



HAL
open science

Jointly estimating subnational mortality for multiple populations

Ameer Dharamshi, Monica Alexander, Celeste Winant, Magali Barbieri

► **To cite this version:**

Ameer Dharamshi, Monica Alexander, Celeste Winant, Magali Barbieri. Jointly estimating subnational mortality for multiple populations. *Demographic Research*, 2025, 52 (3), pp.71-110. 10.4054/DemRes.2025.52.3 . hal-04885695

HAL Id: hal-04885695

<https://hal.science/hal-04885695v1>

Submitted on 14 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



DEMOGRAPHIC RESEARCH

A peer-reviewed, open-access journal of population sciences

DEMOGRAPHIC RESEARCH

**VOLUME 52, ARTICLE 3, PAGES 71–110
PUBLISHED 10 JANUARY 2025**

<http://www.demographic-research.org/Volumes/Vol52/3/>

DOI: 10.4054/DemRes.2025.52.3

Research Article

Jointly estimating subnational mortality for multiple populations

Ameer Dharamshi

Monica Alexander

Celeste Winant

Magali Barbieri

© 2025 Dharamshi, Alexander, Winant & Barbieri.

This open-access work is published under the terms of the Creative Commons Attribution 3.0 Germany (CC BY 3.0 DE), which permits use, reproduction, and distribution in any medium, provided the original author(s) and source are given credit.

See <https://creativecommons.org/licenses/by/3.0/de/legalcode>

Contents

1	Introduction	72
2	Background	74
3	Methods	76
3.1	Principal components models	76
3.2	Model summary	79
3.2.1	Core model	80
3.2.2	Temporal smoothing	84
3.2.3	Overdispersion term	84
3.3	Computation	85
3.4	Simulation study	85
4	Application to age- and sex-specific mortality rates at the county-level in the United States	87
4.1	State-level patterns in sex-specific mortality	88
4.2	Validation	93
5	Discussion	94
6	Acknowledgments	95
	References	96
	Appendices	101

Jointly estimating subnational mortality for multiple populations

Ameer Dharamshi¹

Monica Alexander²

Celeste Winant³

Magali Barbieri⁴

Abstract

BACKGROUND

Understanding patterns in mortality across subpopulations is essential for local health policy decision-making. One of the key challenges of subnational mortality rate estimation is the presence of small populations and zero or near zero death counts. When studying differences between subpopulations, this challenge is compounded as the small populations are further divided along socioeconomic or demographic lines.

OBJECTIVE

We aim to develop a model to estimate subnational age-specific mortality rates that accounts for the dependencies in mortality experiences across subpopulations.

METHODS

We develop a Bayesian hierarchical principal components-based model that shows correlations across subpopulations.

RESULTS

We test this approach in a simulation study and also use the model to estimate age- and sex-specific mortality rates for counties in the United States. The model performs well in validation exercises and the US estimates suggest substantial variation in mortality trends over time across geographic lines.

¹ Department of Biostatistics, University of Washington, Seattle, USA. Email: adharams@uw.edu.

² Departments of Statistics and Sociology, University of Toronto, Toronto, Canada.

³ Department of Demography, University of California, Berkeley, USA.

⁴ Department of Demography, University of California, Berkeley, USA; Institut National d'Études Démographiques, Aubervilliers, France.

CONTRIBUTION

Our proposed model jointly estimates age-specific mortality rates for multiple subpopulations at the subnational level. By sharing information across subpopulations, our model improves on previous approaches that treat subpopulations as independent. Additionally, we demonstrate that ancillary correlation parameters are a useful tool for studying the convergence and divergence of mortality patterns over time.

1. Introduction

Reliable mortality estimates are crucial to support health policy design, implementation, and monitoring. Historically, developments in both mortality modeling and understanding mortality differences across populations have focused on national-level estimates for cross-country comparisons. However, as demonstrated by a growing body of evidence (Ezzati et al. 2008; Kindig and Cheng 2013; Dwyer-Lindgren et al. 2016; Sehgal et al. 2022), patterns in more granular subnational-level mortality rates are needed to facilitate local-level decisions. This insight has led to a recent line of work focused on studying disparities in mortality outcomes within countries at the subnational level. In the United States, for example, researchers are analyzing differences in life expectancy across state jurisdictions (Woolf and Schoemaker 2019; Montez et al. 2020; Harper, Riddell, and King 2021). At the international level, the United Nations has recently moved to producing estimates of key global health indicators, such as the under-five mortality rate, at subnational levels alongside the longstanding national level estimates (Godwin and Wakefield 2021). In addition, in recent years it has become increasingly clear that disparities in mortality outcomes along demographic dimensions, such as sex or gender, race or ethnicity, or socioeconomic status, are growing (Hendi 2015; Crimmins and Zhang 2019; Gutin and Hummer 2021; Masters and Aron 2021). As an example, in the case of child mortality risks, Bhutta (2016) argues that subnational determinants explain a greater portion of mortality than country-level characteristics. In response, estimates of subpopulation mortality risks are needed to identify and understand the mortality patterns of vulnerable groups and track the effects of policy responses.

In this paper, we are motivated by sex differences in mortality and the patterns of this difference across age, time, and geography. Differences in mortality patterns between males and females have long been documented. Women have had longer life expectancy than men in every population and time period where and when female discrimination has not been prevalent. The mortality gap is particularly large at young adult ages, when mortality is mostly determined by external causes (particularly accidents and homicides), but it remains high throughout all working ages. It closes progressively after age 65 to 70 years but never disappears completely. Sex differences in mortality reached a

peak in most high-income countries around 1975–1985 but have declined since, in large part because of the convergence in male and female behavior (e.g., smoking) (Gjonça, Tomassini, and Vaupel 1999). In these countries, the difference in life expectancy between men and women ranges from a low of around 3 years in Northern Europe up to nearly 10 years in Eastern Europe. In the United States, the sex difference in life expectancy increased throughout the 20th century to reach a maximum at 7.74 years in 1975 before declining to a low of 4.74 years in 2012. It has increased again since and currently (2021) reaches 5.75 years (Human Mortality Database 2023). The difference varies across states, ranging from 3.7 years (in Utah) to 7.2 years (in the District of Columbia) (United States Mortality Database 2023). Given that the gap between sexes in life expectancy is highly specific to location, it is of interest to further study the behavior of sex-specific mortality risks at finer geographic resolutions, such as counties. Hence, there is a need for methods that reliably estimate differential mortality risks between males and females (and between other demographic subgroups) in small populations.

While the need for subnational estimates of mortality disaggregated by subpopulation is clear, there are several challenges to obtaining reliable estimates. Although raw death rates can usually be calculated, the natural variations lead to uncertain estimates. It is well understood that subnational mortality estimates are more complex to construct than national-level estimates given the large random fluctuations associated with small numbers (Stevenson and Olson 1993). This problem is further compounded when examining mortality across subgroups within subnational areas, where inferring underlying mortality risks from raw death rates is challenged by the erratic patterns over age.

Producing reliable estimates of mortality rates and measuring the associated uncertainty requires statistical models that take the natural stochasticity in the data into account. In recent years, an increasing number of studies have demonstrated the usefulness of the Bayesian approach in the construction of such models (see for instance Alexander, Zagheni, and Barbieri 2017; Khana et al. 2018; Song and Luan 2022). Building on previous work, in this paper, we propose a general Bayesian model framework to estimate age-specific mortality rates at the subnational level jointly for multiple populations. The model incorporates characteristic shapes of mortality age schedules within a Bayesian hierarchical framework, which allows information on mortality patterns to be shared across populations. We extend previous approaches by accounting for correlation in mortality experiences across subpopulations rather than assuming subpopulations are independent. The model produces estimates and uncertainty for mortality rates for each sex. In addition, it estimates higher-level parameters. These parameters summarize trends in different dimensions of mortality over time and show how similar or dissimilar the trends are across groups. We illustrate the model with an application to estimating sex-specific mortality by county in the United States. While we focus on sex-specific mortality, the modeling framework is generalizable to population groups defined by other characteristics (such as race/ethnicity).

The contributions of our work are twofold. Methodologically, our proposal demonstrates that subpopulation correlation structures are a useful tool for stabilizing small population mortality estimates, and for gaining meaningful insights into the underlying behaviors driving mortality changes. Substantively, the estimates of sex-specific US county-level mortality rates produced in our data application have been released as part of the United States Mortality DataBase (USMDB) (United States Mortality Database 2023) and can now be used by demographers studying small-scale mortality in the United States.

The remainder of the paper is structured as follows. We first give a brief overview of recent developments in subnational mortality estimation. We begin our methods discussion by demonstrating that principal components derived from a set of reference mortality curves capture structural differences in subpopulation-specific mortality patterns, then follow with a formal statement of the model. Results from the model along with validation exercises using simulated and real sex-specific US county-level mortality data are then provided. Finally, we conclude with a discussion of our findings and identify directions for future research.

2. Background

A large body of research exists on small-area estimation issues, and recently demographers have increasingly taken advantage of computational advances that make fitting complex statistical models to small-scale mortality data feasible. In particular, there has been a notable increase in the use of Bayesian methods in demographic estimation, especially for subnational estimation (for example Alexander, Zagheni, and Barbieri 2017; Schmertmann and Gonzaga 2018; Lu, Hanewald, and Wang 2021; Lin and Tsai 2022). Bayesian methods are particularly suited to demographic contexts as they provide a useful framework to incorporate different data sources in the same model, account for various types of uncertainty, and allow for information exchange across time and space (Bijak and Bryant 2016).

One area of previous work has focused on estimating aggregate indicators at the subnational level, such as child mortality and life expectancy (Mercer et al. 2015; Ševčíková and Raftery 2021). Models on aggregate indicators generally involve temporal and spatial smoothing, allowing for information in mortality trends to be shared across these dimensions. In some cases, models rely on covariates (such as education or income) to stabilize mortality rate estimates from noisy data (Wang et al. 2013; Arias et al. 2018; Murray et al. 2006).

In addition to aggregate mortality indicators, research has focused on producing estimates of age-specific mortality rates, which can form the basis of estimating subnational life tables. Recent advances in this area build on classical demographic approaches of model life tables and relational models, which identify key patterns in mortality over age

across a wide range of populations and allow patterns to be shifted based on a reduced set of parameters (Coale, Demeny, and Vaughan 1983). For example, TOPALS (tool for projecting age patterns using linear splines) models consist of a standard age schedule and population-specific deviations away from that standard, which are smoothed using linear splines (de Beer 2012). TOPALS-type models have been used to produce subnational estimates of migration and mortality in varying data quality contexts (Gonzaga and Schmertmann 2016; Schmertmann and Gonzaga 2018; Dyrting 2020). TOPALS models do, however, require selecting a single standard mortality curve. This is challenging when studying multiple populations, as each may have a distinct mortality profile.

