The Impact of Model Misspecification on Parameter Estimates in Mortality Models of the Gompertz Family

(preliminary draft submitted to the 2015 PAA Annual Meeting, San Diego, USA, April 30-May 2)

Trifon I. Missov[1,2] and Laszlo Nemeth[1]

[1]Max Planck Institute for Demographic Research, Rostock, Germany

[2]University of Rostock, Institute of Sociology and Demography, Rostock, Germany

Abstract

We study a class of four mortality models with exponentially increasing hazards of death by senescent causes (the Gompertz family): the Gompertz, the Gompertz-Makeham, the gamma-Gompertz, and the gamma-Gompertz-Makeham. We quantify the bias in parameter estimators if we fit a model from this family that neglects non-zero extrinsic mortality or existing mortality deceleration. We present the resulting distortion in mortality measures imposed by the inaccurate parameter estimates.

The Gompertz Family

Parametric adult-mortality models capture different patterns of death-rate increase over age: exponential, polynomial, logistic, etc. Since Beard (1959) demographers have acknowledged the fact that a given mortality pattern can be generated by a homogeneous population of individuals exposed to one and the same hazard or, equivalently, by a heterogeneous population of individuals subjected to a finite or countably infinite number of different hazards.

Gompertz (1825) suggested an exponential increase of one's risk of dying from senescent causes

$$\mu_G(x) = ae^{bx}\,, \tag{1}$$

where a denotes the level of mortality at the starting adult age (age 0) and b, equal to the relative derivative of $\mu_G(x)$, denotes the rate of aging. If the exponentially increasing hazard of death from senescent causes is complemented at each age by a constant non-aging-related risk c (Makeham 1860), the resulting model is characterized by a force of mortality

$$\mu_{GM}(x) = ae^{bx} + c\,. \tag{2}$$

If the study population is considered heterogeneous with respect to a, i.e. individual mortality patterns have the same rate of aging b, but different starting levels of mortality a, the resulting relative-risk models are known as *frailty models* (Vaupel et al. 1979). Formally, an unobserved random variable Z, called *frailty* (Vaupel et al. 1979), modulates the initial mortality level and the force of mortality for individuals with frailty $Z = z$ is given by

$$\mu_{\Gamma G}(x \,|\, z) = zae^{bx}\,, \tag{3}$$

when risk of dying for all individuals follows a Gompertz pattern, and

$$\mu_{\Gamma GM}(x \,|\, z) = zae^{bx} + c\,, \tag{4}$$

when in addition each individual is subjected to one and the same extrinsic hazard c at all ages x. For empirical (Beard 1959; Vaupel et al. 1979) and theoretical reasons (Finkelstein and Esaulova 2006; Steinsaltz and Wachter 2006; Missov and Finkelstein 2011) frailty is often assumed to be gamma-distributed with an expected value of 1 and a squared coefficient of variation γ, i.e. $Z \sim \Gamma(1/\gamma, 1/\gamma)$ follows a single-parameter gamma distribution. As frailty

is unobserved and human mortality data are aggregated for the entire population by age and year (or age and cohort), researchers fit the marginal models resulting from (3) and (4), namely

$$\mu_{\Gamma G}(x) = \frac{ae^{bx}}{1 + \frac{a\gamma}{b}(e^{bx} - 1)} \tag{5}$$

and

$$\mu_{\Gamma GM}(x) = \frac{ae^{bx}}{1 + \frac{a\gamma}{b}(e^{bx} - 1)}\,. \tag{6}$$

Models (1), (2), (5), and (6) all assume an exponentially increasing hazard of death from senescent causes, and we will address them as the *Gompertz family*. For convenience, we will also abbreviate each of the four models from the Gompertz family: the Gompertz model (1) by G, the Gompertz-Makeham (2) by GM, the gamma-Gompertz (5) by ΓG, and the gamma-Gompertz-Makeham (6) by ΓGM. Vaupel and Missov (2014) present an extensive overview of relationships that hold within each of the four models from the Gompertz family, while Missov (2013) and Missov and Lenart (2013) study the corresponding life expectancies. In this article we study how model misspecification can influence the estimates of model parameters.

A model in the Gompertz family can be misspecified by neglecting either non-zero extrinsic mortality (omitting c) or statistically significant frailty (omitting γ). Disregarding c is usually driven by the fact that estimating a model with a Makeham term (GM or ΓGM) is qualitatively different from estimating a model without it (G or ΓG): the former is a pure optimization problem (likelihood maximization) while the latter reduces to fitting a Poisson regression. Disregarding γ stems usually from one's disbelief in either the heterogeneity of the study population or the existence of mortality deceleration at the oldest ages (Gavrilov and Gavrilova 2011).

We simulate data from models with non-zero extrinsic mortality (GM or ΓGM) and fit

their Gompertz family counterparts that do not contain c (G and ΓG, respectively) to assess the associated bias in a, b, and γ. We perform the same procedure to measure the bias in a, b, and c introduced by ignoring a non-zero γ, i.e. we fit a G and a GM model to data simulated from a ΓG and a ΓGM, respectively. We quantify the resulting bias by using two different measures – the relative absolute bias (Pletcher 1999) and the mean squared error – and provide examples of demographic measures that can be substantially distorted depending on the magnitude of the bias in model parameters.

Model Fitting

The fitting procedure for the models from the Gompertz family is based on the assumption that death counts $D(x)$ at age x are Poisson-distributed, i.e. $D(x) \sim \text{Poisson}(E(x)\mu(x))$ (Brillinger 1986), where $E(x)$ is the exposure at age x and $\mu(x)$ denotes the corresponding hazard from (1), (2), (5) or (6). As a result, we maximize a Poisson log-likelihood

$$L = \sum_x (D(x) \ln \mu(x) - E(x)\mu(x)) \,. \tag{7}$$

Models (1) and (5), the G and the ΓG, are estimated in a qualitatively different way than models (2) and (6) that contain a Makeham term, the GM and the ΓGM. Taking advantage of the fact that the ΓG marginal survival function is equal to

$$s_{\Gamma G}(x) = \left(1 + \frac{a\gamma}{b}(e^{bx} - 1)\right)^{-1/\gamma} , \tag{8}$$

one can represent (5) as (see Vaupel 2002; Vaupel and Missov 2014)

$$\mu_{\Gamma G}(x) = a e^{bx} [s_{\Gamma G}(x)]^{\gamma}. \tag{9}$$

Expressions (1) and (9) aid representing log-mortality in the G and the ΓG setting as

$$\ln \mu_G(x) = \ln a + bx \tag{10}$$

$$\ln \mu_{\Gamma G}(x) = \ln a + bx + \gamma \ln[s_{\Gamma G}(x)]\,, \tag{11}$$

i.e. parameters a, b, and γ can be estimated by a Poisson regression (a generalized linear model), whose solution is calculated analytically by applying iteratively-reweighted least squares (McCullagh and Nelder 1989). Due to the additive Makeham term, expressions (2) and (6) cannot be represented in a regression form, and the only way to estimate model parameters is to maximize the likelihood L in (7). This is an optimization problem over a three- (in the case of GM) or four-dimensional (in the case of ΓGM) space. Finding its solution is not always an easy task as standard gradient-based optimization algorithms might converge to a local maximum instead of the global one. As a result, due to the high sensitivity of parameter estimates to the starting values, many researchers avoid fitting a model with a Makeham term. In this manuscript optimization was carried out by applying differential evolution (Storn and Price 1997) using the `DEoptim` R-package (Mullen et al. 2011).

