establishes a limiting age for the life table. Formally, the following log-quadratic model is fitted by weighted-least squares:

[6] For a detailed analysis see, *e.g.*, Benjamin and Pollard (1993) and Forfar *et al.* (1988).
[7] For a review see, *e.g.*, Buettner (2002) and Pitacco (2004).

$$\ln \hat{q}_x = a + bx + cx^2 + \varepsilon_x, \varepsilon_x \sim N\left(0, \sigma^2\right) \quad (13)$$

to age-specific death probabilities observed at advanced ages. Two restrictions are imposed to equation (13), as to assure a concave configuration to the mortality curve at older ages and restrict a horizontal tangent at the maximum age point considered. The imposed restrictions are:

$$q_{x_{\max}} = 1 \quad (14)$$

$$q'_{x_{\max}} = 0 \quad (15)$$

The inclusion of (14) and (15) into (13) will lead to a new expression of the model equation, given by

$$\ln \hat{q}_x = \left(x_{\max}^2 - 2x\left(x_{\max}\right) + x^2\right)c + \varepsilon_x, \varepsilon_x \sim N\left(0, \sigma^2\right) \quad (16)$$

One of the critical aspects related to this model is determining the adequate age from which to replace the gross mortality probabilities by its correspondent adjusted estimates. The ad hoc method suggested in Denuit and Goderniaux (2005) and Bravo (2007), among others, recommends choosing the age so that the regression coefficient R^2 is maximized, by ranging the initial age of the calibration procedure in the $[50,85]$ interval. Attention is also drawn to the possible need for smoothing the mortality series around the cut age point, so as to avoid abrupt discontinuities between the two series. The suggestion is to apply a geometrical mean to the death probabilities around ages $x = x_0 - 5, ..., x_0 + 5$.

4. APPLICATION TO MORTALITY DATA OF THE LISBON REGION

In this section we describe the results, published by Statistics Portugal, of implementing the approaches described above in constructing life tables and estimating life expectancy for the Portuguese sub-national NUTS II region of Lisbon, with combined sexes. We use aggregate population and death figures corresponding to the three-year period of 2006-2008. These two data sets were published by Statistics Portugal (INE) and are classified by age (ranging from 0 to 100 or older) and sex. They are coherent in that both refer to the Lisbon Region as the place of residence. As the population census takes place every 10 years and during the first year of the ten-year period, the population data for the three-year period of 2006-2008 is based on inter-census estimations calculated by INE. Dating back to December 31, 2006, the estimated population resident in the Lisbon region was of 2794226 individuals.

The first step is to calculate the crude estimates of q_x from the data. The revised methodology developed by Statistics Portugal (INE, 2007) considers life tables based on age specific probabilities of death q_x calculated directly using three-year periods, by pooling deaths and exposures first and then dividing the former by the latter. In the period under study, there were nearly 5.63 million men and women exposed to risk. In the same period, nearly 50.2 thousand men and women died, with the great majority, approximately 91.3%, doing so after the age of 50.

The modelling of q_x has been done through the functions $LGM_\theta(r,s)$, $r \in [0,4]$ and $s \in [2,7]$, using generalized linear models of the binomial family. The first step in applying this graduation procedure is to determine the order (r,s) of the Gompertz-Makeham model that best fits the data. To do so, a total of 30 different combinations are tested by ranging the order (r,s) of the model in the above intervals. The order of $LGM_\theta(r,s)$ is finally determined by selecting the model with highest rank in all (or the majority) of the goodness-of-fit indicators.

Table 1 summarises the results obtained for the log-likelihood function and for the deviance. As expected, as we go along a line or column in the log-likelihood or the deviance table, the values decrease. Our goal is to reach a balance between the number of parameters and the goodness-of-fit of the model. Graduation studies sometimes emphasize the goodness-of-fit without considering the statistical stability of the parameters involved in the regression. The result is an over-parameterization of the model. Recall that the null hypothesis that the difference between the deviance of two adjacent rows/columns can be tested against a chi-squared distribution with one degree of freedom.

