

# GAUSSIAN PROCESS MODELS FOR MORTALITY RATES AND IMPROVEMENT FACTORS

BY

MIKE LUDKOVSKI, JIMMY RISK AND HOWARD ZAIL

## ABSTRACT

We develop a Gaussian process (GP) framework for modeling mortality rates and mortality improvement factors. GP regression is a nonparametric, data-driven approach for determining the spatial dependence in mortality rates and jointly smoothing raw rates across dimensions, such as calendar year and age. The GP model quantifies uncertainty associated with smoothed historical experience and generates full stochastic trajectories for out-of-sample forecasts. Our framework is well suited for updating projections when newly available data arrives, and for dealing with “edge” issues where credibility is lower. We present a detailed analysis of GP model performance for US mortality experience based on the CDC (Center for Disease Control) datasets. We investigate the interaction between mean and residual modeling, Bayesian and non-Bayesian GP methodologies, accuracy of in-sample and out-of-sample forecasting, and stability of model parameters. We also document the general decline, along with strong age-dependency, in mortality improvement factors over the past few years, contrasting our findings with the Society of Actuaries (SOA) MP-2014 and -2015 models that do not fully reflect these recent trends.

## KEYWORDS

mortality, mortality improvement, mortality modeling, Gaussian processes

## 1. INTRODUCTION

Publishing of pension mortality tables and mortality improvement factors for use by actuarial professionals and researchers in longevity risk management is a major endeavor of the actuarial professional organizations. In the United States, the Society of Actuaries (SOA) runs the Retirement Plans Executive Committee (RPEC); its most recent publication is known as the RP-2014 mortality tables and the MP-2015 improvement scales (SOA, 2014a,b). In the United Kingdom, annual tables are released in the form of the Continuous Mortality Investigation reports (Continuous Mortality Investigation, 2015). Being official proposals of the actuarial Societies, such tables enjoy wide use and are also heavily used in

the valuation of pension and post-retirement medical liabilities. For example, in the United States, the SOA tables have been included by the Internal Revenue Service for the purposes of the Pension Protection Act of 2005, or by the Congressional Budget Office for long-term forecasts.

The basic aim in constructing the tables is to convert the raw mortality data into a graduated table of yearly mortality rates and improvement factors, broken down by age and gender. Since the goal is to forecast future mortality from retrospective experience, the process involves two fundamental steps: *smoothing* raw data to remove random fluctuations resulting from finite data sizes; and *extrapolating* future rates. To maximize actuarial credibility of the tables, graduation techniques are applied, in particular for estimating mortality improvement trends based on past experience and then projecting those trends into future years. For example, see the RPEC reports SOA (2014a,b) for the full description of constructing the US tables/scales, as well as more general SOA longevity studies in Purushotham *et al.* (2011) and Rosner *et al.* (2013).

In the present article, we propose a new methodology to graduate mortality rates and generate mortality improvement scales within a single statistical model. More precisely, we advocate the use of Gaussian process (GP) regression, a type of Bayesian nonparametric statistical model. Our aim is to provide a data-driven procedure that produces an alternative to existing methods while enjoying a number of important advantages:

- The GP framework is Bayesian, offering rich uncertainty quantification. The model produces mortality curves smoothed over multiple dimensions, as well as credible intervals that quantify the uncertainty of these curves. This is generated for in-sample smoothing and out-of-sample forecasts. In their basic form, the latter forecasts are Gaussian, allowing for a simple interpretation of the uncertainty by the actuary. Moreover, the GP model is able to generate stochastic *trajectories* of future mortality experience. We demonstrate this projection over both age and calendar year, but the GP model can be consistently applied over higher dimensional data as well. From this, full predictive distributions for annuity values, life expectancies, and other life contingent cash-flows can be produced. Such analyses can provide core components of stress testing and risk management of mortality and longevity exposures.
- Using GPs leads to unified modeling of mortality rates and mortality improvement; one may analytically differentiate the mortality surface to obtain mortality trends (and corresponding credible bands) over calendar years. This structure offers a coherent approach to both tables, jointly quantifying uncertainty in rates and improvement factors.
- Standard graduation techniques are sensitive to edge issues, i.e., the experience in the latest few years. For example, to achieve a better prediction, the MP-2015 method extrapolates rates from 2011 onwards, effectively excluding the last several years of data (as of this writing, Center for Disease Control

(CDC) data go up to 2014). In contrast, our GP approach intrinsically handles the specific shape of the data and is well suited to incorporating missing data. Therefore, dropping the “edge years” is not necessary with GP, with its self-adjusting credible bands.

- The GP approach provides natural “updating” of mortality tables in terms of incorporating the latest mortality experience. The end users can easily update the tables, no longer requiring reliance on official updates.

To recapitulate, the main contribution of this article is to propose the use of GP regression for constructing mortality tables and improvement factors. While being a relatively new “machine learning” paradigm, the underlying statistical methodology and most crucially the software implementation has matured significantly in the past decade. To wit, all of the case studies below have been implemented straightforwardly using publicly available, free, well-documented software and required only basic programming skills. As a companion to the manuscript, we provide an R markdown (.rmd) notebook file containing a concise version of the tables and figures produced in this document, along with the R code used to produce them; please see [github.com/jimmyrisk/GPmortalityNotebook](https://github.com/jimmyrisk/GPmortalityNotebook). Our main focus is on the GP methodology, and accordingly we concentrate on *describing* the important components of this framework in the context of mortality modeling. We illustrate its application over several mortality datasets and show that GPs are competitive with existing methods in terms of their performance and predictive power.

From the empirical direction, our data-driven analysis sheds light on the question of latest mortality experience, whereby mortality improvements appear to have significantly moderated from past trends. Specifically, after implementing the above framework on the latest US mortality experience, we document that as of 2015, mortality improvement factors are (statistically) zero, and possibly *negative* for Ages 55–70 since as early as 2012. These estimates diverge significantly from SOA projections embedded in MP-2015 that continue to bake in past improvements. Lower mortality improvement rates would have a material impact across the pension industry. This paper offers statistical support to the anecdotal demographic evidence of declining US longevity and calls into question traditional backward-looking methods for constructing mortality improvement factors.

### 1.1. Comparison to other approaches

Mortality experience is summarized by a mortality surface, indexed by Age (rows  $i$ ) and calendar Year (columns  $j$ ). Typical data consists of two matrices  $D$  and  $E$  (or  $L$ ), listing the number of deaths  $D$ , exposed-to-risk  $E$  or the mid-year population  $L$ , respectively. In the first step, one postulates a relationship between the individual elements of these matrices,  $D_{ij}$  and  $E_{ij}$ , in terms of the latent (logarithmic) mortality state  $\mu_{ij}$ . In the second step, one estimates  $\mu_{ij}$  through a statistical fitting approach. We may identify two classes of estimation: (i) data-driven models that infer  $\mu_{ij}$  by statistical smoothing techniques; (ii) factor

models that express  $\mu_{ij}$  in terms of several one-dimensional indices. For example, in Age-Period-Cohort (APC) models those factors are additive and correspond to Age, Year and Cohort effects; in Lee-Carter (Lee and Carter, 1992) models they correspond to Age, Year and an Age-Year interaction term. A common distinction is to assume a non-smooth evolution of the mortality surface in time, coupled with a smooth Age effect. The latter Age-modulating terms are then fitted nonparametrically by maximum likelihood, or given a fixed functional form, such as linear or quadratic in Age (Cairns *et al.*, 2006; Hunt and Blake, 2014). Imposing an underlying one-dimensional structure facilitates interpretation of the fitted output, but potentially leads to model risk. In contrast, the data-driven methods, dating back to the classical graduation technique of Whittaker (1922), maintain an agnostic view of mortality experience, and solely focus on removing random fluctuations in observed deaths. Modern frameworks typically work with various types of splines, extending the seminal work by Currie *et al.* (2004) (see also a modern software implementation in Camarda 2012). Here, the main challenge is appropriate smoothing across both Age and Year dimensions; some of the proposed solutions include constrained and weighted regression splines (Hyndman and Ullah, 2007), extensions to handle cohort effects that generate “ridges” (Dokumentov and Hyndman, 2014) and a spatio-temporal kriging approach (Debón *et al.*, 2010). After completing the first version of this work, we learned of the independently executed PhD thesis of Wu (2016) that among other things also considered a GP model for mortality. A mixed strategy of first smoothing the data nonparametrically, and then inferring underlying factor structure was proposed and investigated in Hyndman and Ullah (2007). Finally, we also mention Bayesian approaches (Czado *et al.*, 2005; Girosi and King, 2008) that replace MLE-based point estimates with a posterior distribution of the mortality rate. To date, there is no consensus on which framework is the most appropriate. For example, the influential study by Cairns *et al.* (2009) considered eight different mortality models. Another recent study by Currie (2016) looked at 32 models, nesting the former eight.

A further reason for the large number of models is the use of different link functions (log-Poisson, logit-Poisson, logit-Binomial, etc.), that connect the logarithmic mortality state to deaths and exposures. These choices correspond to using different generalized linear models (GLM) and affect the optimization procedure (usually some variant of maximum likelihood) for model calibration. The Binomial model is defined as  $D_{ij} \sim \text{Bin}(E_{ij}, e^{\mu_{ij}})$  (Hyndman and Ullah, 2007); the Poisson model  $D_{ij} \sim \text{Poisson}(L_{ij}e^{\mu_{ij}})$  (Renshaw *et al.*, 1996; Sithole *et al.*, 2000; Brouhns *et al.*, 2002) and the Gaussian model  $\frac{D_{ij}}{E_{ij}} \sim \mathcal{N}(e^{\mu_{ij}}, \sigma^2 E_{ij})$  (Girosi and King, 2008). A related issue is regularization of the estimated factors that can be achieved via penalization, see Delwarde *et al.* (2007) and Currie (2013).

In terms of forecasting future mortality, a popular strategy is to differentiate the treatment of the Age index, which is incorporated directly into the mortality state and smoothed appropriately, vis-a-vis the Year index, whose

impact is estimated statistically using time-series techniques. This is the basic idea of Lee-Carter models, which construct a time-series process for the Year factor(s) to extrapolate mortality trends and assess forecast uncertainty. More generally, this can be viewed as a principal component approach, expressing the Age-effect as a smooth mortality curve  $\mu_t(x_{ag})$  in the age dimension  $x_{ag}$ , fitted via functional regression or singular value decomposition techniques, and then describing the evolution of this curve over time (Renshaw and Haberman, 2003; Hyndman and Ullah, 2007) as a multivariate time-series. In contrast, in the pure smoothing methods, all covariates are given equal footing, and forecasting is done by extrapolating the fitted *surface* to new input locations.

Precise methods for constructing mortality tables are not without controversy, especially when it comes to extreme age longevity or future forecasts. Ideally one ought to just let the “data speak for itself”. However, this is in fact a very challenging issue, not least because the question of predictive forecasting must acknowledge that any given fixed forecast is only a *point estimate*, and that there is always an element of uncertainty around the prediction. A common paradigm is to specify a stochastic model for mortality that directly prescribes future uncertainty. This is especially relevant for risk management or pricing applications, where the actuary wishes to incorporate (and hopefully manage) mortality risks. However, most stochastic mortality frameworks are “reduced-form” in the sense of specifying a low-dimensional stochastic system with just a few parameters/degrees-of-freedom. For implementation, one “calibrates” the model to data by minimizing e.g., the mean-squared error. In contrast, the RP-2014 mortality table is bottom-up, aiming to directly specify the full mortality experience with minimal a priori specifications. Relative to these two basic strategies, the approach proposed in this article views uncertainty in forecasts as intrinsic to the statistical model, so that all credible bands are obtained simultaneously both in-sample and out-of-sample.

Our proposed approach recasts mortality surface calibration as a *spatial regression* task. Originating in environmental/spatial statistics, GP regression (also known as kriging) takes a functional nonparametric approach to learning the latent response surface  $f$  (Cressie, 2015). Inference of  $f$  is viewed as conditioning on the observed data and fitting the GP concentrates on estimating the spatial dependence. Recently, GP models have gained currency as a machine learning tool for spatio-temporal forecasting thanks to their ability to capture complicated nonlinear dynamics with a high degree of analytic tractability and a minimum of tunable hyperparameters. An introduction to the vast GP landscape can be found in the monograph by Williams and Rasmussen (2006).

## 1.2. Mortality dataset

Our study is US-centric and originated from discussions of the SOA’s MP-2014 and successor tables. There has been some controversy that the scale excluded more recent trends, specifically a slowing of mortality improvement that was not fully reflected in the MP-2014 tables. Indeed, a year later, the SOA updated

TABLE 1

EXCERPT OF CDC MORTALITY DATA TO COMPARE EXPOSURES AND MORTALITY RATES OVER AGES AND GENDER FOR CALENDAR YEAR 2011.

Inputs $x^n$		Log Mortality Rate $y^n$		Mortality Rate $\exp(y^n)$	
Age ( $x_{ag}^n$ )	Year ( $x_{yr}^n$ )	Male	Female	Male	Female
50	2011	-4.931	-5.437	0.00722	0.00435
64	2011	-4.264	-4.707	0.01406	0.00901
74	2011	-3.435	-3.821	0.03222	0.02191
84	2011	-2.408	-2.714	0.08999	0.06625

*Mortality* is the observed proportion  $D^n / L^n$  of the deceased during the Year relative to the mid-year population.

the MP-2014 tables to the MP-2015 tables to include two additional years of mortality experience, and the new tables did in fact reflect a material drop in mortality improvement. In the interim, the CDC has also released new data showing a continued decline in mortality improvement levels.

The mortality data we use comes from CDC. The CDC data covers Ages 0–84 and goes up to 2014 as of the time of writing. For each cell of the table, the CDC data specifies the raw mortality rate for the exposed population. The mid-year exposures  $L_{ij}$  are based on inter-censal estimates interpolated based on the 2000 and 2010 census counts. Thus,  $e^{\mu_{ij}}$  corresponds to central death rates. Table 1 provides a snapshot of the latest year of CDC data (2014). The rapid decrease in sample size causes large variability in reported mortality rates at extreme ages. For a visual representation, two representative years of raw CDC data for Males Aged 60–70 are plotted as the solid lines in Figure 1 in Section 3.1. The figure shows the (super-) exponential increase in mortality with respect to age, along with a clear need for data smoothing.

As our training dataset, we used the CDC database covering Ages 50–84 in years 1999–2014. (Another data source is provided by Social Security Administration (SSA) and was utilized by RPEC.) Since our main aim is to obtain the *present* mortality rates and to forecast short-term calendar trend (through estimating mortality improvement factors) most relevant for actuarial applications, we only consider older ages and recent years. Our analysis targets the early retired group; additional challenges related to handling very young (e.g., infant mortality) and very old 85+ Ages are discussed in Section 4.1. Furthermore, since we approach mortality as a non-stationary surface evolving in Age and Year, we discard most of 20th century data, as distant mortality experience is less influential for our analysis. To understand the impact of excluding some data, we also considered several subsets listed in Table 2.

