Stochasticity, heterogeneity, and variance in longevity in human populations

Nienke Hartemink [a,*], Trifon I. Missov [b,c], Hal Caswell [a,d]

[a] University of Amsterdam, Netherlands
[b] Max Planck Institute for Demographic Research, Germany
[c] University of Rostock, Germany
[d] Woods Hole Oceanographic Institution, United States

ARTICLE INFO

ABSTRACT

Inter-individual variance in longevity (or any other demographic outcome) may arise from heterogeneity or from individual stochasticity. Heterogeneity refers to differences among individuals in the demographic rates experienced at a given age or stage. Stochasticity refers to variation due to the random outcome of demographic rates applied to individuals with the same properties. The variance due to individual stochasticity can be calculated from a Markov chain description of the life cycle. The variance due to heterogeneity can be calculated from a multistate model that incorporates the heterogeneity. We show how to use this approach to decompose the variance in longevity into contributions from stochasticity and heterogeneous frailty for male and female cohorts from Sweden (1751–1899), France (1816–1903), and Italy (1872–1899), and also for a selection of period data for the same countries.

Heterogeneity in mortality is described by the gamma-Gompertz–Makeham model, in which a gamma distributed "frailty" modifies a baseline Gompertz–Makeham mortality schedule. Model parameters were estimated by maximum likelihood for a range of starting ages. The estimates were used to construct an age×frailty-classified matrix model, from which we compute the variance of longevity and its components due to heterogeneous frailty and to individual stochasticity. The estimated fraction of the variance in longevity due to heterogeneous frailty (averaged over time) is less than 10% for all countries and for both sexes. These results suggest that most of the variance in human longevity arises from stochasticity, rather than from heterogeneous frailty.

1. Introduction

Individual variance, especially in fitness components, plays a key role in demography, ecology, and evolutionary biology. From an evolutionary perspective, variance in fitness components is potential material on which natural selection can operate. From a demographic perspective, identifiable differences among individuals are the basis for structured population models (Metz and Diekmann, 1986; Tuljapurkar and Caswell, 1997; Caswell, 2001); differences due to age lead to age-structured models, differences due to size lead to size-structured models, etc.

Longevity (age at death) is a fitness component that varies widely among individuals. This variance arises as a result of two different underlying causes: individual stochasticity and heterogeneity. Individual stochasticity is variance due to random outcomes of probabilistic demographic processes (living or dying, reproducing or not, making or not making a life cycle transition). Even in a completely homogeneous population, in which every individual experienced exactly the same (age-specific) mortality rates, variance due to individual stochasticity would exist (Caswell, 2009). Any calculation of the variance in longevity from an ordinary life table implicitly assumes that every individual is subject to the (age-specific) mortality rates in that life table, and hence that the variance is only due to individual stochasticity.

Variance in longevity can also result from unobserved, or latent, heterogeneity in the properties of individuals. For example, individuals of the same age may differ in their mortality rates due to genetic, environmental, or maternal effects. Such differences are often referred to as heterogeneity in individual frailty (Vaupel et al., 1979). Because more frail individuals are more at risk than others, heterogeneity in frailty leads to changes in cohort composition

with age, due to within-cohort selection. As a cohort ages, the representation of less frail individuals increases, and the average mortality rate in an old cohort will be lower than one would expect based on extrapolation of mortality rates at younger ages. This selection effect has been suggested as an explanation for the mortality plateaus often observed at very old ages (Horiuchi and Wilmoth, 1998; Vaupel et al., 1979; Vaupel, 1985).

The effects of unobserved heterogeneity in survival analysis can be estimated using frailty models (Vaupel et al., 1979; Wienke, 2010). In frailty models, a baseline mortality schedule is modified by a term representing individual frailty. A widely used example is the gamma-Gompertz model, which assumes an exponentially increasing age-specific baseline mortality rate (the Gompertz model), and that frailty acts as a proportional hazard multiplier of the baseline mortality (Vaupel et al., 2014). Frailty, which is fixed over the life of the individual, follows a gamma distribution, the variance of which measures the amount of unobserved heterogeneity.

The variance in longevity in a frailty model is a result of both stochasticity and heterogeneity. Little is known about the relative contribution of each to the total variance in longevity, and how those contributions may depend on species, sex, environmental conditions, etc. Caswell (2014) presented an ad hoc approach to this problem, using an age × frailty-classified matrix model. The variance in longevity was computed from the model and the relative contributions of heterogeneity and stochasticity estimated by reducing the initial variance in frailty to zero and attributing the remaining longevity variance to stochasticity. In an analysis of gamma-Gompertz parameters for a single year Swedish females (obtained from Missov, 2013), the fraction of variance due to heterogeneity was estimated to be only 0.071. Applying the same approach to a gamma-Gompertz–Makeham model for women from Turin (Zarulli et al., 2013) resulted in an even lower estimate of 0.012.