Lisboa 2006-2008, HM						
Log-Liklihood						
r	s = 2	s = 3	s = 4	s = 5	s = 6	s = 7
0	217947.2	216716.8	216523.6	216521.1	216512.1	216489.1
1	217283.4	216689.8	216523.1	216505.9	216504.6	216488.4
2	216767.0	216500.2	216522.0	216505.1	216481.2	216449.7
3	216505.6	216500.1	216498.9	216489.4	216451.4	216441.5
4	216502.2	216498.8	216473.3	216461.0	216450.6	216442.4
Deviance						
r	s = 2	s = 3	s = 4	s = 5	s = 6	s = 7
0	3464.65	1003.89	617.41	612.47	594.35	548.33
1	2136.99	949.78	616.46	582.05	579.41	547.08
2	1104.23	570.63	614.16	508.38	432.61	469.58
3	581.46	570.44	568.02	548.92	473.01	453.20
4	574.62	567.75	516.78	492.13	471.48	455.06

Table 1: Log-likelihood and (unscaled) deviance, Lisbon 2006-2008, sexes combined

Simultaneously, in selecting the optimal model, all the other goodness-of-fit measures should be taken into account. Table 2 presents all relevant overall goodness-of-fit indicators previously discussed. The first column refers to the order of the model. The next two columns show the value of the unscaled χ^2 statistic and the correspondent p-value. Columns four to eight present all p-values for the tests used as a complement to assess model fit and referred to in the previous section (Signs Test, Runs Test, KS Test, Portmanteau Test and Ljung-Box Test). The column denominated as 'Par. Sign.' stand for the significance (Y) or not (N) of the entire model estimated parameters. Obtaining non significant estimates should be interpreted as a first sign of over-parameterization of the proposed model.

In the column 'Conf.', the letter 'Y' ('N') stands for a good (bad) graphical configuration of the adjusted mortality curve, assessed as the ability of the model to generate a plausible mortality curve, particularly the ability of the model to generate increasing probabilities of death at older ages. Finally, 'Int. Conf.' gives information on the goodness-of-fit in the ages where exposure to risk is higher, assessed as the ability of the model to produce small confidence intervals that contain most of the gross probabilities of death.

(r,s)	χ2	p(χ2)	Signs p(+)	Runs p(runs)	KS p(.)	Portm. p(.)	LJ-Box p(.)	Par. Sign	Conf.	Int. Conf.
(0,2)	4075.259	0.112	0.999	0.000	1	0.000	0.000	Y	N	N
(0,3)	980.508	0.546	0.136	0.000	1	0.000	0.000	Y	N	N
(0,4)	606.224	0.531	0.309	0.002	1	0.096	0.040	Y	N	Y
(0,5)	609.297	0.495	0.184	0.001	1	0.182	0.094	Y	N	Y
(0,6)	582.128	0.537	0.242	0.146	1	0.166	0.082	N	N	Y
(0,7)	534.917	0.547	0.460	0.132	1	0.157	0.069	Y	Y	Y
(1,2)	2142.277	0.474	0.382	0.001	1	0.000	0.000	Y	N	N
(1,3)	943.541	0.499	0.618	0.001	1	0.000	0.000	Y	Y	N
(1,4)	602.602	0.543	0.309	0.001	1	0.078	0.031	N	N	Y
(1,5)	568.938	0.543	0.382	0.058	1	0.161	0.077	Y	Y	Y
(1,6)	564.738	0.550	0.242	0.295	1	0.137	0.061	N	Y	Y
(1,7)	534.125	0.545	0.382	0.058	1	0.151	0.065	N	Y	Y
(2,2)	1140.533	0.392	0.097	0.001	1	0.000	0.000	Y	N	N
(2,3)	555.101	0.556	0.242	0.295	1	0.158	0.072	Y	Y	Y
(2,4)	598.726	0.550	0.309	0.001	1	0.051	0.018	N	N	Y
(2,5)	566.304	0.547	0.382	0.058	1	0.134	0.059	N	Y	Y
(2,6)	520.866	0.540	0.618	0.007	1	0.797	0.701	Y	N	Y
(2,7)	454.656	0.566	0.460	0.088	1	0.629	0.507	Y	N	Y
(3,2)	571.059	0.530	0.618	0.484	1	0.192	0.097	Y	Y	Y
(3,3)	554.898	0.555	0.309	0.140	1	0.160	0.074	N	Y	Y
(3,4)	553.648	0.550	0.136	0.329	1	0.200	0.100	N	Y	Y
(3,5)	533.329	0.558	0.460	0.021	1	.0273	0.147	N	Y	Y
(3,6)	456.045	0.577	0.816	0.659	1	0.586	0.468	Y	Y	Y
(3,7)	440.299	0.557	0.540	0.366	1	0.620	0.493	Y	N	Y
(4,2)	563.553	0.533	0.460	0.270	1	0.216	0.113	Y	N	Y
(4,3)	553.782	0.548	0.309	0.283	1	0.205	0.104	N	Y	Y
(4,4)	503.466	0.550	0.618	0.057	1	0.661	0.538	Y	N	Y
(4,5)	477.450	0.561	0.691	0.090	1	0.620	0.495	Y	N	Y
(4,6)	456.482	0.566	0.460	0.917	1	0.650	0.529	N	Y	Y
(4,7)	440.099	0.568	0.758	0.790	1	0.609	0.485	N	Y	Y