In comparison to our dataset, the most recent MP-2015 scales incorporate actual smoothed rates up to 2010 with projections thereafter. However, the CDC already provides actual mortality experience up to 2014. Further results are provided in the appendix based on matching US Female data. Additionally, in

TABLE 2  
DATASETS USED IN ANALYSIS.

Set Name	Training Set	Test Set
All Data	1999–2014, Ages 50–84	N/A: In-Sample
Subset I	1999–2010, Ages 50–84	2011–2014, Ages 50–84
Subset II	1999–2010, Ages 50–84 and 2011–2014, Ages 50–70	2011–2014, Ages 71–84
Subset III	1999–2010, Ages 50–70	2011–2014, Ages 71–84

Mortality data is taken from CDC as described in Section 1.2.

the online supplement, we present results for Japan and UK Males and Females based on Human Mortality Database (HMD) (Wilmoth and Shkolnikov, 2010) datasets. To keep the analysis consistent across countries, in all cases, we worked with the equivalent of All Data, i.e., Ages 50–84 and years 1999–2014 (HMD contains additional years and ages if desired).

## 2. GAUSSIAN PROCESS REGRESSION FOR MORTALITY TABLES

In this paper, we focus on analyzing mortality rates over a two-dimensional input space, namely Age and Year. The mortality data is viewed as a table of  $N$  “cells” (see rows of Table 1), represented by inputs  $x^n$  and outputs or responses  $y^n$ ,  $n = 1, \dots, N$ . In our case,  $x^n$  is in fact a tuple and represents the pair  $(x_{ag}^n, x_{yr}^n)$ . For example,  $x^n = (78, 2016)$  is the input for “78-year old in 2016” cell. We use the logarithmic central mortality rate for  $y^n$ , namely  $y^n = \log(D^n/L^n)$  where  $D^n$  and  $L^n$  represent the annual deaths and midyear count of lives, respectively, for the  $n$ th cell. The overall inputs  $\mathbf{x} = x^{1:N}$  and observations  $\mathbf{y} = y^{1:N}$  are denoted by boldface and aggregated into the mortality dataset  $\mathcal{D} = (\mathbf{x}, \mathbf{y})$ . Superscripts identify individual inputs/outputs, subscripts distinguish coordinates, e.g.,  $x_{ag}^n$ .

**Remark.** *This point of view treats calendar Year as simply another covariate and is easily extendible to further input dimensions, such as Select Period, etc. Also the format easily allows for missing cells, which, for example, is a common issue for dealing with extreme Ages (95+).*

### 2.1. Basics of Gaussian processes

In traditional mortality regression, a parametric function,  $f$ , is postulated that maps the inputs  $\mathbf{x}$  to the noisy measurements of the log-mortality rate,  $\mathbf{y}$ . A cell is modeled as

$$y^i = f(x^i) + \epsilon^i, \tag{1}$$

where  $\epsilon^i$  is the error term. With a GP, the function  $f$  is deemed to be latent and is modeled as a random variable. Consequently, a GP is defined as a set of random variables  $\{f(x)|x \in \mathbb{R}^d\}$  where any finite subset has a multivariate Gaussian distribution with mean  $\mathbf{m}(\cdot)$  and covariance  $\mathbf{C}(\cdot, \cdot)$ . That is for any

$n$ -tuple  $\mathbf{x} = x^{1:n}$ :

$$f(x^1), \dots, f(x^n) \sim \mathcal{N}(\text{mean} = \mathbf{m}(\mathbf{x}), \text{covariance} = \mathbf{C}(\mathbf{x}, \mathbf{x})).$$

In shorthand, we write  $f(\mathbf{x}) \sim GP(\mathbf{m}(\mathbf{x}), \mathbf{C}(\mathbf{x}, \mathbf{x}))$ . An important concept of a GP is that each mortality rate is correlated with every other mortality rate: above,  $\mathbf{C}$  is a  $n \times n$  matrix with entry  $C(x^i, x^j)$  representing the covariance between the  $i$ th and  $j$ th cells.

**Remark.** *We emphasize that the assumption that  $f(\cdot)$  forms a GP is solely a statistical representation of the mortality surface, similar for example to assuming that  $\mu$  can be described in terms of splines. Practically, the assumption is about the shape of  $\mu(\cdot)$ ; for typical kernels the corresponding functional space is dense in the class of continuous functions. It can be compared to the APC models that decompose  $\mu$  as a sum of one-dimensional factors. A more relevant question concerns the mapping from the latent surface to the observed mortality; (1) assumes an additive noise structure like in the classical least-squares framework. Because mortality experience comes from count data, a GLM could be viewed as better suited, see Section 2.2.1.*

Once we collect data  $\mathcal{D}$ , the next step is to determine the posterior distribution for  $f$ , namely  $p(f|\mathcal{D})$ . That is, we want to know the distribution of mortality rates, given the experience data. Using Bayes' rule, we have

$$p(f|\mathcal{D}) \propto p(\mathbf{y}|f, \mathbf{x})p(f) = \{\text{likelihood}\} \cdot \{\text{prior}\},$$

where  $p(\mathbf{y}|f, \mathbf{x})$  is the “likelihood” and  $p(f)$  the “prior”. To complete the definition of the GP, we therefore need to define the “prior”,  $p(f)$ . This is equivalent to setting the initial assumptions for mean function  $\mathbf{m}$  and covariance function  $\mathbf{C}$ .

*The Prior Mean Function:* the prior mean  $m(x)$  stands in for our belief about mortality rate at input  $x$  in the absence of any historic data. We might, for example, define  $m(\cdot)$  as a Gompertz or Makeham curve in the age coordinate  $x_{ag}$ . However, we will show that the choice of  $m(\cdot)$  has little impact on the output of the GP model for purposes of *in-sample* smoothing. Even if we set  $m(x) = 0$  or  $m(x) = \beta_0$  for some constant  $\beta_0$  and for all  $x$ , the results will be largely unaffected, since the posterior mean is largely dominated by the impact of the data. However, for purposes of *out-of-sample* projections, we will conversely show that a more realistic choice of  $m(\cdot)$  is required for long-term mortality projections.

*The Covariance Function:* A core concept of a GP is that for any cells  $i, j$ , if  $x^i$  and  $x^j$  are deemed to be “close”, then we would expect the outputs,  $y^i$  and  $y^j$ , to be “close” too. For example, the mortality rate for a 60-year old in 2016 ( $x^i = (60, 2016)$ ) will be closer to that of a 61-year old in 2017 ( $x^j = (61, 2017)$ ), than that of a 20-year old in 1990 ( $x^j = (20, 1990)$ ). This idea is mathematically encapsulated in  $\mathbf{C}$ : the closer  $x^i$  is to  $x^j$ , the larger the covariance  $C(x^i, x^j)$ . It follows, that if  $x^i$  and  $x^j$  are very close, knowledge of  $y^j$  will greatly affect our



expectations of  $y^j$ . Conversely, if  $x^j$  is far from  $x^i$ , then knowledge of  $y^j$  will have little influence on our expectations of  $y^i$ .

*The Posterior Function:* To project mortality, we evaluate the GP function on new Age and/or Year inputs  $\mathbf{x}_*$ , i.e., evaluate  $\mathbf{f}_* = f(\mathbf{x}_*)|\mathcal{D}$ . We show in the next subsection that when  $m$  is a constant and the likelihood function is Gaussian, then the posterior distribution for  $\mathbf{f}_*$  can be determined analytically. In fact, this posterior itself is a new GP  $\mathbf{f}_*(\mathbf{x}_*)|\mathcal{D} \sim GP(\mathbf{m}_*(\mathbf{x}_*), \mathbf{C}_*(\mathbf{x}_*, \mathbf{x}_*))$  with an updated mean and covariance functions, specified in (7). The posterior mean  $\mathbf{m}_*(\mathbf{x}_*)$  is interpreted as the model prediction for inputs  $\mathbf{x}_*$ , and the posterior covariance  $\mathbf{C}_*(\mathbf{x}_*, \mathbf{x}_*)$  gives a goodness-of-fit measure for this prediction.

The posterior function can be used for both projecting mortality, as well as producing in-sample smoothed mortality curves. For the latter, all we need to do is set  $\mathbf{x}_* = \mathbf{x}$ , namely the training set inputs. In this case, the mean  $\mathbf{m}_*(\mathbf{x})$  of the posterior will produce a smooth set of mortality rates, and the posterior variance  $\mathbf{C}_*(\mathbf{x}, \mathbf{x})$  quantifies the uncertainty around  $\mathbf{m}_*(\mathbf{x})$ . Alternatively, if  $\mathbf{x}_*$  represents inputs of future calendar years, then the posterior will produce an out-of-sample projection of the mortality curves. By fitting a GP, and then analyzing the posterior we are able to achieve the following:

- Estimate the historic smoothed mortality curves by calendar year ( $\mathbf{m}_*(\mathbf{x})$  above).
- Estimate a credible interval around such curves (use the posterior covariance  $\mathbf{C}_*(\mathbf{x}, \mathbf{x})$ ).
- Project the curves forward ( $\mathbf{m}_*(\mathbf{x}_*)$  for future inputs  $\mathbf{x}_*$ ).
- Estimate the credible intervals for such projections ( $\mathbf{C}_*(\mathbf{x}_*, \mathbf{x}_*)$ ).
- Generate stochastic future forecasts (sample from the random vector  $\mathbf{f}_*(\mathbf{x}_*)$  as a future mortality scenario).
- Smooth curves over all dimensions, using automatically determined tuning parameters.

Note that the above projections are about  $\mathbf{f}_*$ . Depending on the context, an actuary might also wish to project future mortality experience  $\mathbf{y}_*$  whose marginal credible intervals are necessarily wider. When the noise  $\epsilon$  is additive and has a Gaussian distribution,  $\mathbf{y}_*$  in fact remains a GP with same mean as  $\mathbf{f}_*$ , and a modified variance due to the variance of  $\epsilon$ . Practically, forecasting realized mortality (for example, in connection with realized annuity payouts) requires also predicting future exposures  $E$ .

**Remark.** In a Lee–Carter framework one first postulates a parametric form for the mortality experience, such as

$$\mu_{ag, yr} = \alpha_{ag} + \beta_{ag}\kappa_{yr} + \epsilon_{ag, yr}, \quad (2)$$

where  $\alpha$  is the Age shape,  $\beta$  is the age-specific pattern of mortality change and  $\kappa$  is the Year trend. In the second step, after fitting  $\alpha, \beta$  by maximum likelihood, one then postulates a time-series model for the  $\kappa$  factor. Relative to a pure

regression model such as ours, the Lee–Carter method treats Age and Year dimensions completely differently; moreover the fit for the Age/Period factors is done globally (i.e., from the full dataset used), so that even spatially distant data directly influences all predictions. Finally, Lee–Carter has no mechanisms for (i) smoothing in-sample experience (beyond model calibration) and (ii) incorporating the uncertainty of the Age/Period factors in out-of-sample forecasts; its forecasts are stochastic only insofar as the time-trend is uncertain.

## 2.2. Mathematical details

GP regression takes a response surface approach, postulating an unknown, non-parametric functional dependence between covariates (inputs)  $\mathbf{x}$  and outputs  $\mathbf{y}$ ,

$$\mathbf{y} = \mathbf{f}(\mathbf{x}) + \epsilon, \quad (3)$$

where  $f$  is the *response surface* (or regression map) and  $\epsilon$  is the mean-zero noise term with observation variance  $\sigma^2(x)$ , independent across  $x$ 's. The meaning of the noise term are the statistical fluctuations that lead to deviations between observed raw mortality rates and the latent “true” rates that are being modeled. The strength of these fluctuations  $\sigma(x)$  is interpreted as the credibility of the corresponding mortality cell exposure. We remind the reader that throughout the paper,  $y^n$  represents log-mortality, and  $x^n = (x_{ag}^n, x_{yr}^n)$  is a two-dimensional Age–Year pair. In GP regression, the map  $f$  is assumed to be a realization of a GP with covariance kernel  $C$  that controls the spatial smoothness of the response surface. The GP model starts with a prior on  $f$ 's over the function space  $\mathcal{M}$  and then computes its posterior distribution conditional on the data  $\mathcal{D}$ . The function space specifying potential  $f$ 's is a reproducing kernel Hilbert space based on the kernel  $C$ . The GP assumption that  $f$  is generated by a GP implies that the posterior distributions are also Gaussian. Hence at any fixed input  $x$ , the marginal posterior is  $f_*(x) \sim \mathcal{N}(m_*(x), C_*(x, x))$ , where  $m_*$  is the predictive mean (also the posterior mode, hence maximum a posteriori (MAP) estimator), and  $C_*(x, x)$  is the posterior uncertainty of  $m_*$ .  $C_*(x, x)$  offers a principled empirical estimate of model accuracy, serving as a proxy for the mean-squared error of  $m_*$  at  $x$ .

A GP model  $GP(\mathbf{m}(\mathbf{x}), \mathbf{C}(\mathbf{x}, \mathbf{x}))$  is specified through its mean function  $m(x^i) = \mathbb{E}[f(x^i)]$  and covariance  $C(x^i, x^j) = \mathbb{E}[(f(x^i) - m(x^i))(f(x^j) - m(x^j))]$ . Specifically, the prior of  $\mathbf{f}(\mathbf{x})$  is  $p(\mathbf{f}|\mathbf{x}) = \mathcal{N}(\mathbf{m}, \mathbf{C})$ , where  $\mathbf{m} = (m(x^i))_{1 \leq i \leq N}$  and  $\mathbf{C} = (C(x^i, x^j))_{i,j}$ . In the standard case, it is further assumed that the noisy observations vector  $\mathbf{y}$  has a Gaussian relationship to the latent  $\mathbf{f}$ , i.e.,  $\epsilon^i \sim \mathcal{N}(0, \sigma^2(x^i))$ , so that

$$p(\mathbf{y}|\mathbf{f}) = \mathcal{N}(\mathbf{y}|\mathbf{f}, \mathbf{\Sigma}), \quad (4)$$

where  $\mathbf{\Sigma} = \text{diag}(\sigma(x^i)^2)$  is the  $N \times N$  noise variance matrix. Certainly, assuming  $\epsilon$  to be Gaussian with a prescribed variance is not realistic for mortality modeling, but as we show this has minimal statistical effect; we return

to this point later. Equation (4) implies that if  $f \sim GP(\mathbf{m}(\mathbf{x}), C(\mathbf{x}, \mathbf{x}))$  then  $\mathbf{y} \sim GP(\mathbf{m}(\mathbf{x}), C(\mathbf{x}, \mathbf{x}) + \Sigma)$ .