A related modeling approach derives a set of ‘principal components’ from reference mortality curves, which are then used as the basis of a regression framework. This allows a large set of plausible mortality curves to be estimated using a reduced set of parameters. This approach is in the spirit of the Lee–Carter model for mortality forecasting and related work though in modern applications, multiple principal components are typically used (Lee and Carter 1992; Renshaw and Haberman 2003). For example, Clark (2019) uses this approach to formulate a new set of model age schedules in data-sparse contexts. In particular, this paper extends an earlier study by Alexander, Zagheni, and Barbieri (2017), which introduces a Bayesian hierarchical framework that builds on principal components derived from national mortality schedules for use at the subnational level. In more recent work, Alexander and Alkema (2022) use a principal-components approach to model mortality in the broader context of subnational population estimation using a cohort component projection framework.

In general, most existing subnational mortality models inherently assume subpopulations are independent. They model each population separately, then combine or interpret after modeling. However, ideally we would model all subpopulations within a single framework as mortality is generally correlated over groups. Existing joint models focus on modeling national or first subnational administrative division-level mortality by sex as it is well understood that changes in male and female mortality are correlated (Noymer and Van 2014). For example, Li and Lee (2005) extend the Lee–Carter model to jointly forecast national-level mortality for both sexes. More recently, national-level sex-specific life expectancy estimates have been produced by the United Nations using gap-based approaches, such as in Raftery, Lalic, and Gerland (2014). Under this approach, female life expectancy is estimated with standard one-sex methods. Estimates of male life expectancy are then constructed by modeling the gap between female and male life expectancy. In the present case, we are interested in estimates of all age-specific mortality rates, not just life expectancy, and a purely gap-based approach to this problem at smaller scales is not appropriate because of the need to select an anchor or reference population, such as females in the case of sex-specific modeling. Due to the small counts in both areas and age groups, standard single-population mortality models may fail to produce reliable estimates for the reference category, which then propagates downstream

into the gap-modeling step. This issue is further compounded when studying variables such as race or ethnicity that have dramatically different compositions across jurisdictions, particularly if the reference population is near absent in certain regions.

At the subnational level, Rau and Schmertmann (2020) jointly model age- and sex-specific mortality in regions across Germany using a Bayesian TOPALS approach. With regards to sex differences, they obtain a typical pattern of age-specific sex differences across aggregated regions of Germany, and then use this information as a basis for a prior for differences at smaller scales, penalizing large deviations away from the observed differences by sex. In this paper, we take a different approach and flip the viewpoint from thinking about sex differences to thinking about covariation in mortality by sex (or any other population subgroups), and explicitly allow for group mortality rates to move together.

3. Methods

To estimate subnational mortality rates while also extracting patterns across key subpopulations, we propose a Bayesian hierarchical model that builds on the principal component-based approach of Alexander, Zagheni, and Barbieri (2017) by incorporating structures that capture subpopulation interactions. Before defining the model equations, we first discuss why principal components are an attractive modeling strategy in this context.

3.1 Principal components models

The use of principal components is motivated by the fact that age-specific mortality rates tend to display strong regularities across different populations (Lee and Carter 1992; Renshaw and Haberman 2003; Clark 2019; Alexander, Zagheni, and Barbieri 2017). This means that systematic variation in mortality rates is well captured by a reduced set of parameters that can be modeled in such a way that information is shared across space and time.

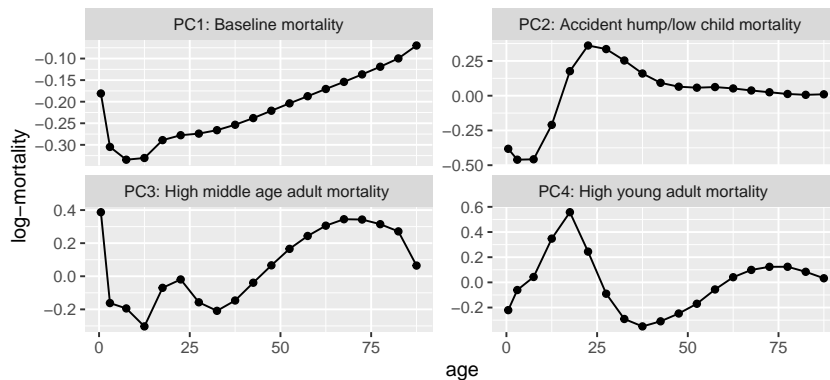
To generate principal components, we compute the singular value decomposition of a collection of regional log-mortality curves. The collection of log-mortality curves, \mathbf{X} , is a $N \times A$ matrix, where N is the number of region-subpopulation-years under consideration and A is the number of age groups. For US sex-specific mortality, $N = 6066$ (60 years of state-level sex-specific data) and $A = 19$ (ages <1, 1–4, 5–9, ..., 75–79, 80–84, 85+). Note that different age groups, such as one-year age groups, could also be considered. The singular value decomposition of \mathbf{X} is then

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}', \tag{1}$$

where \mathbf{U} is the 6066×19 matrix of left singular values, $\mathbf{\Sigma}$ is a 19×19 diagonal matrix of scaling factors, and \mathbf{V} is the 19×19 matrix of right singular values, which we term ‘principal components.’ These extract key patterns in mortality over age. As the \mathbf{X} matrix contains data from both male and female populations, the patterns captured by the principal components are shared across sexes.

The first four principal component curves are plotted in Figure 1. Together, these four principal components explain over 99% of the variation in the state-level data. Since the singular value decomposition is computed using the uncentred data, the first principal component explains the bulk of the variation. The choice to leave the data uncentred is deliberate: Under this approach, the first principal component controls the overall level of mortality, which is useful for capturing systematic shifts in mortality when modeling. Following the first principal component, each successive component explains less variation. Note that to ease interpretation throughout the paper, we have inverted the first two principal components and their associated coefficients.

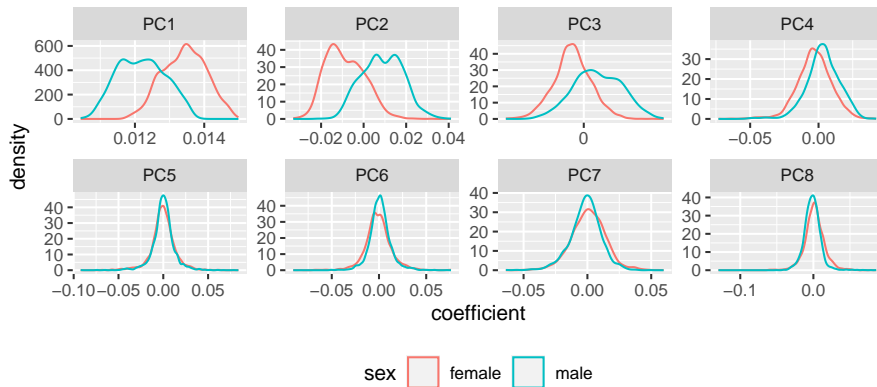
Figure 1: First four state-level log-mortality principal components



In practice, choosing the number of components is a balance between incorporating components that pick up on only systematic patterns, while still allowing for enough flexibility in the model (Sharro, Clark, and Raftery 2014; Clark 2019). The set of principal components that offer useful information on sex-specific mortality can be identified using the \mathbf{U} matrix. The left singular values in each row of the \mathbf{U} matrix represent the contribution of each principal component to the corresponding mortality curve in the \mathbf{X} matrix. If the i th principal component contains material subpopulation-specific differences, we would expect the distributions of the values in the i th column of \mathbf{U} to differ between subpopulations. We thus separate the rows of \mathbf{U} by sex and examine the resulting dis-

tributions. Figure 2 displays the densities of the coefficients by sex for the first eight principal components. We choose eight as the singular values suggest there is limited additional information beyond eight.

Figure 2: Distribution of observed state-level left singular values by sex and principal component



The first four principal components demonstrate clear differences in the sex-specific distributions. In particular, the location of the distributions are noticeably different, suggesting structural differences between sexes. The distributions of the remaining principal components, including those not presented in Figure 2, are largely similar. Beyond a visual inspection, simple t-tests suggest that aside from the fifth, all the first eight principal components have location differences. Balancing these findings with the practical consideration that each additional principal component leads to a significant computational burden during model estimation, we suggest that at least three but ideally four principal components should be used in modeling sex-specific mortality in the United States.

Considering Figure 1 and Figure 2 together offers insight into patterns of interest. The first principal component has the characteristic J shape of log mortality. The higher female coefficient values in the PC1 panel of Figure 2 indicate that female mortality is generally lower than male mortality (to see this, note that the coefficients are multiplied by the first principal component, which as seen in Figure 1 is composed of exclusively negative values). Similarly, the PC3 panel implies that women experience lower mortality through the middle-age adult years. In addition, the distribution of coefficients for PC2 and PC4 indicate higher young adult mortality in the male population than in the female population, which is consistent with the ‘accident hump’ described in the broader mortal-

ity literature, as males are more likely to suffer mortality due to risky behavior (Heligman and Pollard 1980).

There is a substantial body of literature documenting the large and systematic excess mortality of men compared with women (see reviews of this literature in Gjonça, Tomassini, and Vaupel 1999; Barford et al. 2006; Beltrán-Sánchez, Finch, and Crimmins 2015). As previously mentioned, the male disadvantage in the risk of dying is high at all adult ages up to about 65 to 70 years, depending on the population, and it is particularly acute in young adults. The finding that the principal components are recovering these well-understood sex-specific differences motivates their use as the basis for a subpopulation mortality model. Using principal components as a basis for a statistical model can be robust in small population settings as the components allow plausible mortality rates to be estimated even in the presence of high variability.

3.2 Model summary

We frame our proposed model for jointly estimating subnational mortality for multiple populations in the context of US county-level mortality. The basic premise of our model is to leverage the natural relationships between populations of interest in order to improve the quality of mortality estimates for small populations. The model is designed to operate at the state-level, estimating age-specific mortality for all counties within a state, for any number of years and some number of subpopulations of interest. While we are motivated by the United States context, our proposed specification can be easily applied to other countries or administrative areas, substituting states and counties for the appropriate analogs.

We begin by defining $y_{a,s,c,t}$ as the observed number of deaths in age group a , subpopulation s , county c , and year t . Then, following Alexander, Zagheni, and Barbieri (2017), we assume that

$$y_{a,s,c,t} | \lambda_{a,s,c,t} \sim \text{Poisson}(P_{a,s,c,t} \cdot \lambda_{a,s,c,t}), \quad (2)$$

where $P_{a,s,c,t}$ is the population corresponding to age group a , subpopulation s , county c , and year t , and $\lambda_{a,s,c,t}$ is the mortality rate to be estimated for age group a , subpopulation s , county c , and year t .

The mortality rates $\lambda_{a,s,c,t}$ are estimated on the log scale as follows:

$$\log(\lambda_{a,s,c,t}) = \sum_{i=1}^P (\beta_{i,s,c,t} \cdot Y_{i,a}) + \gamma_{a,s,c,t}, \quad (3)$$

where Y_i is the i th principal component, P is the number of principal components, β are the estimated coefficients, and γ is an overdispersion term. The number of principal components, P , can be selected based on their contributions and differences between groups. In the case of sex-specific mortality in US counties, we use the $P = 4$ principal components plotted in Figure 1.

