A Bayesian Model of Later Life Mortality Trends and Implications for Longevity *

Julian Ashwin
Maastricht University

Andrew Scott†
London Business School and CEPR

January 25, 2025

Keywords: Bayesian Methods, Life Expectancy, Lifespan, Lifespan Equality, Longevity, Mortality, State Space Models

Running head: A Bayesian Model of Later Life Mortality Trends

*This research was made possible by funding from ESRC grant T002204. The authors would like to thank Jose Manuel Aburto, Steve Haberman, Andrew Harvey, Rob Hyndman, Ronald Lee, Andrea Nigri, Jay Olshansky, Roland Rau, Andrew Steele, Andres Villegas, Jordan Weiss as well as an Associate Editor and three anonymous referees for extremely valuable comments and feedback.

†Corresponding author: ascott@london.edu

Abstract

Using a novel, flexible and easily interpretable dynamic Bayesian state space model, we analyze historic and future longevity trends across 18 high income countries over the last one hundred years and 16 large population emerging markets from 1950. Our results show the key driver of global life expectancy is now late-life mortality whose importance is projected to increase further. We find no sign of any impending limit to average life expectancy but project a slowdown in future life expectancy gains despite continuing improvement in later-life mortality. Gains to later-life mortality are increasingly driven by the modal age of death with a slowdown in improvements in the speed of ageing and compressions of mortality. The consequence is a projection increase in the upper bound of age at death and a slowdown in lifespan equality improvements. Whereas the 20th century saw widespread cross-country convergence in longevity indicators the projections are for divergence both within high income countries as well as large population emerging markets. A particular outlier is the United States where our model predicts substantial increases in the modal and upper bound for observed age at death but only small improvements in life expectancy and so an increase in lifespan inequality.

1 Introduction

It is well known that a global demographic transition is leading to an ageing society and an increase in the number and proportion of older people (Lee and Mason, 2017). However, with global life expectancy now over 73 years (United Nations, 2024) and exceeding 80 in high income countries, life expectancy gains are increasingly driven by changes in mortality rates at older ages (Eggleston and Fuchs, 2012). That raises the prospect of not just a rising proportion of older people but of ever more people living into older and older ages. In other words, a new stage of a demographic transition is beginning, characterised more by a longevity than an ageing society (Christensen et al., 2009; Scott, 2021). Crucial to this longevity society are two questions - how long can we expect to live for and how well will we age?

Answering those questions requires understanding future longevity trends and in particular those of later-life mortality. If statistical trends in later-life mortality differ from those at other ages, past historical patterns in life expectancy cannot be used to project forward future developments (Olshansky et al., 2024). A focus on later-life mortality also raises deeply profound issues that do not arise when studying mortality at younger ages. Reductions in infant and midlife mortality have propelled global life expectancy to over 70 years and have enabled many to enjoy the long life that was previously the outcome for a minority. But, as emphasised in the "Broken Limits" referred to by Oeppen and Vaupel (2002), continual increases in life expectancy hint at a fundamental change in human existence if life continues to stretch beyond "three score years and ten."

As humanity adapts to the increases in life expectancy achieved in the 20th century a host of 21st century questions are therefore raised. How long can humans live for? (Thatcher, 1999; Dong et al., 2016; Lenart and Vaupel, 2017) and how high can average life expectancy reach? (Cheung et al., 2005) But, as Jonathan Swift remarked in Gulliver's Travels "every man desires to live long, but no man wishes to be old" (Swift, 1726). In other words, it is not just length of life that matters but also its quality. Vaupel (2010) argues longer lives are the result of individuals staying in good health for longer and so delaying the onset of ageing rather than any change in the speed at which we age, consistent with the empirical results of

Zuo et al. (2018). But improvements in the speed at which we age, achieving a compression of mortality and perhaps also morbidity (Fries, 1980), are especially valued from a welfare perspective (Crimmins, 2015; Scott et al., 2021).

As life expectancy increases the concern isn't just that we will outlive our health but also our finances (Milevsky, 2020). That makes the degree of uncertainty around the age of death increasingly important. As emphasised in Colchero (2016), past improvements in life expectancy have led to a reduction in uncertainty regarding age of death. However, depending on trends in later-life mortality, this relationship could be reversed such that future gains might bring about greater uncertainty.

Given considerable international diversity in life expectancy there are also numerous cross-country issues. Will the experience of the large population emerging markets that increasingly define the world population mimic that of the current high income countries? Will the convergence in life expectancy trends amongst high income countries over the past 100 years continue into the future? Not surprisingly given the importance of these issues there has been much research across a range of disciplines providing insights in terms of models, estimation methodologies, data and trends, see *inter alia*, (Aburto et al., 2020; Welsh et al., 2021; Bravo et al., 2021; Heuveline, 2022; Schöley et al., 2022; Levantesi et al., 2023).

In this paper we seek to make two contributions. The first is to propose a model that is flexible and simple enough that it can capture mortality trends at all ages, over long time periods and across multiple countries in a parsimonious and easily interpretable way. The second is to use this model to provide an update on current longevity trends and to make future projections for a wide range of countries.

We proceed as follows. In Section 2 we outline our model and the data we use as well as our estimation methodology. In Section 3 we illustrate our approach with reference to Best Practice Life Expectancy (Oeppen, 2019) whilst the main use of our model is in Section 4 where we apply it to 18 high income countries for whom data exists for more than 100 years and 16 large population emerging economies from 1950. Section 5 uses our approach to make projections into the future enabling us to review whether 20th century trends are likely to continue into the 21st century. A final section concludes.

2 Model, Data and Estimation

In this section we outline our modelling choices, our estimation process and the data used.

2.1 A Dynamic Model of Mortality

Firstly, given the importance of differential mortality trends at different ages and the growing role of improvements in later-life mortality we require a model that allows for different dynamics at each age. Secondly, as we apply the model to a large number of countries and over long periods of time we also require a model that is flexible, robust and parsimonious. Thirdly, we also want the model parameters to be easily interpretable in order to shed light on key longevity issues.

Therefore, we base our analysis on a time varying Siler mortality function (Siler, 1979), as parameterised in Bergeron-Boucher et al. (2015). That is, denoting expected mortality rates at age a in period t by $\mu_{a,t}$, we assume:

$$\mu_{a,t} = b_t \exp(-b_t(a + B_t)) + c_t \exp(c_t(a - C_t)) + d_t, \tag{1}$$

where each of the parameters (b_t, B_t, c_t, C_t, d_t) are non-negative. The parameters b_t and B_t reflect infant mortality whilst d_t denotes age-independent mortality. The parameters c_t and C_t are of particular interest as they capture later-life mortality. C_t reflects the modal age at death whilst increasing c_t slows the increase in mortality with age below the modal age of death but increases it for ages above, bringing about a "compression of mortality". Thus c_t reflects the "speed" of aging and C_t postponements in its timing (Horiuchi et al., 2013; Missov et al., 2015). Figure 1 illustrates this by comparing the change in survival rates due to increases to c_t (in red) and C_t (in green) that both generate an additional 5 years of life expectancy at birth, relative to a baseline survival curve (in blue).

Our choice of this Siler model is based on its relative parsimony (we found more heavily parameterised models either computationally expensive or, given the variety of periods and countries examined, not always stable or convergent) and its relative success in terms of out of sample predictions compared to other popular low-dimensional parameterisations (see Appendix A.1).

Figure 1: Illustration of how c_t and C_t change survival rates

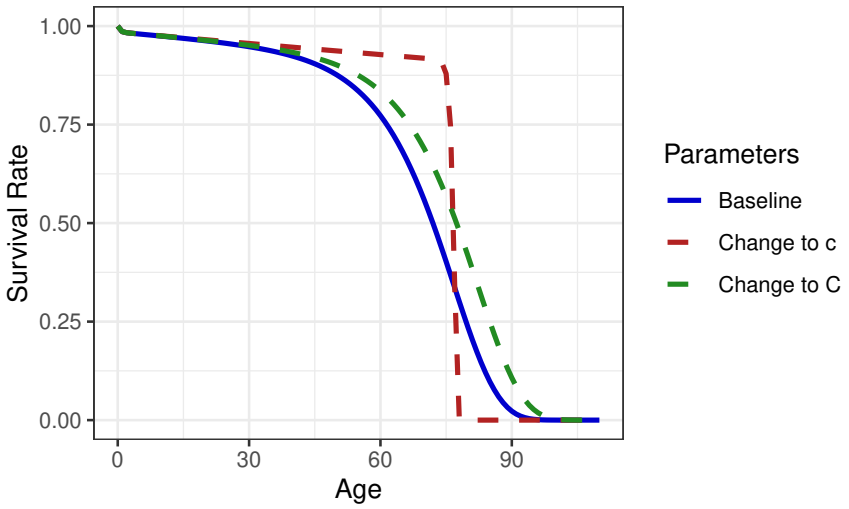

For estimation purposes we assume observed mortality rates, $m_{a,t}$, are log-normally distributed[1] with mean $\ln(\mu_{a,t})$ and variance σ_t:

$$\ln(m_{a,t}) = \ln(\mu_{a,t}) + \nu_{a,t} \quad \text{where} \quad \nu_{a,t} \sim \mathcal{N}(0, \sigma_t) \tag{2}$$

An alternative assumption would be to model death counts (Basellini and Camarda, 2019) as a Poisson distribution whilst accounting for the size of the total population, although at the cost of having to model a population process.[2]

Our model has six time-varying parameters - b_t, B_t, c_t, C_t, d_t and σ_t. The Siler functional form requires these to be (weakly) positive, so we allow their log value to vary according to a random walk with time-varying drift. A parameter θ_t (each of d_t, b_t, B_t, c_t, C_t, σ_t) thus evolves according to Equation 3

$$\ln(\theta_t) = \alpha_{\theta,t} + \ln(\theta_{t-1}) + \epsilon_{\theta,t} \quad \text{where} \quad \epsilon_t \sim \mathcal{N}(\mathbf{0}, \Sigma_\epsilon), \tag{3}$$

where the drift terms α_t is also assumed to follow a random walk :

$$\alpha_{\theta,t} = \alpha_{\theta,t-1} + \xi_{\theta,t} \quad \text{where} \quad \xi_t \sim \mathcal{N}(\mathbf{0}, \Sigma_\xi). \tag{4}$$

[1]This assumption is also implicit in the Lee and Carter (1992) model, see Basellini and Camarda (2019).

[2]Forecasting exercises in Appendix A.1 provide support for our approach with its performance often beating alternatives using population data in this way.

As shown in Harvey (1989), our specification is flexible enough to approximate a wide range of ARIMA processes. Importantly, by modelling both the level and drift as random walks we allow for considerable variation in mortality trends across ages, countries and over long periods of time.

A substantial literature exists on the statistical modelling of life expectancy and mortality, much of it focused on Lee and Carter (1992) and variants. In contrast, we specify a non-linear functional form for mortality at each age with separate random walk processes for both the level and growth of these parameters and estimate using a Bayesian state space approach. Such a parametric approach has several potential advantages over more reduced-form approaches such as Lee and Carter (1992), in particular it offers greater flexibility over longer samples and more easy interpretation given the parameteric representation.

A number of other studies have also used Bayesian state space methods i.e De Jong and Boyle (1983); Pedroza (2006); De Jong and Tickle (2006). Fung et al. (2017) incorporate the approach into Lee-Carter specifications , whilst (Aliverti et al., 2022) utilize mixtures of skewed distribution functions in estimation, Alexander et al. (2017) uses the approach to estimate small area mortality and Dodd et al. (2018) to smooth mortality data. A novel feature of our approach is to apply these methods to Siler functional forms and allow for dynamic effects (see Raferty et al. (2013) for a Bayesian model with a unit root and non-constant drift but applied to life expectancy rather than mortality).

Closest to our approach is the body of work that estimates parameterised mortality functions with time-varying parameters (McNown and Rogers, 1989, 1992; Congdon, 1993; Li and Liu, 2019; Tang et al., 2023).[3] In particular, Dellaportas et al. (2001) who estimate a time invariant Heligman-Pollard parameterisation of mortality, which Alexopoulos et al. (2019) extends for time varying parameters. Similarly, Cairns et al. (2006) estimate a time-varying two parameter logistic model. These papers focus on relatively short sample periods (1983-1992) and forecasting at medium run horizons (up to 21 years ahead). In contrast, the flexibility of our model, in particular the time-varying drift terms, allow us to model the evolution of mortality over considerably longer sample periods and a wide vari-

[3]We show in Appendix A.2 that as well as allowing for a more internally consistent approach with well-defined estimation uncertainty, our dynamic model also produces more accurate forecasts than a static version.

ety of countries. As described in Appendix A.3, where we compare the Siler function to the Gompertz-Makeham and Heligman-Pollard functions, we chose the Siler function as it strikes a balance between parsimoniously capturing mortality changes across the age distribution whilst being flexible enough to model a wide range of countries across multiple time periods.

2.2 Data

We use mortality data from the Human Mortality Database (HMD) as well as the United Nations World Population Prospects for a total of 34 countries. The available sample and data source for each country is show in in Appendix A.5. We include in our sample countries with at least 100 years of data available from HMD and/or they are one the world's 20 most populous countries (data taken from UN WPP where not available from HMD). For each country we use any available data from 1900 onwards. This gives us a sample roughly balanced between high income and large population emerging markets countries, while still ensuring each country has at least a 70 year sample (with the exception of Germany).[4] We use five year averages as our basic data point so that our analysis focuses on medium and long run trends.[5] The results required considerably more computational time but provided little different in the way of conclusions.

2.3 Priors and Estimation

We use a Bayesian approach to directly estimate posterior (in-sample) and predictive (out-of-sample) distributions taking into account joint parameter uncertainty. We define our priors as follows:

- For the initial values of the Siler function parameters θ which evolve according to Equation 3, we select log-normal prior distributions in order to be consistent with our random walk processes. Specifically: $\ln(B_0) \sim \mathcal{N}(\ln(5), 2)$, $\ln(b_0) \sim \mathcal{N}(\ln(1), 1)$, $\ln(C_0) \sim \mathcal{N}(\ln(80), 2)$, $\ln(c_0) \sim \mathcal{N}(\ln(0.1), 1)$, $\ln(d_0) \sim \mathcal{N}(\ln(0.01), 1)$, $\ln(\sigma_0) \sim \mathcal{N}(\ln(0.01), 1)$.

[4] Data quality will vary across countries and over time, adding a note of caution to interpreting results earlier in our sample.

[5] For robustness purposes we also estimate our model for a number of countries using annual data (see Appendix A.4).