Thanks to the Gaussian assumption, determining the posterior distribution  $p(f|\mathbf{y})$  reduces to computing the predictive mean  $\mathbf{m}_*$  and covariance  $C_*$ . Combining the above likelihoods and denoting by  $\Theta$  the hyperparameters of the GP model, the log-likelihood is

$$\log p(\mathbf{y}|\mathbf{x}, \Theta) = -\frac{1}{2}\mathbf{y}^T(C + \Sigma)^{-1}\mathbf{y} - \frac{1}{2}\log|C + \Sigma| - \frac{N}{2}\log(2\pi), \tag{5}$$

where  $\mathbf{y}^T$  denotes vector transpose.

The basic GP model treats the prior mean function  $m$  as given (i.e., known and fixed). In Section 2.3.1, we discuss the more relevant case where we simultaneously infer a parametric prior mean  $m(x)$  and the kernel hyperparameters, which is known as Universal Kriging. For now, by de-trending via  $\mathbf{f} - \mathbf{m}(\mathbf{x})$ , we may assume without loss of generality that  $\mathbf{f}$  is centered at zero and  $m \equiv 0$ . The resulting posterior distribution  $f_*(\mathbf{x}_*)$  at a vector of inputs  $\mathbf{x}_*$  is multivariate Gaussian (Roustant *et al.*, 2012) with mean/covariance:

$$f_*(\mathbf{x}_*|\mathbf{x}, \mathbf{y}) \sim GP(\text{mean} = C(\mathbf{x}, \mathbf{x}_*)^T(C + \Sigma)^{-1}\mathbf{y}, \tag{6}$$

$$\text{covariance} = C(\mathbf{x}_*, \mathbf{x}_*) - C(\mathbf{x}, \mathbf{x}_*)^T(C + \Sigma)^{-1}C(\mathbf{x}_*, \mathbf{x})), \tag{7}$$

where  $C^T$  is the transpose of  $C$ .

The effect of (7) is that if we have new inputs  $\mathbf{x}_*$ , then draws from the posterior distribution of  $f_*$  at  $\mathbf{x}_*$  will be primarily influenced by historic data that have inputs close to  $\mathbf{x}_*$ . Marginally at a single cell  $x_*$ , and similar to kernel regression, the predicted value  $m_*(x_*)$  is a linear combination of observed  $y^j$ 's, capturing the idea of the GP model nonparametrically smoothing the raw mortality data. The covariance kernel  $C$  quantifies the relative contribution of different  $y^j$ 's in terms of the distance of their  $x^j$ 's to  $x_*$ , see Section 2.3 below.

2.2.1. *Observation model.* The observation noise matrix  $\Sigma$  represents the credibility of the corresponding observations  $y$ 's and is used by the GP to automatically determine how much of interpolation versus smoothing to carry out; in the limiting case  $\sigma = 0$ , the posterior mean exactly interpolates the observation  $y^j$  at  $x^j$ :  $m_*(x^j) = y^j$ .

In reality, the credibility of mortality experience is non-constant because of the different number of exposed-to-risk in different age brackets. Indeed, in the existing literature it is common to replace the additive noise structure of (4) with a GLM approach to match the fact that observed mortality is based on the counts  $D_x, L_x$ . A popular choice is a (log-link) Poisson GLM model that replaces (4) (equivalent to  $D_x = e^{f(x)+\epsilon(x)} L_x$ ) with

$$D_x \sim Poi(e^{f(x)} \cdot L_x), \tag{8}$$

and constructs a linear model

$$f(x) = \sum_i \beta_i B_i(x), \quad (9)$$

see Brouhns *et al.* (2002) and Czado *et al.* (2005). Because mortality data tend to exhibit *over-dispersion*, other approaches like Negative binomial GLM have been proposed instead of (8). In Currie *et al.* (2004), a Poisson GLM model was adjusted by fitting an age-dependent overdispersion factor with a spline.

Conceptually, it is straightforward to combine a GLM link function with a GP model: one simply adjusts the log-likelihood function in (5) and proceeds to fit the GP hyperparameters. This is equivalent to working with (8), where  $D_x$ ,  $L_x$  are observations and  $f(\mathbf{x}) \sim \text{GP}(\mathbf{m}(\mathbf{x}), \mathbf{C}(\mathbf{x}, \mathbf{x}))$  is the latent GP. A small caveat is that the non-Gaussian observations (8) ruins the Bayesian conjugacy, so that the posterior  $f_*$  is no longer Gaussian. The typical solution is a Laplace approximation that constructs a Gaussian distribution for  $f_*$  around the posterior mode. Such details are gracefully handled by the software packages and do not pose practical difficulties. Because the exposed counts  $E_x$  are very large (on the order of  $10^5$  or more), the Gaussian likelihood approximation to  $D_x/E_x$  is very close, see Table 4.

To capture the non-uniform credibility of the different cells one may take the noise level  $\sigma(x^i)$  to be state-dependent.

Specifically, the mortality table structure can be used to estimate  $\sigma(x^n)$ :  $L^n \cdot \exp\{y^n\}$  is expected to be binomially distributed with parameters  $p^{(n)} := D^n/E^n$ , and size  $E^n \simeq L^n + D^n/2$ . We then have  $\text{Var}(\exp\{y^n\}) = p^{(n)}(1 - p^{(n)})/E^n$ , and large population  $L^n$  implies the delta-method estimate

$$\sigma^2(x^n) = \text{Var}(y^n) \simeq \frac{(1 - p^{(n)})}{p^{(n)} E^n}. \quad (10)$$

We find however that (10) does not perform well, partly due to the mentioned over-dispersion effect and partly because the computed  $p^{(n)}$  is not the true mortality rate. In fact, our experience has been that a precise estimate of  $\sigma(x^n)$  is not important for GP performance, because  $\sigma$  is only used for smoothing.

Specifically, in our main analysis we take  $\sigma^2(x^n) \equiv \sigma^2$  to be an unknown constant, estimated as part of fitting the model. We return to this issue in Section 3.

### 2.3. Covariance kernels and parameter estimation

Given the covariance kernel  $C$ , (7) fully specifies the posterior distribution  $f_*(\mathbf{x}_*)|\mathcal{D}$  conditional on the dataset  $\mathcal{D}$ . GP inference is thus reduced to simply applying the above formulas, akin to the ordinary least-squares equations that specify the coefficients of a linear regression model. Of course in practice the

kernel  $C$  is not known and must be inferred itself. This corresponds to fitting the hyperparameters  $\Theta$ .

Our examples use the separable, spatially stationary kernel of the squared-exponential family, which written out explicitly takes

$$C(x^i, x^j) = \eta^2 \exp \left( -\frac{(x_{ag}^i - x_{ag}^j)^2}{2\theta_{ag}^2} - \frac{(x_{yr}^i - x_{yr}^j)^2}{2\theta_{yr}^2} \right). \quad (11)$$

In (11), covariance between  $y^i$  and  $y^j$  is determined by the distance between inputs of the respective cells, measured through the (squared) difference in Ages and Years between  $x^i, x^j$ , and modulated by the  $\theta$ 's. This use of spatial dependence can be straightforwardly extended to incorporate other dimensions, such as year-of-birth cohorts to conduct an APC allocation, or to include duration, to create a select and ultimate mortality table in the context of life insurance mortality analysis.

The hyperparameters  $\theta_\ell$  are called characteristic length-scales and their effect on the model is quite subtle. Informally, larger  $\theta$ 's result in smoother mortality curves, i.e., correlation dissipates slower. Smaller lengthscales reduce smoothing and lead to "rougher" curves. (The form of (11) implies that the mortality curves are infinitely differentiable both in Age and Year dimensions.) Note that the two lengthscale parameters  $\theta_\ell$  for Age and Year are different, so that the covariance kernel is anisotropic.

The lengthscales also determine the speed at which the latent process reverts back to its prior outside the dataset. For example, considering the Year coordinate and the question of projecting mortality rates into the future, the GP prediction will automatically blend smoothed mortality rates derived from the experience data and the specified Year trend. Indeed,  $\mathbf{m}_*(\mathbf{x}_*)$  is a weighted average of observed experience  $\mathbf{y}$ , and  $\mathbf{m}(\mathbf{x}_*)$ , with the weights determined by the lengthscale parameters  $\theta_{yr}$  and  $\theta_{ag}$ . We contrast this to APC-type models where such blending is ad hoc based on user-defined parameters.

Two further GP parameters are the process variance  $\eta^2$ , which controls the natural amplitude of  $f$  and the observation noise  $\sigma^2$  in (1) which is viewed as a constant to be estimated. Thus, the overall hyperparameter set is  $\Theta \doteq (\theta_{ag}, \theta_{yr}, \eta^2, \sigma^2)$ .

The classical method for inferring  $\Theta$  is obtained by optimizing the marginal likelihood  $p(\mathbf{y}|\mathbf{x}, \Theta) = \int p(\mathbf{y}|\mathbf{f}, \Theta)p(\mathbf{f}|\mathbf{x}, \Theta)d\mathbf{f}$  that can be written out explicitly since all the integrands are Gaussian. This leads to a nonlinear optimization problem of simultaneously fitting  $\theta_\ell$ 's and variance terms  $\eta^2, \sigma^2$ . Details on this procedure can be found in Section 3.2 of Picheny and Ginsbourger (2013). Alternatively, it is possible to directly specify  $C$ , for example, from expert knowledge regarding the expected correlation in mortality rates. Given  $\theta$ 's, the MLEs for  $\eta$  and  $\sigma$  can be analytically inferred (Picheny and Ginsbourger, 2013). This approach increases interpretability of the final smoothing/prediction and makes the GP model less of a black-box.

2.3.1. *Fitting the mean function.* A generalized version of (3) incorporates a parametric prior mean of the form  $m(x) = \beta_0 + \sum_{j=1}^p \beta_j h_j(x)$ , where  $\beta_j$  are constants to be estimated, and  $h_j(\cdot)$  are given basis functions. The coefficient vector  $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$  is obtained in parallel with computing  $\mathbf{m}_*$ ,  $\boldsymbol{\Sigma}$ . Letting  $\mathbf{h}(x) \doteq (h_1(x), \dots, h_p(x))$  and  $\mathbf{H} \doteq (\mathbf{h}(x^1), \dots, \mathbf{h}(x^N))$ , the posterior mean and variance at cell  $x$  are (Roustant *et al.*, 2012)

$$\left\{ \begin{array}{l} \hat{\boldsymbol{\beta}} = (\mathbf{H}^T(\mathbf{C} + \boldsymbol{\Sigma})^{-1}\mathbf{H})^{-1} \mathbf{H}^T(\mathbf{C} + \boldsymbol{\Sigma})^{-1}\mathbf{y}, \\ m_*(x_*) = \mathbf{h}(x_*)\hat{\boldsymbol{\beta}} + \mathbf{c}(x_*)^T(\mathbf{C} + \boldsymbol{\Sigma})^{-1}(\mathbf{y} - \mathbf{H}\hat{\boldsymbol{\beta}}), \\ s_*^2(x_*) = C(x_*, x_*) + (\mathbf{h}(x_*)^T - \mathbf{c}(x_*)^T(\mathbf{C} + \boldsymbol{\Sigma})^{-1}\mathbf{H})^T (\mathbf{H}^T(\mathbf{C} + \boldsymbol{\Sigma})^{-1}\mathbf{H})^{-1} \\ \cdot (\mathbf{h}(x_*)^T - \mathbf{c}(x_*)^T(\mathbf{C} + \boldsymbol{\Sigma})^{-1}\mathbf{H}), \end{array} \right. \tag{12}$$

where  $\mathbf{c}(x_*) = (C(x_*, x^j))_{1 \leq j \leq N}$ .

Note that (12) reduces to (7) when  $\mathbf{h} \equiv \mathbf{0}$ . We also see that the fitted coefficients  $\boldsymbol{\beta}$  are in analogue to the classical least-squares linear model.

A non-constant mean function is important for imposing structural constraints about the shape of the mortality curve, as well as the long-term improvement trends in mortality rates. Appropriate choices for parameterizing  $m$  are needed to be able to give reasonable out-of-sample projections, which corresponds to extrapolating in Age, or in calendar Year.

Use of a mean function for the GP via (12) combines the idea of parametrically de-trending the raw data through a fitted Age shape, and then modeling the residual fluctuations into a single step. We note that as  $m$  is assigned more and more structure, the residuals necessarily decrease and becomes less correlated. This calls to attention the typical over-fitting concern. Standard techniques, such as cross-validation or information criteria could be applied as safeguards, but their precise performance within the GP framework is not yet fully analyzed. We, therefore, confine ourselves to a qualitative comparison regarding the impact of the prior mean  $m(\cdot)$  on the GP model output.

2.3.2. *Bayesian GP and Markov chain Monte Carlo.* One can also consider a fully Bayesian GP model, where the mean and/or covariance parameters have a prior distribution, see Williams and Rasmussen (2006). Bayesian GP implies that there is additional, intrinsic uncertainty about  $C$ , which is propagated through to the predictive distributions  $f_*$ . Starting from the hyper-prior  $p(\Theta)$ , the posterior distribution of the hyperparameters is obtained via  $p(\Theta|\mathcal{D}) \propto p(\Theta)p(\mathbf{y}|\mathbf{x}, \Theta)$ . This hierarchical posterior distribution is typically not a GP itself. Practically, this means that one draws realizations  $\Theta^m, m = 1, 2, \dots$  from the posterior hyperparameters and then applies (7) to each draw to compute  $m_*(\mathbf{x}_*|\Theta^m), C_*(\mathbf{x}_*, \mathbf{x}_*|\Theta^m)$ .

In general, sampling from  $p(\Theta|\mathcal{D})$  requires approximate techniques, such as Markov Chain Monte Carlo (MCMC). The output of MCMC is a sequence  $\Theta^1, \Theta^2, \dots, \Theta^M$  of  $\Theta$  values that can be used as an empirical approximation for the marginal distribution of  $\Theta$ , namely  $p(\Theta|\mathbf{y}, \mathbf{x})$ . From this sequence, it is possible to calculate means and modes of the model parameters or use the  $\Theta$  sequence directly to conduct posterior predictive inference. A hybrid approach first specifies hyperparameter priors but then simply uses the MAP estimates of  $\Theta$  for prediction (thus bypassing the computationally intensive MCMC steps). This idea is motivated by the observation that under a vague prior  $p(\Theta) \propto 1$ , the posterior of  $\Theta$  is proportional to the likelihood, so that the MAP estimator  $\hat{\Theta}$ , which optimizes  $p(\Theta|\mathbf{y}, \mathbf{x})$  becomes identical to the MLE maximizer above.