Equation (3) can be intuitively understood as constructing a log-mortality curve for each age-sex-county-year as a linear combination of the four principal components plus some additional age-specific variation. Different linear combinations of the principal components (that is, different estimated values of the β coefficients) lead to different plausible log-mortality curves. The overdispersion term accounts for the possibility that deaths will be overdispersed relative to the patterns captured by the principal components.

3.2.1 Core model

For many counties, population and death counts are small and highly variable, which would lead to uncertain estimates of β if each county and year were estimated independently. As such, we propose a hierarchical model for β based on the nesting structure of counties within states. Specifically, we assume that counties within each state are more likely to share similar mortality patterns than counties in different states, and allow high-data counties to share information with low-data counties in the same state. Similar hierarchical structures have been used successfully in the mortality modeling literature (Alexander, Zagheni, and Barbieri 2017; Lu, Hanewald, and Wang 2021; Lin and Tsai 2022). We note that while an analogous argument could be made to justify an additional hierarchy where states are nested within the United States as a whole, we opt not to take this approach as US states have large enough populations to produce reliable age-specific mortality, implying that pooling across states will be of limited benefit.

Similar to the dynamic of counties within a state, while subgroups of the population are expected to experience specific drivers of mortality due to their unique experience within each county, they are also expected to jointly experience more general drivers, whether those are due to policy, the environment, or other factors. For instance, regulations regarding maternity leaves have an impact on women's health, but natural catastrophes like heat waves or hurricanes affect everyone in the same area (albeit to various extents). This structure suggests that β coefficients, and thus by extension log-mortality curves, for all subgroups within a county should be modeled jointly, thereby allowing for distinct mortality estimates in each group while also exploiting the dependence between groups.

To capture geographic and subgroup dependence, we propose to model the vector of β coefficients for all subgroups within each county jointly as multivariate normal with a common state-level mean vector and covariance matrix:

$$\begin{pmatrix} \beta_{i,1,c,t} \\ \dots \\ \beta_{i,S,c,t} \end{pmatrix} = \begin{pmatrix} \mu_{\beta_{i,1,t}} \\ \dots \\ \mu_{\beta_{i,S,t}} \end{pmatrix} + \begin{pmatrix} \omega_{i,1,c,t} \\ \dots \\ \omega_{i,S,c,t} \end{pmatrix}, \quad i = 1, \dots, P, \quad (4)$$

$$\begin{pmatrix} \omega_{i,1,c,t} \\ \dots \\ \omega_{i,S,c,t} \end{pmatrix} \left| \sigma_{\beta_{i,t}}, L_{i,t}^{(\beta)} \right. \sim \mathcal{N} \left(0_S, \sigma_{\beta_{i,t}} 1_S L_{i,t}^{(\beta)} L_{i,t}^{(\beta)\top} 1_S \sigma_{\beta_{i,t}} \right), \quad (5)$$

$$L_{i,t}^{(\beta)} L_{i,t}^{(\beta)\top} \sim \text{LKJ}(1), \quad (6)$$

$$\sigma_{\beta_{i,t}} \sim \mathcal{N}^+(0, 1), \quad (7)$$

where S is the number of subpopulations under consideration, $\mu_{\beta_{i,s,t}}$ is the state-level coefficient for the i th principal component, subpopulation s , and year t , and $\omega_{i,s,c,t}$ is the county deviation for the i th principal component, subpopulation s , county c , and year t .

We discuss the two components of β in turn, starting with the state-level mean vector, μ_{β} . Information-rich observations from county subgroups with large populations are the primary contributors to estimates of μ_{β} for the corresponding subpopulation. Estimates of β for small county subgroups with less informative observed death counts are then partially informed by the high-data counties through the shared state-level means. This pooling effect stabilizes estimates of β for small populations by pulling the estimates toward the state mean. It is important to note that this effect occurs at the county-subgroup level and not uniformly within a county: A large population county with an uneven subgroup composition can experience stronger pooling effects in its smaller subgroups. This is not particularly important for the application to sex-specific mortality as the populations are roughly balanced, but may be critical in other applications involving subgroups defined by other demographic variables.

The ω vector models the specific deviations in the β vector for each county from the state-level mean vector. By jointly modeling the deviations for all subpopulations as in Equation (4), the model captures patterns in the dispersion of β by the principal component in the correlation matrix $L_{i,t}^{(\beta)} L_{i,t}^{(\beta)\top}$, where $L_{i,t}^{(\beta)}$ refers to the Cholesky factor of the correlation matrix. This enables information sharing across subpopulations. In counties where there is an imbalance in the number of nonzero death observations across subpopulations, jointly generating β coefficients allows higher data groups to support lower data groups. An example of such a county is given in Figure A-3 in Appendix B.

Regarding the prior specification for the covariance of the ω vector, we adopt a separation strategy whereby the variance terms and correlation matrix are assigned independent priors (Barnard, McCulloch, and Meng 2000). For each principal component, we assume that the variance is constant across subpopulations but varying in time. The

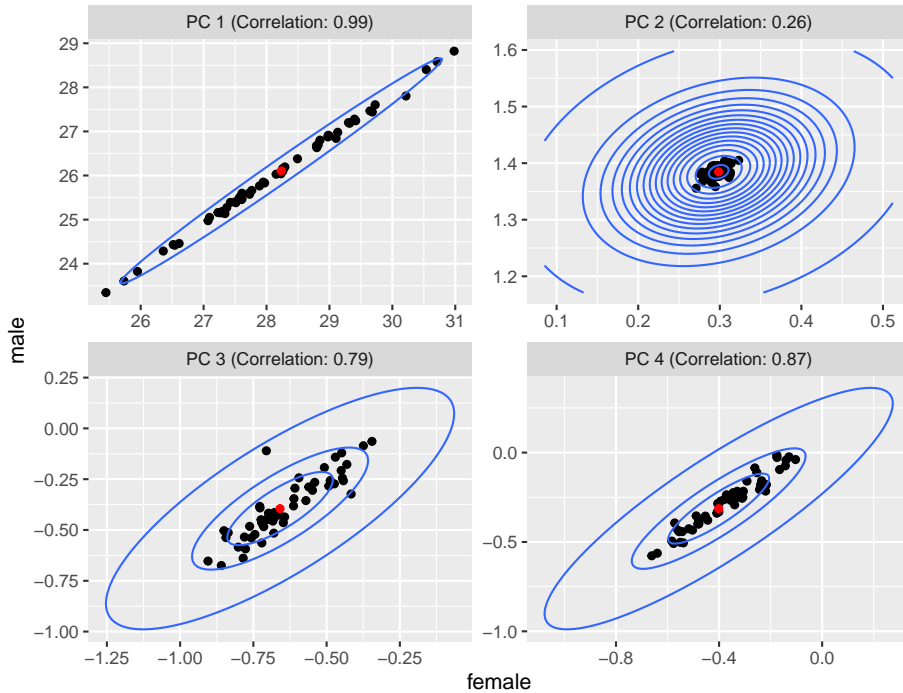
constant variance assumption is a simplifying choice predicated on the relative symmetry of male and female subpopulations in our motivating context. In applications where such an assumption may not be preferred, one can simply add an additional index for subpopulation in Equation (7). For the correlation component, our goal is to recover patterns from the data, not impose patterns on the data. As such, we choose to assign the uninformative LKJ(1) prior to the subpopulation correlation Cholesky factors (Lewandowski, Kurowicka, and Joe 2009). The LKJ(1) prior is uniform over the space of all valid $S \times S$ correlation matrices. We note that for $S > 2$, this is different from a marginal uniform prior on each entry of the correlation matrix as the positive semidefinite constraint leads to a higher marginal prior density on smaller correlations. However, in our context, where S is expected to be fairly small, this difference is limited, and we can understand the LKJ(1) prior as a generic uninformative choice for the correlation structure.

By modeling correlations in principal component coefficients between sexes, it is possible for subgroups to experience joint variation in certain parts of their mortality curves but not in others. For example, the coefficients corresponding to the first principal component (i.e., baseline mortality) could experience strong correlation, leading to a common trend in baseline mortality levels, but the coefficients on the second principal component (i.e., the accident hump) may not, leading to differences by sex in young adult mortality.

An illustration of the correlated principal component coefficients is provided in Figure 3. Each plot captures the patterns of one of the four principal component coefficients for California in 2017. The red point represents the state-level value, the black points represent each county, and the blue contours are the distribution of ω centered at the state means. High correlation in the first principal component is expected, since when the general conditions in a region improve, both sexes are positively impacted, leading to highly correlated changes in baseline mortality. Similarly, the low correlation in the second principal component is consistent with males experiencing an elevated level of mortality in young adult ages (i.e., the accident hump) whereas females tend not to. High correlations are recovered for the third and fourth principal components, suggesting that changes in young adult and middle-age adult mortality experiences are similar across sexes.

The β specification is intended to be flexible. If there is only a small number of years of data available or there is prior information suggesting that correlation structures are time-invariant, one could reasonably omit the time index or share correlations across a short time interval. Similarly, if certain correlations are known in advance, the correlation matrices themselves could be constrained accordingly.

Figure 3: Median posterior county-level principal component coefficients (β) and the corresponding state-level mean (μ_β) for California in 2017



Notes: The contour plots display the recovered correlation patterns between the male and female populations. Each panel corresponds to one principal component. The red and black points are the point estimates of the state-level and county-level principal component coefficients, respectively. The blue contours indicate the recovered correlation patterns between the two populations.

3.2.2 Temporal smoothing

State-level means are smoothed over time by penalizing the second-order differences to produce gradual changes (Alexander, Zagheni, and Barbieri 2017). In other words, we assign a random walk 2 prior on the state-level means with flat priors on the first two time points. Smoothing occurs at the state-level as opposed to the county-level to allow the needed flexibility for counties to experience irregular mortality patterns driven by localized events such as a natural disaster or a pandemic:

$$\mu_{\beta_{i,s,t}} | \mu_{\beta_{i,s,t-1}}, \mu_{\beta_{i,s,t-2}}, \sigma_{\mu_{\beta_i}} \sim \mathcal{N} \left(2 \cdot \mu_{\beta_{i,s,t-1}} - \mu_{\beta_{i,s,t-2}}, \sigma_{\mu_{\beta_i}} \right), \quad (8)$$

$$\sigma_{\mu_{\beta_i}} \sim \text{Log-Normal}(-1.5, 0.5). \quad (9)$$

3.2.3 Overdispersion term

Finally, the $\gamma_{a,s,c,t}$ term that allows for overdispersion of the log-mortality rate is modeled similarly to $\omega_{i,s,c,t}$ in a S -dimensional multivariate normal setup:

$$\left(\begin{array}{c} \gamma_{a,1,c,t} \\ \dots \\ \gamma_{a,S,c,t} \end{array} \right) \Bigg| \sigma_a, L_{a,t}^{(\gamma)} \sim \mathcal{N} \left(0_S, \sigma_a 1_S L_{a,t}^{(\gamma)} L_{a,t}^{(\gamma)\top} 1_S \sigma_a \right), \quad (10)$$

$$L_{a,t}^{(\gamma)} L_{a,t}^{(\gamma)\top} \sim \text{LKJ}(1), \quad (11)$$

$$\sigma_a \sim \mathcal{N}^+(0, 0.25). \quad (12)$$

Relationships between subpopulation γ 's are captured using the age-year correlation matrix $L_{a,t}^{(\gamma)} L_{a,t}^{(\gamma)\top}$. An intuitive way of understanding this component of the model is that the principal components produce the expected mortality derived from aggregate and local patterns, while γ captures additional age-specific deviations. These deviations from the expectation are often correlated across subpopulations, though this correlation may differ across ages and over time. As with the principal component coefficient correlations, here we also use an LKJ(1) prior. Given that the log-mortality values for different age groups occur in substantially different parts of the log curve, we estimate a separate scaling factor σ_a for each age group.