- For the initial values of the α_θ drift terms which evolve according to Equation 4, we select independent normal prior distributions of $\alpha_1 \sim \mathcal{N}(0, 0.02)$ for each θ.

- We assume the covariance matrices Σ_ϵ and Σ_ξ are diagonal and select independent inverse-gamma prior distributions of $\mathcal{IG}(2, 0.1)$ for the diagonal elements of both.

We estimate posterior and predictive distributions using the No-U-Turn Sampler (NUTS), a Hamiltonian Monte Carlo (HMC) method (Neal et al., 2011; Hoffman et al., 2014; Betancourt, 2017), as provided by the Turing.jl package in *Julia*. A more detailed description of the estimation method is given in Appendix A.6.

3 Best Practice Life Expectancy

To illustrate our approach we first apply it to Best Practice Life Expectancy (BPLE), see Oeppen and Vaupel (2002). BPLE is defined by the mortality curve of the country with the highest life expectancy at birth at a particular moment in time.[6] Based on this definition, Figure 2 shows mortality and survival curves for BPLE from 1900 to 2019, our constructed BPLE[7] and a measure of lifespan equality (defined by the inverse of a measure of dispersion at age of death - see Appendix A.7 for definitions).

Posterior distributions for the time varying parameters in Eq. 1 for BPLE are shown in Figure 3. In Figure 4 we use these and a linear approximation for life expectancy, based on our Siler function, to perform a historical decomposition of past increases in life expectancy.

Consistent with previous results, Figure 4 shows large declines (with a brief spike during the Great Influenza Pandemic) in age-independent mortality d_t to a level which allows for only limited further progress. Up until 1960 it was this decline in d_t, in conjunction with falling infant mortality, that were the main drivers of increasing BPLE.

Focusing on later-life mortality our results show that in the first half of the 20th century c_t increased, leading to a compression of mortality, but since then has shown little change. By comparison the modal age of death, C_t, has shown a relatively constant rate of trend

[6] An alternative approach is to take the minimum mortality rate at each age across all countries in a given period (Vallin et al., 2008; Liu and Li, 2019; Canudas-Romo et al., 2019; Li et al., 2021).

[7] As is customary this is constructed using female mortality data - to enable comparison we use female mortality data throughout this paper.

Figure 2: Best practice mortality and life expectancy (1900-2019)

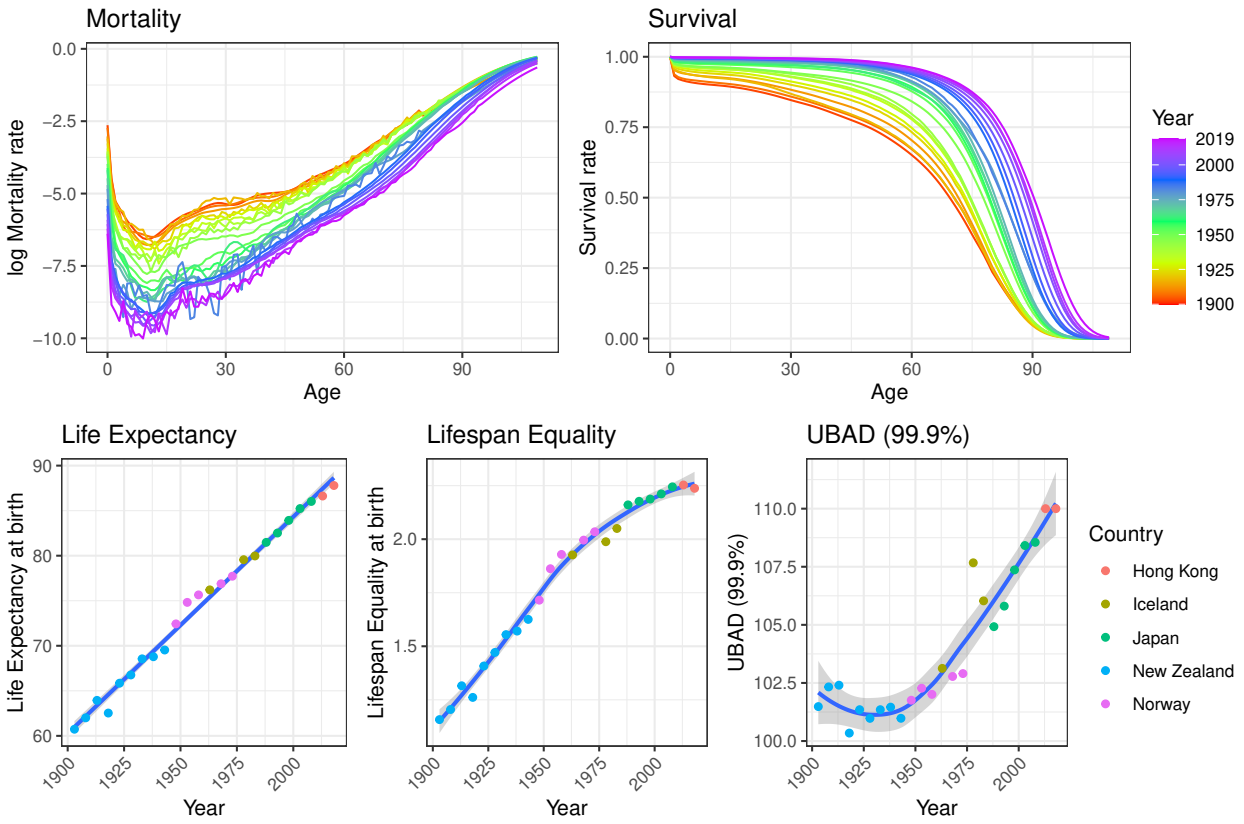

Note: Best practice is defined as the country with the highest life expectancy at birth in a given period. The colour of each point indicates the identity of the country defining best practice in each period. Data plotted are five year averages.

improvement. As a result, since 1960, the substantial majority (86%) of increases in BPLE are explained by increases in C_t (Canudas-Romo, 2010). This is a product of two effects. The first is the trend increase in C_t, shown in Figure 3, and the stalling of change in other parameters. The second is the rise in the sensitivity of life expectancy to changes in C_t. In particular, with a slowing in the speed of ageing the derivative of life expectancy (based on our Siler specification) with respect to shifts in the modal age of death increases. This narrative is confirmed by the posterior distributions for the α_θ drift terms shown in Figure A.9. These show that while the drift term for c_t has fallen to close to zero, it has remained at around 0.01 for C_t (suggesting an expected increase of 1% every five years).

Figure 3: Siler parameter estimates for best practice mortality (1900-2019)

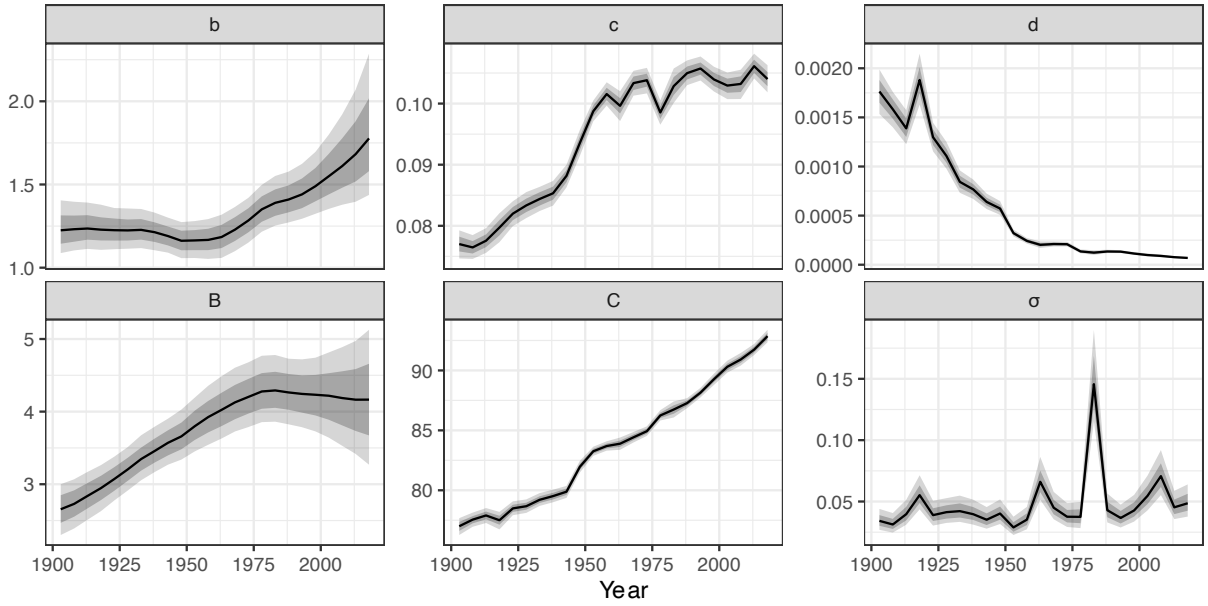

Note: Panels show posterior distributions for the five parameters in Eq 3 and the variance of the error term in Eq 2 over time, for best practice mortality. The solid line indicates the median with the shaded areas indicating the 70% and 95% credible intervals. Posterior distributions for the drift and variance terms are shown in Appendix A.8.

Figure 4: Historical decomposition of changes (1900-2019)

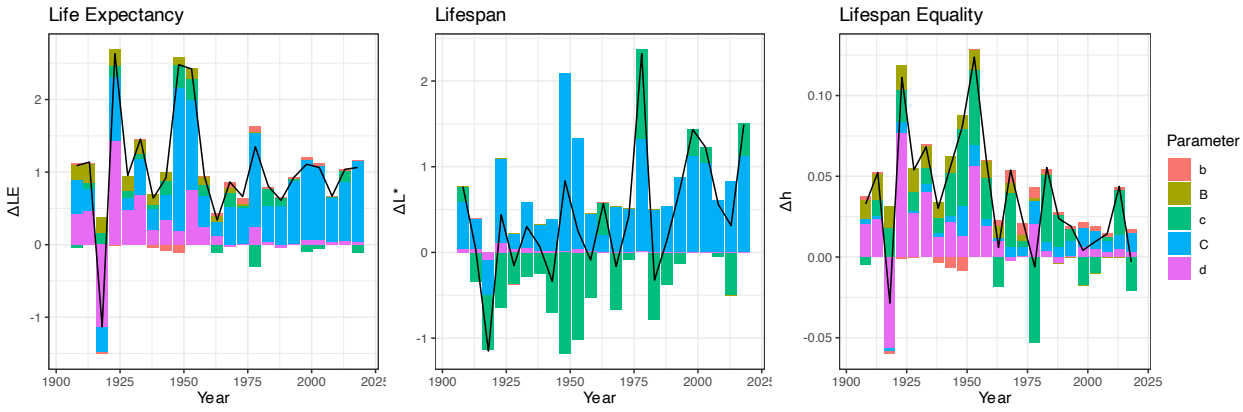

Note: The decomposition shown by the colored bars indicate the part of the period on period change due to each parameter. Black lines indicates the actual change. The median of the posterior distribution for each parameter is used throughout.

4 International Results

BPLE is a useful way of framing improvements in life expectancy but it is a synthetic variable constructed from multiple countries and so may possess different dynamics than any specific

nation. It also provides no insight into the mortality dynamics of countries who do not define best practice, especially for emerging markets far from this frontier.

For these reasons we turn our attention to understanding long-run mortality trends in our 34 countries which vary considerably in terms of income and population size. In 1950, the lowest life expectancy amongst our sample was in Nigeria (40.2 years) and the highest in New Zealand (72.6). By 2023, the lowest was again Nigeria (55.7) and the highest Japan (86.9). There are clear signs of convergence in life expectancy amongst our countries over this period. As shown in Figure 5, taking 1950 as the starting point, the correlation coefficient between initial life expectancy and subsequent increases is -0.85 (-0.97 amongst the high income economies and -0.40 amongst the large population countries). Thus, countries with the lowest life expectancy at the beginning of the sample tend to show the fastest subsequent growth therefore pointing towards convergence. Going beyond these life expectancy trends to understand the underlying drivers in terms of mortality changes, both past and future, across such a range of countries and time periods is the key contribution of this paper.

Figure 5: Convergence of parameters across countries

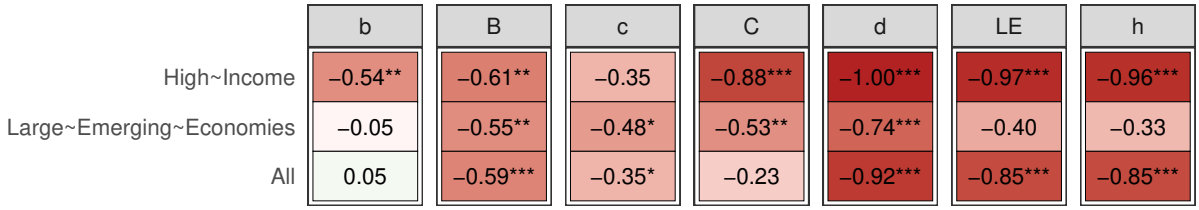

Note: Figure shows the correlation of the 1950 value with the change until the latest period for each Siler parameter, life expectancy at birth and lifespan equality at birth. If this correlation is negative, it indicates convergence over the time period as countries with more positive starting points have a less positive change over the period. In each case the Pearson correlation coefficient with statistical significance indicated by stars (*p<0.1; **p<0.05; ***p<0.01).

4.1 Drivers of Change

Figure 6 shows the decomposition of changes in life expectancy across our various mortality components. Focusing on the high income countries, initially (1918-43 and 1943-1968) the most important driver was reductions in age-independent mortality (d_t) followed by improvements in infant mortality (b_t and B_t). However, since 1968, it is later-life mortality (c_t and

C_t) that plays the most dominant role, accounting for over 80 per cent of overall changes in life expectancy. Within this the overwhelmingly important role is changes in C_t, with c_t explaining relatively little. In other words, shifts in the modal age of death not changes in the speed of ageing explain life expectancy gains in high income countries. While compressions in mortality are occurring they are relatively small and there is considerable variation across countries in whether the speed of ageing is improving or deteriorating. The final column lists the current estimate of the modal age of death which exceeds 85 years (an earlier limit suggested by Fries (1980) and Olshansky and Carnes (2019)) in all but one (Russia) of our high income countries. Japan, at 92 years, is estimated to have the highest modal age of death.