Other fitting procedures can also be applied, but only when their underlying assumptions hold. For instance, one can fit a linear regression for the logarithm of the death rates, given that they are normally distributed. The latter, however, is not always fulfilled in human life tables (HMD 2014). Another alternative would be to maximize a quasi-Poisson likelihood for the mortality rates if the sample mean and variance differ substantially. This is also rarely the case for HMD data. Finally, one can maximize a binomial likelihood for the probability of death $q(x)$ if age-specific exposures $E(x)$ are unknown. The binomial log-likelihood is given by

$$\ln L = \sum_x \left[D(x) \ln q(x) + (N(x) - D(x)) \ln (1 - q(x)) \right] ,$$

where $N(x)$ denotes the population at age x. Maximizing a Poisson likelihood for the death

counts is equivalent to maximizing a binomial likelihood for the corresponding probabilities of dying (for further discussion, see Lenart 2014).

Fitting Misspecified Models to Simulated Data

We simulate individual lifespans from the GM and the ΓGM by inverting the corresponding survival functions. For $y \sim \text{Unif}(0,1)$ we solve for x numerically

$$\exp\{bx\} + \frac{bcx}{a} + \frac{b}{a}\ln y - 1 = 0$$

for the GM, and

$$\frac{a\gamma}{b}\exp\{bx\} + \gamma cx + 1 - \frac{a\gamma}{b} - y^{-\gamma} = 0$$

for the ΓGM. For the other two models of the Gompertz family we have an exact simulation formula resulting from the inversion of the corresponding survival function:

$$x = \frac{1}{b}\ln\left(1 - \frac{b}{a}\ln y\right)$$

for the G, and

$$x = \frac{1}{b}\ln\left(\frac{b}{a\gamma}y^{-\gamma} - \frac{b}{a\gamma} + 1\right)$$

for the ΓG. Death counts and exposures are aggregated age-wise from individual lifespans. For each of GM, ΓG and ΓGM we generate 500 samples of 5000 individual lifespans. Parameter values for the simulation have been chosen in accordance with the range of estimates (see Missov 2013; Missov and Lenart 2013; Missov et al. 2014) obtained from fitting models from the Gompertz family to period mortality data from the HMD (2014).

Influence of Neglecting Frailty

Neglecting frailty means that a G or a GM is fitted to a dataset with a non-zero γ. This implies that we do not account for mortality deceleration. Figure 1 shows observed logarithmic death rates for France in 1999 (triangles for females and circles for males). Estimating a G model starting from age 80 results in a (brown) line on a log-scale that substantially overestimates observed mortality after age 95 for both sexes. Estimating a ΓG provides a much better fit (the blue curves) as it captures mortality deceleration. Moreover, if mortality levels off (Gampe 2010), incorporating frailty is essential (Missov and Vaupel 2014).

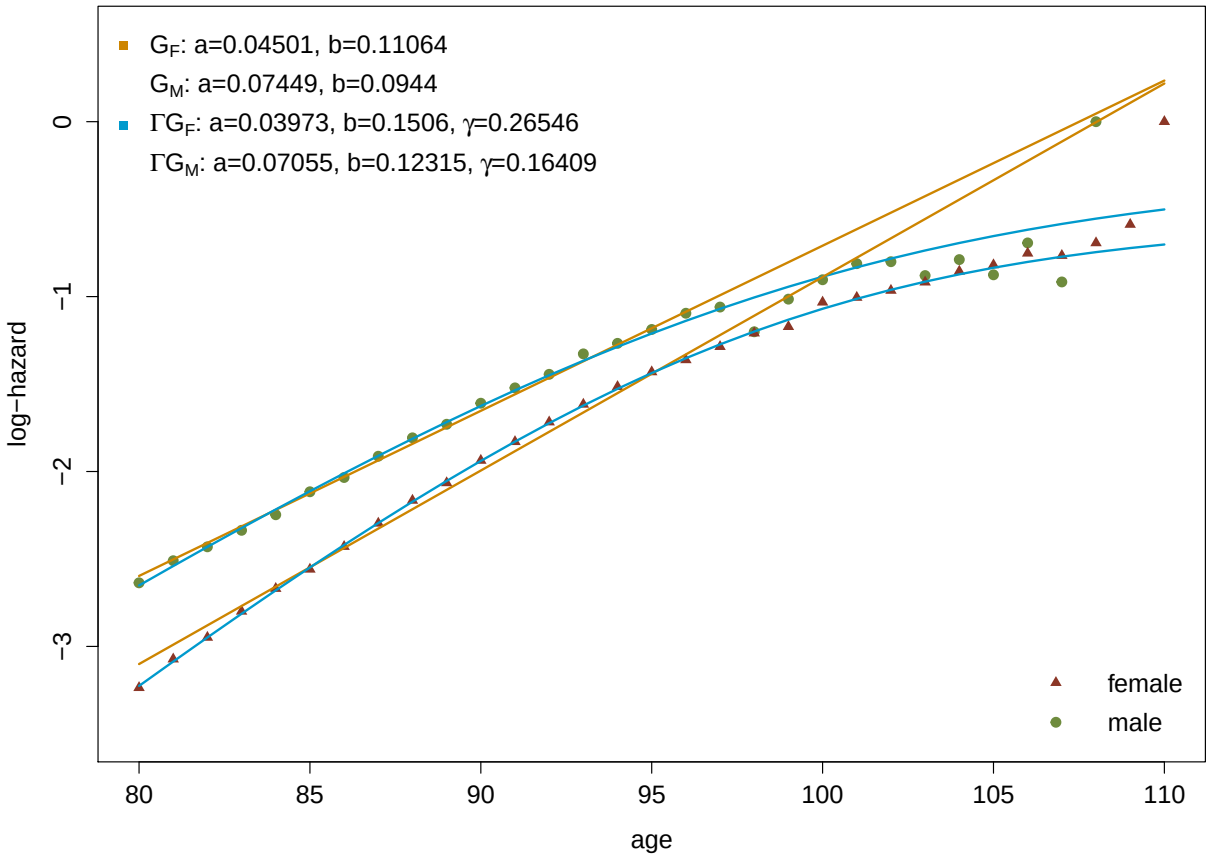

Figure 1: Observed vs fitted mortality in France, 1999, ages 80–110. Observed death rates are represented by circles for males and triangles for females (Source: KTD 2014). Fitted mortality is presented in brown for the G model and in blue for the ΓG model.

Neglected non-zero frailty results in overestimated a, underestimated b (Figures 2 and

3), and underestimated c (Figure 3). The reverse direction in which the misspecification-generated estimates of a and b shift with respect to the true values of the G parameters is not surprising: the maximum-likelihood estimators of $\hat{a}$ and $\hat{b}$ exhibit an almost perfect negative correlation (for further discussion, see Lenart and Missov 2014; Missov et al. 2014).