3.3 Computation

The model described here is fit using a Bayesian framework. Posterior samples are drawn using the No-U-turn sampling Hamiltonian Monte Carlo algorithm (Hoffman and Gelman 2014; Neal 2011) implemented in the Stan R package (Stan Development Team 2022). In the US county-level mortality context, we execute the model once for each state using 4 chains with 500 iterations of burn-in and 2,500 iterations of samples each. All 10,000 samples for each state are retained. Convergence is diagnosed using trace plots, effective sample sizes, and the Gelman and Rubin diagnostic (Gelman and Rubin 1992).

The runtime of the model is a function of both the computing setup and the number of county-year-subpopulations of interest. Our analysis is run on a virtual machine (Xen 4) with memory ranging from 100GB up to 500GB and CPU cores ranging from 4 to 20 depending upon project demand. The VM runs Ubuntu 20.04TLS and is not directly accessible from the internet. The underlying server is a Supermicro server with 768GB of memory and 36 CPU dual cores (Xeon E5-2695 v4) (72 virtual CPUs). In the US county-level mortality context, the model runs for Delaware, the state with the fewest counties, and Texas, the state with most counties, took approximately 1 hour and 1.5 days, respectively.

3.4 Simulation study

We conducted a simulation study to test the ability of the model to simultaneously estimate mortality rates and extract patterns in subgroup mortality. We generated 10 years of population data for 25 simulated counties of various sizes. Each population is composed of five subgroups ranging in size from 10% to 50% of the total population. We then generated deaths for all subgroups using log-mortality rates constructed as the linear combination of a pair of standard curves with coefficients subject to a variety of correlation patterns. The goal of the simulation study was to recover the true correlations and age-specific mortality rates when fitting the proposed model to the simulated data. Specific details on the data-generating process are provided in Appendix A.

The model is run on the simulated data using the standard curves as the principal components. To validate the results, we compare the estimated correlation matrices and log-mortality rates to the corresponding true values. Figure 4b plots estimated correlation matrices extracted from the model. When we compare corresponding facets against Figure 4a, it appears that the model has successfully identified and extracted the patterns observed in the data.

Figure 4: Simulation study true and estimated posterior median correlation matrices

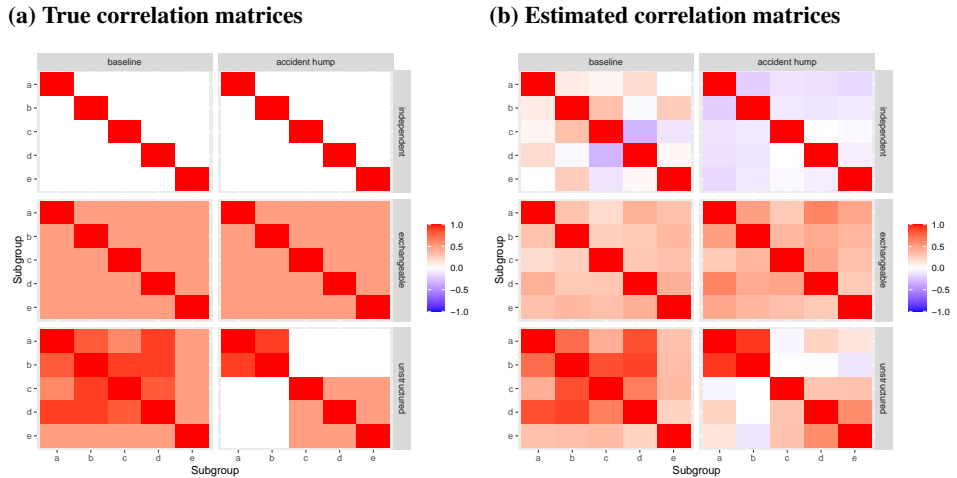


Table 1 presents the coverage values for correlation and log-mortality rate parameters at the 80%, 90%, and 95% nominal levels. Coverage is defined for correlation and log-mortality rates as $\frac{1}{n} \sum_{i=1}^n 1_{l_i \leq \theta_i \leq u_i}$, where n is the number of parameters, θ_i refers to an individual parameter, and l_i and u_i are, respectively, the lower and upper bounds of the posterior credible intervals. For the correlation matrices, we compute entrywise coverage values for the off-diagonal entries of the lower triangles of the correlation matrices.

Table 1: Simulation study coverage values for correlation and log-mortality rates

	Coverage (80%)	Coverage (90%)	Coverage (95%)
Correlations	0.78	0.90	0.94
Log-mortality rates	0.83	0.92	0.96

For both sets of parameters, the model’s coverage values are in line with the nominal values at all levels, suggesting that the model is well calibrated.

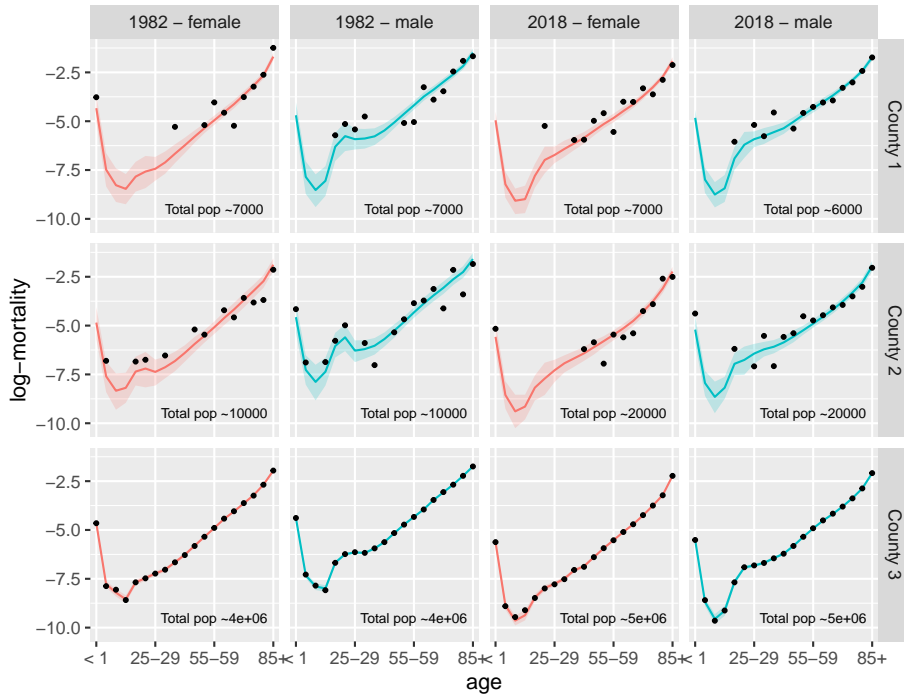
4. Application to age- and sex-specific mortality rates at the county-level in the United States

We applied the proposed model to estimate sex- and age-specific mortality at the county-level in the United States over the period 1982–2019. In the United States, relatively high quality data are collected through vital registration systems. The National Center for Health Statistics (NCHS) publishes detailed mortality records for each year included in our series from which we can compile the mortality tabulations (National Center for Health Statistics 2023). For the purposes of the present analysis, we obtained access to the necessary protected data with individual death records from 1989 to 2019 from the NCHS under a Data User Agreement as the publicly available data do not include geographic information. The United States Census Bureau publishes midyear population estimates by county of residence, year, sex, and age (United States Census Bureau 2022). To comply with our data agreement with the NCHS, counties are anonymized in all figures.

As stated previously, the model is applied to each state individually. Across all states, we use the first four principal components plotted in Figure 1. Note that estimates based on this model are now published as part of the USMDB and can be accessed at <https://usa.mortality.org/> (United States Mortality Database 2023).

As an illustration of the model outputs, in Figure 5 we plot the estimated sex-specific log-mortality curves in 1982 and 2019 for a subset of US counties along with associated uncertainty intervals. The observed log-mortality rates for age groups with non-zero deaths are presented as points, and the female and male county-level estimates and 95% credible intervals are presented in red and blue, respectively. The selected counties represent settings with small, medium, and large populations (of approximately 13,000, 40,000, and 10,000,000 people, respectively). As expected, smaller counties, such as County 1, have greater uncertainty as compared to larger ones, such as County 3, due to the higher levels of noise caused by low or zero death counts. For County 3, modeled estimates follow the data exactly.

Figure 5: Examples of US county-level sex-specific log-mortality plots by county-year-sex for three counties of varying sizes



Note: Black dots represent observed values, coloured curves and regions indicate posterior medians, and 95% credible intervals.

4.1 State-level patterns in sex-specific mortality

In addition to producing mortality rates for subnational areas, the proposed model specification contains a number of parameters that are useful in observing broad patterns in mortality at the state-level. Specifically, changes in μ_β describe broad, state-level trends in mortality patterns over time, and the correlation matrices offer insight into county-level trends.

In Figures 6 and 7, we plot posterior medians and 95% credible intervals of the estimated state-level coefficients for the first and second principal components for male and female populations in all states and years. Recall that higher values for the first and second principal component coefficients correspond to lower overall mortality and

a larger accident hump, respectively. States are organized roughly along their relative geographic location.

The most notable finding is that in the plot for the first principal component, the gap in the baseline mortality coefficients between men and women is shrinking over time. This suggests a convergence in mortality across the sexes (Seligman, Greenberg, and Tuljapurkar 2016). Geographically, we see that baseline mortality patterns differ substantially by state, suggesting evidence for a divergence in mortality across US states due to stagnation or regression in some states and continuous improvement in others (Fenelon 2013). For example, states such as California and New Jersey experienced significant declines in baseline mortality in both sexes. Others, such as Louisiana, experienced limited improvement and have stagnated in the 2010s. Some states, such as Ohio, saw a deterioration in baseline mortality in the 2010s. This trend reversal occurs across the eastern half of the United States.

For the second principal component, increases are experienced across the United States in both sexes with female values transitioning from negative to positive. As the second principal component contains an accident hump shape, this suggests increasing young adult mortality relative to the baseline. Our findings suggest that this trend accelerated in the 2010s for males in New Jersey, West Virginia, Ohio, and more broadly across the Northeast and Midwest. This is consistent with the sharp increase in opioid overdose deaths in these regions of the United States (Alexander, Kiang, and Barbieri 2018).