Turning to the large population emerging economies reveals a somewhat different pattern. In these the most important driver of life expectancy improvements between 1943 and 1993 was reductions in infant mortality, explaining nearly half of all improvements, followed by reductions in age independent mortality. However, over the most recent period (1993-2018) it is declines in later-life mortality that are most important in all but four countries. Once again it is delays in the modal age of death rather than changes in the speed of ageing that dominate. In all but three countries (Congo, Nigeria and Pakistan) the modal age of death is now estimated as above 80 years of age.

While Figure 7 decomposes life expectancy gains into proportions, we can also calculate the absolute number of years of life expectancy attributable the each parameter. This reveals a similar pattern of an increasing importance of C_t across both high income and large population emerging economies, as shown in Appendix A.9. This also reveals that the number of extra years attributable to increasing C_t has been fairly constant over time in high income economies (2.78 years from 1968-1993 and 2.77 years from 1993-2018) while gains from other parameters have decreased over time. In contrast, for the emerging economies in our sample the number of extra years attributable to increasing C_t has increased continuously (2.49 from 1968-2993 and 3.07 from 1993-2018).

Comparing estimates across countries shows the strongest evidence for convergence (in terms of the correlations between initial levels and subsequent growth shown in Figure 5) in age-independent mortality, d, modest convergence in C and B and relatively little in c and b

Figure 6: Proportion of life expectancy gains due to c and C across countries over time

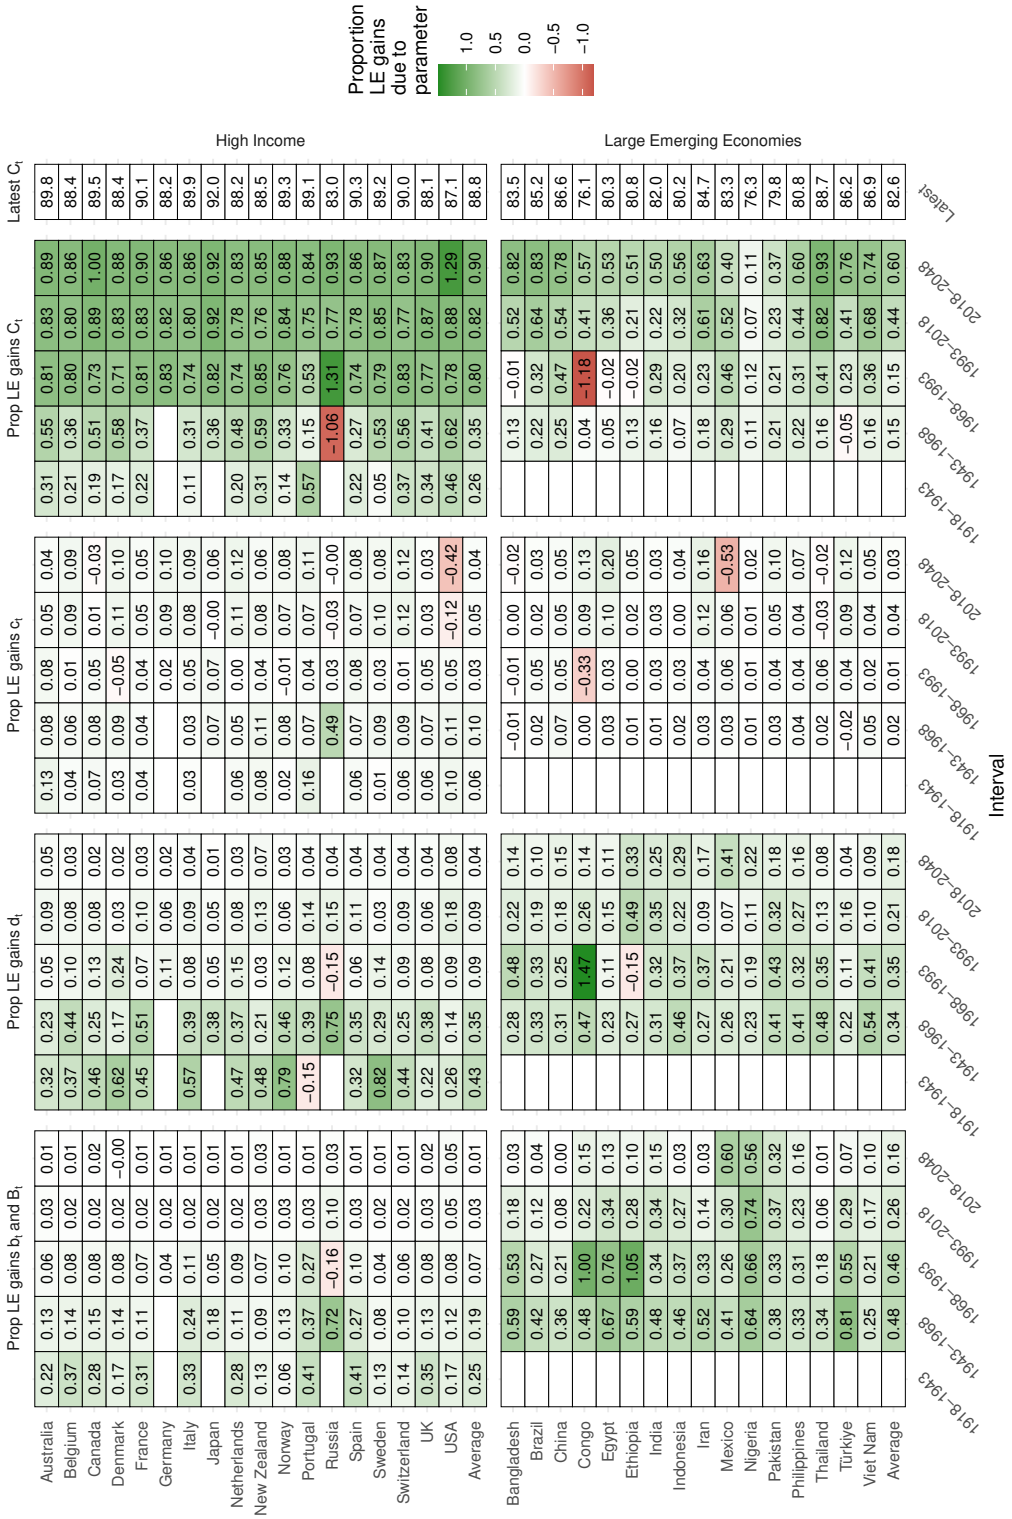

Note: Figure shows the proportion of increases in life expectancy at birth due to each parameter in Equation 2. Where data are unavailable for a given period, cells are left blank. The final column shows the value of C in the latest in-sample period. The posterior median for each parameter is used throughout.

Figure 7: Evolution of c and C over time

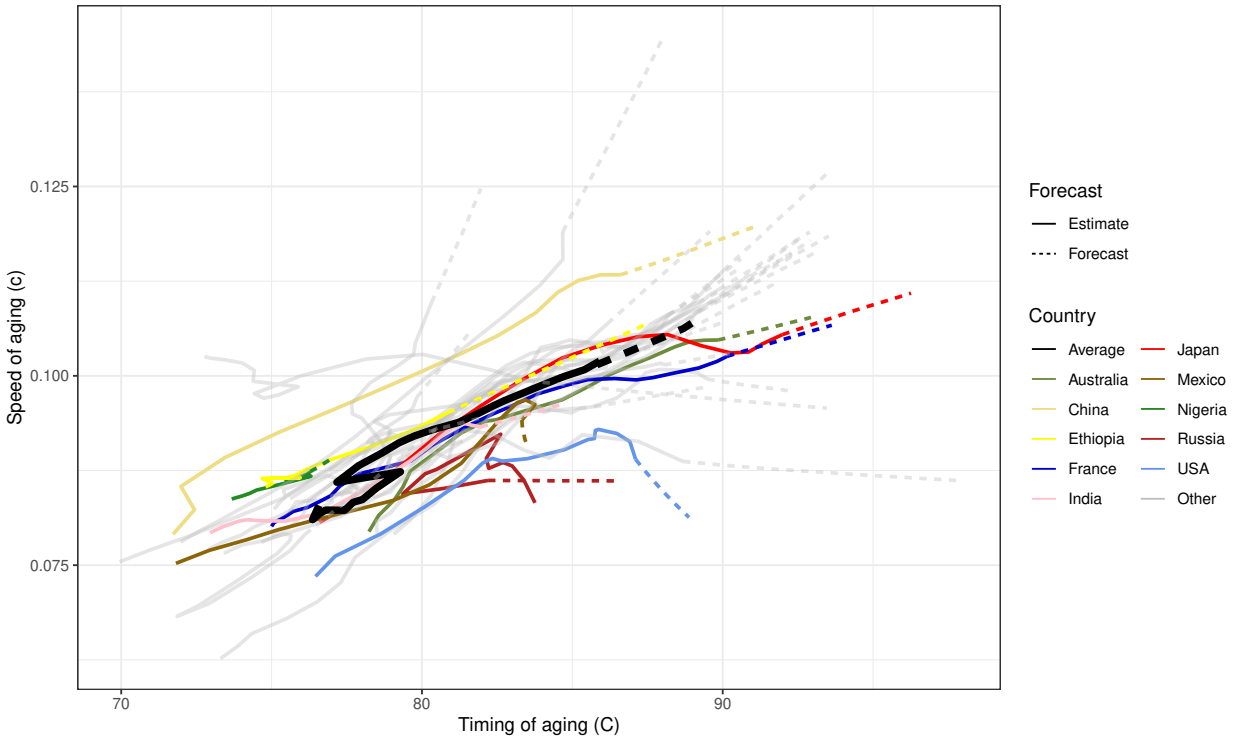

Note: Figure shows the path of c and C for each country over time in (C, c) space. Selected countries are shown in color with all other countries appearing in grey. Solid lines indicate estimated values of the parameters and dashed lines an out-of-sample forecast. The median of the posterior/predictive distribution for each parameter is used throughout.

(Aksan and Chakraborty, 2023). Convergence always tends to be strongest within the high income countries rather than the large population nations or the whole sample.

Figure 7 shows the co-evolution of c_t and C_t for each country with solid lines showing the estimated historical dynamics and dashed lines the forecast. This shows a clear historical positive correlation, broadly following a common international pattern. Consistent with Figure 6, the percentage increases in C are larger than for c and in the most recent periods the general tendency is for greater improvements in C than c.

A positive correlation between c_t and C_t is also found in the cross section across countries so that in general countries which have postponed the age of death the most also have achieved the largest compression of mortality. However, there are important differences which has relevance for BPLE. Japan has the highest estimated C_t but not for c_t. Recent improvements in Japanese life expectancy are driven much more by C_t with relatively little

improvement in c_t. The dynamics of BPLE differ from general international experience. To the extent compressions of mortality reflects improvements in how we age questions whether "Best Practice Life Expectancy" really does denote best practice.

4.2 Lifespan Equality and Lifespan

Given relatively few people die at exactly average life expectancy, measures of lifespan equality are an important demographic variable. With more of life occurring after retirement age this uncertainty about age at death has important implications for financial planning.

To measure lifespan equality we follow Aburto et al. (2020) and use the negative of the logarithm of the normalised entropy of the survival rate, see Appendix A.7. Figure 8 shows our estimates (and projections) for lifespan equality at birth $h(0)$ and clear evidence of a global trend towards improving lifespan equality. Every single one of our countries in the long sample experiences an improvement from 1900 to 1950 and all our countries show a further increase in the shorter sample. There is evidence of convergence across the whole group (e.g. the negative correlation coefficient of -0.84 between initial lifespan equality and subsequent growth shown in Figure 5) although once again this is strongest amongst the high income countries.

A key longevity concept is the maximal possible length of life, sometimes referred to as "lifespan". Whilst the existence of such a limit, its magnitude and even its definition is the subject of much controversy (Dong et al., 2016; Barbi et al., 2018; Dang et al., 2023). We follow Olshansky et al. (1990) and construct a statistical measure of the upper bound of age at death (UBAD) based on the age at which only 0.1 percent of the population can be expected to survive. This serves as an empirical measure of the oldest old in society. For some countries and time periods the Human Mortality Database interpolates a logistic model to infer deaths above 90 adding potential noise and misspecification to our estimates. For robustness purposes we therefore compute our results for 1, 5 and 10 percent survival probabilities as well to compare with these 0.1 percent outcomes (see Figure A.11 for a direct comparison). Whilst we show results for 0.1 percent similar qualitative findings held at all levels.

Figure 8: Evolution of $h(0)$ over time

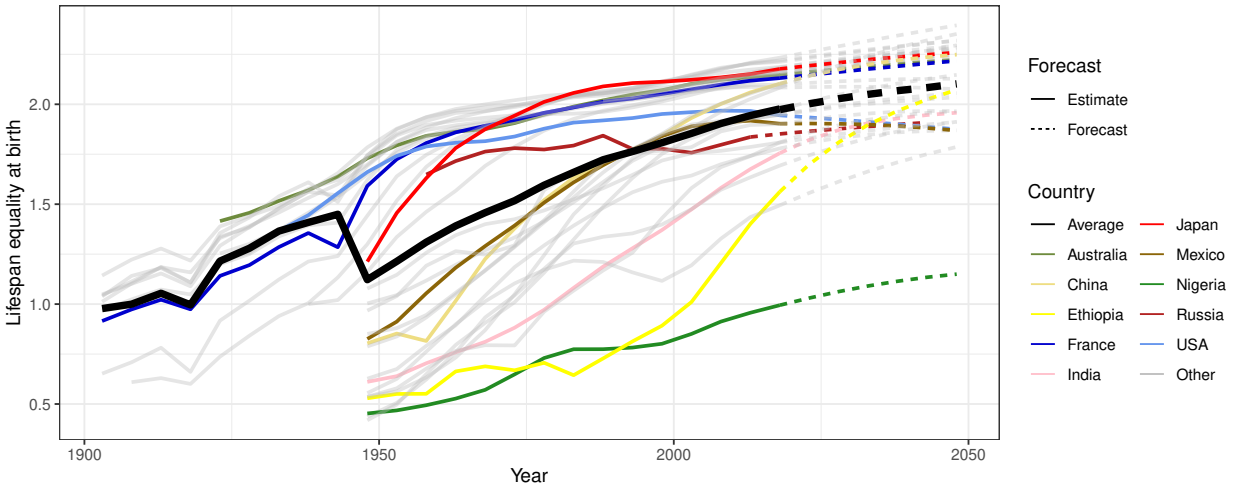

Note: This Figure shows lifespan equality at birth for each country over time. A higher value of $h(0)$ indicates more equality in expected lifespan. Selected countries are shown in color with all other countries appearing in grey. Solid lines indicate estimated values of the parameters while dashed lines indicate that the path is an out-of-sample forecast. The median of the posterior(predictive) distribution for each parameter is used throughout.