We note that standard MCMC techniques are not well suited for GP as the components of  $\Theta$  tend to be highly correlated resulting in slow convergence of the MCMC chains. One solution is to use Hamiltonian Monte Carlo (HMC) (Brooks *et al.*, 2011), which is better equipped for managing correlated parameters.

**2.3.3. Setting priors for the Bayesian model.** To improve the efficiency of the MCMC routines, we first standardize the input covariates, for example,  $x_{ag, std}^i := (x_{ag}^i - \text{mean}(\mathbf{x}_{ag}))/\text{sd}(\mathbf{x}_{ag})$ . We then set priors relative to this standardized data model. Note that for comparative purposes with non-Bayesian models, the resulting posteriors of  $\beta$  and  $\Theta$  then need to be transformed back to the original scale.

Priors are taken to be weakly informative, accounting for the specifics of each hyperparameter. For the lengthscale,  $\theta_\ell$  should be below the scale of the input  $x_\ell$ , otherwise the resultant model will be essentially linear in the  $\ell$ th input dimension (Carpenter *et al.*, 2016). Thus, a prior that curtails values much beyond the data scale is appropriate. After standardization, we found that  $\log \theta_\ell \sim \mathcal{N}(0, 1)$  is reasonable. The  $\eta$  parameter plays a role similar to that of the prior variance for linear model weights in a standard linear regression, and we found  $\log \eta^2 \sim \mathcal{N}(0, 1)$  prior to be reasonable for the linear and quadratic-mean models. The prior for  $\sigma$  should reflect the noise in the data. For the CDC data, we set the prior  $\sigma^2 \sim \mathcal{N}_+(0, 0.2)$ , restricted to be positive. When including trend, priors for the  $\beta$  parameters are also required. These are set similarly to standard regression coefficients. In our analysis, we tested both Cauchy priors of  $\text{Cauchy}(0, 5)$  or Gaussian priors of  $\mathcal{N}(0, 5)$  and found both to be reasonable. For the intercept coefficient, we chose  $\beta_0 \sim \mathcal{N}(-4, 5)$  to reflect log-mortality, whereby  $\exp(y) \simeq 2\% = \exp(-3.9)$ .

**Remark.** *The Bayesian hierarchical approach for determining the parameters of the covariance matrix is also coined “automatic relevance determination”. The Bayesian model will automatically select the values of  $\theta_\ell$  and  $\eta$  without the need for using cross-validation or other approaches to set the parameter levels. Smaller values of  $\theta$  amplify the effect of the difference calculation in the covariance*

matrix, hence determining the relevance of an input dimension. Thus the Bayesian approach automatically sets the level of covariance among the  $y$ -values.

## 2.4. Software

There are several software suites that implement GP modeling and can be used for our application. The software is complementary in terms of its capabilities and approaches, in particular for inferring the covariance kernel  $C$  and for handling extensions of GPs discussed in Section 4 below.

To implement Bayesian GP models, we built models in Stan (Carpenter *et al.*, 2016). Stan is a probabilistic programming language and is a descendant of other Bayesian programming languages, such as BUGS and JAGS. In its default setting, Stan's engine utilizes MCMC techniques, and in particular a version of HMC (Brooks *et al.*, 2011). Stan also allows the option of working with the MAP estimate  $\hat{\Theta}$  or the incorporation of non-conjugate priors, and implementation of idiosyncratic features within a model. Stan automatically infers the GP hyperparameters, specifically the lengthscales  $\theta$ 's, that determine the smoothness of the mortality curves. This allows for a more data-driven approach compared to traditional graduation that a priori imposes the degree of smoothing to apply to raw data.

Within the R environment, we utilized the package "DiceKriging" (Roustant *et al.*, 2012). *DiceKriging* can fit both standard and parametric trend (12) models, and works with several different kernel families (Gaussian, exponential, Matérn). Moreover, *DiceKriging* can handle non-constant observation noise and has multiple options regarding the underlying nonlinear optimization setup. It estimates hyperparameters through maximum likelihood (but does not do MCMC).

## 3. RESULTS

We implemented a GP model for CDC mortality rates using a squared-exponential (11) covariance structure. To analyze and compare the different choices available within the GP framework, we have experimented with

1. other covariance kernel families, in particular Matern-5/2,
2. MLE and Bayesian approaches to inference of hyperparameters  $\Theta$ ,
3. a variety of mean function specifications,
4. choice of inhomogeneous noise variance  $\sigma(x)$ .

For easier reading, the figures and tables show the results for the Males; in the appendix, we report the corresponding figures and tables for Females. Most of the conclusions are identical for both genders; where appropriate we make further remarks.

We tested both the *DiceKriging* and Stan models as described in Section 2.4. Table 3 reports the MLE and MAP hyperparameter estimates for the



TABLE 3

HYPERPARAMETER ESTIMATES BASED ON MAXIMUM LIKELIHOOD (DiceKriging) AND MAXIMUM A POSTERIORI PROBABILITY (Stan), ALONG WITH MCMC SUMMARY STATISTICS.

	DiceKriging	Stan		
	MLE	MAP	MCMC Mean	MCMC 80% Posterior CI
$\theta_{ag}$	15.8384	14.7988	11.0401	(6.3369, 17.0395)
$\theta_{yr}$	15.5308	15.7910	25.8306	(14.6287, 39.4763)
$\eta^2$	1.8468	1.2365	1.9920	(0.8744, 3.3930)
$\sigma^2$	2.808e-04	2.753e-04	2.760e-04	(2.536e-04, 2.998e-04)
$\beta_0$	-3.8710	-3.8003	-3.8302	(-4.7305, -2.9350)

The GP is fitted to all data and uses squared-exponential covariance kernel (11) with prior mean  $m(x) = \beta_0$ . Stan hyper-priors (on standardized data) were  $\log \theta_{ag}, \log \theta_{yr}, \log \eta^2 \sim \mathcal{N}(0, 1)$  i.i.d.,  $\sigma^2 \sim \mathcal{N}_+(0, 0.2)$ ,  $\beta_0 \sim \mathcal{N}(-4, 5)$ .

intercept-only models fitted with All data (Males Aged 50–84, years 1999–2014, see Table 2). All of the MLEs are quite close to the MAP estimates and both fall in the 80% credible intervals for the MCMC runs. Closer analysis of the Stan output revealed that the hyperparameter posteriors are reasonably uncorrelated, justifying the use of the MAP estimates and corresponding marginal credible intervals.

Comparing both methods showed the resulting posterior distributions for the GP to be near identical, with the posterior means  $m_*$  on average within 0.3% (relative error) of each other, and the credible bands within 1.2% of each other. This indicates stability of the GP estimates given slightly different hyperparameters.

Consequently, the rest of the analysis in this paper is done using the simpler DiceKriging model which is quicker to fit and produces a convenient Gaussian posterior for the log-mortality (the fully Bayesian model built in Stan can be viewed as a mixture-of-Gaussians). Similarly, there was no major difference in prediction and smoothing when picking different covariance kernels. In general, picking a kernel is like picking a basis family for linear regression; basic caveats apply, but it is mostly a secondary effect. Below we focus on the squared-exponential kernel. One benefit of this choice is that the resulting scenarios  $f_*$  are guaranteed to be infinitely differentiable, which enables analytic treatment of instantaneous mortality improvement  $\partial_{yr} f_*$ , see Section 3.5.

For the observation noise, estimating a constant noise variance led to MLE of  $\hat{\sigma}^2 = 2.808 \times 10^{-4}$ . Figure A1 in Appendix B gives a descriptive analysis of the resulting residuals; we observe that both the Gaussian assumption and the i.i.d assumptions are statistically plausible. As a further check, we tried to work with a non-constant  $\sigma^2(x)$  by plugging-in the delta method estimate in (10). However, this led to credible bands that are too narrow in terms of coverage ratios due to the aforementioned over-dispersion effect. Manual calibration found that  $\check{\sigma}^2(x) = 4(1 - p_x)/(p_x E_x)$ , i.e., an overdispersion factor of 2, works fine. The resulting estimated  $\check{\sigma}^2$  values ranged over  $[1.066 \times 10^{-4}, 1.304 \times 10^{-3}]$

TABLE 4  
 POSTERIOR HYPERPARAMETER ESTIMATES FROM MCMC RUNS IN STAN  
 FOR THE POISSON GP MODEL.

	MCMC Mean	MCMC 80% Posterior CI
$\theta_{ag}$	10.5955	(5.8652, 16.4272)
$\theta_{yr}$	26.1259	(14.4609, 40.1992)
$\eta^2$	1.8643	(0.7938, 3.2357)
$\sigma^2$	2.170e-04	(1.945e-04, 2.400e-04)
$\beta_0$	-3.7912	(-4.6851, -2.9317)

The GP is fitted to all data and uses squared-exponential covariance kernel (11) with prior mean  $m(x) = \beta_0$ . Stan hyper-priors (on standardized data) were  $\log \theta_{ag}, \log \theta_{yr}, \log \eta^2 \sim \mathcal{N}(0, 1)$  i.i.d.,  $\sigma^2 \sim \mathcal{N}_+(0, 0.2)$ ,  $\beta_0 \sim \mathcal{N}(-4, 5)$ .

with a mean of  $4.36 \times 10^{-4}$ . This is close to the constant- $\sigma^2$  MLE estimate and the respective projections were very close, confirming that with a GP model the whole question of capturing observation errors is a “higher order” concern. For ease of interpretation, we thus used a constant  $\sigma^2$ , estimated via MLE, for the remainder of the analysis.

As an alternative, we also implemented a Poisson GP model in Stan. The resulting parameter estimates are reported in Table 4. As can be seen, the resulting GP covariance structure is very similar to Table 3. We find that the actual outputs of the two models are also essentially identical. Namely, the mean percent error between the Gaussian and Poisson GP was  $-8.14e-05$ , indicating no systematic discrepancy between the two models and the root mean squared error was  $4.99e-4$ , which is not material for the ultimate actuarial use. Our conclusion is that there is no statistical difference between using additive Gaussian noise or a Poisson link function.

### 3.1. Retrospective analysis

We begin with a retrospective look at smoothed mortality experience over the recent past. Traditionally, this is done using actuarial graduation techniques; for the GP framework smoothing is simply the in-sample prediction  $m_*(\mathbf{x})$ . Specifically, we fit a model using all the data, and investigate the mortality during the last 5 years of the period. Figure 1 shows the estimated mortality rates as a function of Age, specifically Males Aged 60–70. The left panel compares the raw and GP-smoothed rates for 2010 and 2014, while the right panel shows the overall yearly trend for years 2010–2014. As a complement to above, Figure 2 provides a preliminary analysis of mortality improvement by plotting mortality rates against time. We show the observed and smoothed mortality rates against calendar years 1999–2014 for Males and Females Aged 60, 70, and 84, along with the forecasted rates up to 2016. From the figure, we clearly observe the

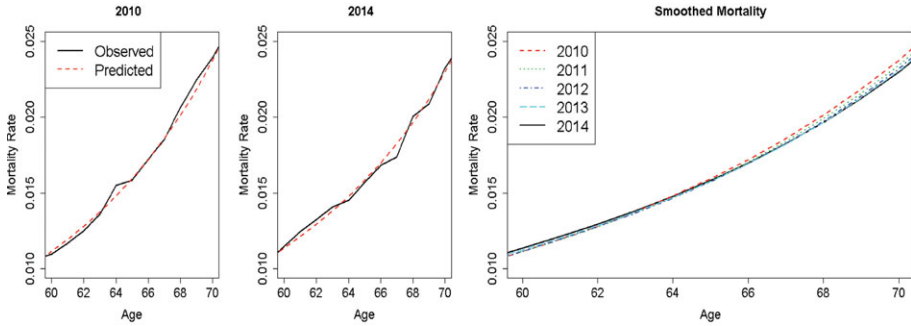


FIGURE 1: Mortality rates for US Males Aged 60–70 during the years 2010–2014. Raw (solid) vs. smoothed (dashed) mortality curves. Models are fit to 1999–2014 CDC data for Ages 50–84 (all data). Mean function  $m(x)$  is intercept-only,  $m(x) = \beta_0$ . (Color online)

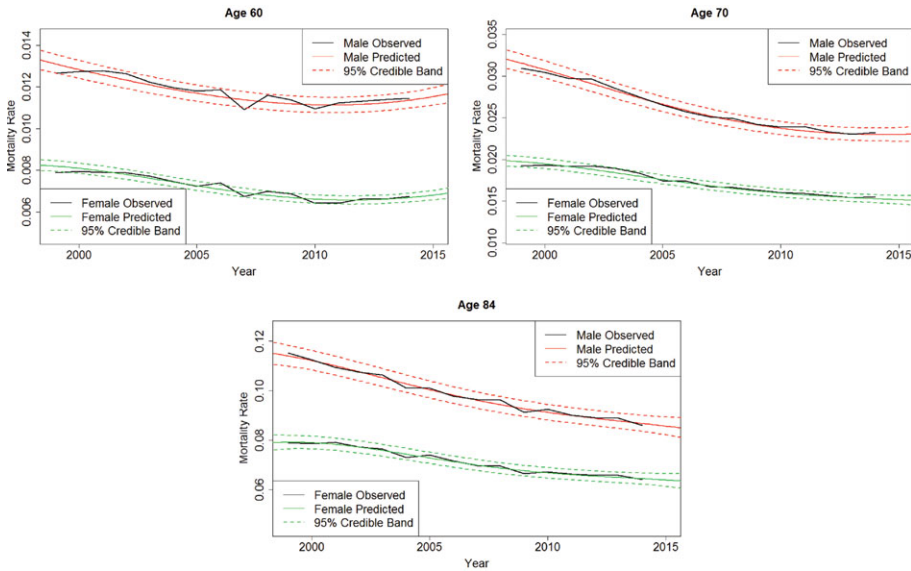


FIGURE 2: Mortality rates for US Males (top) and US Females (bottom) Aged 60, 70 and 84 over time. The plots show raw mortality rates (solid black) for years 1999–2014, as well as predicted mean of the smoothed mortality surface (solid red) and its 95% credible band, for 1999–2016. Models are fit to the 1999–2014 CDC data for Ages 50–84 (All data). Mean function is intercept-only,  $m(x) = \beta_0$ . (Color online)

decrease of mortality at older ages which is, however, slowing down in the last few years.