For the final two principal components (shown in Appendix B) there are again clear geographic clusters of patterns over time. For the third principal component, which has relatively higher mortality in middle-age adults, trends are generally declining, but stagnating more in the South, and there is evidence of a trend reversal in Kentucky and West Virginia. Unlike the rest of the states, Alaska does not show any significant gap between the two sexes in the third principal component, suggesting that the deviation to baseline mortality associated with this component contributes similarly to male and female mortality. However, it is important to recognize that the population in Alaska is small, leading to large uncertainty in the extracted trends.

For the fourth principal component, a transition from positive to negative coefficient values is found across the United States. The interpretation of this trend is similar to that of the second principal component. An inverted fourth principal component implies relatively higher early adult mortality, which is again consistent with the opioid crisis.

Figure 6: Posterior medians and 95% credible intervals for the state-level coefficients for the first principal component (μ_β)

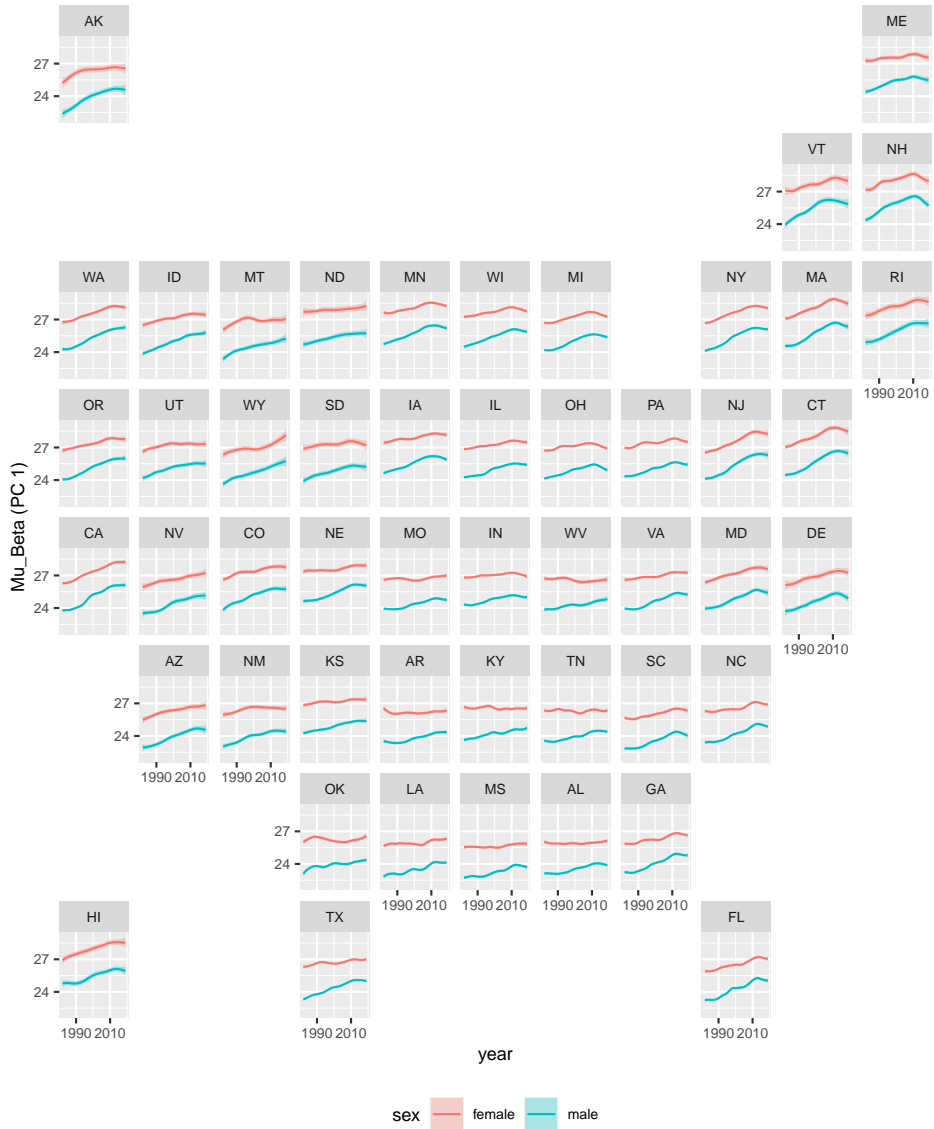


Figure 7: Posterior medians and 95% credible intervals for the second principal component (μ_{β})

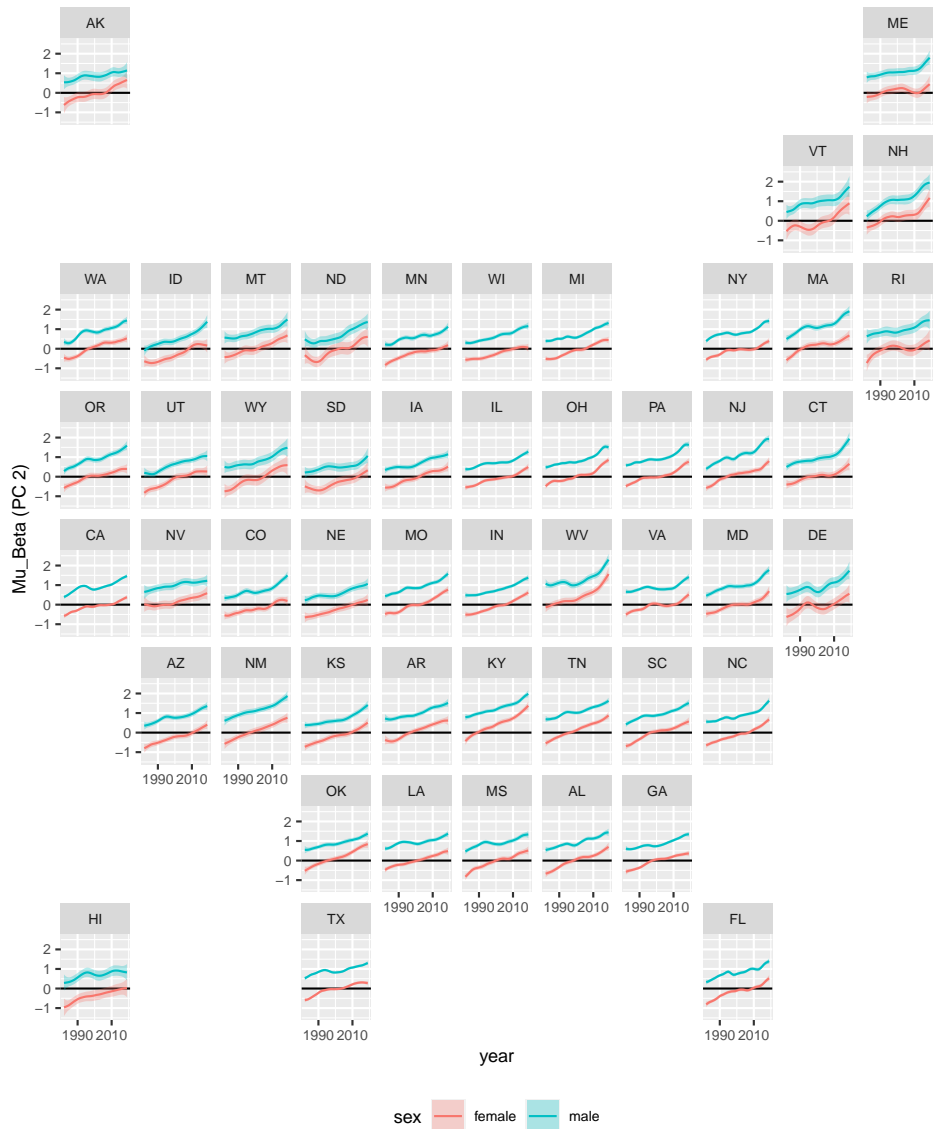
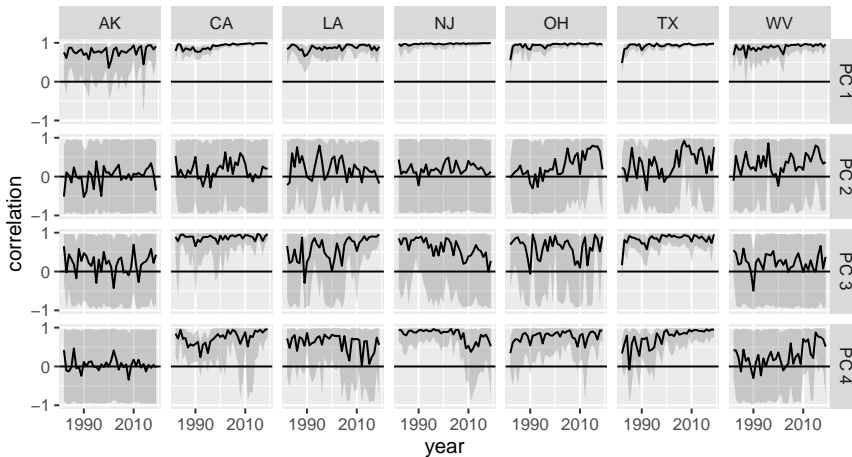


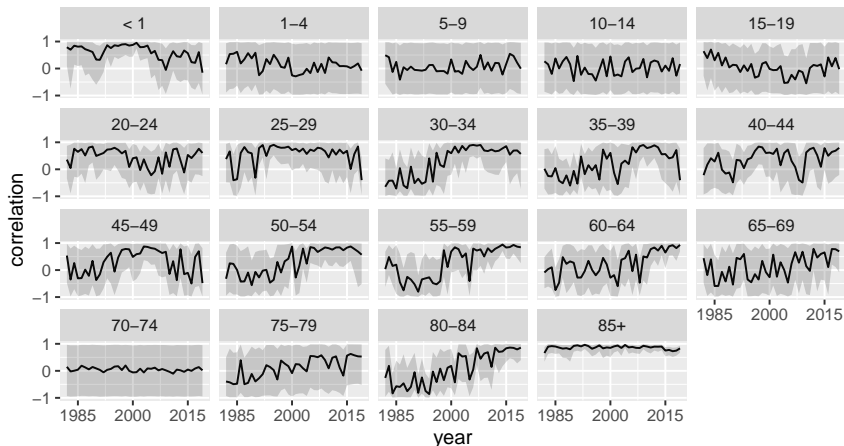
Figure 8 plots the posterior values of the between-sex principal component coefficient correlations extracted from $L_{i,t}^{(\beta)} L_{i,t}^{(\beta)\top}$ for a subset of states over time. Plots for all states are given in Appendix B. Each series includes the associated uncertainty intervals along with a horizontal line at zero. Across states, the first principal component has high correlation between the two sexes, reinforcing the notion that the baseline mortality for both sexes moves together. In contrast, the posterior distributions of the second principal component correlation terms mimic the corresponding uniform priors: They cover nearly the whole $(-1, 1)$ range for all states and years, suggesting that the data contain little information to support a relationship between sexes for this component. This is consistent with the idea that the accident hump is a predominantly male phenomenon, which would imply weak correlation between sexes. The third and fourth principal component correlations deviate by state. For Alaska, the relationships are weak at best; for California and Texas, there are strong or growing correlations; and New Jersey has declining correlations.