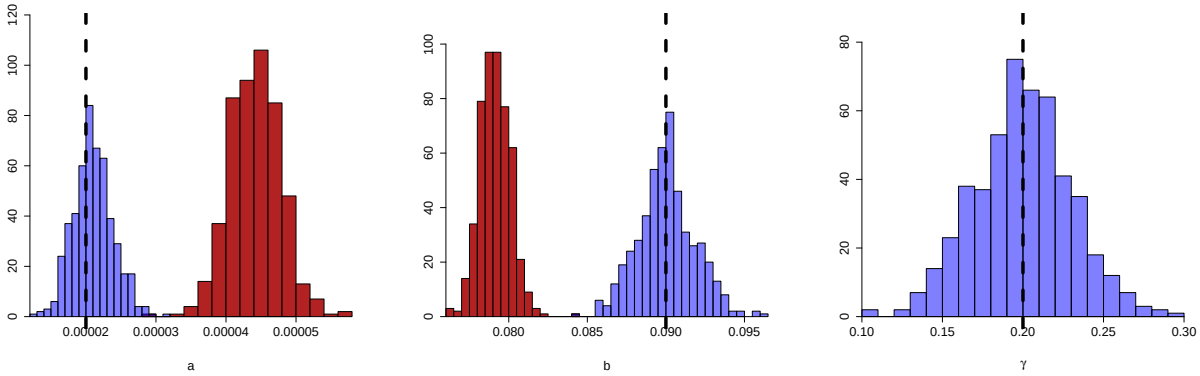

Figure 2: Histograms of estimated a and b of the Gompertz model (red bars) and estimated a, b and γ of the gamma-Gompertz model (blue bars). The G and the ΓG were fitted to data simulated from a ΓG model with parameters $a = 0.00002$, $b = 0.09$ and $\gamma = 0.2$.

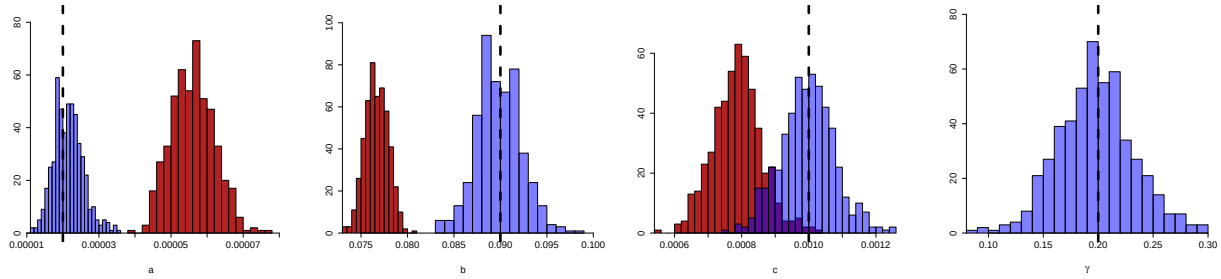

Figure 3: Histograms of estimated a,b and c of a Gompertz-Makeham model (red bars) and estimated a, b and γ of a gamma-Gompertz-Makeham model (blue bars). The GM and the ΓGM were fitted to data simulated from a ΓG model with parameters $a = 0.00002$, $b = 0.09$, $c = 0.001$ and $\gamma = 0.2$.

Influence of Neglecting the Makeham Term

Neglecting the Makeham term means that a G or a ΓG is fitted to a dataset with a non-zero c. Empirical evidence for the presence of the latter can be detected if plotted logarithmic

death rates at adult ages have a "tail" to the right of the log-Gompertz line, i.e., $c \neq 0$ if logarithmic death rates grow at a decreasing pace until some adult age, after which they start increasing linearly (see Figure 4). Fitting a G instead of a GM model (or a ΓG instead of a ΓGM) can sometimes be misleading about mortality dynamics. For example, in Figure 4 we fitted a G and a GM to mortality data for Italian males and females in 2004, ages 25-80. We restricted ourselves to the latter age range to avoid the effect of non-zero frailty γ in the subsequent ages, i.e. in this example we are interested solely in the effect of neglecting c. If we neglect the Makeham term, the estimated bs suggest that the rate of aging for females exceeds the one for males. If we include c, though, the results are exactly the opposite – the rate of aging for males is higher than the one for females.

Neglecting c shifts the estimates of the Gompertz parameters in the same direction as in the case when frailty was neglected: a is overestimated, while b is underestimated (see Figures 5 and 6). The estimates for the frailty parameter γ are close to 0, i.e., neglecting c can result in substantial underestimation of the "amount" of unobserved heterogeneity in the data.

The Bias in Parameter Estimators Induced by Model Misspecification

The bias of an estimator $\hat{\theta}$ is defined as

$$\text{Bias}(\hat{\theta}) = \text{E}\hat{\theta} - \theta\,,$$

where θ is the true value of the study parameter. Other related measures are, for instance, the relative absolute bias [1] (Pletcher 1999)

[1] Pletcher calls it "absolute bias", but we think it sounds better if we add "relative" to this name as we divide by the true value of the parameter.

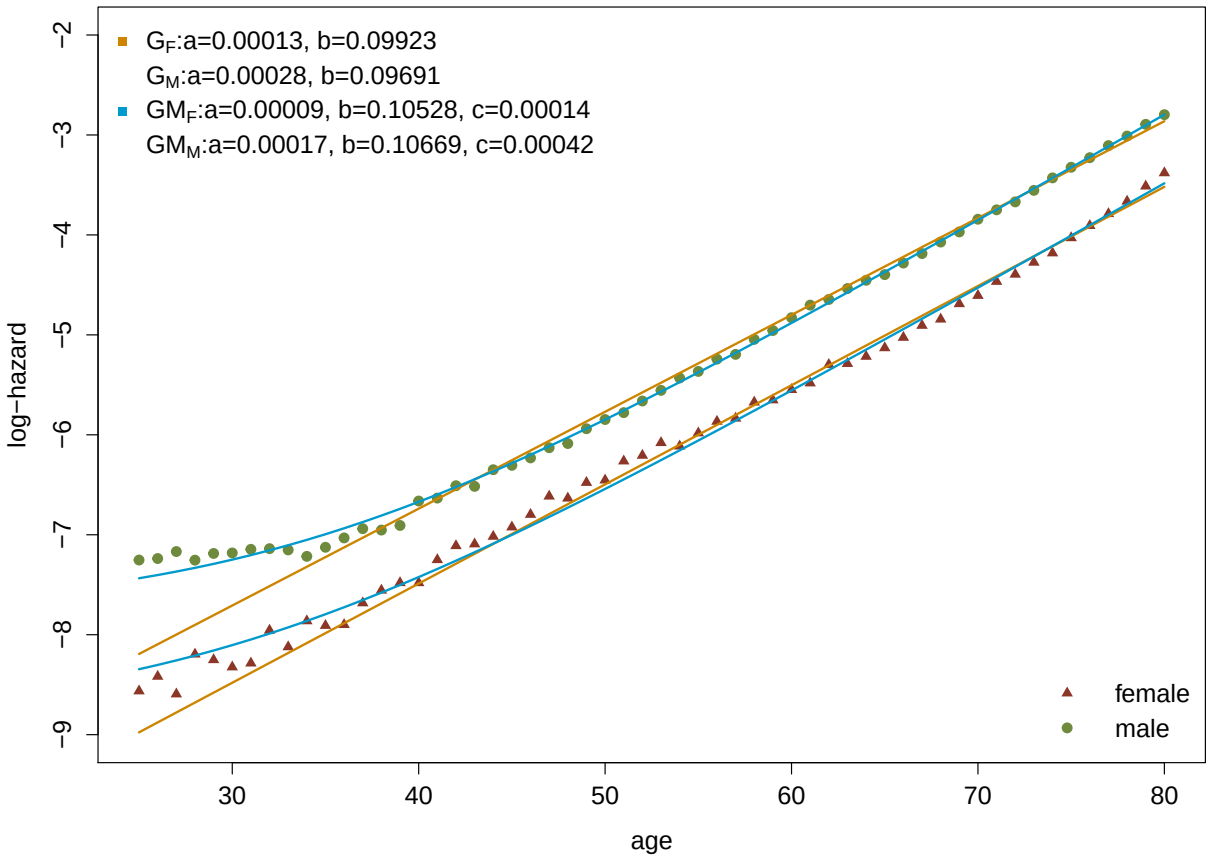

Figure 4: Observed vs fitted mortality in Italy, 2004, ages 25–80. Observed death rates are represented a green curve for males and by a red curve for females (Source: HMD 2014). Fitted mortality is presented in brown for the G model and in blue for the ΓG model.