The country with the highest current UBAD is estimated to be Japan at 110.8.[8] With gains in life expectancy increasingly driven by improvements in later-life mortality, c_t and C_t, the gradient of UBAD relative to life expectancy is increasing. Currently, for most countries, every one year increase in life expectancy at birth is translates into an approximate 8 month increase in UBAD.

5 Future Projections

We now use our model to make projections to see if past trends are expected to continue. Obviously this approach is entirely statistical and takes no account of any biological factors that may constrain improvements in later-life mortality and the potential length of human life. Its validity also depends on the out-of-sample performance of our model. Evidence on this is given in Appendix A.1 where we perform various out of sample forecasting exercises from 1965 onwards. Overall our model performs well, often outperforming other popular models (such as a variety of Lee-Carter models).

[8]Using alternative thresholds, Japan also has the highest UBAD which values of 99.86 (10%), 102.36 (5%) or 104.44 (1%).

We utilise our Bayesian approach to quantify the uncertainty around our projections arising from uncertainty regarding the Siler parameters. We do so by taking into account the posterior distribution of the in-sample parameters and the covariance matrices of shocks. Let M be the number of draws from the in-sample posterior distribution provided by the MCMC sampling. We thus have M draws from the posterior for $\alpha_{\theta,T}$, θ_T, $\sigma_{\epsilon,\theta}$ and $\sigma_{\xi,\theta}$ enabling us to take into account any potential covariance between the different parameters i.e if a higher C_t is likely only when c_t is low. In order to approximate the predictive distribution we simulate forward from each of the particles of this sampled posterior drawing from the distribution of the shocks $\epsilon_{\theta,T+1}$ and $\xi_{\theta,T+1}$, which we assume independent. We draw these N times, so our predictive distribution has $N \times M$ particles. With this sampled predictive distribution we can quantify the forecast distributions for future values of the Siler parameters (and for life expectancy and other derived statistics).

5.1 Best Practice Life Expectancy

Figure 9 shows our model forecast for BPLE along with 70% and 95% percent predictive intervals. Because our model assumes both the level of mortality rates and their rate of change follow random walks the predictive intervals are large and increasing. As such the projections inevitably provide support for "futurists", "optimists" and "realists" to use the nomenclature of Carnes and Olshansky (2007) as well as "pessimists". However most striking is that the central forecast is for a continuation of current trend increases in best practice of 1.94 years per decade between now and 2050, remarkably close to the historical trend noted in Oeppen and Vaupel (2002).

It is important to emphasise that the linearity in Figure 8 is not a built in assumption of our approach. We model mortality rates directly rather than life expectancy and so the projection is not a simple linear extrapolation. Further, we assume mortality rates depend non-linearly on our core parameters and by allowing for different trends at each age as well as considerable time variation in each of these trends we allow for a rich range of potential future trajectories.

The reason why, despite these modelling choices, we end up with a simple linear projection is the growing importance of changes in later-life mortality and C_t in particular. Our estimate

Figure 9: Forecast LE for Best Practice case

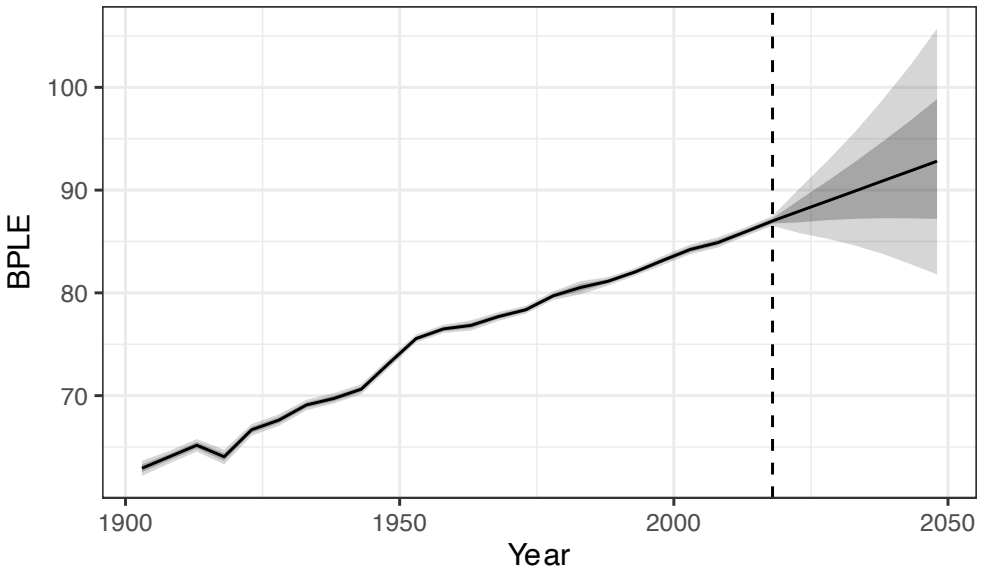

Note: Best Practice Life Expectancy at birth is calculated for each particle sampled from the joint posterior (predictive) distributions of each parameter of the dynamic Siler model, thus approximating the posterior (predictive) distribution of life expectancy based on the model. The solid line indicates the median with the shaded areas indicating the 70% and 95% credible intervals.

of the trend growth in C_t, α_C, has settled at a stable 0.01, implying a 1% increase in C (modal age of death) every five years or roughly 2 years every decade. With little projected change in the other BPLE mortality parameters the result is a relatively simple linear projection.

5.2 International Evidence

Looking at individual countries again reveals a different pattern. Whereas the projection for BPLE is for a constant rate of improvement there is broad evidence of a slowdown in life expectancy gains at a country level. Across all countries (see Table A.2 in Appendix A.1) the average rate of growth slows from around two years a decade over the period 1988 to 2018 to a projected 1.3 years for the period 2018 to 2048. Amongst high income countries, growth shows a modest decline, from 1.3 years a decade to nearly 1.1.[9] In large population emerging economies the rate of increase goes from just under 3 years to just over 1.5. Out of our total of 34 countries, only four see a projected increase (Netherlands and Norway to a small extent, as well as Russia and Congo).

[9]Although in Japan, which currently, defines BPLE the rate of increase remains broadly unchanged

However, whilst trend growth in life expectancy is projected to slow it remains positive in all countries (except for Mexico which is predicted to see life expectancy effectively remain unchanged). This slowdown is not surprising given that Olshansky and Carnes (2019) show that maintaining a constant trend growth requires an acceleration in the rate of improvements in later-life mortality. Examining our model projections reveals that the rate of improvement in mortality at the oldest ages is indeed increasing but not at a fast enough rate to deliver no change in the trend rate of growth.

Also impacting the trend rate of growth is the changing sensitivity of life expectancy with respect to later-life mortality (see Figure 10, evaluated for BPLE). The left hand panel shows the sensitivity of remaining life expectancy at different ages to C_t. This has been increasing substantially at older ages precisely because the modal age itself is now above 80 in so many of our countries. The result is that for a constant trend increase in C_t gains to remaining life expectancy at higher ages is increasing. Combining this with a relatively constant rate of increase in C_t helps support future longevity gains.

Figure 10: Gradient of LE with respect to later-life mortality

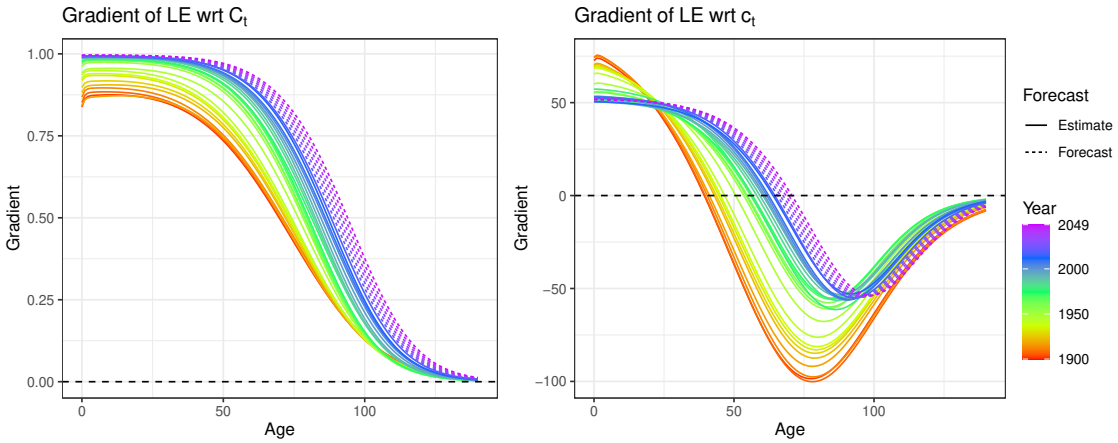

Note: Figures shows the gradient of remaining life expectancy at each age with respect to the parameters c and C, over time for the best practice country. The color of the lines indicates the year, solid lines indicate estimated values of the parameters while dashed lines indicate that the path is an out-of-sample forecast. The median of the posterior(predictive) distribution for each parameter is used throughout.

The sensitivity of life expectancy to c_t is more complicated (see Aburto et al. (2020) and the right hand panel in Figure 10). As the modal age of death increases, pre-modal older ages benefit more from changes in the speed of ageing than previously and so the negative impact of increases in c_t on life expectancy is lessened. Therefore a shift towards C_t as a

driver of mortality improvements and these changes in derivatives combine to magnify the impact of improvements in mortality at older ages on life expectancy.

As shown in Figure 6, the projection is across nearly all countries for a growing importance of later life mortality and C in particular. In all but 3 countries (Mexico, Nigeria and Pakistan) the projection is for later life mortality to explain the majority of changes in life expectancy. Across high income countries improvements in C are projected to account for 80 percent or more of all life expectancy gains and in the large population emerging markets it accounts for the majority of gains except in Nigeria, Pakistan and Phillipines.

Whilst life expectancy is projected to continue to rise, albeit at slower rates, the projections also point to a major shift. Whereas life expectancy and most of the mortality parameters showed broad signs of convergence in earlier periods this is not the expectation going forward. That holds both across all countries and within both high income countries (as noted by Nigri et al. (2022)) and the large population emerging markets. Amongst both groups the projection is that those countries with the highest life expectancy will see the largest future increases, with this being especially pronounced amongst high income countries. The U.S, U.K, Canada, New Zealand, Germany and Netherlands are all forecast to have the slowest life expectancy increase and all sit in the bottom half in terms of their current life expectancy amongst high income nations. Amongst the large population emerging markets the evidence of divergence is somewhat weaker although Mexico, Egypt, Indonesia and the Phillippines are projected to have the slowest future growth and all sit in the bottom half in terms of current life expectancy. Meanwhile Thailand and China are predicted to have the third and fifth fastest growth in life expectancy despite having the first and second highest current life expectancy amongst this group.

The United States in particular is a strong outlier showing continued trend improvements in the modal age of death but declining c_t, indicative of higher mortality rates below this modal age (consistent with (Case and Deaton, 2015) and their evidence of rising mortality rates in midlife).

5.3 Lifespan Equality and Lifespan

Figure 8 shows projections for lifespan equality at birth $h(0)$. In the majority of countries lifespan equality continues to improve whilst three countries (Canada, Mexico and the United States) are projected to see modest declines.

However, again projections point to shifting trends. The first is a substantial slowdown in improvements in lifespan equality. The reason for this diminishing trend is the increasing dominance of C_t relative to c_t in driving later-life mortality. Increases in c_t lead to a compression of mortality around the mode but as trend growth in c_t has declined so the projection is for a slowing in improvements in equality.

The second changing trend is a switch from convergence to divergence amongst our sample, once more particularly amongst high income countries. The largest improvements in lifespan equality are predicted to occur in those countries with already high levels of equality.

Our model also predicts, in the main, continued increases in UBAD. Based on our Siler specification the derivative of UBAD with respect to C_t is almost exactly one but is negative with respect to c_t. With a general tendency across our countries for continued projected increases in C_t and an attenuation of improvements in c_t the relationship between UBAD and C is getting stronger. If there is no strong compression of mortality but continued increase in the modal age of death then the age of the oldest person will continue to increase.

Once again though the United States is a marked outlier. The American combination of rising C_t and declining c_t leads to a steeper slope than other countries and the projection of limited life expectancy gains but a sharp increase in UBAD-C_t. Growing inequality in US life expectancy is leading to very different trends in the mean, the modal and the upper bound of age at death.

6 Conclusion

Our focus is on a robust but flexible low parameter model which allows for differential mortality trends at different ages and across countries. This allows us to provide a comprehensive and integrated assessment of past and likely future trends in later-life mortality, we well as

their implications for future longevity across a wide range of countries.

Our findings show that a) later-life mortality is expected to continue to be the main driver of life expectancy, both in our sample of high income countries and now also in the majority of our large population emerging markets; b) increasingly, the dominant factor driving later-life mortality is changes in the modal age of death rather than any change in the speed of ageing/compression of mortality; c) Best Practice Life Expectancy displays different characteristics from life expectancy in most countries as it is currently driven only by shifts in the modal age of death; d) although we find evidence of an increase in the rate of improvement in later-life mortality it is not enough to prevent a slowdown in future life expectancy gains; e) because of the shifting modal age of death and limited changes in the speed of ageing we find no evidence of any approaching limit to the upper bound of age at death; f) we find a reversal away from convergence in key longevity indicators to one of divergence both across countries and within high income nations especially as good performers continue to outperform poor performers; and g) the United States is a marked outlier and is expected to see only modest gains in average life expectancy, continued growth in the upper bound of age at death but widening lifespan inequality.

Our results and insights are of course limited by our empirical approach. Firstly they inevitably depend on our chosen Siler function and assumptions on parameter dynamics. Support for this approach is given by the ability of our model to track past changes, conform with past findings and in particular its credible out of sample forecast performance.

Because the limits to average life expectancy and the upper bound for age of death are biologically determined any future projections based on statistical analysis of past data has obvious limitations. However, our finding that there are no indications of sharp changes in major historical trends but instead more smooth adjustments is indicative that whatever these biological limits are they are not yet exerting a dramatic effect on demographic trends.