A key output of official tables are the mortality improvement scales, such as the MP-2015 rates  $MI_{back}^{MP}(x_{ag}, yr)$ , where we distinguish the common indexing by Age, keeping Year fixed. These are intuitively the smoothed version of the raw

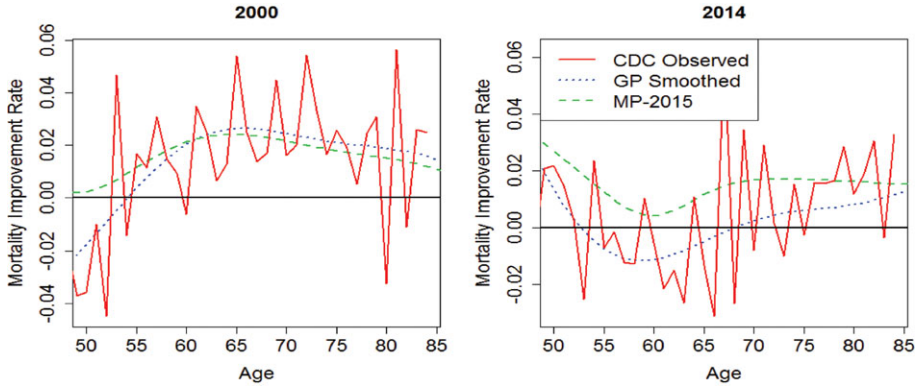


FIGURE 3: Mortality improvement factors for US Males using All Data. Solid lines indicate the empirical mortality experience  $MI_{back}^{obs}(\cdot; yr)$  for years  $yr \in \{2000, 2014\}$ , the dotted and dashed lines are  $\partial m_{back}^{GP}(\cdot; yr)$  from (13) and (14), and the MP-2015 improvement scale  $MI_{back}^{MP}(\cdot; yr)$ , respectively. (Color online)

annual percentage mortality improvement which is empirically observed via

$$MI_{back}^{obs}(x_{ag}; yr) \doteq 1 - \frac{\exp(\mu(x_{ag}, yr))}{\exp(\mu(x_{ag}, yr - 1))}, \tag{13}$$

with  $\mu(x_{ag}, yr)$  the raw log-mortality rate for  $(x_{ag}, yr)$ . In analogue to above, we can obtain the predicted mean improvement by replacing  $\mu$ 's by the GP model posteriors  $f_*$ 's and integrating over their posterior distributions:

$$\partial m_{back}^{GP}(x_{ag}, yr) := \mathbb{E} [MI_{back}^{GP}(x_{ag}, yr)] \doteq \mathbb{E} \left[ 1 - \frac{\exp(f_*(x_{ag}, yr))}{\exp(f_*(x_{ag}, yr - 1))} \right]. \tag{14}$$

Figure 3 shows these different improvement scales for Ages 50–85 and two sample years, 2000 and 2014; the MP-2015 curves are from the published SOA reports (SOA, 2015). We observe that the raw mortality improvements  $MI_{back}^{obs}$  are extremely noisy, which is not surprising since they are based on the relative difference of two very similar raw mortality rates. Figure 3 also indicates that the MP-2015 estimates are significantly higher than either the actual experience (which has moderated a lot in the past decade) or our fit  $\partial m_{back}^{GP}$ , with differences of as much as 2% p/a in improvement factors. Figure 4 emphasizes that there is a downward trend in mortality improvement, and moreover non-uniform behavior across ages. This throws into question the MP-2015 concept of a sustained, age-uniform projected long-term mortality improvement trend.

### 3.2. Mean function modeling

We tested three models for the prior mean function: an intercept-only model  $m(x^n) = \beta_0$ , a linear model,  $m(x^n) = \beta_0 + \beta_1^{ag} x_{ag}^n + \beta_1^{yr} x_{yr}^n$ , and a quadratic Age

TABLE 5

FITTED MEAN FUNCTION AND COVARIANCE PARAMETERS USING SUBSET III (AGES 50–70 AND YEARS 1999–2009) FOR MALES.

	Mean Function Parameter MLE's				GP Hyperparameter MLE's			
	$\beta_0$	$\beta_1^{ag}$	$\beta_2^{ag}$	$\beta_1^{yr}$	$\eta^2$	$\sigma^2$	$\theta_{ag}$	$\theta_{yr}$
Intercept	-4.526	-	-	-	6.213e-01	3.428e-04	8.384	12.746
Linear	18.737	0.081	-	-1.397e-02	8.521e-04	1.761e-04	3.610	3.543
Quadratic	19.641	0.064	1.459e-04	-1.417e-02	1.403e-03	2.998e-04	3.629	3.475

The mean functions are  $m(x^n) = \beta_0$  for Intercept,  $m(x^n) = \beta_0 + \beta_1^{ag} x_{ag}^n + \beta_1^{yr} x_{yr}^n$  for Linear and  $m(x^n) = \beta_0 + \beta_1^{ag} x_{ag}^n + \beta_1^{yr} x_{yr}^n + \beta_2^{ag} (x_{ag}^n)^2$  for Quadratic.

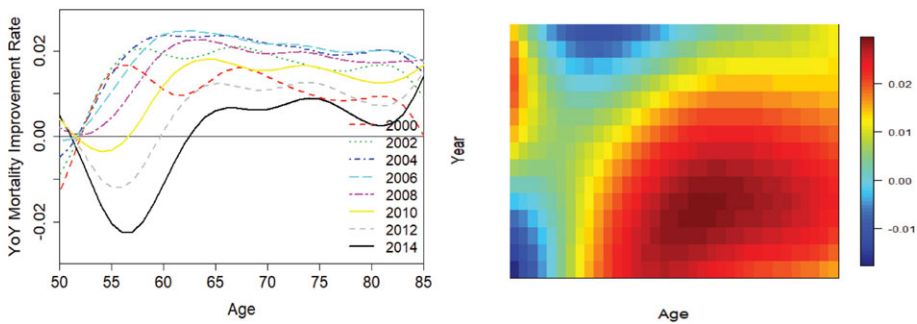


FIGURE 4: Smoothed yearly mortality improvement factors  $\partial m_{back}^{GP}(ag, yr)$  from (14) for US Males using All data. Left: Age factors for  $yr = 2000, \dots, 2014$ . The curves for 2000 and 2014 are the same as in Figure 3. Right: Heatmap of the estimated YoY improvement factors. (Color online)

model,  $m(x^n) = \beta_0 + \beta_1^{ag} x_{ag}^n + \beta_1^{yr} x_{yr}^n + \beta_2^{ag} (x_{ag}^n)^2$ . Thus, the linear model has the log mortality increasing linearly in Age and decreasing linearly in calendar year. The quadratic model then adds a convexity component to the Age axis.

The coefficients of these functions were estimated concurrently with fitting the GP models using (12). A summary of the models and the coefficient estimates is shown in Table 5. finding is that the fitted year-trend coefficient  $\beta_1^{(yr)}$  is consistently estimated by both the linear and quadratic model and indicates a linear improvement in log mortality rates of about 1.4% per calendar year in both of these models regardless of assumptions on age shapes. Since this model is fitted to Ages 50–70, these results are consistent with the long-term trend of improving mortality. As expected, the table also indicates a strong Age effect; we note that the fitted coefficient  $\beta_2^{(ag)} = 1.459 \times 10^{-4}$  for the quadratic age component confirms a significant convexity of log-mortality in Age.

Intuitively, the mean function provides a fundamental explanation of mortality rates by age and year, while the covariance structure captures deviations from this postulated relationship based on nearby observed experience (with the influence depending on the lengthscale). Consequently, the choice of the mean function affects the covariance structure; a stronger trend/shape lowers

the spatial dependence of the residuals. We observe this effect in Table 5, where the intercept-only model has length-scales of  $\theta_{ag} \approx 8.5$ ,  $\theta_{yr} \approx 12.5$ , while for the linear and quadratic models the range of the length-scales is much smaller  $\theta_\ell \approx 3.5$ . Another effect of the mean function is on the hyperparameter  $\eta$  that can be viewed as the variance of the model residuals. If the mean function fits well then we expect smaller  $\eta$ . In turn, smaller  $\eta$  translates into tighter credible intervals around in-sample smoothing and out-of-sample forecasts. Table 5 shows that the values for  $\eta$  and  $\sigma$  are similar across linear and quadratic models, while the intercept-only model has uniformly larger values across parameters.

Figure 5 illustrates these three models fit to Subset III, which emulates deep out-of-sample extrapolation. As discussed, out-of-sample forecasting by the GP model can be viewed as blending the data-driven prediction with the estimated trend encapsulated by  $m$ . Specifically, as  $x_{*,ag}^n$  moves beyond the Age range of  $\{50, 51, \dots, 70\}$  in Subset III we have  $m_*(x_*) \rightarrow m(x_*)$ . In the case of an intercept-only model, this implies that  $m_*(x_*) \rightarrow \beta_0$ , i.e., the projected mortality is independent of either Age or Year. In Figure 5, the asymptotic projected rate was  $\exp(\hat{\beta}_0) = 1.082\%$ . A similar issue pertains to the linear-mean model whose long-range forecasts imply exponential Age dependence which is not appropriate for Ages above 80. This discrepancy is successfully resolved by the quadratic  $m(x_*)$  model. The lengthscales  $\theta$  control this transition; roughly speaking extrapolating more than  $\theta$  distance away reduces to  $m_*(x_*) \simeq m(x_*)$ . This can also be seen in Figure 5: Since the training data includes up to 2010, the forecast for 2011 is much more driven by past data compared to the one for 2014. As a result, for the intercept-only model with  $\theta_{yr} = 11.461$ , the forecast is acceptable in 2011, but deteriorates dramatically for 2014. This effect is also present but less apparent in the trend models; due to smaller values of  $\theta_{yr}$  the latter forecasts already rely more heavily on their mean functions for extrapolation.

As an additional comparator, Figure 5 plots the fit of a Poisson GLM model using (9) and a quadratic age-trend. The respective fitted coefficients  $\beta$  are listed in Table 6. We observe that the estimated coefficients using a Poisson link function are very similar to those in Table 5 (confirmed visually in the figure), which is consistent with our earlier discussion that adding a link function does not materially modify the results. At the same time, the goodness-of-fit of the GP model is significantly better than a parametric GLM fit, confirming the complex spatial structure that cannot be captured in a simple GLM.

Another way for model comparison is to look at the widths of the respective credible intervals. For example, for year 2010 Age 84, the observed mortality rate was 8.999% and the intercept, linear and quadratic models generated 95% credible intervals of (0.783%, 10.254%), (7.100%, 8.562%) and (6.379%, 11.188%), respectively. Certainly, the first interval is too wide (partly due to the large  $\eta$  and  $\theta$ 's of the intercept-only model), while the second interval is too narrow and does not even contain the raw data point (the linear model apparently underestimates  $\eta$ ). On the other hand, for Age 71 in year 2014, the raw rate was 2.489% and the respective 95% credible intervals were (2.258%, 2.927%), (2.378%, 2.703%)

TABLE 6

FITTED POISSON GLM PARAMETERS FOR (9) WITH THREE DIFFERENT AGE-YEAR PARAMETERIZATIONS (INTERCEPT-ONLY  $f(x^n) = \beta_0$ , LINEAR  $f(x^n) = \beta_0 + \beta_1^{ag} x_{ag}^n + \beta_1^{yr} x_{yr}^n$ , AND QUADRATIC-AGE  $f(x^n) = \beta_0 + \beta_1^{ag} x_{ag}^n + \beta_1^{yr} x_{yr}^n + \beta_2^{ag} (x_{ag}^n)^2$  BASED ON US MALE CDC SUBSET III DATA.

Poisson GLM Mean Function Parameter MLE's				
	$\beta_0$	$\beta_1^{ag}$	$\beta_2^{ag}$	$\beta_1^{yr}$
Intercept	-4.442	-	-	-
Linear	23.264	0.080	-	-1.62e-02
Quadratic	24.218	0.0403	3.24e-04	-1.608e-02




FIGURE 5: Comparison of mean function choices in extrapolating mortality rates at old ages. Models are fit to years 1999–2010 and Ages 50–70 (Subset III) for US Males, with estimates made for Age 50–85 in 2011 and 2014. The vertical line indicates the boundary of the training dataset in  $x_{ag}$ . The mean functions are given in Table 5. Also shown is the fit based on a quadratic-age Poisson GLM model (9), cf. Table 6. (Color online)

and (2.346%, 2.798%). While all models now contain the observed rate, the linear model again has the tightest credible interval, which might indicate poor goodness-of-fit.

Returning to in-sample smoothing and looking again at Figure 5, we observe that all three models generate very similar forecasts for Ages 55–70. This confirms that in-sample  $m_*$  is data-driven and the choice of  $m(\cdot)$  is secondary. To summarize, the most important criterion in including a mean function is whether the goal is to predict out-of-sample and if so, how far out-of-sample. In general, mean modeling is crucial, but the precise choice of the mean function is not as clear. In Section 4.1, we discuss one further method for mean-modeling based on Age-grouping.

### 3.3. Predictive accuracy

Figure 5 can also be viewed as a first glimpse into the predictive accuracy of a GP mortality model. Recall that in the figure we fit to mortality data from 1999 to 2010 and then forecast 1 year out (2011) and 4 years out (2014). The figure then compares these projections to the observed mortality experience in 2011 and 2014. As discussed, these projections are highly sensitive to the choice of

TABLE 7

GP MODEL PREDICTIONS FOR US MORTALITY IN 2014 AND AGE 70/80 WHEN FITTED TO VARIOUS DATA SUBSETS, CF. TABLE 2, INDICATED BY SUPERSCRIPTS.

Intercept-Only $m(x) = \beta_0$							
$x_{ag}$	Fit to Subset III		Fit to Subset I		Fit to All Data		Observed $\mu$
	$m_*^{(III)}$	$(s_*^{(III)})$	$m_*^{(I)}$	$(s_*^{(I)})$	$m_*^{(All)}$	$(s_*^{(All)})$	
70	-3.7520	(0.0580)	-3.7380	(0.0427)	-3.7702	(0.0169)	-3.7630
80	-3.7177	(0.4565)	-2.8416	(0.0463)	-2.8579	(0.0170)	-2.8531
Quadratic $m(x) = \beta_0 + \beta_1^{ag} x_{ag} + \beta_1^{yr} x_{yr} + \beta_2^{ag} x_{ag}^2$							
$x_{ag}$	Fit to Subset III		Fit to Subset I		Fit to All Data		Observed $\mu$
	$m_*^{(III)}$	$(s_*^{(III)})$	$m_*^{(I)}$	$(s_*^{(I)})$	$m_*^{(All)}$	$(s_*^{(All)})$	
70	-3.7507	(0.0419)	-3.7711	(0.0332)	-3.7671	(0.0163)	-3.7630
80	-2.8774	(0.1046)	-2.8546	(0.0333)	-2.8553	(0.0164)	-2.8531

We report the predictive mean  $m_*(x)$  and the predictive standard deviation  $s_*^2(x) = C_*(x_*, x_*)$ .

$m(x)$ , especially in terms of the Age-structure because the models are only given experience up to Age 70 and hence have zero information about how mortality behaves for higher ages.

To provide a more “fair” comparison, Table 7 shows projections for other input datasets. Overall, we observe excellent predictive power for making projections 4-years out (fit using Subset I, forecast for 2014), confirming the competitive performance of the GP fitted models.