$$\mathrm{RABias}(\hat{\theta}) = \frac{|\mathrm{E}\hat{\theta} - \theta|}{\theta}$$

or the mean squared error

$$\mathrm{MSE}(\hat{\theta}) = \mathrm{Var}\hat{\theta} + [\mathrm{Bias}(\hat{\theta})]^2 .$$

We simulated datasets with fixed $a = 0.00002$, $b = 0.09$ and varying c or γ and fitted the corresponding misspecified models. Figures 7-10 show that the relative absolute bias of all parameter estimators increases as a function of the neglected parameter (c or γ). If we

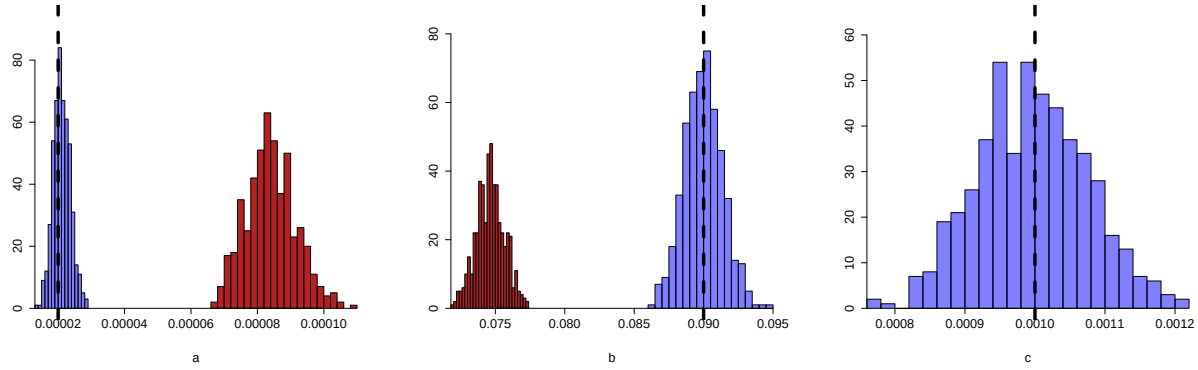

Figure 5: Histograms of estimated a and b of the Gompertz model (red bars) and estimated a, b and c of the Gompertz-Makeham model (blue bars). The G and the GM were fitted to data simulated from a GM model with parameters $a = 0.00002$, $b = 0.09$ and $c = 0.001$.

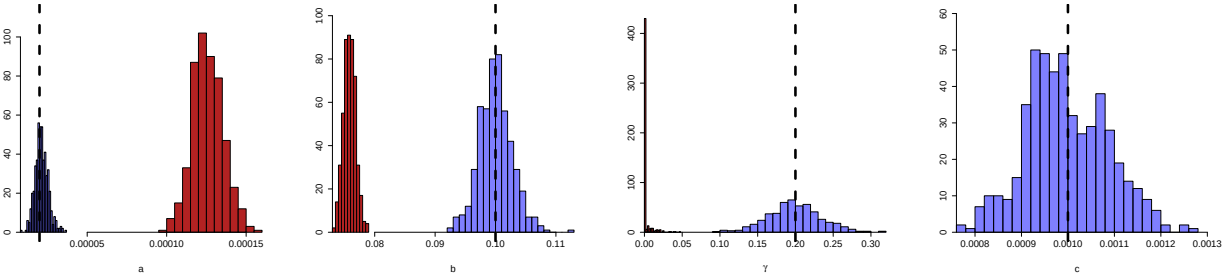

Figure 6: Histograms of estimated a, b and γ of the gamma-Gompertz model (red bars) and estimated a, b, c and γ of the gamma-Gompertz-Makeham model (blue bars). The ΓG and the ΓGM were fitted to data simulated from a ΓGM model with parameters $a = 0.00002$, $b = 0.1$, $c = 0.001$ and $\gamma = 0.2$.

disregard a non-zero c, the relative absolute bias of $\hat{a}$ increases convexly, while the one of $\hat{b}$ grows at a decreasing rate (Figures 7 and 8). Note that when c increases, the estimated γ tends to zero (Figure 8), and moreover, in the presence of extrinsic mortality c that is comparable in magnitude with a, the fitted ΓG does not detect frailty at all. Neglecting non-zero γ results in an exponentially increasing relative absolute bias of $\hat{a}$ and an almost linearly increasing one for $\hat{b}$ and $\hat{c}$ (Figures 9 and 10).

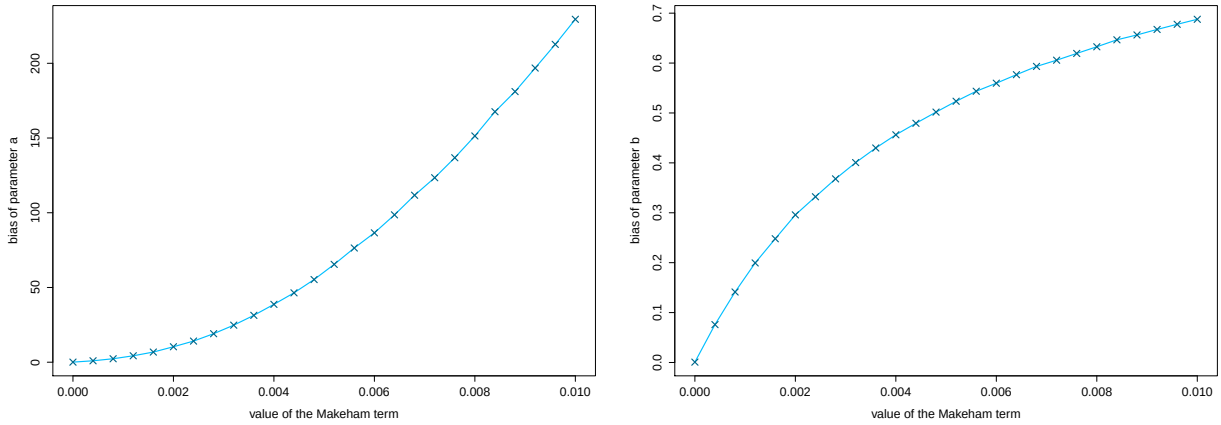

Figure 7: The relative absolute bias of G parameter estimators as a function of Makeham's c. In each case we generated 100 samples of size 5000 from the GM with $a = 0.00002$, $b = 0.09$, and a c corresponding to the marks on the horizontal axis. A G model was fitted to the generated dataset.