Our analysis can shed no light on the impact of Covid-19 due to the fact this occurred outside our sample and our use of five year intervals. Clearly Covid-19 will have impacted the short run behaviour of our model parameters, and the short term impact of the pandemic on mortality has been widely studied (Hardy et al., 2021; Leggat-Barr et al., 2021; Morwinsky et al., 2021; Aburto et al., 2022; Schöley et al., 2022). The duration of Covid-19 and any

potential impact on all-cause mortality resulting from lockdowns, pressures on health systems or long Covid will determine whether its impact extends beyond a specific five year interval.

In predicting further gains to life expectancy and lifespan for the majority of countries our paper clearly points to the continued need for society to adapt to longevity. In identifying divergent trends across countries our results also point to the notion that progress is not automatic. For those countries experiencing a deterioration, especially in the United States, it is important to identify and correct the factors behind this divergence.

In detecting a diminishing future role for c_t and a rising one for C_t our results suggest that future gains to life expectancy will be driven by shifts in the timing rather than speed of ageing. To the extent that compressions of mortality bring about compressions of morbidity (Fries, 1980) this is concerning. Our predicted combination of rising life expectancy, increases in the oldest age at death and stalling improvements in lifespan equality and the speed of ageing all point to the urgency of ensuring that life isn't just longer but also healthier for longer. The danger is that life expectancy is being extended by interventions that keep people alive but not healthier for longer. A number of studies show the value of compressing morbidity relative to further gains in life expectancy (Goldman et al., 2013; Olshansky, 2016; Scott et al., 2021). These results and our model forecasts point to the importance of potential developments in geroscience and the biology of ageing (Campisi et al., 2019) or broader public health and social policy programs (Wong et al., 2023) aimed at achieving healthy longevity as urgent priorities.

References

Aburto, J. M., Schöley, J., Kashnitsky, I., Zhang, L., Rahal, C., Missov, T., Mills, M., Dowd, J. and Kashyap, R. (2022), 'Quantifying impacts of the covid-19 pandemic through life-expectancy losses: a population-level study of 29 countries', *International Journal of Epidemiology* **51**(1), 63–74.

Aburto, J. M., Villavicencio, F., Basellini, U., Kjærgaard, S. and Vaupel, J. W. (2020), 'Dynamics of life expectancy and life span equality', *Proceedings of the National Academy of Sciences* **117**(10), 5250–5259.

Aksan, A.-M. and Chakraborty, S. (2023), 'Life expectancy across countries : Convergence, divergence and fluctuations', *World Development* **168**.

Alexander, M., Zagheni, E. and Barbieri, M. (2017), 'A flexible bayesian model for estimating subnational mortality', *Demography* **54**(6), 2025–2041.

Alexopoulos, A., Dellaportas, P. and Forster, J. J. (2019), 'Bayesian forecasting of mortality rates by using latent gaussian models', *Journal of the Royal Statistical Society Series A: Statistics in Society* **182**(2), 689–711.

Aliverti, E., Mazzuco, S. and Scarpa, B. (2022), 'Dynamic modelling of mortality via mixtures of skewed distribution functions', *Journal of the Royal Statistical Society Series A: Statistics in Society* **185**(3), 1030–1048.

Barbi, E., Lagona, F., Marsili, M., Vaupel, J. and Wachter, K. (2018), 'The plateau of human mortality : Demography of longevity pioneers', *Science* **360**, 1459–1461.

Basellini, U. and Camarda, C. G. (2019), 'Modelling and forecasting adult age-at-death distributions', *Population studies* **73**(1), 119–138.

Bergeron-Boucher, M.-P., Ebeling, M. and Canudas-Romo, V. (2015), 'Decomposing changes in life expectancy: Compression versus shifting mortality', *Demographic Research* **33**, 391–424.

Betancourt, M. (2017), 'A conceptual introduction to hamiltonian monte carlo', *arXiv preprint arXiv:1701.02434* .

Booth, H., Maindonald, J. and Smith, L. (2002), 'Applying lee-carter under conditions of variable mortality decline', *Population Studies* **56**(3), 325–336.

Bravo, J. M., Ayuso, M., Holzmann, R. and Palmer, E. (2021), 'Addressing the life expectancy gap in pension policy', *Insurance: Mathematics and Economics* **99**, 200–221.

Cairns, A. J., Blake, D. and Dowd, K. (2006), 'A two-factor model for stochastic mortality with parameter uncertainty: theory and calibration', *Journal of Risk and Insurance* **73**(4), 687–718.

Campisi, J., Kapahi, P., Lithgow, G. J., Melov, S., Newman, J. C. and Verdin, E. (2019), 'From discoveries in ageing research to therapeutics for healthy ageing', *Nature* **571**(7764), 183–192.

Canudas-Romo, V. (2010), 'Three measures of longevity: Time trends and record values', *Demography* **47**(2), 299–312.

Canudas-Romo, V., Booth, H. and Bergeron-Boucher, M.-P. (2019), 'Minimum death rates and maximum life expectancy: The role of concordant ages', *North American Actuarial Journal* **23**(3), 322–334.

Carnes, B. and Olshansky, S. (2007), 'A realist view of aging, mortality, and future longevity', *Population and Development Review* **33**(2), 367–381.

Case, A. and Deaton, A. (2015), 'Rising morbidity and mortality in midlife among white non-hispanic americans in the 21st century', *Proceedings of National Academy of Sciences* **112**(49), 15078–15083.

Cheung, S. L. K., Robine, J.-M., Tu, E. J.-C. and Caseli, G. (2005), 'Three dimensions of the survival curve: Horizontalization, verticalization, and longevity extension', *Demography* **42**(2), 243–258.

Christensen, K., Doblhammer, G., Rau, R. and Vaupel, J. W. (2009), 'Ageing populations: the challenges ahead', *The Lancet* **374**(9696), 1196–1208.

Colchero, F., Aburto, J. M., Archie, E. A., Boesch, C., Breuer, T., Campos, F. A., Collins, A., Conde, D. A., Cords, M., Crockford, C. et al. (2021), 'The long lives of primates and the 'invariant rate of ageing'hypothesis', *Nature Communications* **12**(1), 1–10.

Colchero, F. e. a. (2016), 'The emergence of longevous populatons', *Proceedings of National Academy of Sciences* **113**(48), E7681–E7690.

Congdon, P. (1993), 'Statistical graduation in local demographic analysis and projection', *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **156**(2), 237–270.

Crimmins, E. (2015), 'Lifespan and healthspan : Past, present, and promise', *Gerontologist* **55**(6), 901–11.

Dang, L. H. K., Camarda, C. G., Meslé, F., Ouellette, N., Robine, J.-M. and Vallin, J. (2023), 'The question of the human mortality plateau: Contrasting insights by longevity pioneers', *Demographic Research* **48**, 321–338.

De Jong, P. and Boyle, P. P. (1983), 'Monitoring mortality: a state-space approach', *Journal of Econometrics* **23**(1), 131–146.

De Jong, P. and Tickle, L. (2006), 'Extending lee-carter mortality forecasting', *Mathematical Population Studies* **13**(1), 1–18.

Dellaportas, P., Smith, A. F. and Stavropoulos, P. (2001), 'Bayesian analysis of mortality data', *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **164**(2), 275–291.

Demetrius, L. (1978), 'Adaptive value, entropy and survivorship curves', *Nature* **275**(5677), 213–214.

Dodd, E., Forster, J. J., Bijak, J. and Smith, P. W. (2018), 'Smoothing mortality data: the english life tables, 2010–2012', *Journal of the Royal Statistical Society Series A: Statistics in Society* **181**(3), 717–735.

Dong, X., Milholland, B. and Vijg, J. (2016), 'Evidence for a limit to human lifespan', *Nature* **538**, 257–259.

Eggleston, K. N. and Fuchs, V. R. (2012), 'The new demographic transition : Most gains in life expectancy now realized late in life', *Journal of Economic Perspectives* **26**(3), 137–156.

Fries, J. F. (1980), 'Aging, natural death, and the compression of morbidity', *New England Journal of Medicine* **303**(3), 130–135.

Fung, M. C., Peters, G. W. and Shevchenko, P. V. (2017), 'A unified approach to mortality modelling using state-space framework: characterisation, identification, estimation and forecasting', *Annals of Actuarial Science* **11**(2), 343–389.

Gelman, A. and Rubin, D. B. (1992), 'Inference from iterative simulation using multiple sequences', *Statistical science* **7**(4), 457–472.

Goldman, D. P., Cutler, D., Rowe, J. W., Michaud, P.-C., Sullivan, J., Peneva, D. and Olshansky, S. J. (2013), 'Substantial health and economic returns from delayed aging may warrant a new focus for medical research', *Health affairs* **32**(10), 1698–1705.

Hardy, O. J., Dubourg, D., Bourguignon, M., Dellicour, S., Eggerickx, T., Gilbert, M., Sanderson, J.-P., Scohy, A., Vandael, E. and Decroly, J.-M. (2021), 'A world apart', *Demographic Research* **45**, 1011–1040.

Harvey, A. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press.

Heligman, L. and Pollard, J. H. (1980), 'The age pattern of mortality', *Journal of the Institute of Actuaries* **107**(1), 49–80.

Heuveline, P. (2022), 'Global and national declines in life expectancy: An end-of-2021 assessment', *Population and development review* **48**(1), 31–50.

Hoffman, M. D., Gelman, A. et al. (2014), 'The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo.', *J. Mach. Learn. Res.* **15**(1), 1593–1623.

Horiuchi, S., Ouellette, N., Cheung, S. L. K. and Robine, J.-M. (2013), 'Modal age at death: lifespan indicator in the era of longevity extension', *Vienna Yearbook of Population Research* pp. 37–69.

Hyndman, M. R. J., Booth, H., Tickle, L. and Maindonald, J. (2012), 'Package 'demography".

Hyndman, R. J. and Ullah, M. S. (2007), 'Robust forecasting of mortality and fertility rates: A functional data approach', *Computational Statistics & Data Analysis* **51**(10), 4942–4956.

Keyfitz, N. and Caswell, H. (2005), *Applied mathematical demography*, Vol. 47, Springer.

Lee, R. and Carter, L. (1992), 'Modelling and forecasting u.s mortality', *Journal of the American Statistical Association* **87**(419), 659–671.

Lee, R. and Mason, A. (2017), 'Cost of aging', *Finance and Development)* **54**(1).

Leggat-Barr, K., Uchikoshi, F. and Goldman, N. (2021), 'Covid-19 risk factors and mortality among native americans', *Demographic Research* **45**, 1185–1218.

Lenart, A. and Vaupel, J. W. (2017), 'Questionable evidence for a limit to human lifespan', *Nature* **546**(7660), E13–E14.

Levantesi, S., Nigri, A., Piscopo, G. and Spelta, A. (2023), 'Multi-country clustering-based forecasting of healthy life expectancy', *Quality & Quantity* **57**(Suppl 2), 189–215.

Li, J., Lee, M. and Guthrie, S. (2021), 'A double common factor model for mortality projection using best-performance mortality rates as reference', *ASTIN Bulletin: The Journal of the IAA* **51**(2), 349–374.

Li, J. and Liu, J. (2019), 'A logistic two-population mortality projection model for modelling mortality at advanced ages for both sexes', *Scandinavian Actuarial Journal* **2019**(2), 97–112.

Liu, J. and Li, J. (2019), 'Beyond the highest life expectancy: Construction of proxy upper and lower life expectancy bounds', *Journal of Population Research* **36**, 159–181.

McNown, R. and Rogers, A. (1989), 'Forecasting mortality: A parameterized time series approach', *Demography* **26**(4), 645–660.

McNown, R. and Rogers, A. (1992), 'Forecasting cause-specific mortality using time series methods', *International Journal of Forecasting* **8**(3), 413–432.

Milevsky, M. (2020), 'Calibrating gompertz in reverse: what is your longevity-risk-adjusted global age?', *Insurance: Math. Econ.* **92**, 147–161.

Missov, T. I., Lenart, A., Nemeth, L., Canudas-Romo, V. and Vaupel, J. W. (2015), 'The gompertz force of mortality in terms of the modal age at death', *Demographic Research* **32**, 1031–1048.

Morwinsky, S., Nitsche, N. and Acosta, E. (2021), 'Covid-19 fatality in germany', *Demographic Research* **45**, 1355–1372.

Neal, R. M. et al. (2011), 'Mcmc using hamiltonian dynamics', *Handbook of Markov Chain Monte Carlo* **2**(11), 2.

Nigri, A., Barbi, E. and Levantesi, S. (2022), 'The relay for human longevity: country-specific contributions to the increase of the best-practice life expectancy', *Quality & Quantity* **56**(6), 4061–4073.

Oeppen, J. (2019), 'Life expectancy convergence among nations since 1820: separating the effects of technology and income', *Old and New Perspectives on Mortality Forecasting* pp. 197–219.

Oeppen, J. and Vaupel, J. W. (2002), 'Broken limits to life expectancy', *Science* **296**(5570), 1029–1031.

Olshansky, S. J. (2016), 'Articulating the case for the longevity dividend', *Cold Spring Harbour Perspectives Medicine* **6**.

Olshansky, S. J. and Carnes, B. A. (2019), 'Inconvenient truths about human longevity', *The Journals of Gerontology: Series A* **74**(Supplement_1), S7–S12.

Olshansky, S. J., Carnes, B. and Cassel, C. (1990), 'In search of methuselah : Estimating the upper limits to human longevity', *Science* pp. 634–640.

Olshansky, S. J., Wilcox, B., Demetrius, L. and Sanchez, H. (2024), 'Implausibility of radical life extension in humans in the 21st century', *Nature Ageing* .

Pedroza, C. (2006), 'A bayesian forecasting model: predicting us male mortality', *Biostatistics* **7**(4), 530–550.

Raferty, A., Chunn, J., Gerland, P. and Sevcikova, H. (2013), 'Bayesian probabilistic projections of life expectancy for all countries', *Demography* **50**, 777–801.

Sacher, S., Battaglia, L. and Hansen, S. (2021), 'Hamiltonian monte carlo for regression with high-dimensional categorical data', *arXiv preprint arXiv:2107.08112* .

Schöley, J., Aburto, J. M., Kashnitsky, I., Kniffka, M. S., Zhang, L., Jaadla, H., Dowd, J. B. and Kashyap, R. (2022), 'Life expectancy changes since covid-19', *Nature human behaviour* **6**(12), 1649–1659.

Schöley, J., Aburto, J., Kashnitsky, I., Kniffa, M., Zhang, L., Jaadla, J., Dowd, J. and Kashyap, R. (2022), 'Life expectancy changes since covid-19', *Nature Human Behaviour* pp. 1649–1659.

Scott, A. J. (2021), 'The longevity society', *Lancet Healthy Longevity* **2**(12), e820–827.