Beyond the predictive mean  $m_*(x_*)$ , we also report the corresponding posterior marginal variance  $s_*^2(x_*) := C_*(x_*, x_*)$  which is a proxy for the confidence the model assigns to its own prediction. This provides a valuable insight: for example, if fitted to Ages 50–70 (Subset III) and projecting for Age 80 in year 2014:  $\tilde{x} := (x_{ag}, x_{yr}) = (80, 2014)$ , the intercept-only model reports minimal predictive power that is reflected in the very large  $s_*^{(III)}(\tilde{x}) = 0.4565$ , in light of which the poor prediction  $m_*^{(III)}(\tilde{x}) = -3.7177$  is simply a “shot in the dark”. Indeed, the model predicts mortality rate of 2.43%, which is nowhere the realized 5.78%, but is still within its 95% credible interval of (0.98%, 6.03%). Including more ages (Subset I) gives a more reasonable and much more confident forecast of  $m_*^{(I)}(\tilde{x}) = -2.8416$  and  $s_*^{(I)}(\tilde{x}) = 0.0463$ , and including more years (which makes  $\tilde{x}$  to be right at the edge of observed data) raises credibility even further,  $m_*^{(All)}(\tilde{x}) = -2.8579$  and  $s_*^{(All)}(\tilde{x}) = 0.0170$ . Table 7 also quantifies the gains from using a more sophisticated  $m$ —the quadratic trend allows to shrink  $s_*^{(I)}(\tilde{x})$  from 0.0463 to 0.0333, and brings the prediction  $m_*^{(I)}(\tilde{x})$  closer to the eventually realized experience.

For another angle on forecasting with GP models, Figure 6 shows that the intercept-only model still performs well when predicting only slightly out-of-






FIGURE 6: Mortality rate prediction for years 2011 and 2014 and Ages 71–84. Model is fit with Subset II data with intercept-only mean functions and squared-exponential kernel. “Simulated paths of  $f$ ” refers to simulated trajectories of the latent  $\mathbf{f}_*$ . Credible bands are for the mortality surface  $\mathbf{f}_*$ ; vertical intervals are for predicted observable mortality experience  $\mathbf{y}_*$ . (Color online)

sample. In this figure, we fitted mortality curves using the “notched” Subset II: years 1999–2010 and Ages 50–84, plus 2011–2014 with Ages 50–70, and then predicted out-of-sample for mortality rates for 2011–2014 and Ages 71–85. This differs from the previous setup where the model had no prior information on Ages 71–84. We observed that in this setup the uncertainty from the intercept-only model is only slightly worse (wider interval) relative to the quadratic trend model, confirming the reasonableness of using the simpler  $m(x) = \beta_0$ .

Figure 6 also plots the marginal credible bands for  $\mathbf{f}_*(\mathbf{x}_*)$  and intervals for future observations  $\mathbf{y}_*$ . As expected, the prediction uncertainty increases for the oldest ages and for later calendar years (compare credible intervals in Figure 6 for 2014 vis-à-vis 2011). Also note that the intervals for  $\mathbf{y}_*$  are always a fixed distance away from the pointwise bands of  $\mathbf{f}_*$  regardless of Age/Year due to the assumed constant noise variance  $\sigma^2$ ; this is much more noticeable when in-sample, where posterior variance  $C_*(\mathbf{x}_*, \mathbf{x}_*)$  is negligible relative to  $\sigma^2$ .

As discussed, the GP model automatically generates credible intervals around any prediction, giving a principled approach for assessing uncertainty in forecasts. Moreover, since GP considers the full covariance structure of mortality curves, one can analytically evaluate the joint predictive uncertainty of any number of mortality rates. In particular, one can use the posterior predictive distribution to generate conditional trajectories of mortality rates for any collection of cells. For instance, fixing a calendar year, we may sample from the multivariate normal distribution of  $\mathbf{f}_*(\mathbf{x}_*)$  across Ages  $\mathbf{x}_*$  to obtain a stochastic scenario of the respective mortality age structure. This is illustrated in Figure 6 that shows several such scenarios of log-mortality rates for calendar year 2014 (which is in-sample up to Age 70 and extrapolating for Ages 71–84), along with the overall credible band. Note that in contrast to factor models like Lee–Carter that force the log-mortality curve  $\mu_{\cdot, yr}$  to be confined to a low-dimensional space (e.g., one degree of freedom in classical Lee–Carter), within a GP framework, the shape of  $f_*(\cdot, yr)$  remains nonparametric and infinite-dimensional. Alternatively, we could sample potential evolutions of mortality at selected age into a desired future projection interval. Sampling such trajectories is crucial for




FIGURE 7: Observed and predicted mortality rates for 1994–2040 for three representative datasets. GP model uses quadratic mean function  $m(x) = \beta_0 + \beta_1^{ag} x_{ag} + \beta_1^{yr} x_{yr} + \beta_2^{ag} x_{ag}^2$ , and the APC model is as in Equation (15). Models fit to HMD data for 1994–2009 and Ages 50–84. (Color online)

quantifying aggregate mortality risk in a portfolio (say in a pension plan or life insurance context).

### 3.4. Comparison of GP and APC forecasts

To provide a brief comparison of the popular stochastic mortality models, we fit a cohort extension of the Lee–Carter model in (2), introduced by Renshaw and Haberman (2006), which is as follows:

$$\mu_{ag, yr} = \alpha_{ag} + \frac{1}{n_a} \kappa_{yr} + \frac{1}{n_a} \gamma_{yr-ag} + \epsilon_{ag, yr}, \tag{15}$$

where  $\gamma_{yr-ag}$  is the cohort effect and  $n_a$  is the number of years in the dataset. Using the StMoMo software suite (Villegas *et al.*, 2015) on our data yielded a random walk with drift for  $\kappa$ . and ARIMA(0,1,2) model for  $\gamma$ . Cairns *et al.* (2011) showed that this model performed well in US Male data analysis.

Figure 7 compares the predictions from a GP model against those of an APC model (15) for three representative ages and across three different country/gender datasets. We use years 1994–2040, i.e., both in-sample and up to 30 years into the future. We observe that relative to the GP model, the APC model generates both volatile in-sample projections (as it is not designed with smoothing in mind), and erratic short-term projections due to the underlying time series fitted to the  $\kappa$ . and  $\gamma$ . factors. Recall that the APC framework tries to average out trends via a parametric model that makes the projections dependent even on distant historical experience, while the GP effectively uses the history to learn the spatial dependence structure and then makes data-driven projections based on recent experience. Also, the APC has generally tighter predictive intervals (i.e., it is more confident in its forecast) compared to the GP predictions, although the final uncertainty band in 2040 is about the same. We note that there is no general pattern—sometimes the two models agree on the likely mortality trend, sometimes APC indicates larger improvements, sometimes GP does.

### 3.5. Forecasting mortality improvement

To focus more precisely on mortality *improvement*, we proceed to analyze changes in  $\mu(x_{ag}, \cdot)$  over time. Section 3.1 discussed already backward-looking annual (YoY) improvements  $MI_{back}^{obs}$  and  $\partial m_{back}^{GP}$  as defined in Equations (13) and (14). For a more prospective analysis, one could consider a centered difference

$$1 - \left( \frac{\exp(f_*(x_{ag}, yr + h))}{\exp(f_*(x_{ag}, yr - h))} \right)^{1/2h} \approx - \frac{f_*(x_{ag}, yr + h) - f_*(x_{ag}, yr - h)}{2h}, \tag{16}$$

which is possible to compute for any  $h$  since the GP model for  $f_*$  yields an entire mortality surface spanning over all  $(x_{ag}, x_{yr}) \in \mathbb{R}^+ \times \mathbb{R}^+$ . Note that since  $f_*$  is a GP, the right-hand side of (16) remains Gaussian. We may also take the limit  $h \rightarrow 0$  that gives the instantaneous rate of change of mortality in terms of calendar time. As an analogue to (16), we term the negative of the above differential as the instantaneous mortality improvement process

$$MI_{diff}^{GP}(x_{ag}; x_{yr}) \doteq - \frac{\partial f_*}{\partial x_{yr}}(x_{ag}, yr). \tag{17}$$

A remarkable property of the GP is that  $MI_{diff}^{GP}$  is once again a GP with explicitly computable mean and covariance functions (Williams and Rasmussen, 2006).

**Proposition 1.** *For the GP  $f_*$  with a twice differentiable covariance kernel  $C$ , the limiting random variables*

$$\frac{\partial f_*}{\partial x_{yr}}(x_{ag}, yr) \doteq \lim_{h \rightarrow 0} \frac{f_*(x_{ag}, yr + h) - f_*(x_{ag}, yr)}{h}, \tag{18}$$

*exist in mean square and form a GP  $\frac{\partial f_*}{\partial x_{yr}} \sim GP(\partial m_{diff}, s_{diff})$ . Given the training set  $\mathcal{D} = (\mathbf{x}, \mathbf{y})$ , the posterior distribution of  $\frac{\partial f_*}{\partial x_{yr}}(x_*)$  has mean and variance*

$$\partial m_{diff}(x_*) = \mathbb{E} \left[ \frac{\partial f_*}{\partial x_{yr}}(x_*) \mid \mathbf{x}, \mathbf{y} \right] = \frac{\partial C}{\partial x'_{yr}}(\mathbf{x}, x_*)(\mathbf{C} + \Sigma)^{-1} \mathbf{y}, \tag{19}$$

$$\begin{aligned} s_{diff}^2(x_*) &= \text{Var} \left( \frac{\partial f_*}{\partial x_{yr}}(x_*) \mid \mathbf{x}, \mathbf{y} \right) = \frac{\partial^2 C}{\partial x_{yr} \partial x'_{yr}}(x_*, x_*) \\ &\quad - \frac{\partial C}{\partial x'_{yr}}(\mathbf{x}, x_*)(\mathbf{C} + \Sigma)^{-1} \frac{\partial C}{\partial x_{yr}}(x_*, \mathbf{x}), \end{aligned} \tag{20}$$

where  $\frac{\partial C}{\partial x'_{yr}}(\mathbf{x}, x_*) = [\frac{\partial C}{\partial x'_{yr}}(x^1, x_*), \dots, \frac{\partial C}{\partial x'_{yr}}(x^N, x_*)]$  and each component is computed as the partial derivative of  $C(x, x')$ .




FIGURE 8: Estimated Male mortality improvement using the differential GP model (instantaneous improvement) and the YoY improvement from the original GP model. We show the means and 80% credible bands for  $MI_{diff}^{GP}$  and  $MI_{back}^{GP}$  for Males Aged 50–84 and years 2000 and 2014. Models used are fit to All Data with  $m(x) = \beta_0$ . (Color online)

See Theorem 2.2.2 in Adler (2010) for more details. By analogy, Proposition 1 can also be extended to consider the differential of mortality to age or other covariates. Note that the squared exponential kernel in (11) is infinitely differentiable with derivatives

$$\frac{\partial C}{\partial x'_{yr}}(x, x') = -C(x, x') \frac{\eta^2}{\theta^2_{yr}} (x_{yr} - x'_{yr}), \tag{21}$$

$$\frac{\partial^2 C}{\partial x_{yr} \partial x'_{yr}}(x, x') = C(x, x') \frac{\eta^2}{\theta^2_{yr}} \left( 1 - \frac{1}{\theta^2_{yr}} (x_{yr} - x'_{yr})^2 \right). \tag{22}$$

Observe that the mean  $\partial m_{diff}(x_*)$  mortality improvement is equal to the derivative of the predicted mortality surface,  $\frac{\partial}{\partial x_{yr}} m(x_*)$ , a desirable self-consistency property. However, Proposition 1 goes much further, providing also analytic credible bands around  $\partial m_{diff}(x_*)$  and even the full predictive distribution of the mortality improvement process. Compare these features to a non-Bayesian smoothing model, such as P-splines, that only models  $m(x_*)$  and therefore beyond direct differentiation provides no uncertainty quantification for  $\frac{\partial f_*}{\partial x_{yr}}$ .

To sum up the previous discussion, the GP framework yields a probabilistic estimate of the *instantaneous* mortality improvement that is analytically consistent with the projected mortality rates. Figure 8 shows mortality improvement estimates  $\partial m_{back}^{GP}$ ,  $\partial m_{diff}^{GP}$  and MP-2015 improvement factors for US Ages 50–85 in years 2000 and 2014. The 80% credible bands of  $MI_{back}^{GP}$  and  $MI_{diff}^{GP}$  are also shown. The bands for  $MI_{diff}^{GP}$  were produced from (20), while for  $MI_{back}^{GP}$  they were generated from empirical sampling from (14). While we observe similar overall structure (in terms of similar predicted values and

similar predicted uncertainty), we also note that there are some differences that indicate the changing rate of mortality improvement. Thus, in 2000, mortality improvement was accelerating, leading to  $\partial m_{diff}^{GP}(\cdot; 2000) > \partial m_{back}^{GP}(\cdot; 2000)$ . In contrast, the fact that  $\partial m_{diff}^{GP}(\cdot; 2014) < \partial m_{back}^{GP}(\cdot; 2014)$  suggests that mortality improvement continues to decelerate as of 2014, so that the gap with the level improvement scale embedded in MP-2015 is likely to grow. In our analysis, we find that this deceleration started around 2010, so that in the past 5–6 years mortality evolution over time has been convex, generating a growing wedge against the MP-2014/15 forecasts.

**Remark.** *In our analysis, we concentrate on modeling the log mortality surface, obtaining the mortality improvement factors as a by-product. An alternative is to first directly calculate observed mortality improvement  $MI_{back}$  and then model it with a GP. This would effectively replace the  $\beta_1^{yr}$  component of the mean function with a richer structure. This procedure is similar to that of Mitchell et al. (2013) where mortality improvement itself is modeled in a Lee–Carter framework.*

### 3.6. Further datasets

As further empirical evidence, Appendix B presents analysis for US Females, listing the equivalents of Figures 1 and 3–7 and Table 5. Additional results for four more datasets—UK/Japan Males/Females—can be found in the online supplement. These were generated using the provided R notebook on the respective HMD datasets.

Overall, the results are consistent, yielding similar covariance structure estimates. This suggests the possibility of building a hierarchical model that can improve credibility through cross-national borrowing of (statistical) information. The lengthscales in Age are all in the range [8, 16] and in Year are around 10. The estimated  $\sigma^2$  values and  $\eta^2$  are very similar throughout. The one exception is the UK Males dataset, which produced an exceptionally high  $\theta_{ag}$  and an exceptionally low  $\theta_{yr}$ . The low  $\theta_{yr}$  causes out-of-sample forecasts to mean-revert too quickly, leading to poor prediction. This can be observed in the equivalent of Figures 6/B5 for the UK and Japan data sets, where the GP model trained on Subset II data under-estimates the mortality experience in 2014 for UK and Japan Males, but does well for Females. This is partly mitigated by the wider uncertainty bands for Males, i.e., the models are aware that its forecasts are less accurate.