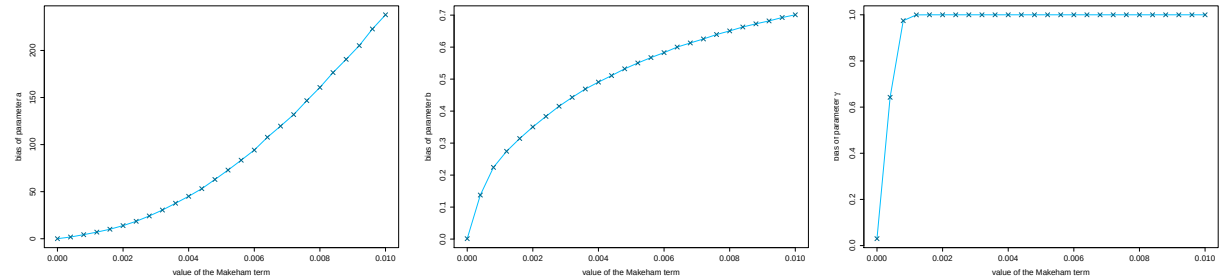

Figure 8: The relative absolute bias of ΓG parameter estimators as a function of Makeham's c. In each case we generated 100 samples of size 5000 from the ΓGM with $a = 0.00002$, $b = 0.09$, $\gamma = 0.2$ and a c corresponding to the marks on the horizontal axis. A ΓG model was fitted to the generated dataset.

Distorted Mortality Measures in the Presence of Biased Parameter Estimators

Misspecified models, i.e., models that disregard non-zero Makeham term c or non-zero frailty parameter γ, lead to biased parameter estimators. Even if the fit of the observed death rates by the estimated force of mortality is satisfactory, single-parameter estimates $\hat{\theta}$ that are not capturing the true θ can lead to wrong calculations of mortality measures in which parameters θ are involved. Examples of such are the life-table aging rate (LAR), the modal

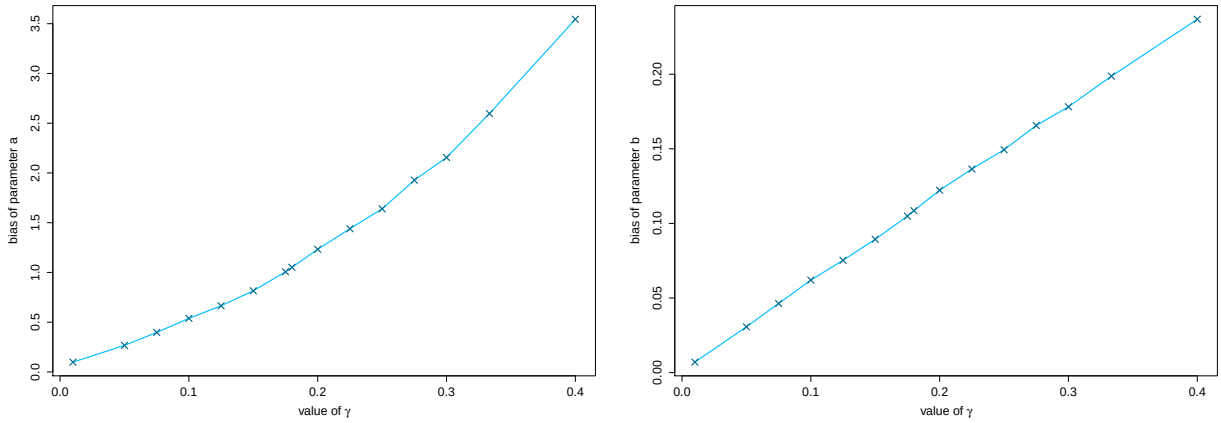

Figure 9: The relative absolute bias of G parameter estimators as a function of γ. In each case we generated 100 samples of size 5000 from the ΓG with $a = 0.00002$, $b = 0.09$, and a γ corresponding to the marks on the horizontal axis. A G model was fitted to the generated dataset.

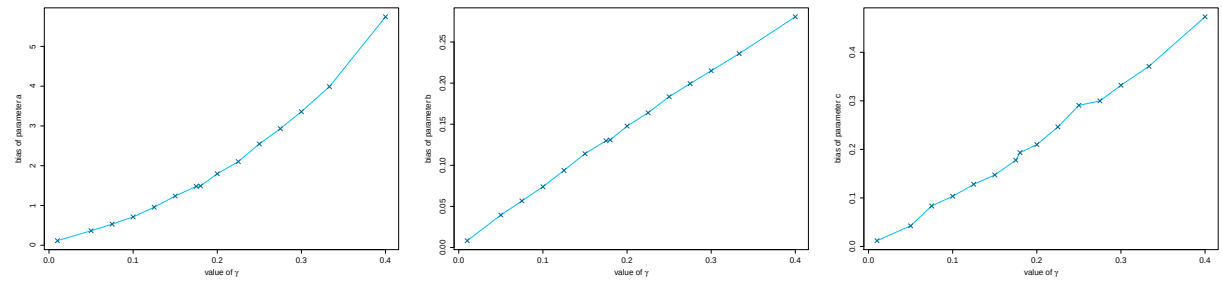

Figure 10: The relative absolute bias of GM parameter estimators as a function of γ. In each case we generated 100 samples of size 5000 from the ΓG with $a = 0.00002$, $b = 0.09$, $c = 0.001$ and a γ corresponding to the marks on the horizontal axis. A GM model was fitted to the generated dataset.

age at death, and the cross-sectional average length of life (CAL).

Life-Table Aging Rate (LAR)

The life-table aging rate (LAR) measures the age-specific pace of mortality change. It is defined as (Horiuchi and Coale 1990)

$$\mathrm{LAR}(x) = \frac{\frac{d}{dx}\,\mu(x)}{\mu(x)}\,,$$

where $\mu(x)$ denotes the hazard of death for the population. For contemporary populations it has a bell-shaped pattern over the adult age range (Horiuchi and Wilmoth 1997) and the peak of the curve signifies the age of mortality deceleration. Vaupel and Zhang (2010) present a formula for LAR in a ΓGM setting:

$$\mathrm{LAR}(x) = b\left(1 - \frac{c}{\mu_{\Gamma GM}(x)}\right) - \gamma\left(1 - \frac{c}{\mu_{\Gamma GM(x)}}\right)(\mu_{\Gamma GM}(x) - c)\ .$$

The age of mortality deceleration x^* in this case is given by (Missov and Ribeiro 2014)

$$x^* = \frac{1}{b}\ \ln\left(\frac{c(b+c\gamma)}{2ab} + \frac{\sqrt{c\gamma(b+c\gamma)[c(b+c\gamma) - 4b(a\gamma - b)]}}{2ab\gamma}\right)\ .$$

Neglecting the Makeham term results in a LAR curve that decreases monotonically, i.e. does not capture the age of mortality deceleration, and provides an inadequate fit to the empirical LAR (Figure 11).

The bigger the size of the Makeham term the more the LAR curve based on ΓG model deviates from the empirical values.