Here, we present a more rigorous variance decomposition, which does not require a hypothetical reduction of frailty variance to zero. We apply it to large cohort mortality data sets for three different countries: Sweden, France and Italy, over a long time period. This will enable us to see whether any patterns in variance can be generalized across countries, time periods, or sexes.

The paper is organized as follows. Section 2 describes the gamma-Gompertz–Makeham mortality model. Section 3 presents the construction of the age × frailty matrix model, and Section 3.2 provides the methods used to calculate longevity statistics and decompose the variance. Section 4 gives details about the mortality data and estimation of the gamma-Gompertz–Makeham parameters. Section 5 presents the implementation of the age × frailty matrix model. Section 6 presents the results for the cohort and period data, and Section 7 discusses the interpretation of the results.

2. Frailty in the gamma-Gompertz–Makeham model

Unobserved heterogeneity in mortality risk, or frailty, can be included in mortality models by assuming that this frailty acts to modify a baseline mortality rate shared by all individuals. The gamma-Gompertz–Makeham mortality model has been shown to give a good fit to human mortality data (Manton et al., 1986; Yashin et al., 1994). It is an extension of the Gompertz model (Gompertz, 1825), in which mortality at adult and older ages is an exponentially increasing function of age. The Gompertz mortality function has a baseline mortality parameter a and a parameter b that determines the steepness of the exponential increase with age. The Makeham model an age-independent component c to the mortality (Makeham, 1860). The Makeham term has been shown to be essential to prevent distorted parameter estimates (Missov

and Németh, 2016). In the Gompertz–Makeham model, the hazard at age x equals:

$$\mu(x) = ae^{bx} + c. \tag{1}$$

In the gamma-Gompertz–Makeham model (hereafter called ΓGM), the frailty of an individual is included as a (gamma-distributed) random effect that is fixed over the lifetime. Dynamic frailty, that can change with age or with health-related events, has been included in other models (e.g., Vaupel and Yashin, 2006; Le Bras, 1976; Gavrilov and Gavrilova, 1991; Yashin et al., 1994, 2000). The matrix analysis we develop here also applies to such dynamic frailty models (Caswell, 2014); see Section 3.1. Frailty in the ΓGM affects the (age-dependent part of) mortality as a proportional hazard; the hazard $\mu(x, z)$ for an individual with frailty z at age x is

$$\mu(x, z) = zae^{bx} + c. \tag{2}$$

The initial frailty distribution in the cohort is gamma-distributed, $Z \sim \Gamma(\kappa, \lambda)$, with shape parameter κ and scale parameter λ. The mean and variance of this distribution are $E(Z) = \kappa/\lambda$ and $V(Z) = \kappa/\lambda^2$. The mean is set equal to 1, so that the cohort starts life with an average frailty of 1. When this is the case, i.e. $E(Z) = 1$, $\lambda = \kappa$ and the variance $V(Z) = 1/\lambda := \gamma$. The marginal hazard function, which gives the unconditional population hazard Manton et al. (1981) and Missov and Vaupel (2015), is a sigmoid function

$$\mu(x) = \frac{ae^{bx}}{1 + \frac{a\gamma}{b}(e^{bx} - 1)} + c. \tag{3}$$

Heterogeneity is described by the variance γ of frailty at the starting age of analysis; the higher the variance, the greater the heterogeneity between individuals. In a completely homogeneous population, $V(Z) = 0$ and every individual experiences the same age-dependent hazard.

Using (3) and applying maximum likelihood yields estimates for the baseline mortality parameters a, b, c and for γ, the parameter that describes the heterogeneity in frailty. This optimization is described in more detail in Section 4.2.

3. An age × frailty matrix model

We incorporated the ΓGM mortality function into an age × frailty-classified matrix model (for a more general description of age–stage classified matrix models see Caswell (2009, 2012)). Age is described by a set of ω discrete age classes and frailty by a set of g frailty classes that discretize the gamma distribution of frailty. Vector $\boldsymbol{\mu}_0$ of dimension ω contains the baseline age-specific part of the mortality rates:

$$\boldsymbol{\mu}_0 = \begin{pmatrix} ae^{0b} \\ \vdots \\ ae^{(\omega-1)b} \end{pmatrix}. \tag{4}$$

If z_i is the frailty for the ith group, then the mortality vector for frailty group i is

$$\boldsymbol{\mu}_i = z_i\boldsymbol{\mu}_0 + c \quad i = 1, \dots, g. \tag{5}$$

3.1. Cohort projection

The state of the cohort at age t is given by a vector $\tilde{\mathbf{n}}(t)$, which is derived from an array

$$\mathcal{N}(t) = \begin{pmatrix} n_{11} & \cdots & n_{\omega 1} \\ \vdots & & \vdots \\ n_{1g} & \cdots & n_{\omega g} \end{pmatrix} \tag{6}$$