Comparing Male and Female mortality, Female mortality is always lower, but the Age-shapes are mostly the same. In United States, we observe a smaller mortality deterioration for Females around Ages 50–65, and a slightly lower improvement  $\partial m$  overall (compare Figures 4 and B3). For smoothing, Figures 1 and B1 are nearly identical in shape. The curve in Figure B2 for 2014 is slightly different in shape compared to the Male Figure 3 around Ages 50–65 due to the observed mortality improvement declining in this range.

Comparing Table B1 with Table 5 and Figure B4 with Figure 5, we see that the trend model comparison results are near identical; the only noticeable differences are that the quadratic model is a much better fit on the test set for US Females, and that the  $\theta$  values for the intercept-only model are larger. We do notice differences across genders in the quadratic trend function parameters, see Table B2 vis-à-vis Table 8. In particular, the intercept terms are different in magnitude, and some of the higher order terms differ in sign. Thus, the trend curves differ in shape between Males and Females in their respective age groups, which is unsurprising since the age group endpoints were chosen to match the Male dataset.

Some interesting features can be glimpsed by comparing the mortality improvement rate plots across datasets. First, raw YoY improvements are extremely noisy, even more than in United States, so smoothing is essential for detecting trends. Second, relative to reported low and deteriorating mortality improvements in United States, *MI*'s for Japan and United Kingdom are higher and are all positive, except around Age 50 in United Kingdom in 2014. Japanese Males are experiencing the highest improvement rates, although their improvement pattern fluctuates a lot over the years. All populations indicate present mortality “deceleration” manifested in  $\partial m_{\text{diff}} < \partial m_{\text{back}}$  in 2014, suggesting a trend of declining mortality improvement going forward as discussed in Section 3.5. Third, the age-shapes of  $x_{ag} \mapsto \partial m(\cdot, x_{ag})$  are quite volatile over time and frequently “rotate” (perhaps due to cohort effects). Of interest is that the improvement in the 50–60 Age range is very steady over the years for all Females (1.5–2% p/a), while it is all over the place for Males or older ages. There is often a noticeable acceleration/deceleration of mortality, i.e., a clear trend in  $x_{yr} \mapsto \partial m(x_{yr}, ag)$  across years. For instance, decreasing improvements in Japan and UK Females in their 60s, or increasing improvement in Japan Males at ages 50–60. These observations suggest diverging mortality experiences across different sub-groups. Finally, we clearly observe the difficulty in accurately learning mortality improvement rates: while the models are fairly confident about *MI* back in year 2000, they are much less so for the “edge” year 2014, generating much wider relative uncertainty bands.

## 4. EXTENSIONS OF GP MODELS

### 4.1. Inhomogeneous GP models

Basic GP assumes a stationary covariance structure that may not be appropriate. If the spatial dependence in mortality experience is state-dependent, i.e.,  $C(x^i, x^j)$  depends on  $x^i, x^j$  (and not just  $|x^i - x^j|$ ), this would introduce model misspecification and lead to poor model performance (i.e., too much or too little smoothing).

To test for inhomogeneous correlation, we consider a GP model segmented by age. This means that we introduce a piecewise setup, fitting three different

TABLE 8

GP MODELS FITTED BY AGE GROUPS. ALL MODELS ARE FITTED TO YEARS 1999–2014 AND USING A SQUARED-EXPONENTIAL KERNEL WITH A QUADRATIC MEAN FUNCTION

$$m(x^n) = \beta_0 + \beta_1^{ag} x_{ag}^n + \beta_1^{yr} x_{yr}^n + \beta_2^{ag} (x_{ag}^n)^2.$$

Ages Fit	$\beta_0$	$\beta_1^{ag}$	$\beta_2^{ag}$	$\beta_1^{yr}$	$\eta^2$	$\sigma^2$	$\theta_{ag}$	$\theta_{yr}$
Extended [1, 84]	-23.533	-0.005	8.402e-04	7.797e-03	1.904e-01	1.184e-03	3.966	12.795
Younger [50, 69]	10.521	0.084	-3.336e-05	-9.908e-03	2.633e-03	2.964e-04	4.501	4.196
Older [70, 84]	26.806	-0.016	7.113e-04	-1.635e-02	1.489e-03	1.517e-04	14.709	6.661
All [50, 84]	19.336	0.041	3.324e-04	-1.367e-02	1.760e-03	2.336e-04	4.543	3.825

The reported hyperparameter values are maximum likelihood estimates from DiceKriging.

GP models depending on  $x_{ag}$ . The age grouping was done manually according to (younger)  $x_{ag} \in \{50, \dots, 69\}$ , (older)  $x_{ag} \in \{70, \dots, 84\}$ , as well as the full model  $x_{ag} \in \{50, \dots, 84\}$ , and an extended model considering all ages  $x_{ag} \in \{1, \dots, 84\}$ . Table 8 presents the fitted trend and hyperparameters for each group using a model fitted to all years 1999–2014 and quadratic mean function.

Table 8 shows that the Extended age group trend/shape parameter estimates differ from the remaining groups, likely due to the fact that infant and adolescent mortality produce a non-quadratic mortality shape in age. Furthermore, the respective positive coefficient of the Extended  $\beta_1^{yr}$  parameter contradicts the idea of mortality improvement and possibly indicates poor goodness-of-fit.

Segmenting the older ages does generate some reasonable differences in fitted models: log-mortality is linear in the younger group, so that the  $\beta_2^{ag}$  coefficient is negligible; it is larger in the older age group due to the rapid increase of mortality in age; combining the two as was done originally yields an average of the two estimates. The estimates of  $\beta_1^{yr}$  also support the claim of Older mortality improving faster than Younger mortality: log-mortality decreases annually at 1% for the Younger group and at 1.6% for the Older group. The  $\theta_{yr}$  values are all similar across groups, except for the Extended group that needs to compensate for its poor trend fit. The Younger and Extended fits share similar  $\theta_{ag}$  values. We attribute the larger  $\theta_{ag}$  for Older ages to fitting issues due to a complicated age dependence and only 15 ages worth of data (it could also suggest that mortality rates of older ages are more correlated). There is further evidence of this when comparing with Table B2 for Females in the appendix that also produces an unreasonably large value of  $\theta_{ag}^{Fem} = 44.118$  for Older ages.

In sum, this preliminary investigation suggests that a single model that includes all ages is inappropriate and both the mean and covariance structures have further age-dependence. More detailed “change-point” analysis may be warranted to determine the best segmentation of data, and whether the lower cutoff at Age 50 is appropriate. We remark that there exist hierarchical GP models (Gramacy and Taddy, 2012) that attempt to automatically carry out such

data splitting. See also Li and O'Hare (2015) for a discussion about “local” versus “global” approaches to mortality.

#### 4.2. Modeling cause of death scales

The raw CDC data are classified by cause of death and hence it is in fact possible to build a comprehensive mortality improvement model that is broken down beyond the basic Male/Female distinction. Understanding the different trends in cause-of-death can be important as there has been uneven progress (and in some situations reversal) of longevity improvements by cause. For example, the large improvement in mortality from coronary artery disease has not been matched by improvements in mortality from cancer. Different causes of death affect different ages, creating multiple “cross-currents” that drive mortality, a fact which is important for long-term projections.

Thus, mortality improvement models can benefit from analyzing by-cause data. Building such models would need to balance the risk of over-specification with the benefit of incorporating additional data. Key issues and concepts in building a by-cause model are as follows:

- The mean function,  $m$ , would need to be fit to each cause.
- The covariance function controlling spatial correlation would also likely differ by cause.
- This paper focuses on modeling the log mortality rate. A by-cause model would benefit instead from modeling the force of mortality from each cause, as the total force of mortality is simply a sum of the underlying by-causes forces of mortality. However, this additive structure does not match the log transformation applied in this paper.
- Bayesian models with informative priors for mean function and other coefficients would provide a degree of protection against overfitting the models.
- A hierarchical model that builds in a relationship between the by-cause trend coefficients could be tested.

Such analysis is left for further research.

#### 4.3. Model updating

The GP model is convenient for analysis when new data becomes available. This is in contrast to methods, such as splines, which require a full model refit. With GPs, once the correlation structure is fit (and assuming it did not change), the Gaussian posterior  $\mathbf{f}_*$  allows for an updated  $\mathbf{m}_*$  and  $\mathbf{C}_*$ , see Ludkovski (2015, Section 5.1) for details. These formulas showcase the explicit impact of additional data, both for smoothing past experience, or projecting forward in time.

To illustrate the effect of a new year of data, we compute the predicted mean  $\mathbf{m}_*$  and standard deviation  $\mathbf{C}_*$  for Age 65 and years 1999, 2013 and 2016, first based on data for all ages and calendar years 1999–2013, and then updated with year-2014 data. The results are listed in Table 9.



TABLE 9

GP MODEL UPDATING:  $\mathbf{x}$ ,  $\mathbf{y}$  REFERS TO OBSERVED MORTALITY FOR AGES 50–84, YEARS 1999–2013;  $\tilde{\mathbf{x}}$ ,  $\tilde{\mathbf{y}}$  IS THE SAME DATA AUGMENTED WITH YEAR-2014 EXPERIENCE.

$x_{yr}$	Before Updating (1999–2013)		After Updating (1999–2014)	
	$\mathbb{E}[f(65, x_{yr}) \mathbf{x}, \mathbf{y}]$	$s_*(65, x_{yr})$	$\mathbb{E}[f(65, x_{yr}) \tilde{\mathbf{x}}, \tilde{\mathbf{y}}]$	$\tilde{s}_*(65, x_{yr})$
1999	−3.8845	0.0174	−3.8849	0.0173
2013	−4.1497	0.0174	−4.1502	0.0170
2016	−4.1197	0.0266	−4.1248	0.0208

The mean function is intercept-only,  $m(x) = \beta_0$ ;  $s_*$  is posterior standard deviation.

The additional year of credibility decreases posterior standard deviations  $s_*$ . Unsurprisingly, the impact on 1999-prediction is negligible since it is so far in the past. The standard deviation for 2013 has a slight decrease after updating, while 2016 has a much larger reduction: the original model was initially predicting 3 years out-of-sample, while the updated one does for just 2 years out-of-sample. Similarly, the in-sample means change only slightly, while the out-of-sample 2016 has a larger adjustment. The overall decrease in updated posterior means is consistent with the fact that the observed log-mortality for Age 65 in 2014 was −4.1543, lower than the predicted −4.1443 using the 1999–2013 model.

**4.4. Other extensions**

A standard assumption is that mortality curves are increasing in Age, i.e.,  $x_{ag} \mapsto f(x_{ag}, \cdot)$  is monotone. The basic GP framework does not impose any monotonicity restriction. Such structural constraints on  $f$  can help in improving mortality projection in terms of  $m_*$  (especially for long-range forecasts), as well as reduce predictive uncertainty measured by  $s_*^2$ . At the same time, constraints are at odds with the underlying Gaussian random field statistical paradigm, introducing additional complexity in fitting and making inference from the constrained posterior.

One promising recent solution was proposed in Riihimäki and Vehtari (2010) who suggested incorporating monotonicity by adding virtual observation points  $\tilde{x}_i, \tilde{m}_i$  for the derivative of  $f(x_i)$ . Because the derivative  $f'$  also forms a GP, one can explicitly write down the joint covariance structure of  $(f, f')$  (for example, the posterior mean of  $f'$  is the derivative of  $m_*$ ). Monotonicity is then implied by requiring the derivative to be positive at the given  $\tilde{x}_i$ 's. As the size of the latter collection increases, the resulting estimate is more and more likely to be increasing *everywhere* in the domain. This strategy circumvents the direct monotonicity restriction while maintaining computational tractability through linear constraints. Riihimäki and Vehtari (2010) give a recipe for adaptively placing such virtual derivative points by iteratively adding new  $\tilde{x}_i$ 's where the current

$m_*$  violates monotonicity. Further constraints, such as expert opinions about mortality at extreme Ages (100+) could be beneficially added.

An additional extension involves use of multiple data sets; there are many instances where mortality data from one source might be more up-to-date than from other sources, for example, CDC data provides at least three more years of information than SSA data. The use of co-kriging models or the use of CDC data as an input to a GP used to model SSA data is another avenue of possible future research. Such co-kriging models might also be helpful when using population improvement data to supplement a GP analysis of a specific insurance company's or pension fund's mortality experience.

## 5. CONCLUSION

We have proposed and investigated the use of GP models for smoothing and forecasting mortality surfaces. Our approach takes a unified view of the mortality experience as a statistical *response surface* that is noisily reflected in realized mortality experience. A statistical procedure is then used to calibrate the spatial dependence among the latent log-mortality rates. The GP model provides a consistent, nonparametric framework for uncertainty quantification in *both* the mortality surface itself, as well as mortality improvement, which corresponds to relationship between  $f$  and  $x_{yr}$ . This quantification can be done in-sample, by retrospectively smoothing raw mortality counts, or out-of-sample, by building mean forecasts, uncertainty bands, and full scenarios for future mortality/mortality improvement evolution. In contrast, traditional actuarial techniques for graduating data commonly and currently in use (e.g., the Whittaker–Henderson model used by RPEC) focus on smoothing noisy data but fail to provide measures of uncertainty about the fit.

We have focused on population data and smoothing over age and year. The model can be easily extended to additional dimensions, such as duration and net worth in the context of life insurance, or year-of-birth cohort for pension mortality analysis. Adding covariates to the definition of the covariance kernel  $C(x, x')$  is straightforward, with the main challenge lying in interpreting the resulting GP parameters that would reflect a modified concept of spatial distance.

Perhaps the most useful application of our model is for analyzing the latest mortality data, i.e., at the “edge” of the mortality surface. Here, we find and document the statistical evidence that US mortality improvements have materially moderated across a large swath of ages. In particular, for Ages 55–70, US mortality has been effectively flat, or possibly even increasing in the 2010's. This points to a large divergence from the MP-2015 improvement scales that continue to assume significant mortality gains for all Ages and would seem to be overstated at least in the near-term. Moreover, by explicitly computing the differential mortality improvement  $MI_{diff}^{GP}$ , our model gives the most current, instantaneous forecast on mortality improvement, in contrast to the traditional year-over-year estimates.

On a related note, our analysis quantifies the apparent correlation in observed mortality experience across Age and calendar Year. Thus, the obtained estimates of length-scales  $\theta_{yr}$ , imply that studies with very long historical analysis (e.g., going back to 1950 or even 1900) may not add much value to our understanding of current or future projected trends in mortality improvement. Similarly, long-term projections of future mortality improvement (e.g., MP-2015 which is used for projecting mortality often 40–60 years into the future) contain a higher degree of uncertainty than is typically recognized in actuarial analyses. Indeed, our results suggest that projections more than a decade into the future are entirely based on the assumed prior calendar trend and hence have almost no credibility based on observed experience.