Scott, A. J., Ellison, M. and Sinclair, D. A. (2021), 'The economic value of targeting aging', *Nature Aging* **1**(7), 616–623.

Siler, W. (1979), 'A competing-risk model for animal mortality', *Ecology* **60**(4), 750–757.

Swift, J. (1726), *Gulliver's Travels*, Benjamin Motte.

Tang, S., Li, J. and Tickle, L. (2023), 'A hermite spline approach for modelling population mortality', *Annals of Actuarial Science* **17**(2), 243–284.

Thatcher, A. R. (1999), 'The long-term pattern of adult mortality and the highest attained age', *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **162**(1), 5–43.

United Nations (2024), 'World population prospects 2024: Summary of results', *UNDESA/POP/2024/TR/NO.9* .

Vallin, J., Meslé, F. and Divinagracia, E. (2008), 'Minimum mortality: A predictor of future progress?', *Population* **63**(4), 557–590.

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

Vehtari, A., Gelman, A., Simpson, D., Carpenter, B. and Bürkner, P.-C. (2021), 'Rank-normalization, folding, and localization: An improved r for assessing convergence of mcmc (with discussion)', *Bayesian analysis* **16**(2), 667–718.

Villegas, A., Kaishev, V. K. and Millossovich, P. (2015), Stmomo: An r package for stochastic mortality modelling, *in* '7th Australasian Actuarial Education and Research Symposium'.

Welsh, C. E., Matthews, F. E. and Jagger, C. (2021), 'Trends in life expectancy and healthy life years at birth and age 65 in the uk, 2008–2016, and other countries of the eu28: An observational cross-sectional study', *The Lancet Regional Health–Europe* **2**.

Wong, J., Fried, L. and Dzau, V. (2023), 'The global roadmap for healthy longevity: United states of america national academy of medicine consensus study report, 2022', *Journal of the Economics of Ageing* **24**.

Zuo, W., Jiang, S., Guo, Z., Feldman, M. W. and Tuljapurkar, S. (2018), 'Advancing front of old-age human survival', *Proceedings of the National Academy of Sciences* **115**(44), 11209–11214.

Online Appendices for *A Bayesian Model of Later Life Mortality Trends and Implications for Longevity*

A Appendices

A.1 Forecasting Appendix

To construct the out-of-sample forecasts we estimate models using data up to and including 1965-1969, 1975-1979, 1985-1989, 1995-1999, 2005-2009 and 2015-2019. We then forecast 6 periods (i.e. 30 years) ahead in each case and compare forecasts with outcomes where observed. We find our model outperforms a range of benchmarks when estimated on all available data and its performance is comparable when estimated using rolling windows.

To compare forecast performance we use the R packages *demography* (Hyndman et al., 2012) and *StMoMo* (Villegas et al., 2015).

1. LC (demo): standard Lee-Carter, under default parameters of *demography* package.
2. LC (SMM): standard Lee-Carter, under default parameters of *StMoMo* package.
3. LC (dt): Lee-Carter with the adjustment suggested by Lee and Carter (1992), as implemented in *demography* package.
4. LC (dxt): Lee-Carter with the adjustment suggested by Booth et al. (2002), as implemented in *demography* package.
5. LC (e0): Lee-Carter adjusted to fit in-sample life expectancy at birth, as implemented in *demography* package.
6. FDM: Basis function model for demographic data proposed by Hyndman and Ullah (2007), as implemented in *demography* package.
7. CBD: the two factor Cairns-Blake-Dowd (CBD) model introduced byCairns et al. (2006), as implemented in the *StMoMo* package.

We estimate each model on both the entire available sample and on a rolling window of the most recent 50 years of data (the Siler model is always estimated on all available data). The benchmark models generally have better out-of-sample performance with the rolling window approach. In contrast, the performance of the Siler model does not worsen if a longer sample is used.

A.1.1 Best Practice

Figure A.1 compares the out-of sample errors for life expectancy at birth in the BPLE case, measured as the root mean squared error (RMSE) at each forecasting horizon. Compared to models estimated on the full available sample, in the left panel, we see that from 10 years ahead onward, our model outperforms the benchmarks in forecasting BPLE. Compared to the models estimated on a 50 year rolling sample, in the right panel, we see that the Siler model outperforms the benchmarks in forecasting BPLE from 20 years ahead onward.

Figure A.1: Best practice errors for LE(0)

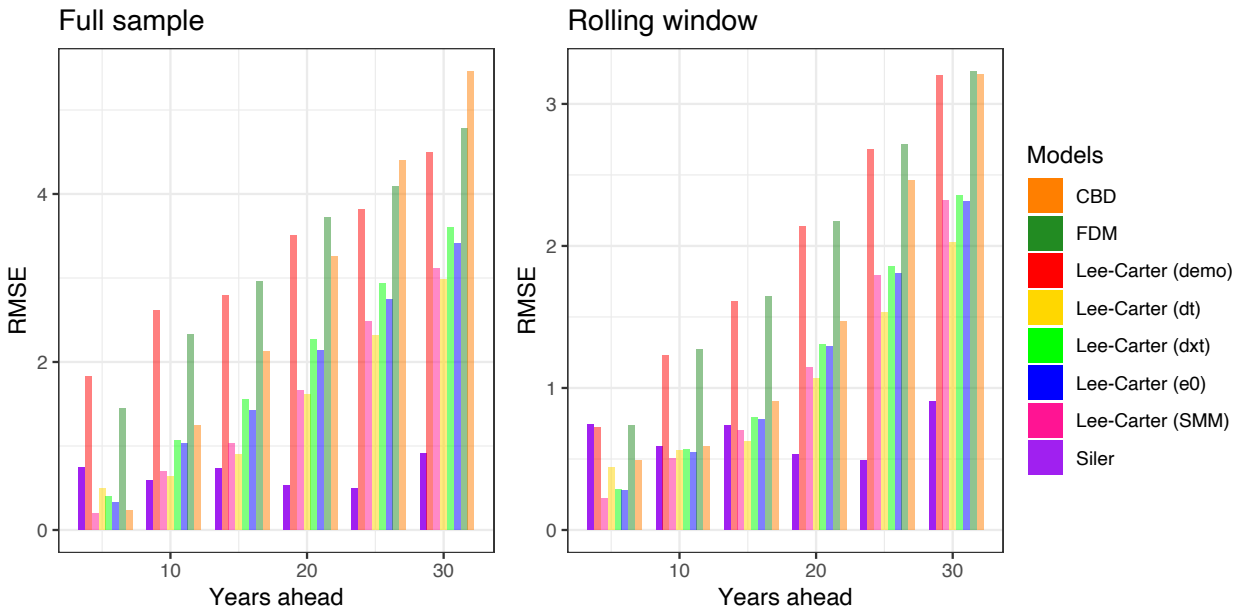

Note: Figure shows the root mean squared forecast error of a range of models when predicting BPLE. Models are either estimated on the full available sample (left panel) or on a 50 year rolling window (right panel). Performance is averaged over models estimated on data up to and including 1965-1969, 1975-1979, 1985-1989, 1995-1999, 2005-2009 and 2015-2019.

A.1.2 Country-level

Figure A.2 shows the historic and forecast changes in life expectancy across 33 countries in our sample.

Figure A.2: Increases in Life expectancy at birth

Note: Tables show the distribution across countries of increase in life expectancy at birth for historical and future decades, for high income and other countries. Columns show average increase in years per decade and cells number of countries in that range. As models are estimated in five year intervals, periods are referred to by their last year, i.e. 1989 refers to average mortality from 1985-1989.

Figure A.3: Country level forecast errors

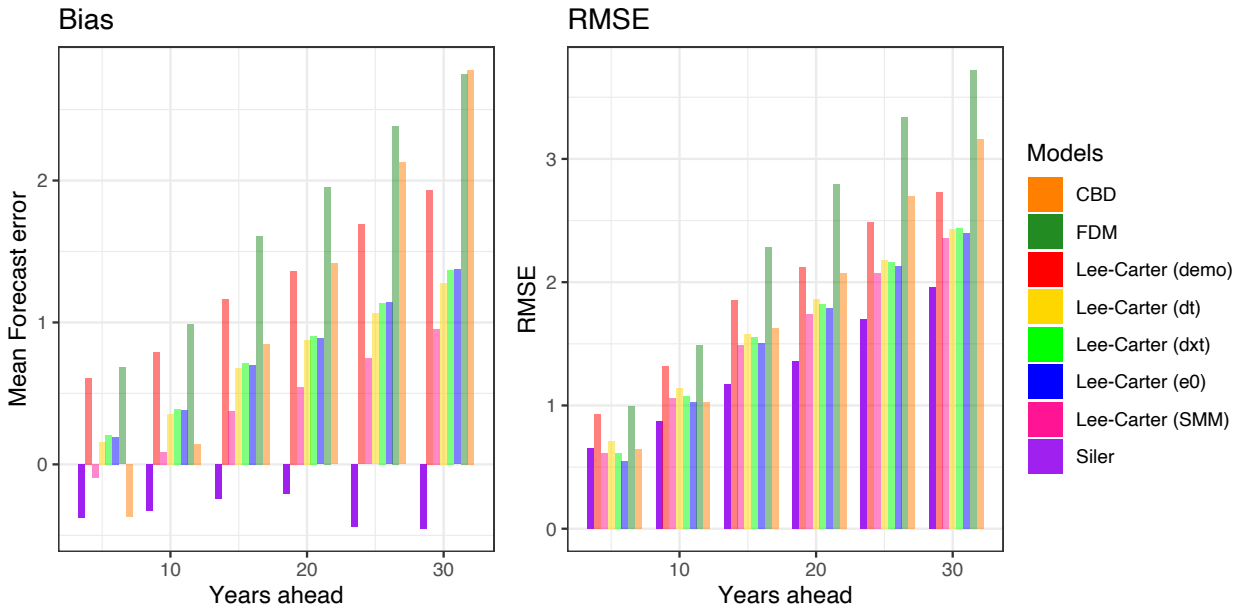

Note: Figure shows the mean forecast error (left panel) and root mean squared forecast error (right panel) of a range of models. The benchmark models are estimated on a 50 year rolling window (right panel). Performance is averaged over models estimated on data up to and including 1965-1969, 1975-1979, 1985-1989, 1995-1999, 2005-2009 and 2015-2019, and averaging across all countries with sufficient data availability.

Figure A.3 shows the forecast errors for each model across all countries, estimated on 50 year rolling windows.[10] We conduct this exercise on countries where at least 10 periods (50 years) worth of training data is available. The left hand panel shows the *mean* forecast error by horizon (across all countries and time periods). Here we see that the benchmark models have

[10]Results for benchmark models estimated on full available samples are qualitatively similar, but with larger errors.

a tendency to under-predict life expectancy and so generally have positive forecast errors. The Siler model has a smaller bias but slightly over-predicts at longer horizons. However, this over prediction is only driven by the models estimated in 1965-1969 and 1975-1979, after that the Siler model also has small positive mean forecast errors.

In the right hand panel we see the RMSE by horizon across all countries and time periods. The Siler model is competitive with the benchmark models and, as in the Best-Practice case, it sees less of a decrease in performance over longer horizons than its competitors.

Figure A.4 shows the out-of-sample RMSE for each country across the models estimated on a 50 year rolling window and averaging across all forecast horizons and estimation periods. For each country the best performing model, in terms of out-of-sample RMSE is highlighted in red, demonstrating that the Siler model is the best performing in more countries than any other model.

Figure A.4: Country level forecast errors by income group

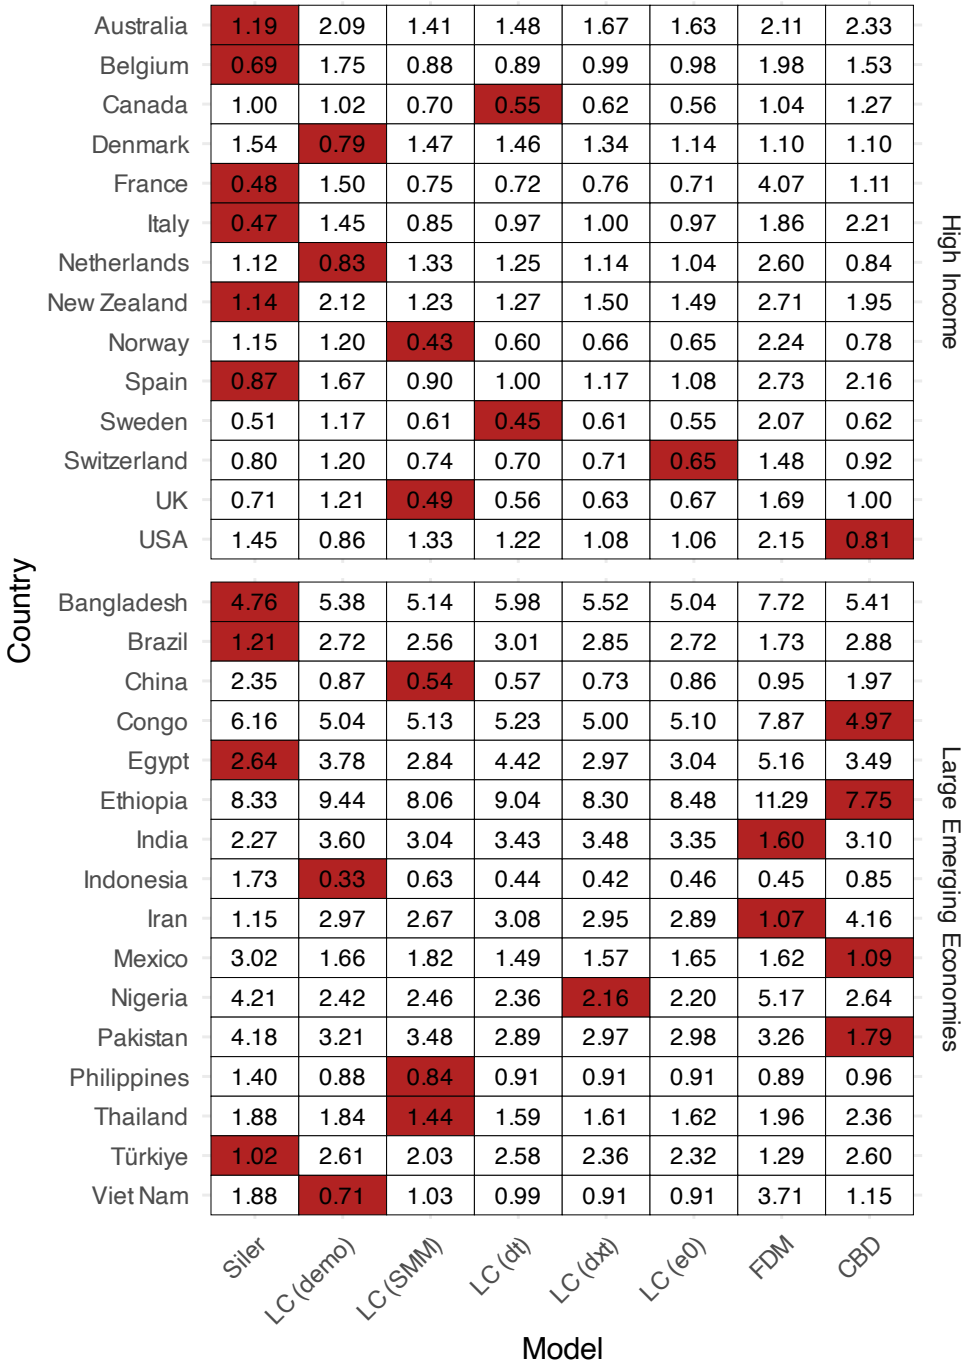

Note:

A.2 Comparing dynamic and static model results

To assess the value of our dynamic Bayesian approach, we compare it to a "static" model in which the Siler model is fitted to each period independently. Figure A.5 shows the in-

sample parameter estimates from the full dynamic version (in red) and a version in which each period is estimated separately. In most cases the parameter estimates are very similar with the dynamic model resulting in considerably smoother distributions for b_t and B_t in particular. We also see considerably more precise posterior distributions with the dynamic model, highlighting the the fact that nearby periods contain useful information.