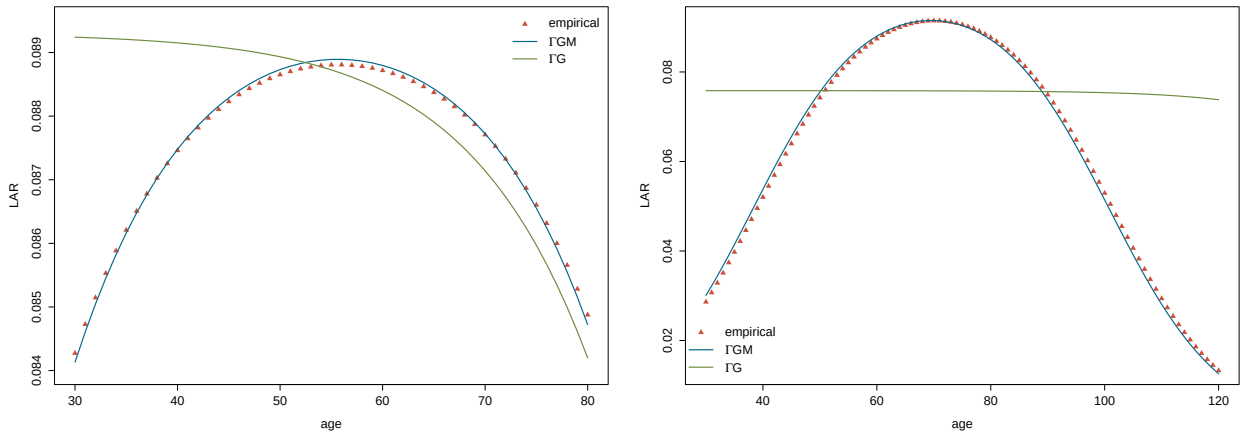

Figure 11: Empirical, ΓG and ΓGM fitted LAR values. Data simulated from a ΓGM models with parameters $a = 0.00002$, $b = 0.09$, $\gamma = 0.2$, $c = 0.00002$ (left) and $c = 0.001$ (right).

Modal Age at Death

The old-age mortality mode M is the age at which the largest number of senescent deaths occur. Unlike life expectancy e_0, i.e., the average lifespan, the modal age at death M is not influenced by infant and young-adult mortality. As a result, usually $M > e_0$. Knowing M is important because this is the age around which hospitals, nursing homes and public health in general spend most resources.

The Gompertz function ae^{bx} can be equivalently expressed via M as $be^{b(x-M)}$ (Gumbel 1958), i.e. M can be estimated directly from any model of the Gompertz family by substituting $a = be^{-bM}$ in (1), (2), (5) and (6) (for further details on the estimation procedure, see Missov et al. 2014). Note than in all four cases M denotes the Gompertz mode, i.e. the mode of the distribution of deaths by senescent causes.

The modal age at death can be inaccurately estimated if we fit a G model instead of a ΓG (Figure 12). The G model overestimates M, and the size of the overestimation depends on the magnitude of the neglected frailty. Results are similar if we fit a GM model instead of a ΓGM (Figure 13).

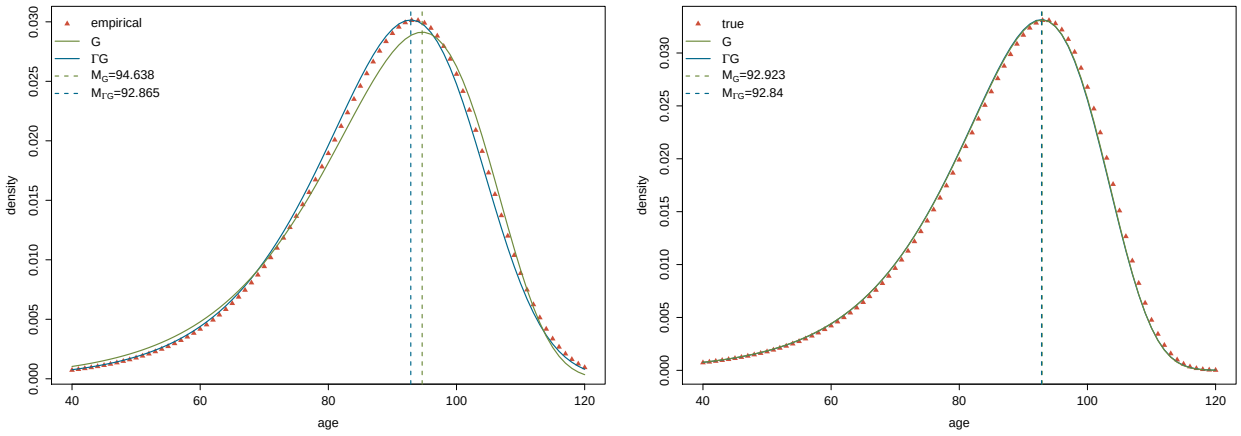

Figure 12: Distribution of deaths of estimated G and ΓG models. Data simulated from ΓG models with parameters $a = 0.00002$, $b = 0.09$, $\gamma = 0.2$ (left) and $\gamma = 0$ (right).

Neglecting the Makeham term also results in biased estimates for the modal age at death. As a consequence of fitting a G model instead of a GM (Figure 14) M is underestimated.

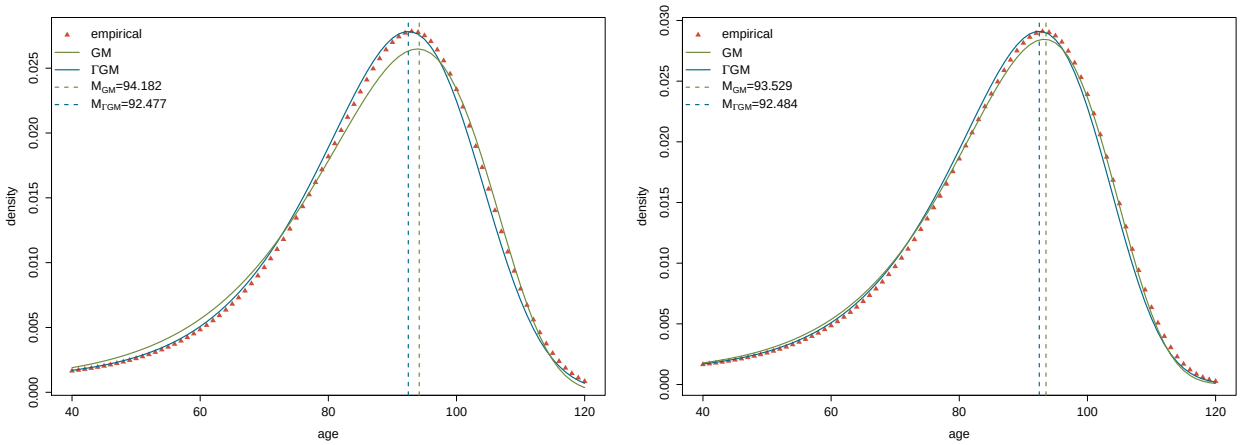

Figure 13: Distribution of deaths of estimated GM and ΓGM models. Data simulated from ΓGM models with parameters $a = 0.00002$, $b = 0.09$, $c = 0.001$, $\gamma = 0.2$ (left) and $\gamma = 0.1$ (right).

The magnitude of underestimation depends on the size of the neglected Makeham term.