Table 2: LGM(r,s) - Goodness-of-fit measures, Lisbon, 2006-2008, sexes combined

Considering the last three indicators, the only models where all estimates are significant and the curve configuration is acceptable both at younger and older ages are models $LGM(0,7)$, $LGM(1,5)$, $LGM(3,2)$ and $LGM(3,6)$. From these candidate solutions, if we take into account all other goodness-of-fit measures, the best model is $LGM(3,6)$. The coefficients of the optimal model are shown in Table 3 and are all statistically significant.

As a complement for the information discussed previously, it can be helpful to observe some selected graphical representations not only of the overlap between the gross and estimated mortality curve, but also of the residuals, so as to better understand their distribution. Figures 1 to 3 help visualising the fit obtained in the selected model.

	Coef.	**se**	***t*-ratio**	***p*-value**
α0	0,003332	0,00009	35,743	< 0.0001
α1	0,009357	0,00031	30,308	< 0.0001
α2	0,006380	0,00027	23,997	< 0.0001
β0	-7,895357	0,15654	-50,436	< 0.0001
β1	5,667404	0,27884	20,325	< 0.0001
β2	10,428748	0,71621	14,561	< 0.0001
β3	-8,786856	0,94184	-9,329	< 0.0001
β4	-9,507530	0,87421	-10,876	< 0.0001
β5	10,509087	0,98526	10,666	< 0.0001

Table 3: Coefficients of model LGM(3,6), Lisbon, 2006-2008, sexes combined

In figure 1, the gross and graduated probabilities of deaths are represented, as well as the 95% confidence intervals. In figure 2, the Pearson residuals by age for the optimum model are represented, as well as a comparison between the quantiles of a standardized normal distribution and the obtained residual distribution.

Some observations can be seen to fall into the tails of the standardized normal distribution, mainly at older ages. Apart from this fact, the overall model seems to result in an acceptable fit, having most of observations within the estimated confidence interval.

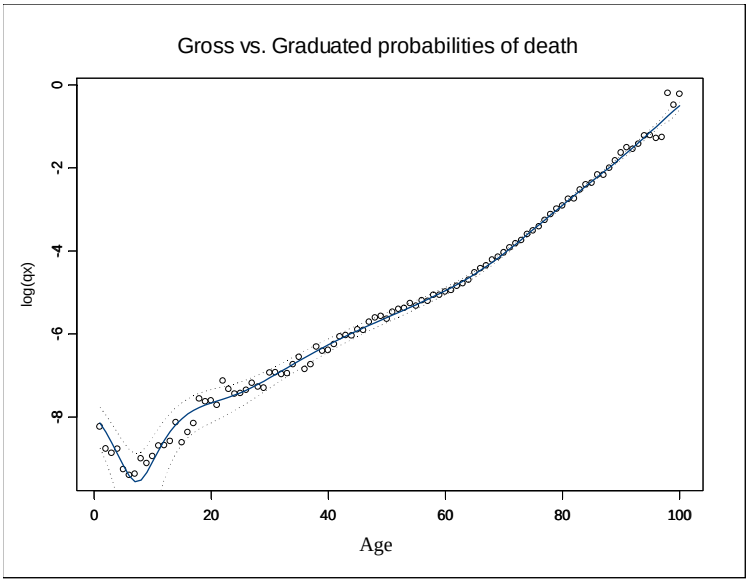

Figure 1 – Adjusted mortality curve, and confidence intervals, Lisbon, 2006-2008, sexes combined

It is important to remember that population estimates at older ages are less reliable, meaning also that graduated probabilities of death suffer a considerable bias. As explained earlier, applying the Denuit-Goderniaux method for estimating mortality probabilities at older ages will result in a more adequate curve shape, and consequently reflect the known mortality laws in a more adequate way. Note that Statistics Portugal established as the life table limit the age of 115, meaning that at this age, the death probability is restricted to 1.