Figure A.5: Dynamic and independent parameter estimates for best practice mortality (1900-2019)

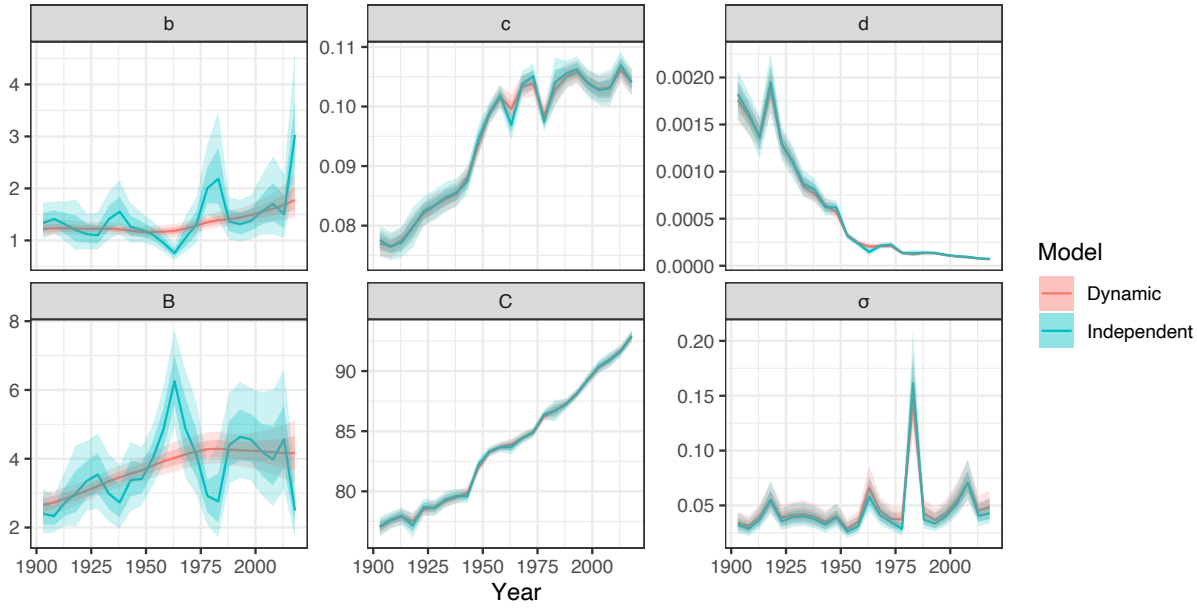

Note: Panels compare the posterior distributions for the latent parameters in Eq 1, in the dynamic model (in red) and in repeated independent estimations of a static model (in blue). The red distributions are thus the same as those in Figure 3. In both cases, the solid line indicates the median with the shaded areas indicating the 70% and 95% credible intervals.

A more extensive exploration of our model's forecasting performance is shown in Appendix A.1, but here we briefly compare the out-of-sample forecasting performance of the static and dynamic models. To construct the out-of-sample forecasts we estimate models using data up to and including 1965-1969, 1975-1979, 1985-1989, 1995-1999, 2005-2009 and 2015-2019. We then forecast 6 periods (i.e. 30 years) ahead in each case. We can then compare the forecasts to the true values, where we observe them.

For the dynamic model, this process is described in Section 5. For the independent model, we project forward the posterior medians independently estimated parameters assuming a constant trend of the log value of the parameters, i.e. we estimate $\ln \theta_t = \beta_0 + \beta_1 t + u_t$. This highlights a major short-coming of the static model as it is not clear how to quantify uncertainty around these forecasts.

As shown in Figure A.6, the static model performs worse in terms of out-of-sample RMSE across all horizons up and including to 25 years ahead.

Figure A.6: Comparing dynamic versus static model out-of-sample RMSE for LE(0)

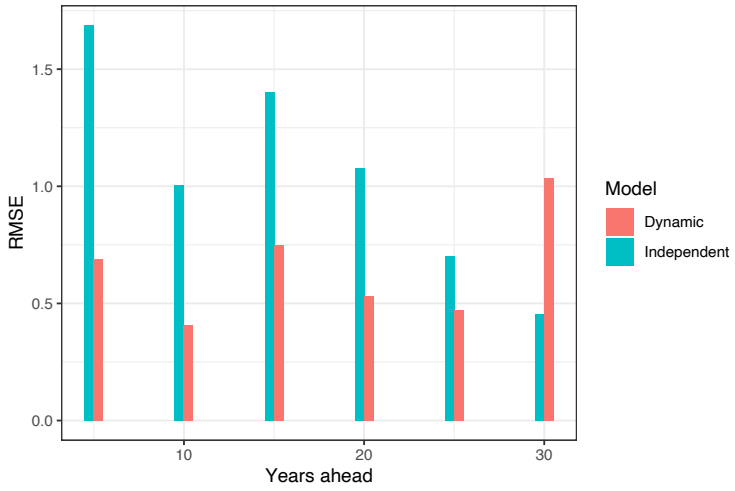

A.3 Alternative Parametric Mortality Functions

The Siler function used in this paper is only one of several possible parametric mortality functions. In this Section we explain our choice of Siler rather two commonly used alternatives - Gompertz-Makeham (GM) and Heligman-Pollard (HP). Figure A.7 compares the three functions graphically.

Figure A.7: Comparing parameteric mortality functions

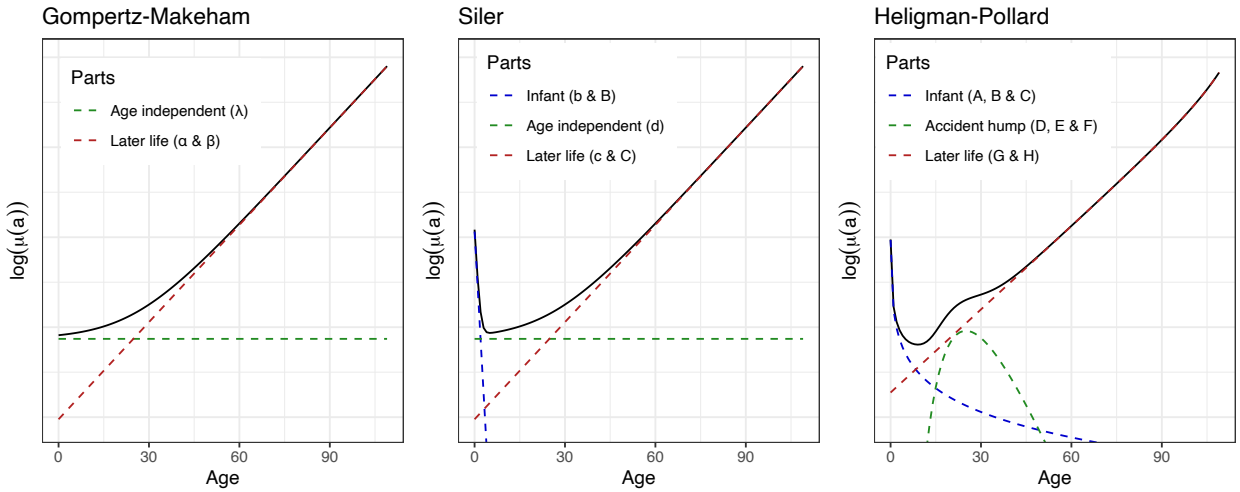

Note: Figure compares three different parametric mortality functions. In each case, age is shown along the horizontal axis and the log mortality rate at that age on the vertical. Parameters are chosen for each function to give plausible and roughly comparable mortality profiles.

Defining $\mu(a)$ as the mortality rate at age a (i.e. the probability of dying within a year

for a person of age a), the GM model is :

$$\mu(a) = \lambda + \alpha \exp(\beta a) \tag{A.1}$$

which is nested by the Siler function in Equation 1 through the addition of the infant mortality terms, b and B terms. Given we want to explain changes in life expectancy across a wide range of countries (including low income) and allowing for different trends at different ages the inclusion of an infant mortality term makes the Siler function preferable.

The HP model (Heligman and Pollard, 1980) is an eight-parameter parameterisation function of the form:

$$\frac{\mu(a)}{1-\mu(a)} = A^{(a+B)^C} + D \exp\left(-E(\ln(a) - \ln(F))^2\right) + GH^a \tag{A.2}$$

The function has three parts. The first reflects infant mortality, where A measures the level of mortality, B age displacement, and C the rate of decline in childhood mortality. The second part represents the "accident hump" of increased mortality in early adulthood due to accidents and maternal mortality. F sets the location of this hump, D its severity and E its spread. The third part, is a Gompertz term reflecting exponentially rising mortality in later life.

The difference between HP and Siler is basically the three parameters allowing for a hump in mortality during early 20s in lieu of an age-independent mortality term. Whilst these additional terms will provide a better fit to mortality for early adult years it does add a computational cost which can lead to numerical issues in estimation (which we found for some countries and some years given the demands of our model). For that reason we chose the Siler model as it strikes a balance between capturing mortality changes at all ages whilst being flexible enough to model a wide range of countries across multiple time periods. Given our focus on female mortality in which Heligman and Pollard (1980) themselves find that the accident hump is less pronounced the cost of this omission is likely to be less.

A.4 Estimation over 1 year age intervals

Figure A.8 shows the parameter estimates of the dynamic Siler model in the best practice case with one-year rather than five-year periods. The results are broadly similar to the five-year case (see Figure 2), although considerably more volatile.

Figure A.8: Parameter estimates under best practice with 1 year intervals

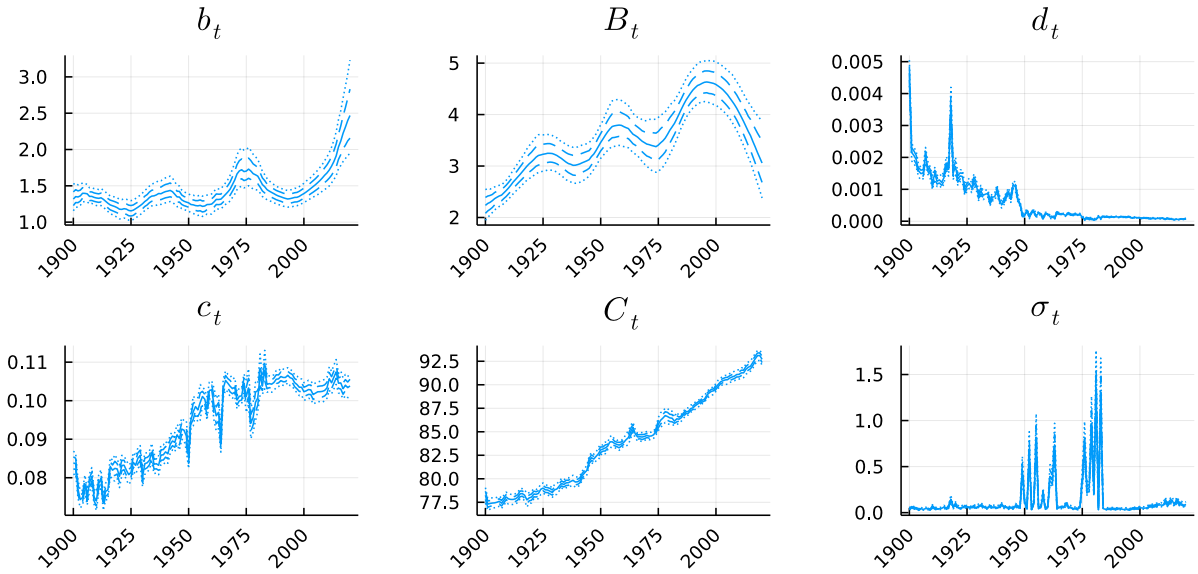

A.5 Data Appendix

All mortality data is taken from the Human Mortality Database (HMD) where available, otherwise it is taken from the United Nation's World Population Prospects database (UN WPP). Mortality rates are obtained at a single-age, single-year level and then aggregated to five year intervals. The use of five year averages means our focus is on medium and long run mortality trends. The use of quinquennial averages also considerable speeds up estimation. The length of sample period varies across countries due to data availability.

Table A.1 shows the sample period for each variable at the country level. In all cases, the mortality models are estimated on all available data.