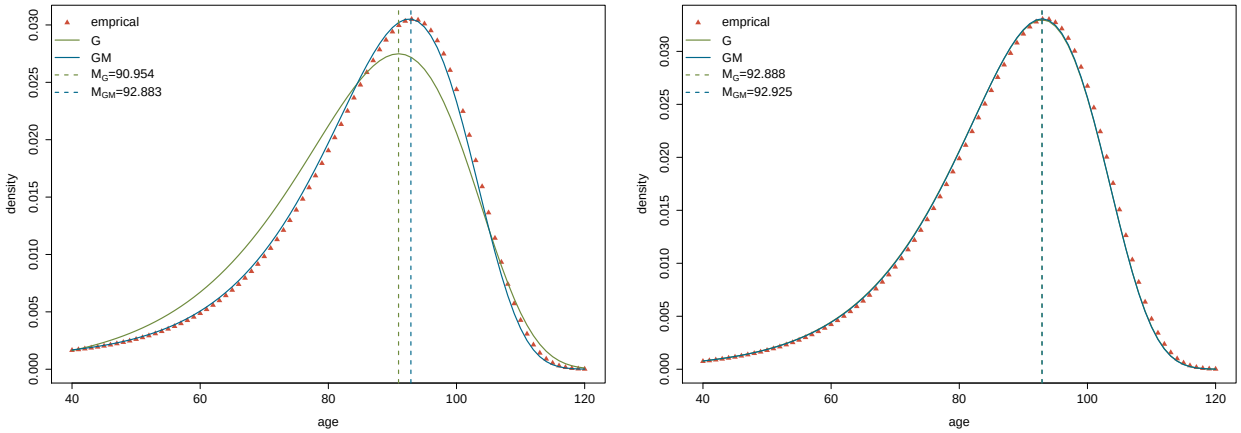

Figure 14: Distribution os deaths of estimated G and GM models. Data simulated from GM models with parameters $a = 0.00002$, $b = 0.09$, $c = 0.001$ (left) and $c = 0.00002$ (right).

Fitting a ΓG model instead of a ΓGM estimates M correctly if the population is heterogeneous "enough" (to "absorb" the bias resulting from having a high value of the Makeham term. But if we converge to GM data, we see the similar results like G fits instead of GM.). However, the higher the Makeham term, the more deaths are predicted for ages lower than M.

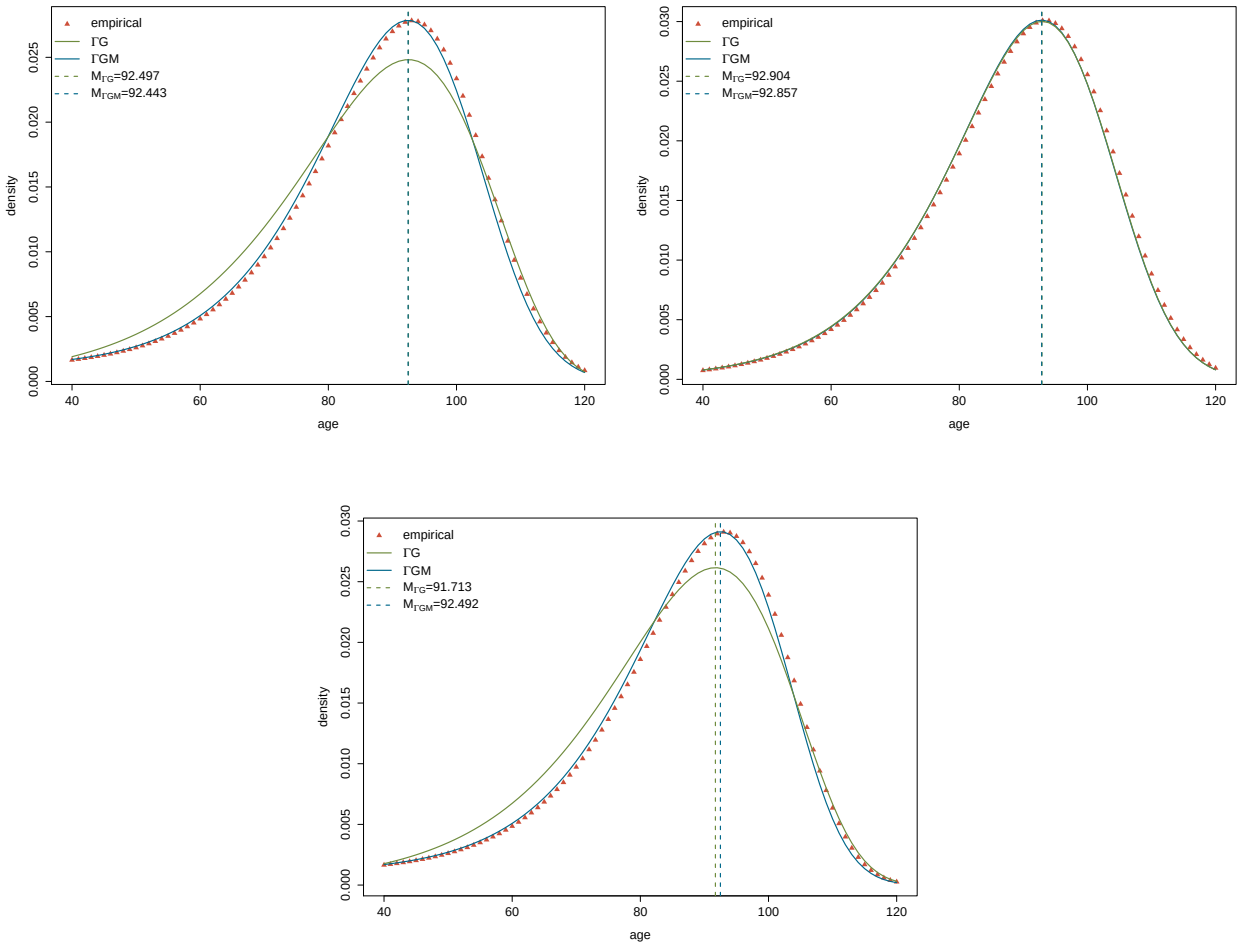

Figure 15: Distribution of deaths of estimated ΓG and ΓGM models. Data simulated from ΓGM models with parameters $a = 0.00002$, $b = 0.09$, $\gamma = 0.2$, $c = 0.001$ (topleft), $c = 0.00002$ (topright) and $\gamma = 0.1$, $c = 0.001$ (bottom).

Life Expectancy

Life expectancy is conventionally calculated from a life table which characterizes the distribution of deaths. If we assume a parametric structure for the latter, we can estimate model parameters and calculate model-based life expectancy (the integral of the explicitly given survival function). One can use model-based life expectancy to reconstruct exposures (for discussion and examples, see Missov and Lenart 2013) or forecast mortality.

In the aforementioned parametric setting, the life table is characterized by the true values of model parameters. If the maximum-likelihood estimators of the parameters are biased, the resulting model-based (remaining) life expectancy will not match the conventional one (the $e(x)$ column of the life table) and, as a result, reconstructed exposures and mortality forecasts can be distorted.

Tables 1-3 show the effect of neglecting γ on remaining life expectancy at different ages. In all cases the ΓGM model-based life expectancy is the closest to the conventional one. This not surprising as, among all models in the Gomeprtz family, the ΓGM captures best the observed S-shaped mortality pattern. Neglecting frailty does not seem to affect life expectancy much. In the next draft of this manuscript we will study how these small discrepancies affect exposures and forecasts.