Figure 8: Time-series of posterior principal component correlations for seven states



In Figure 9, we plot the posterior medians and 95% credible intervals for the between-sexes correlations in the overdispersion terms captured by $L_{a,t}^{(\gamma)} L_{a,t}^{(\gamma)\top}$ for California as an example. We find that patterns vary substantially by age group. For the 85+ category, consistent high correlation is found. However, in the youth age groups, the posteriors again cover the whole range, suggesting that there is no meaningful evidence of correlation in the data. The most interesting patterns are those in the middle-age adult groups, where transitions from limited correlation to strong positive correlation are found.

Figure 9: Posterior medians and 95% credible intervals for γ correlation over time by age group for California



4.2 Validation

To formally evaluate the model, we perform out-of-sample validation exercises comparing the present model with the model proposed by Alexander, Zagheni, and Barbieri (2017). By comparing these two models, we can isolate the contributions of the between-subpopulation correlation matrices introduced. The validation exercises focus on five states (Alaska, California, Louisiana, New Jersey, and Texas) that capture the diversity in population size and number of counties across the United States.

For each state, we leave out 20% of the observed age-sex-year death counts for each county in the years 1982–2019 as an out-of-sample dataset, then execute the model on the remaining in-sample dataset. We then generate a distribution of deaths for all left-out observations using the posterior log-mortality rates for the corresponding county-age-sex-years. Finally, we calculate coverage at the 80%, 90%, and 95% nominal levels for the death uncertainty intervals as well as mean squared errors (MSE) and mean absolute deviations (MAD) between the median death estimates and the observed deaths. In Table 2, the results of this exercise by state are presented. Note that the model introduced in this paper is labeled as the ‘joint’ model and the version in Alexander, Zagheni, and Barbieri (2017) with independent sex modeling is denoted ‘independent’.

Table 2: Out-of-sample coverage and errors

State	Model	Cov80	Cov90	Cov95	MAD	MSE
Alaska	independent	0.884	0.943	0.973	1.931	14.447
Alaska	joint	0.887	0.946	0.974	1.899	12.614
California	independent	0.843	0.920	0.961	9.747	2004.650
California	joint	0.846	0.923	0.963	8.568	1194.914
Louisiana	independent	0.873	0.939	0.970	3.054	40.795
Louisiana	joint	0.874	0.940	0.972	2.921	30.257
New Jersey	independent	0.843	0.925	0.963	8.388	298.082
New Jersey	joint	0.850	0.930	0.964	7.838	238.458
Texas	independent	0.878	0.940	0.970	3.232	92.449
Texas	joint	0.879	0.941	0.970	3.030	61.618

The joint model consistently outperforms the independent model on the out-of-sample set across metrics. The level of outperformance is most pronounced in the larger states, notably California. The outperformance of the joint model on the out-of-sample set suggests that the subpopulation correlation structures materially improve subnational mortality estimation in this context. Further, as the out-of-sample exercises replicate an incomplete registration setting, it suggests that our model may have uses beyond the US county context for estimating subnational mortality by subpopulation in jurisdictions without complete data.

5. Discussion

In this paper we extended principal component-based methods to jointly estimate subnational mortality across subpopulations. This approach leverages the inherent structural mortality patterns associated with the individual principal components and extracts correlations between groups to offer insight into the joint movements of mortality trends across groups. The model centers on a regression-based framework with four principal components that allow core patterns in age-specific mortality to be captured. The principal component coefficients are modeled hierarchically to allow for information exchange within states. County-specific effects on each component are assumed to be correlated across subgroups within a county. Our approach is validated using both a simulation study and with out-of-sample exercises using real subnational mortality data. We find that the proposed model is well calibrated and that its errors are smaller than those of the independent model in out-of-sample exercises.

We illustrate the model through estimating US county-level sex-specific mortality rates. An investigation into the parameters of the model highlights general trends in US mortality as well as state specific patterns. Notably, while movements in baseline mortality tend to manifest in both sexes simultaneously, specific features of the mortality curve, such as elevated young adult mortality, appear independent by sex. These results can offer direction to those working to reduce mortality pressures. For example, when correlations are high, aggregate policies may be effective. However, in jurisdictions where correlations are low or zero, more targeted policies may be necessary. State-level parameter estimates also highlight clear geographic clustering of mortality patterns over time. Notably, there is evidence of stagnation in parts of the country with increasing early adult mortality particularly in the eastern states, and no improvement in middle-age mortality in the southern states.

We identify two directions for future extensions to this work. The first is to study how the model performs in real-world circumstances similar to the simulation study where there are more than two subpopulations and the composition of the population is not approximately balanced across groups – for example, race/ethnicity-based mortality data in US counties. Many counties are dominated by one subgroup, and thus the frequency of low or zero death data is significantly higher than with aggregate or sex-specific data. The second direction is to apply the model in contexts where there are differences in mortality data collection or death registration coverage. In many countries, complete mortality data, such as the US county data used here, are not available and death registration coverage can vary substantially by sex, geographical area, or time (Peralta et al. 2019; Basu and Adair 2021). Typically, male coverage is higher than female coverage. By studying relationships between sexes in the years and regions with higher quality data, one may be able to support estimation of female mortality rates by using the male data and the modeled correlations.

The code used for the simulation study and data analysis is available at <https://github.com/AmeerD/joint-mortality>. As we cannot share the US county-level mortality data, an illustrative sample dataset is provided.

6. Acknowledgments

This study was made possible by an award from the Center on the Economics and Demography of Aging (National Institute on Aging grant P30-AG012839) at the University of California, Berkeley. However, any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors alone and do not necessarily represent the official views of the National Institute on Aging. The authors would like to thank three anonymous reviewers and the journal's editor for helpful comments on an earlier version of this manuscript.

References

- Alexander, M. and Alkema, L. (2022). A Bayesian cohort component projection model to estimate women of reproductive age at the subnational level in data-sparse settings. *Demography* 59(5): 1713–1737. doi:10.1215/00703370-10216406.
- Alexander, M., Kiang, M.V., and Barbieri, M. (2018). Trends in black and white opioid mortality in the United States, 1979–2015. *Epidemiology* 29(5): 707–715. doi:10.1097/EDE.0000000000000858.
- Alexander, M., Zagheni, E., and Barbieri, M. (2017). A flexible Bayesian model for estimating subnational mortality. *Demography* 54(6): 2025–2041. doi:10.1007/s13524-017-0618-7.
- Arias, E., Escobedo, L.A., Kennedy, J., Fu, C., and Cisewski, J. (2018). U.S. small-area life expectancy estimates project: Methodology and results summary. *Vital and Health Statistics* 181: 1–40. https://www.cdc.gov/nchs/data/series/sr_02/sr02_181.pdf.
- Barford, A., Dorling, D., Smith, G.D., and Shaw, M. (2006). Life expectancy: Women now on top everywhere. *BMJ* 332(7545): 808. doi:10.1136/bmj.332.7545.808.
- Barnard, J., McCulloch, R., and Meng, X.L. (2000). Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica* 10(4): 1281–1311.
- Basu, J.K. and Adair, T. (2021). Have inequalities in completeness of death registration between states in India narrowed during two decades of civil registration system strengthening? *International Journal for Equity in Health* 20(195). doi:10.1186/s12939-021-01534-y.
- Beltrán-Sánchez, H., Finch, C.E., and Crimmins, E.M. (2015). Twentieth century surge of excess adult male mortality. *Proceedings of the National Academy of Sciences* 112(29): 8993–8998. doi:10.1073/pnas.1421942112.
- Bhutta, Z.A. (2016). Mapping the geography of child mortality: A key step in addressing disparities. *The Lancet Global Health* 4(12): e877–e878. doi:10.1016/S2214-109X(16)30264-9.
- Bijak, J. and Bryant, J. (2016). Bayesian demography 250 years after Bayes. *Population Studies* 70(1): 1–19. doi:10.1080/00324728.2015.1122826.
- Clark, S.J. (2019). A general age-specific mortality model with an example indexed by child mortality or both child and adult mortality. *Demography* 56(3): 1131–1159. doi:10.1007/s13524-019-00785-3.

- Coale, A.J., Demeny, P., and Vaughan, B. (1983). *Regional model life tables and stable populations: Studies in population (second edition)*. London: Academic Press. doi:10.1016/C2013-0-07295-7.
- Crimmins, E.M. and Zhang, Y.S. (2019). Aging populations, mortality, and life expectancy. *Annual Review of Sociology* 45: 69–89. doi:10.1146/annurev-soc-073117-041351.
- de Beer, J. (2012). Smoothing and projecting age-specific probabilities of death by TOPALS. *Demographic Research* 27(20): 543–592. doi:10.4054/DemRes.2012.27.20.
- Dwyer-Lindgren, L., Bertozzi-Villa, A., Stubbs, R.W., Morozoff, C., Kutz, M.J., Huynh, C., Barber, R.M., Shackelford, K.A., Mackenbach, J.P., van Lenthe, F.J., Flaxman, A.D., Naghavi, M., Mokdad, A.H., and Murray, C.J.L. (2016). US county-level trends in mortality rates for major causes of death, 1980–2014. *JAMA* 316(22): 2385–2401. doi:10.1001/jama.2016.13645.
- Dyrting, S. (2020). Smoothing migration intensities with P-TOPALS. *Demographic Research* 43(55): 1607–1650. doi:10.4054/DemRes.2020.43.55.
- Ezzati, M., Friedman, A.B., Kulkarni, S.C., and Murray, C.J.L. (2008). The reversal of fortunes: Trends in county mortality and cross-county mortality disparities in the United States. *PLOS Medicine* 5(4): 0557–0568. doi:10.1371/journal.pmed.0050066.
- Fenelon, A. (2013). Geographic divergence in mortality in the United States. *Population and Development Review* 39(4): 611–634. doi:10.1111/j.1728-4457.2013.00630.x.
- Gelman, A. and Rubin, D.B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science* 7(4): 457–472. doi:10.1214/ss/1177011136.
- Gjonça, A., Tomassini, C., and Vaupel, J. (1999). Male-female differences in mortality in the developed world. Rostock: Max Planck Institute for Demographic Research (MPIDR working paper WP-1999-009). doi:10.4054/MPIDR-WP-1999-009.
- Godwin, J. and Wakefield, J. (2021). Space-time modeling of child mortality at the Admin-2 level in a low and middle income countries context. *Statistics in Medicine* 40(7): 1593–1638. doi:10.1002/sim.8854.
- Gonzaga, M.R. and Schmertmann, C.P. (2016). Estimation of mortality rates by age and sex for small areas with topals regression: An application for Brazil in 2010. *Revista Brasileira De Estudos De População* 33(3): 629–652. doi:10.20947/S0102-30982016c0009.
- Gutin, I. and Hummer, R.A. (2021). Social inequality and the future of US life expectancy. *Annual Review of Sociology* 47: 501–520. doi:10.1146/annurev-soc-072320-100249.