Table A.1: Data availability by country

Country	Mortality	Source
Australia	{1921,1925}:{2016,2019}	HMD
Bangladesh	{1946,1950}:{2016,2019}	UN WPP
Belgium	{1901,1905}:{2016,2019}	HMD
Brazil	{1946,1950}:{2016,2019}	UN WPP
Canada	{1921,1925}:{2016,2019}	HMD
China	{1946,1950}:{2016,2019}	UN WPP
Congo	{1946,1950}:{2016,2019}	UN WPP
Denmark	{1901,1905}:{2016,2019}	HMD
Egypt	{1946,1950}:{2016,2019}	UN WPP
Ethiopia	{1946,1950}:{2016,2019}	UN WPP
France	{1901,1905}:{2016,2019}	HMD
Germany	{1986,1990}:{2016,2019}	HMD
India	{1946,1950}:{2016,2019}	UN WPP
Indonesia	{1946,1950}:{2016,2019}	UN WPP
Iran	{1946,1950}:{2016,2019}	UN WPP
Italy	{1901,1905}:{2016,2019}	HMD
Japan	{1946,1950}:{2016,2019}	HMD
Mexico	{1946,1950}:{2016,2019}	UN WPP
Netherlands	{1901,1905}:{2016,2019}	HMD
Nigeria	{1946,1950}:{2016,2019}	UN WPP
Norway	{1901,1905}:{2016,2019}	HMD
Pakistan	{1946,1950}:{2016,2019}	UN WPP
Philippines	{1946,1950}:{2016,2019}	UN WPP
Portugal	{1936,1940}:{2016,2019}	HMD
Russia	{1956,1960}:{2016,2019}	HMD
Spain	{1906,1910}:{2016,2019}	HMD
Sweden	{1901,1905}:{2016,2019}	HMD
Switzerland	{1901,1905}:{2016,2019}	HMD
Thailand	{1946,1950}:{2016,2019}	UN WPP
Türkiye	{1946,1950}:{2016,2019}	UN WPP
UK	{1921,1925}:{2016,2019}	HMD
USA	{1931,1935}:{2016,2019}	HMD
Viet Nam	{1946,1950}:{2016,2019}	UN WPP

A.6 Estimation and Convergence

HMC is a class of Markov Chain Monte Carlo (MCMC) algorithms that use information on the gradient of a joint distribution in order to sample from it. HMC has several key advantages (Sacher et al., 2021), not least its flexibility, computational efficiency and its quality of inference compared to other gradient-based and variational inference approaches.

The goal of HMC is to draw samples from a posterior distribution $q(\Theta)$ over the M parameters of interest Θ. In our case these are the values of the Siler model parameters over

time and the parameters defining their evolution. HMC only requires calculating the gradient of q with respect to each element of Θ rather than closed form expressions of conditional distributions.

To generate a stable "orbit" around the mode of the posterior distribution, HMC introduces M auxiliary momentum variables $\mathbf{r} \sim \mathcal{N}(\mathbf{0}, I_M)$. These momentum variables allow the algorithm to explore the parts of the posterior distribution that account for most of its mass, rather than remaining close to the posterior mode. To obtain a news sample for Θ, we first re-sample $\mathbf{r}$ and then follow a path described by the Hamiltonian function:

$$H(\Theta, \mathbf{r}) = -\log(q(\Theta)) + \sum_{i=1}^{M} \frac{r_i^2}{2} \tag{A.3}$$

The evolution of Θ is then given by

$$\frac{d\Theta_i}{dt} = \frac{\partial H}{\partial r_i} \tag{A.4}$$

and the evolution of $\mathbf{r}$ is given by

$$\frac{dr_i}{dt} = -\frac{\partial H}{\partial q_i} = \frac{\partial q_i(\Theta)}{q(\Theta)} \tag{A.5}$$

As in the well known Metropolis-Hastings algorithm, HMC accepts or rejects proposed samples from the Hamiltonian dynamics based on the value of the joint density.

The NUTS algorithm is a specific variance of HMC which endogenously determines the number of steps followed in the path for generating samples to increase efficiency. The algorithm follows the Hamiltonian dynamics until the path begins to circle back towards its starting point. This generates proposals that are relatively far from one another while keeping correlation between draws low. We initialise the sampler with the maximum a posteriori (MAP) estimate, and run 4 separate chains in parallel.

We assess the convergence of our HMC sampler based on the potential scale reduction statistic $\hat{R}$ (Gelman and Rubin, 1992) which measures the degree to which samples from different chains are drawn from the same distribution.[11] Each estimated parameter (which in the case of our time varying parameters is the value of that parameter in a single period) has its own $\hat{R}$ statistic. Vehtari et al. (2021) recommend only using a sample if the $\hat{R}$ statistics for the parameters of interest is 1.1 or lower. We see in all cases these statistics are very close to 1, with $\hat{R}$ statistics well below 1.05 for all estimated parameters.

A.7 Definitions

Given the mortality rate (Eq. 1), the period survival rate from birth until age a under the Siler parameterisation is

$$S(a) = \exp(-da + (\exp(-b(a+B)) - \exp(-bB)) - (\exp(c(a-C)) - \exp(-cC))) \tag{A.6}$$

[11] Roughly speaking, $\hat{R}$ is the ratio of the variance of a parameter when the data is pooled across all of the chains to the within-chain variance

Given this we can derive expressions for period life expectancy, lifespan equality and lifespan. Defining $S(a,\tau) = S(\tau)/S(a)$ as the probability of an individual aged a surviving until age τ, remaining life expectancy at age a is

$$LE(a) = \int_a^\infty S(a,\tau)\mathrm{d}\tau \tag{A.7}$$

Lifespan inequality is the normalised entropy of the survival rate (Demetrius, 1978; Keyfitz and Caswell, 2005).

$$H(a) = -\frac{1}{LE(a)}\int_a^\infty S(a,\tau)\ln\left(S(a,\tau)\right)\mathrm{d}\tau \tag{A.8}$$

Lifespan equality is then defined, following Aburto et al. (2020), as

$$h(a) = -\log H(a) \tag{A.9}$$

Upper bound of age at death (UBAD) is defined by Olshansky et al. (1990) as the age at which only 0.1% of a cohort can expect to live so that $S(UBAD) = 0.001$.

To compute the sensitivity of both life expectancy and lifespan inequality to each parameter, we need to calculate the derivative of S with respect to each parameter.

$$S_B(a) = b\left(\exp(-bB) - \exp(-b(a+B))\right)S(a) \tag{A.10}$$

$$S_b(a) = \left(B\exp(-bB) - (a+B)\exp(-b(a+B))\right)S(a) \tag{A.11}$$

$$S_C(a) = c\left(\exp(c(a-C)) - \exp(-cC)\right)S(a) \tag{A.12}$$

$$S_c(a) = \left(C\exp(-cC) - (a-C)\exp(c(a-C))\right)S(a) \tag{A.13}$$

$$S_d(a) = -aS(a) \tag{A.14}$$

By Leibniz's rule, the derivative of life expectancy with respect to parameter θ is

$$LE_\theta(a) = \int_a^\infty S_\theta(a,\tau)\mathrm{d}\tau \tag{A.15}$$

For lifespan inequality, Colchero et al. (2021) show that.

$$H_\theta(a) = -\frac{1}{LE(a)}\left(LE_\theta(a)(1+H(a)) + \int_a^\infty S_\theta(a,\tau)\ln S(a,\tau)\mathrm{d}\tau\right) \tag{A.16}$$

from which the gradient of $h(a)$ with respect to θ follows from Equation A.9

The above expressions refer to *period* measures of survival rate, life expectancy and lifespan inequality. To calculate cohort versions define $\tilde{S}(a,t)$, the cohort survival rate for a person born in period t until age a, as :

$$\tilde{S}(a,t) = \exp(-\int_0^a \mu(a,t+a)da) \tag{A.17}$$

Given $\mu(a, t+a)$ is now time varying, this expression is less analytically tractable than the period survival rate, but can be straightforwardly calculated numerically.

A.8 Additional BPLE results

Figure A.9 shows the posterior distribution of the time-varying drift terms in Eq 4, estimated on best practice mortality. Particularly of note here is that:

1. The drift term for c_t has fallen from around 0.02 in 1950 (meaning a 2% increase per period) gone from positive to very close to zero by the end of the sample.

2. The drift term for C_t has remained roughly constant around 0.01 throughout the sample.

Figure A.9: Siler parameter drift estimates for best practice mortality

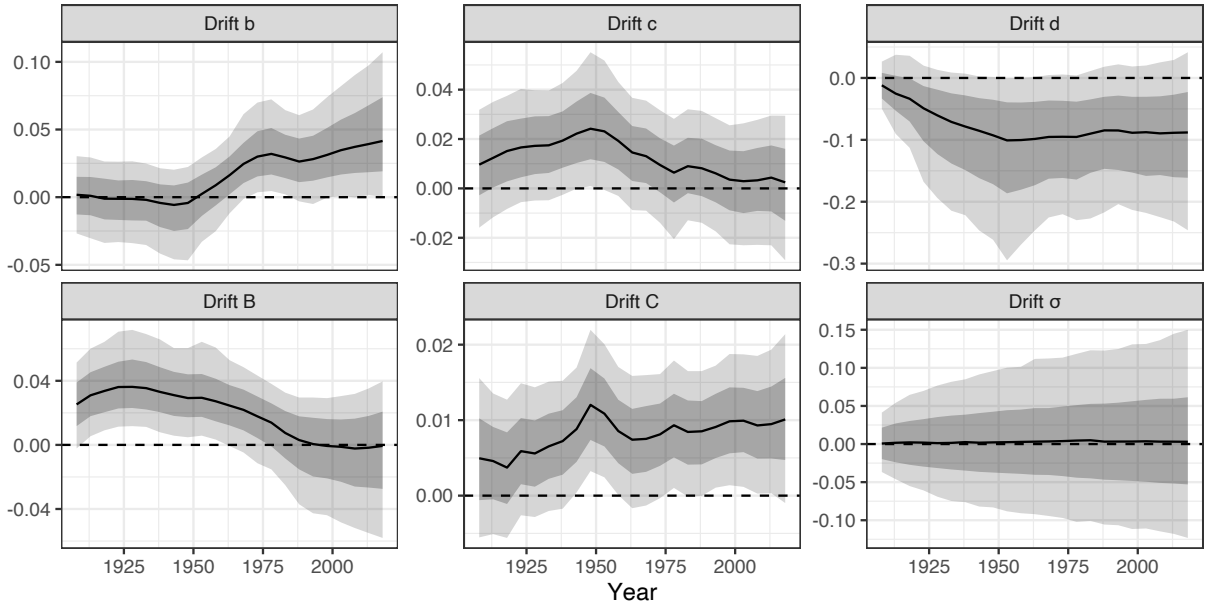

Note: These panels show posterior distributions of the drift terms for the Siler model parameters estimated on best practice mortality. The parameters evolve according to Eq 3, in which the drift terms themselves evolve according to Eq 4. The posterior distributions for parameters themselves are shown in Figure 2. The solid line indicates the median with the shaded areas indicating the 70% and 95% credible intervals.

A.9 Additional international results

Figure A.10 shows the absolute number of years gained the life expectancy at birth across each country over time. It has the same format as Figure 6 in the main text, but expressed in years rather than proportions.

Figure A.11 shows the estimated UBAD for thresholds at the 10%, 5%, 1% and 0.1% survival probabilities. This illustrates that results are fairly consistent across the choices of these thresholds and so not driven by the imprecisely measured mortality rates at very old ages.

Figure A.10: Absolute LE gains (in years) from each parameter

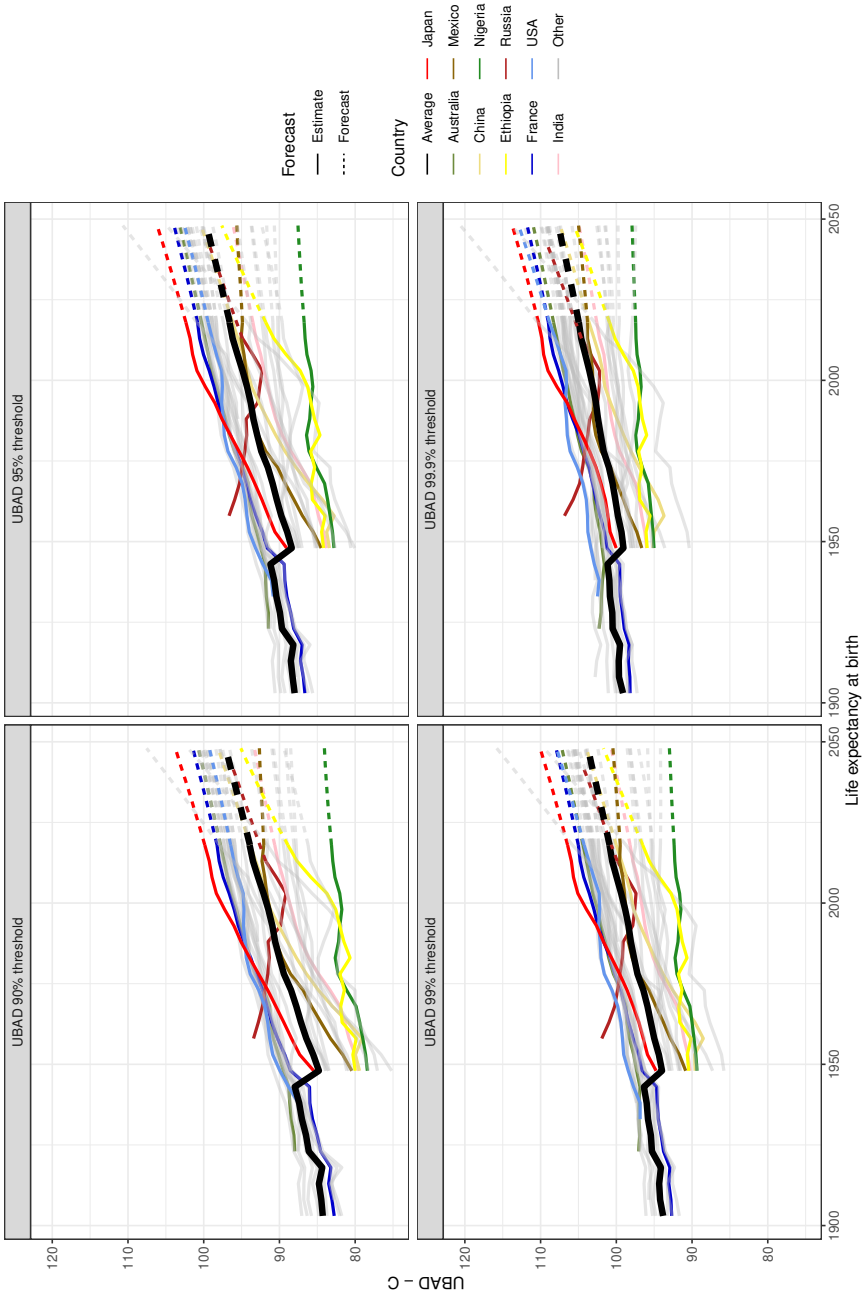

Figure A.11: Value of UBAD over time across countries, for different thresholds