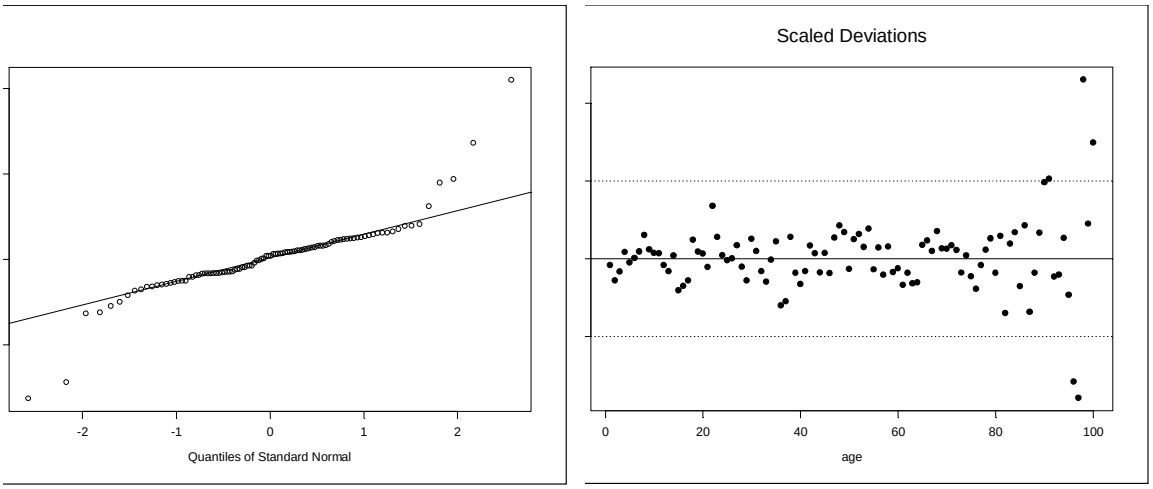

Figure 2 – Residuals from LGM(3,6) model, Lisbon, 2006-2008, MF

The DG models tested led us to apply an estimation from age 75, and to consider age 81 as the point of junction of the two estimated series, with no need for smoothing procedures. Figure 3 represents the profile of the three mortality curves.

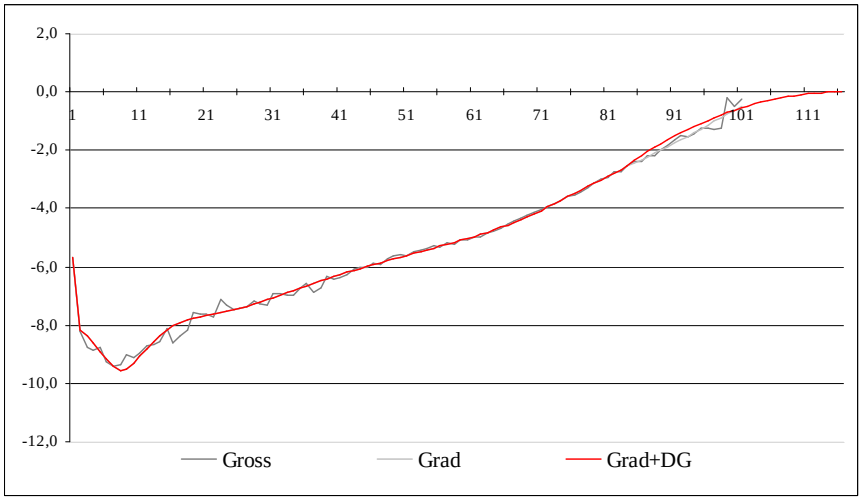

Figure 3 – Comparison between crude and fitted death probabilities

5. CONCLUSIONS

In this paper we discussed the use of graduation methods in smoothing mortality data as a feasible solution for estimating life expectancy indicators for small population areas. In particular, we give a comprehensive presentation of the parametric method developed by the Continuous Mortality Investigation Bureau (CMIB) and its extension to generalized linear models, recently adopted by Statistics Portugal. The methodology is empirically tested using data for the Portuguese sub-national region of Lisbon and for the period of 2006-2008. Our results show that the methodology is robust and can be used to construct life tables and estimate life expectancy. The flexibility of the method makes it applicable to mortality data for a wide range of ages from any geographical conditions.