- Harper, S., Riddell, C.A., and King, N.B. (2021). Declining life expectancy in the United States: Missing the trees for the forest. *Annual Review of Public Health* 42: 381–403. doi:10.1146/annurev-publhealth-082619-104231.
- Heligman, L. and Pollard, J.H. (1980). The age pattern of mortality. *Journal of the Institute of Actuaries* 107(1): 49–80. doi:10.1017/S0020268100040257.
- Hendi, A.S. (2015). Trends in US life expectancy gradients: The role of changing educational composition. *International Journal of Epidemiology* 44(3): 946–955. doi:10.1093/ije/dyv062.
- Hoffman, M.D. and Gelman, A. (2014). The No-U-Turn Sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* 15(47): 1593–1623. <http://jmlr.org/papers/v15/hoffman14a.html>.
- Human Mortality Database (2023). Max Planck Institute for Demographic Research (Germany), University of California, Berkeley (USA), and French Institute for Demographic Studies (France). <http://www.mortality.org>.
- Khana, D., Rossen, L.M., Hedegaard, H., and Warner, M. (2018). A Bayesian spatial and temporal modeling approach to mapping geographic variation in mortality rates for subnational areas with R-INLA. *Journal of Data Science* 16(1): 147–182.
- Kindig, D.A. and Cheng, E.R. (2013). Even as mortality fell in most US counties, female mortality nonetheless rose in 42.8 percent of counties from 1992 to 2006. *Health Affairs* 32(3): 451–458. doi:10.1377/hlthaff.2011.0892.
- Lee, R.D. and Carter, L.R. (1992). Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association* 87(419): 659–671. doi:10.2307/2290201.
- Lewandowski, D., Kurowicka, D., and Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis* 100(9): 1989–2001. doi:10.1016/j.jmva.2009.04.008.
- Li, N. and Lee, R. (2005). Coherent mortality forecasts for a group of populations: An extension of the Lee–Carter method. *Demography* 42(3): 575–594. doi:10.1353/dem.2005.0021.
- Lin, T. and Tsai, C.C.L. (2022). Hierarchical Bayesian modeling of multi-country mortality rates. *Scandinavian Actuarial Journal* 2022(5): 375–398. doi:10.1080/03461238.2021.1979639.
- Lu, Q., Hanewald, K., and Wang, X. (2021). Subnational mortality modelling: A Bayesian hierarchical model with common factors. *Risks* 9(11): 1–21. doi:10.3390/risks9110203.

- Masters, R. and Aron, L. (2021). Changes in US life expectancy in the wake of COVID-19: Differences by race/ethnicity and relative to other high-income countries [unpublished manuscript].
- Mercer, L.D., Wakefield, J., Pantazis, A., Lutambi, A.M., Masanja, H., and Clark, S. (2015). Space-time smoothing of complex survey data: Small area estimation for child mortality. *The Annals of Applied Statistics* 9(4): 1889–1905. doi:10.1214/15-AOAS872.
- Montez, J.K., Beckfield, J., Cooney, J.K., Grumbach, J.M., Hayward, M.D., Koytak, H.Z., Woolf, S.H., and Zajacova, A. (2020). US state policies, politics, and life expectancy. *The Milbank Quarterly* 98(3): 668–699. doi:10.1111/1468-0009.12469.
- Murray, C.J.L., Kulkarni, S.C., Michaud, C., Tomijima, N., Bulzacchelli, M.T., Iandiorio, T.J., and Ezzati, M. (2006). Eight Americas: Investigating mortality disparities across races, counties, and race-counties in the United States. *PLOS Medicine* 3(9): 1513–1524. doi:10.1371/journal.pmed.0030260.
- National Center for Health Statistics (2023). Detailed mortality files 1982–2019, as compiled from data provided by the 57 vital statistics jurisdictions through the Vital Statistics Cooperative Program. Hyattsville: National Center for Health Statistics.
- Neal, R.M. (2011). MCMC using Hamiltonian dynamics. In: *Handbook of Markov Chain Monte Carlo*. New York: Chapman and Hall/CRC: 113–162. doi:10.1201/b10905.
- Noymer, A. and Van, V. (2014). Divergence without decoupling: Male and female life expectancy usually co-move. *Demographic Research* 31(51): 1503–1524. doi:10.4054/DemRes.2014.31.51.
- Peralta, A., Benach, J., Borrell, C., Espinel-Flores, V., Cash-Gibson, L., Queiroz, B.L., and Marí-Dell’Olmo, M. (2019). Evaluation of the mortality registry in Ecuador (2001–2013) – social and geographical inequalities in completeness and quality. *Population Health Metrics* 17(3): 1–12. doi:10.1186/s12963-019-0183-y.
- Raftery, A.E., Lalic, N., and Gerland, P. (2014). Joint probabilistic projection of female and male life expectancy. *Demographic Research* 30(27): 795–822. doi:10.4054/DemRes.2014.30.27.
- Rau, R. and Schmetmann, C.P. (2020). Lebenserwartung auf Kreisebene in Deutschland. *Deutsches Ärzteblatt International* 117(29–30): 493–499. doi:10.3238/arztebl.2020.0493.
- Renshaw, A.E. and Haberman, S. (2003). Lee–Carter mortality forecasting with age-specific enhancement. *Insurance: Mathematics and Economics* 33(2): 255–272. doi:10.1016/S0167-6687(03)00138-0.

- Schmertmann, C.P. and Gonzaga, M.R. (2018). Bayesian estimation of age-specific mortality and life expectancy for small areas with defective vital records. *Demography* 55(4): 1363–1388. doi:10.1007/s13524-018-0695-2.
- Sehgal, N.J., Yue, D., Pope, E., Wang, R.H., and Roby, D.H. (2022). The association between COVID-19 mortality and the county-level partisan divide in the United States. *Health Affairs* 41(6): 853–863. doi:10.1377/hlthaff.2022.00085.
- Seligman, B., Greenberg, G., and Tuljapurkar, S. (2016). Convergence in male and female life expectancy: Direction, age pattern, and causes. *Demographic Research* 34(38): 1063–1074. doi:10.4054/DemRes.2016.34.38.
- Ševčíková, H. and Raftery, A.E. (2021). Probabilistic projection of subnational life expectancy. *Journal of Official Statistics* 37(3): 591–610. doi:10.2478/jos-2021-0027.
- Sharrow, D.J., Clark, S.J., and Raftery, A.E. (2014). Modeling age-specific mortality for countries with generalized HIV epidemics. *PLoS One* 9(5): e96447. doi:10.1371/journal.pone.0096447.
- Song, I. and Luan, H. (2022). The spatially and temporally varying association between mental illness and substance use mortality and unemployment: A Bayesian analysis in the contiguous United States, 2001–2014. *Applied Geography* 140: 102664. doi:10.1016/j.apgeog.2022.102664.
- Stan Development Team (2022). RStan: The R interface to Stan. R package version 2.21.3. <https://mc-stan.org>.
- Stevenson, J.M. and Olson, D.R. (1993). Methods for analysing county-level mortality rates. *Statistics in Medicine* 12(3–4): 393–401. doi:10.1002/sim.4780120320.
- United States Census Bureau (2022). Annual estimates of the resident population by single year of age and sex: April 1, 2020 to July 1, 2022 (SC-EST2022-SYASEX). Washington: Population Estimates and Projections Unit. <https://www.census.gov/data/datasets/time-series/demo/pepopest/2020s-state-detail.html>.
- United States Mortality Database (2023). University of California, Berkeley (USA). <http://usa.mortality.org>.
- Wang, H., Schumacher, A.E., Levitz, C.E., Mokdad, A.H., and Murray, C.J. (2013). Left behind: Widening disparities for males and females in US county life expectancy, 1985–2010. *Population Health Metrics* 11(8). doi:10.1186/1478-7954-11-8.
- Woolf, S.H. and Schoemaker, H. (2019). Life expectancy and mortality rates in the United States, 1959–2017. *JAMA* 322(20): 1996–2016. doi:10.1001/jama.2019.16932.

Appendices

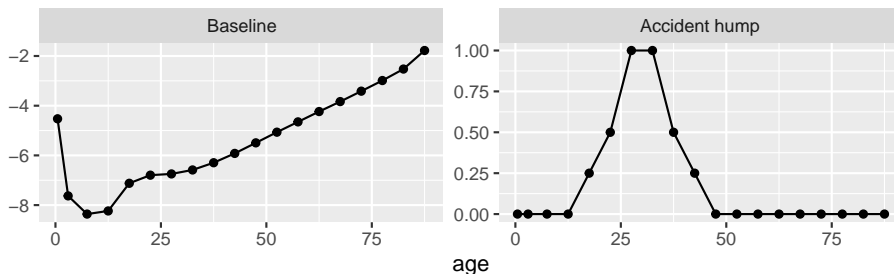
A. Simulation study details

To illustrate the model's ability to estimate mortality rates and extract mortality patterns across subgroups, we simulate population and death data for populations composed of five subgroups of varying sizes. The simulation study data includes population and mortality data for 10 years, 25 subnational areas, and 5 population subgroups. For each year-area-age subgroup, we construct the population data as follows:

1. The total population in year 1 in each area is set as 100,000 times the area index, which ranges from 1 to 25.
2. The total population in each area in subsequent years is computed by increasing the total population in each area by 1% annually.
3. Total populations in each year-area are allocated to the 19 age groups using a jittered version of the age distribution of Los Angeles County in California.
4. Each year-area-age population is then disaggregated into the 5 subgroups using the following proportions: 50% in Group A, 20% in Group B, and 10% in Groups C, D, E.

To generate corresponding mortality data for the populations in each year-area-age subgroup, we first construct true log-mortality curves for each year-area subgroup. The mortality curves are constructed using a linear combination of two standard curves: the first representing a baseline mortality curve, and the second containing an accident hump. Plots of the standard curves are given in Figure A-1.

Figure A-1: Simulation study mortality curve components



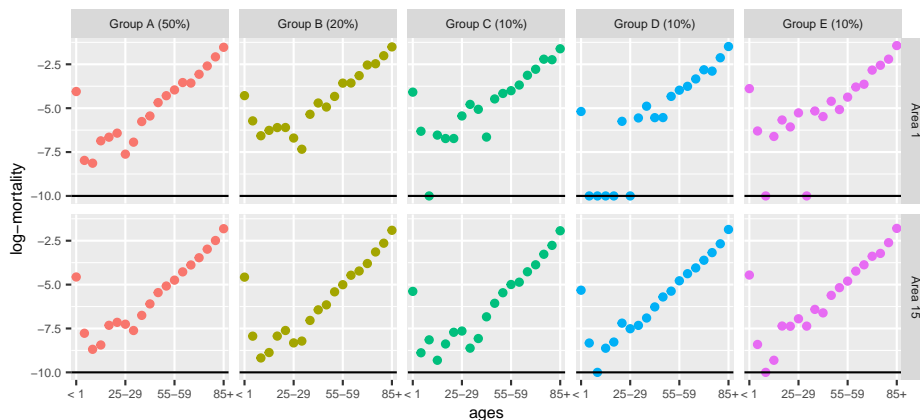
The log-mortality curves are then constructed with the following steps:

1. The baseline mortality curve is multiplied by a random coefficient sampled from a 5-dimensional multivariate normal distribution with mean 1_5 and a year-dependent covariance matrix. The covariance matrix describes the dependence between the 5 subgroups. It is assigned a constant variance of 0.1^2 and one of the correlation structures plotted in Figure 4a. In years 1 to 3, the independent correlation matrix is used. In years 4 to 6, the exchangeable correlation matrix (i.e., constant off-diagonal values) is used. In years 7 to 10, the unstructured correlation matrix based on state-level observed race/ethnicity-based US mortality is used.
2. An accident hump is added to the mortality curve by adding the accident hump standard curve multiplied by random coefficients sampled from a 5-dimensional multivariate normal distribution with mean 0_5 , and covariance matrix with a constant variance of 0.5^2 and one of the correlation matrices plotted in Figure 4a. The matrices used in each year mirror those used for the baseline mortality standard.