	LT	GM	ΓGM
${}_{1900}e^{f}_{30}$	39.23	38.96423	38.87021
${}_{1900}e^{m}_{30}$	37.14	37.26261	37.26079
${}_{1900}e^{f}_{20}$	46.51	46.65176	46.65493
${}_{1900}e^{m}_{20}$	44.26	44.40381	44.40485
${}_{2010}e^{f}_{50}$	34.54	34.52609	34.54378
${}_{2010}e^{m}_{50}$	31.17	31.19214	31.19471
${}_{2010}e^{f}_{30}$	53.99	53.95806	53.95827
${}_{2010}e^{m}_{30}$	50.32	50.30969	50.31780

Table 1: Female and male remaining life expectancy estimates for the Swedish cohorts of 1900 and 2010. Conventional life expectancy is represented in the second column, while model-based life expectancies are listed in columns 3-6.

	LT	G	ΓG
80	9.43	9.38	9.44
85	6.58	6.76	6.68
90	4.44	4.67	4.67
95	3.05	3.10	3.36
100	2.14	1.98	2.58
105	1.63	1.23	2.17
110	0.50	0.74	1.96

Table 2: Remaining life expectancy for French females, ages 80-110. Source: KTD. Conventional life expectancy is represented in the second column, while model-based life expectancies are listed in columns 3-4.

	LT	G	ΓG
80	7.30	7.39	7.41
85	5.14	5.36	5.33
90	3.51	3.79	3.82
95	2.49	2.60	2.82
100	1.83	1.74	2.19
105	1.55	1.14	1.81
110	-	0.74	1.60

Table 3: male remaining life expectancy, ages 80-110, France, 1999, source: KTD

Conclusion

Neglecting non-zero extrinsic mortality or unobserved heterogeneity distorts parameter estimates of the models from the Gompertz family: a is overestimated, while b is underestimated. When c is neglected, $\hat{\gamma}$ tends to zero. When γ left out of the model, parameter c is underestimated. The bigger the magnitude of the neglected c or γ, the higher the bias in the estimators of the model parameters.

Mortality measures that rely on estimated parameters of the Gompertz family get distorted if the parameters are inaccurately estimated. Among them are the life-table aging rate (LAR), the age at mortality deceleration, the modal age at death, but not life expectancy.

References

Beard, R.E. 1959. *Note on Some Mathematical Mortality Models*, pp. 302–311. Woolstenholme G.E.W., O'Connor M. (eds.). Boston: Little, Brown and Company.

Brillinger, D.R. 1986. "The Natural Variability of Vital Rates and Associated Statistics." *Biometrics* pp. 693–734.

Finkelstein, M.S. and V. Esaulova. 2006. "Asymptotic Behavior of a General Class of Mixture Failure Rates." *Advances in Applied Probability* 38:242–262.

Gampe, J. 2010. "Human Mortality Beyond Age 110." In *Supercentenarians*, edited by H. Maier, J. Gampe, B. Jeune, J.-M. Robine, and J.W. Vaupel, number 7 in Demographic Research Monographs, chapter III, pp. 219–230. Heidelberg [et al.]: Springer.

Gavrilov, L.A. and N.S. Gavrilova. 2011. "Mortality Measurement at Advanced Ages: A Study of the Social Security Administration Death Master File." *North American Actuarial Journal* 15:442–447.

Gompertz, B. 1825. "On the Nature of the Function Expressive of the Law of Human Mortality, and on a New Mode of Determining the Value of Life Contingencies." *Philosophical Transactions of the Royal Society of London* 115:513–583.

Gumbel, E.J. 1958. *Statistics of Extremes*. Columbia University Press.

HMD. 2014. "The Human Mortality Database." http://www.mortality.org/.

Horiuchi, S. and A.J. Coale. 1990. "Age Patterns of Mortality for Older Women: An Analysis Using the Age-Specific Rate of Mortality Change with Age." *Mathematical Population Studies* 2:245–267.

Horiuchi, S. and J.R. Wilmoth. 1997. "Age Patterns of the Life Table Aging Rate for Major Causes of Death in Japan, 1951–1990." *Journal of Gerontology, Biological Sciences* 52A:B67–B77.

KTD. 2014. "The Kannisto-Thatcher Database." http://www.demogr.mpg.de/databases/ktdb/.

Lenart, A. 2014. "The moments of the Gompertz distribution and maximum likelihood estimation of its parameters." *Scandinavian Actuarial Journal* 2014:255–277.

Lenart, A. and T.I. Missov. 2014. "Goodness-of-fit tests for the Gompertz distribution." *Communications in Statistics - Theory and Methods (in press)* .

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

McCullagh, P. and J. A. Nelder. 1989. *Generalized Linear Model*. Monographs on Statistics Applied Probability. London: Chapman & Hall, 2nd edition.

Missov, T.I. 2013. "Gamma-Gompertz Life Expectancy at Birth." *Demographic Research* 28:259–270.

Missov, T.I. and M. Finkelstein. 2011. "Admissible Mixing Distributions for a General Class of Mixture Survival Models with Known Asymptotics." *Theoretical Population Biology* 80:64–70.

Missov, T.I. and A. Lenart. 2013. "Gompertz-Makeham Life Expectancies: Expressions and Applications." *Theoretical Population Biology* 90:29–35.

Missov, T.I., A. Lenart, L. Nemeth, V. Canudas-Romo, and J.W. Vaupel. 2014. "The Gompertz Force of Mortality in Terms of the Modal Age of Death." *Demographic Research (under review)* .

Missov, T.I. and F. Ribeiro. 2014. "Do We Age at the Same Rate? Evidence from Cause-of-Death Data." *Demography (under review)* .

Missov, T.I. and J.W. Vaupel. 2014. "Mortality Implications of Mortality Plateaus." *SIAM Review (in press)* .

Mullen, K.M., D.L. Gil D. Ardia, D. Windover, and J. Cline. 2011. "DEoptim: An R Package for Global Optimization by Differential Evolution." *Journal of Statistical Software* 40:1–26.

Pletcher, S.D. 1999. "Model fitting and hypothesis testing for age-specific mortality data." *Journal of Evolutionary Biology* 12:430–439.

Steinsaltz, D.R. and K.W. Wachter. 2006. "Understanding Mortality Rate Deceleration and Heterogeneity." *Mathematical Population Studies* 13:19–37.

Storn, R. and K. Price. 1997. "Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces." *Journal of Global Optimization* 11:341–359.

Vaupel, J.W. 2002. "Life Expectancy at Current Rates vs Current Conditions: A Reflexion Stimulated by Bongaarts and Feeney's "How Long Do We Live?"." *Demographic Research* 7:365–378.

Vaupel, J.W., K.G. Manton, and E. Stallard. 1979. "The Impact of Heterogeneity in Individual Frailty on the Dynamics of Mortality." *Demography* 16:439–454.

Vaupel, J.W. and T.I. Missov. 2014. "Unobserved Population Heterogeneity: A Review of Formal Relationships." *Demographic Research (in press)* .

Vaupel, J.W. and Z. Zhang. 2010. "Attrition in heterogeneous cohorts." *Demographic Research* 23:737–748.