6. REFERENCES

Benjamin, B. and Pollard, J. (1992). The Analysis of Mortality and Other Actuarial Statistics, 6th edition. London: Butterworth-Heinemann.

Benjamin, B. and Pollard, J. (1993). The Analysis of Mortality and other Actuarial Statistics. Third Edition. The Institute of Actuaries and the Faculty of Actuaries, U.K.

Brass, W (1971). On the Scale of Mortality. Brass. W. Biological Aspects of Demography. London: Taylor & Francis.

Bravo, J. M. (2007). Tábuas de Mortalidade Contemporâneas e Prospectivas: Modelos Estocásticos, Aplicações Actuariais e Cobertura do Risco de Longevidade. Tese de Doutoramento, Universidade de Évora.

Buettner, T. (2002). Approaches and experiences in projecting mortality patterns for the oldest-old. North American Actuarial Journal, 6(3), 14-25.

Chiang C, (1984). The life table and its application. Malabar, Fl., Robert E. Krieger Publishing Co.

Chiang, C. (1979). Life table and mortality analysis. World Health Organization, Geneva.

Coale, A. and Demeny, P. (1966). Regional Model Life Tables and Stable Populations (Princeton, N. J., Princeton University Press.

Delwarde, A., Kachakhidze, D., Olié, L., & Denuit, M. (2004). Modèles linéaires et additifs généralisés, maximum de vraisemblance local et méthodes relationnelles en assurance sur la vie. Bulletin Français d'Actuariat 6, 77-102.

Denuit, M. and Goderniaux, A. (2005). Closing and projecting life tables using log-linear models. Bulletin of the Swiss Association of Actuaries, 29-49.

Forfar, D., McCutcheon, J. and Wilkie, D. (1988). On Graduation by Mathematical Formula. Journal of the Institute of Actuaries 115, 1-135.

Gompertz, B. (1825). On the nature of the function of the law of human mortality and on a new mode of determining the value of life contingencies. Philosophical Transactions of The Royal Society, 115, 513-585.

Haberman, S. and Renshaw, A. (1996). Generalized linear models and actuarial science. The Statistician, 45, 407-436.

Heligman, L. and Pollard, J. (1980). The age pattern of mortality. Journal of the Institute of Actuaries, 107, 49-80.

Irwin, R. (1980). Methodology for Estimates of the Population of Counties by Age and Sex: July 1, J975. Current Population Reports N. 103, U.S Bureau of Census, Washington.

Makeham, W. (1860). On the law of mortality. Journal of the Institute of Actuaries, 13, 325-358.

Mazur, D. (1969). Expectancy of Life at Birth in 36 Nationalities of the Soviet. Union: 1958-60. Population Studies, 23, 225-46.

Perks W. (1932). On some experiments in the graduation of mortality statistics. Journal of the Institute of Actuaries, 63, 12-57.

Pitacco, E. (2004). Survival models in a dynamic context: A survey. Insurance: Mathematics & Economics 35, 279-298.

Renshaw, A. (1991). Actuarial graduation practice and generalised linear models. Journal of the Institute of Actuaries, 118, 295-312.

Renshaw, A. and Hatzopoulos, P. (1996). On the graduation of amounts. British Actuarial Journal, 2, 185-205.

Renshaw, A., Haberman, S., and Hatzopoulos, P. (1997). On the duality of assumptions underpinning the construction of life tables. Astin Bulletin, 27, 5-22.

Siler, W. (1979). A competing risk model for animal mortality, Ecology, 60, 750-757.

Verrall, R. (1996). A unified framework for graduation. Actuarial Research Paper, 91, 2-25.

Wang, J. (2005). Encyclopedia of Biostatistics, Chap. Smoothing Hazard Rates. Willey, 2nd edition.