Finally, observed deaths for each year-area-age subgroup are generated from the log-mortality curves using a Poisson likelihood with rate equal to the corresponding population multiplied by the exponential of the corresponding log-mortality rate.

Examples of simulated data in two subnational areas are given in Figure A-2. Points correspond to log-mortality observations (i.e., the log of observed deaths divided by population). Note that zero death observations are encoded as -10 on the log-scale.

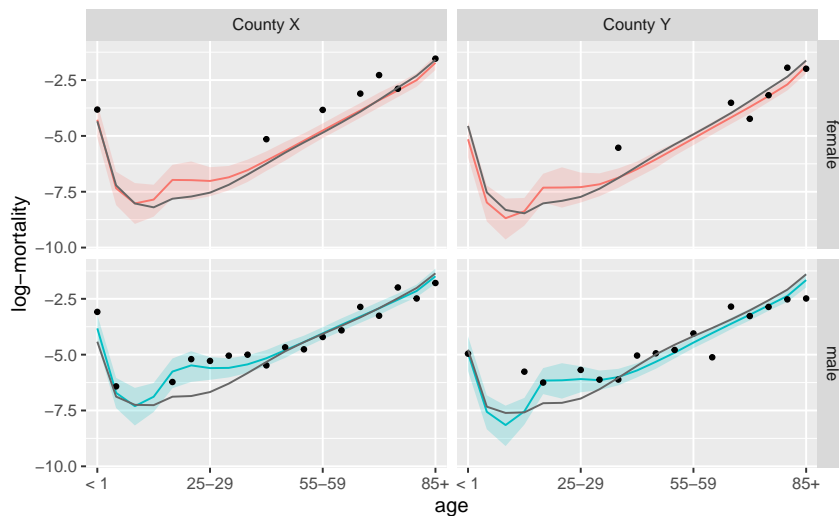
Figure A-2: Examples of simulated log-mortality observations



B. Additional results

Figure A-3 plots two county-years where the amount of nonzero death data varies by sex. The observed log-mortality rates for age groups with nonzero deaths are presented as points, the estimated state means for each sex are denoted by the gray lines, and the male and female county-level estimates and 95% credible intervals are presented in blue and red, respectively. In both county-years, the majority of age groups had zero female deaths in the year whereas the male data is nearly complete. As opposed to defaulting to the female state-level means, estimates for female mortality are also informed by the behavior of male mortality relative to the male state-level means.

Figure A-3: Examples of US counties with limited female but sufficient male data



Figures A-4 and A-5 plot maps of the posterior medians and 95% credible intervals of the estimated state-level coefficients for the third and fourth principal components for male and female populations for all years.

Figures A-6, A-7, A-8, and A-9 plot the posterior medians and 95% credible intervals for the principal component coefficient correlations between sexes over time. As expected, strong correlation in the first principal component is found in most states whereas there is negligible correlation in the second principal component. Patterns in the third and fourth principal component coefficient correlations vary regionally with some evidence for stronger correlations in the fourth principal component coefficients in the Great Lakes region and surrounding states.

Figure A-4: Posterior medians and 95% credible intervals for the third principal component (μ_β)



Figure A-5: Posterior medians and 95% credible intervals for the fourth principal component (μ_{β})

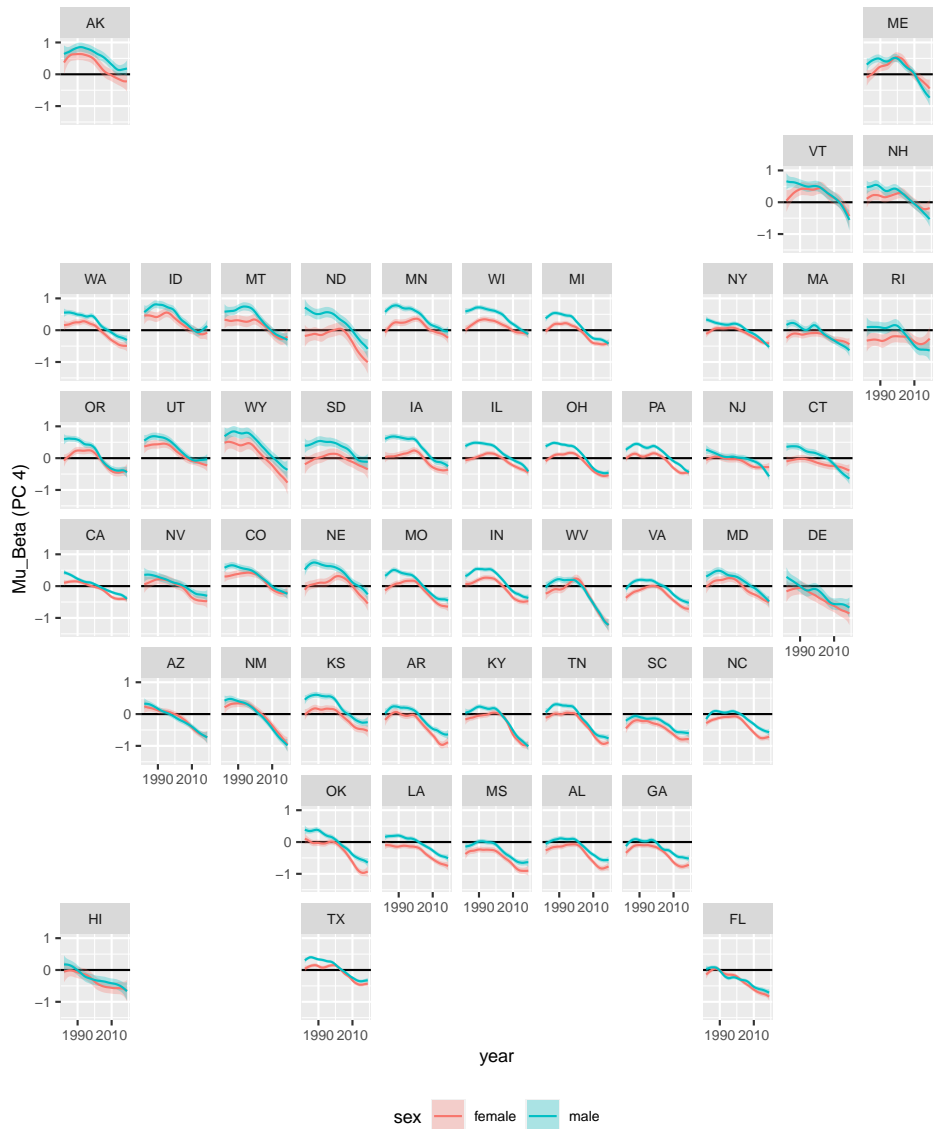


Figure A-6: Time-series of posterior medians and 95% credible intervals for the first principal component correlations



Figure A-7: Time-series of posterior medians and 95% credible intervals for the second principal component correlations

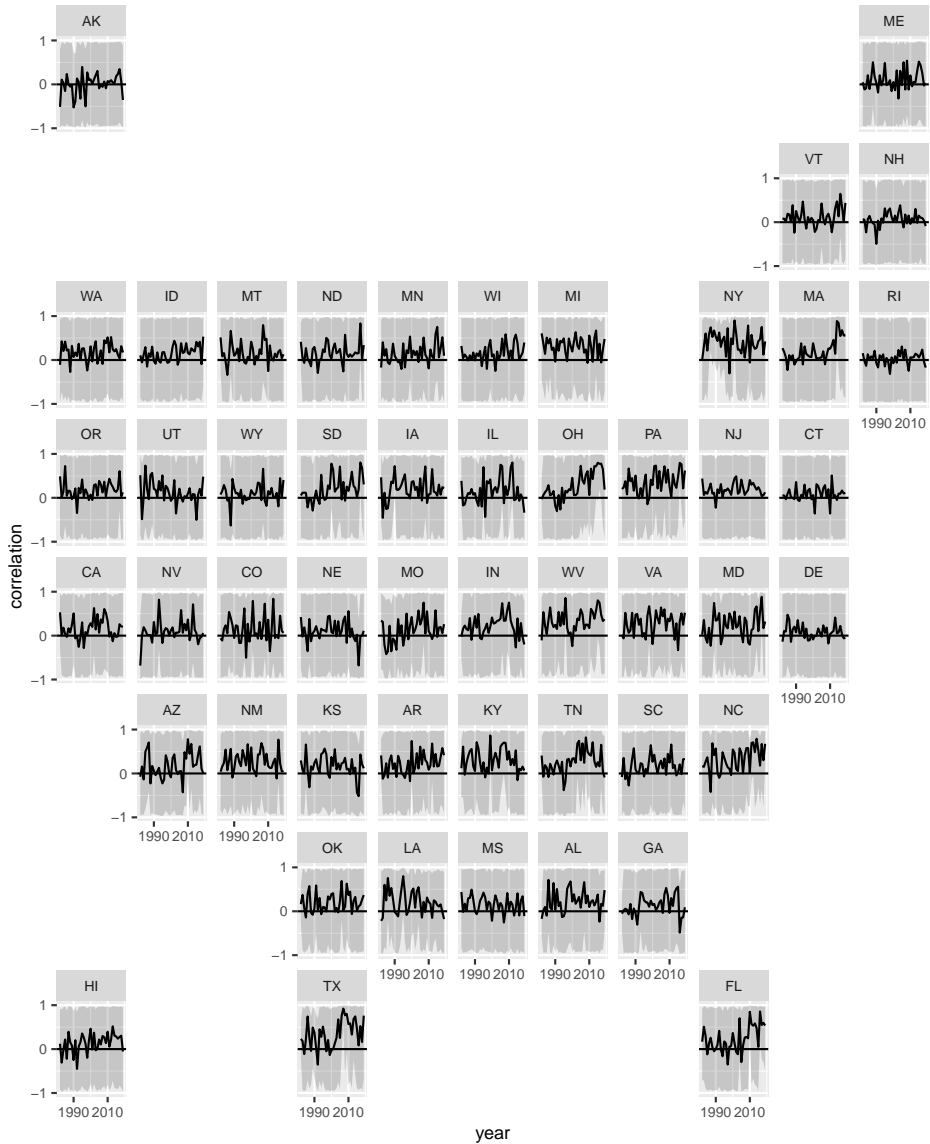


Figure A-8: Time-series of posterior medians and 95% credible intervals for the third principal component correlations



Figure A-9: Time-series of posterior medians and 95% credible intervals for the fourth principal component correlations