Our results show that even a “vanilla” implementation of a GP model already produces useful statistical description of the mortality experiences that is competitive with existing methods in terms of its probabilistic richness and accuracy. We therefore see an enormous potential for further works in this direction, in particular to resolve some further methodological challenges. Mean function modeling which is typically not an important component of GP models in other contexts, is critical for actuaries when projecting out-of-sample. Also, constrained GP models that structurally enforce the age-shape of mortality could be promising in creating better future forecasts. Yet another challenge is better blending of the data-influenced prediction and the prior mean for extrapolation that can be achieved with other Gaussian field specifications or other techniques (Salemi *et al.*, 2013; Lee and Owen, 2015). A different challenge consists in creating meaningful backtesting analyses which would test not just predictive accuracy of  $m_*$ , but also the quality of the generated credibility intervals (both for mortality rates and mortality improvements), and the assumption of Age- and Year-stationary covariance structure. On that point, it would be worthwhile to investigate data from other countries to infer commonalities in mortality correlations.

#### SUPPLEMENTARY MATERIALS

For supplementary material for this article, please visit <http://dx.doi.org/10.1017/asb.2018.24>

#### REFERENCES

- ADLER, R.J. (1981) *The Geometry of Random Fields*, Classics in Applied Mathematics, volume 62. SIAM. Chichester, NY.
- BROOKS, S., GELMAN, A., JONES, G. and MENG, X.-L. (2011) *Handbook of Markov Chain Monte Carlo*. CRC press. Boca Raton, FL.
- BROUHNS, N., DENUIT, M. and VERMUNT, J.K. (2002) A Poisson log-bilinear regression approach to the construction of projected life tables. *Insurance: Mathematics and Economics*, **31**(3), 373–393.

- 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.
- CAIRNS, A.J., BLAKE, D., DOWD, K., COUGHLAN, G.D., EPSTEIN, D. and KHALAF-ALLAH, M. (2011) Mortality density forecasts: An analysis of six stochastic mortality models. *Insurance: Mathematics and Economics*, **48**(3), 355–367.
- CAIRNS, A.J., BLAKE, D., DOWD, K., COUGHLAN, G.D., EPSTEIN, D., ONG, A. and BALEVICH, I. (2009) A quantitative comparison of stochastic mortality models using data from England and Wales and the United States. *North American Actuarial Journal*, **13**(1), 1–35.
- CAMARDA, C.G. (2012) Mortalitysmooth: An R package for smoothing Poisson counts with P-splines. *Journal of Statistical Software*, **50**(1), 1–24.
- CARPENTER, B., LEE, D., BRUBAKER, M.A., RIDDELL, A., GELMAN, A., GOODRICH, B., GUO, J., HOFFMAN, M., BETANCOURT, M. and LI, P. (2016) Stan: A probabilistic programming language. *Journal of Statistical Software*, **76**(1). doi 10.18637/jss.v076.i01
- Continuous Mortality Investigation (2015) The CMI mortality projections model, CMI 2015. Technical report, CMI Working Paper 84.
- CRESSIE, N. (2015) *Statistics for Spatial Data*. Canada: John Wiley & Sons.
- CURRIE, I.D. (2013) Smoothing constrained generalized linear models with an application to the Lee-Carter model. *Statistical Modelling*, **13**(1), 69–93.
- CURRIE, I.D. (2016) On fitting generalized linear and non-linear models of mortality. *Scandinavian Actuarial Journal*, **2016**(4), 356–383.
- CURRIE, I.D., DURBAN, M. and EILERS, P.H. (2004) Smoothing and forecasting mortality rates. *Statistical Modelling*, **4**(4), 279–298.
- CZADO, C., DELWARDE, A. and DENUIT, M. (2005) Bayesian poisson log-bilinear mortality projections. *Insurance: Mathematics and Economics*, **36**(3), 260–284.
- DEBÓN, A., MARTÍNEZ-RUIZ, F. and MONTES, F. (2010) A geostatistical approach for dynamic life tables: The effect of mortality on remaining lifetime and annuities. *Insurance: Mathematics and Economics*, **47**(3), 327–336.
- DELWARDE, A., DENUIT, M. and EILERS, P. (2007) Smoothing the Lee-Carter and Poisson log-bilinear models for mortality forecasting a penalized log-likelihood approach. *Statistical Modelling*, **7**(1), 29–48.
- DOKUMENTOV, A. and HYNDMAN, R.J. (2014) Bivariate Data with Ridges: Two-Dimensional Smoothing of Mortality Rates. Technical report, Working paper series, Monash University.
- GIROSI, F. and KING, G. (2008) *Demographic Forecasting*. Princeton, NJ: Princeton University Press.
- GRAMACY, R. and TADDY, M. (2012) Tgp, an R package for treed Gaussian process models. *Journal of Statistical Software*, **33**, 1–48.
- HUNT, A. and BLAKE, D. (2014) A general procedure for constructing mortality models. *North American Actuarial Journal*, **18**(1), 116–138.
- 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.
- LEE, M.R. and OWEN, A.B. (2015) Single nugget kriging. Technical report, arXiv preprint arXiv:1507.05128.
- LEE, R.D. and CARTER, L.R. (1992) Modeling and forecasting US mortality. *Journal of the American Statistical Association*, **87**(419), 659–671.
- LI, H. and O'HARE, C. (2015) Mortality forecast: Local or global? Technical report, Available at SSRN 2612420.
- LUDKOVSKI, M. (2015) Kriging metamodels for bermudan option pricing. arXiv:1509.02179.
- MITCHELL, D., BROCKETT, P., MENDOZA-ARRIAGA, R. and MUTHURAMAN, K. (2013) Modeling and forecasting mortality rates. *Insurance: Mathematics and Economics*, **52**(2), 275–285.
- PICHENY, V. and GINSBOURGER, D. (2013) A nonstationary space-time Gaussian process model for partially converged simulations. *SIAM/ASA Journal on Uncertainty Quantification*, **1**(1), 57–78.
- PURUSHOTHAM, M., VALDEZ, E. and WU, H. (2011) Global mortality improvement experience and projection techniques. Technical report, Society of Actuaries.
- RENSHAW, A., HABERMAN, S. and HATZOPOULOS, P. (1996) The modelling of recent mortality trends in united kingdom male assured lives. *British Actuarial Journal*, **2**(2), 449–477.

- RENSHAW, A.E. and HABERMAN, S. (2003) Lee–Carter mortality forecasting with age-specific enhancement. *Insurance: Mathematics and Economics*, **33**(2), 255–272.
- RENSHAW, A.E. and HABERMAN, S. (2006) A cohort-based extension to the Lee–Carter model for mortality reduction factors. *Insurance: Mathematics and Economics*, **38**(3), 556–570.
- RIIHIMÄKI, J. and VEHTARI, A. (2010) Gaussian processes with monotonicity information. *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, in *PMLR*, **9**:645–652.
- ROSNER, B., RAHAM, C., ORDUÑA, F., CHAN, M., XUE, L., ZAK, B. and YANG, G. (2013) Literature review and assessment of mortality improvement rates in the US population: Past experience and future long-term trends. Technical report, Society of Actuaries.
- ROUSTANT, O., *et al.* (2012) Dicekriging, Diceoptim: Two R packages for the analysis of computer experiments by Kriging-based metamodeling and optimization. *Journal of Statistical Software*, **51**(1), 1–55.
- SALEMI, P., STAUM, J. and NELSON, B.L. (2013) Generalized integrated Brownian fields for simulation metamodeling. In *Proceedings of the 2013 Winter Simulation Conference*, IEEE Press, pp. 543–554.
- SITHOLE, T.Z., HABERMAN, S. and VERRALL, R.J. (2000) An investigation into parametric models for mortality projections, with applications to immediate annuitants and life office pensioners data. *Insurance: Mathematics and Economics*, **27**(3), 285–312.
- SOA (2014a) Mortality improvement scale MP-2014 report. Technical report, Retirement Plans Experience Committee. <https://www.soa.org/Research/Experience-Study/Pension/research-2014-mp.aspx>.
- SOA (2014b) RP-2014 mortality tables. Technical report, Society of Actuaries Pension Experience Study. <https://www.soa.org/Research/Experience-Study/pension/research-2014-rp.aspx>.
- SOA (2015) Mortality improvement scale MP-2015. Technical report, Retirement Plans Experience Committee. <https://www.soa.org/Research/Experience-Study/Pension/research-2015-mp.aspx>.
- VILLEGAS, A.M., KAISHEV, V.K. and MILLOSVOVICH, P. (2015) StMoMo: An R package for stochastic mortality modelling. Technical report, SSRN documentation at [papers.ssrn.com/sol3/papers.cfm?abstract-id=2698729](https://papers.ssrn.com/sol3/papers.cfm?abstract-id=2698729).
- WHITTAKER, E.T. (1922) On a new method of graduation. *Proceedings of the Edinburgh Mathematical Society*, **41**, 63–75.
- WILLIAMS, C.K. and RASMUSSEN, C.E. (2006) *Gaussian Processes for Machine Learning*. Massachusetts: The MIT Press.
- WILMOTH, J.R. and SHKOLNIKOV, V. (2010) *Human Mortality Database*. University of California.
- WU, R. (2016) *Gaussian Process and Functional Data Methods for Mortality Modelling*. PhD thesis, Department of Mathematics University of Leicester.

MIKE LUDKOVSKI

*Department of Statistics & Applied Probability*

*University of California, Santa Barbara CA 93106-3110, USA*

*E-Mail: ludkovski@pstat.ucsb.edu*

JIMMY RISK (Corresponding author)

*Department of Mathematical Sciences*

*Cal Poly Pomona, Pomona, CA 91768, USA*

*E-Mail: jrisk@cpp.edu\**

HOWARD ZAIL

*Actuary and Founder of Elucidor, LLC, NY 10016, USA*

*E-Mail: hzail@elucidor.com*

\*Jimmy Risk's email address has been corrected. A corrigendum notice detailing this change was also published (DOI: 10.1017/asb.2018.39).

## APPENDIX A. SUPPLEMENTARY PLOTS

### A1. GP MODEL RESIDUALS

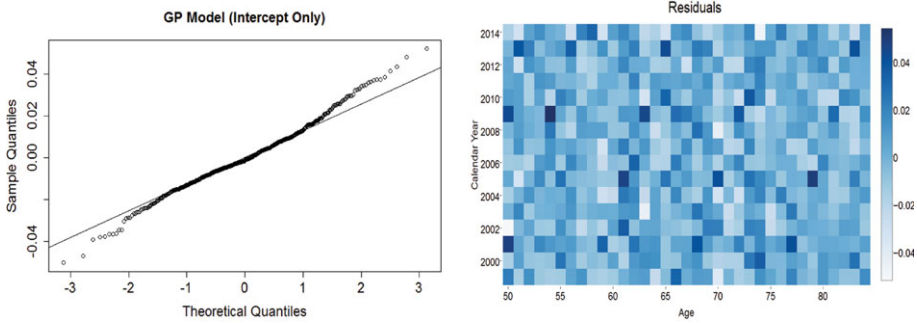


FIGURE A1: Left: Q-Q Plots for residuals of a fitted GP model with mean function  $m(x) = \beta_0$ . We use Male All Data to test the normality assumption of  $\epsilon$  in (3). We observe that the GP residuals are reasonably Gaussian with mildly heavy tails. Right: heatmap of  $\epsilon(x)$  as a function of the two-dimensional input  $x = (x_{ag}, x_{yr})$ . We observe no apparent correlation in the fitted residuals. (Color online)

## APPENDIX B. TABLES AND FIGURES FOR US FEMALE DATA




FIGURE B1: Mortality rates for US Females Aged 60–70 during years 2010–2014. Raw vs. estimated smoothed mortality curves. Models are fit to All Female data. (Color online)




FIGURE B2: Mortality improvement factors for US Females using All data. Solid red lines indicate the empirical mortality experience; dotted blue lines are the smoothed estimates using a GP. (Color online)




FIGURE B3: Comparison of yearly mortality improvement factors for US Females using All data. The curve for 2014 is the same as in Figure B2. (Color online)

TABLE B1

MEAN FUNCTIONS AND FITTED COVARIANCE PARAMETERS USING SET I US FEMALE DATA (AGES 50–70 AND YEARS 1999–2010).

	Trend Parameter MLE's				GP Hyperparameter MLE's			
	$\beta_0$	$\beta_1^{ag}$	$\beta_2^{ag}$	$\beta_1^{yr}$	$\eta^2$	$\sigma^2$	$\theta_{ag}$	$\theta_{yr}$
Intercept	-5.101	-	-	-	4.444e-01	2.968e-04	7.363	10.882
Linear	4.484	0.083	-	-7.167e-03	2.802e-03	3.682e-04	4.432	4.505
Quadratic	11.207	0.054	2.712e-04	-1.014e-02	2.053e-03	2.911e-04	4.464	4.384

The mean functions are  $m(x) = \beta_0$  for Intercept,  $m(x) = \beta_0 + \beta_1^{ag} x_{ag} + \beta_1^{yr} x_{yr}$  for Linear, and  $m(x) = \beta_0 + \beta_1^{ag} x_{ag} + \beta_1^{yr} x_{yr} + \beta_2^{ag} x_{ag}^2$  for Quadratic.

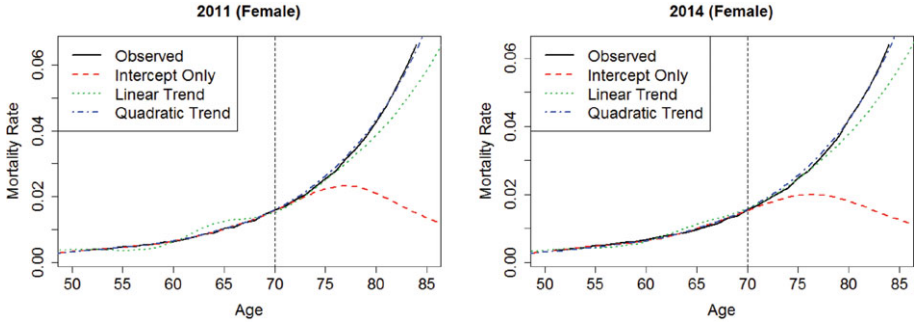


FIGURE B4: Comparison of mean function choices in extrapolating mortality rates at old ages for US Females. Models are fit to years 1999–2010 and Ages 50–70 (Subset III), with estimates made for Age 50–85 in 2011 and 2014. The vertical line indicates the boundary of the training dataset in  $x_{\text{age}}$ . The mean functions are given in Table B1. (Color online)

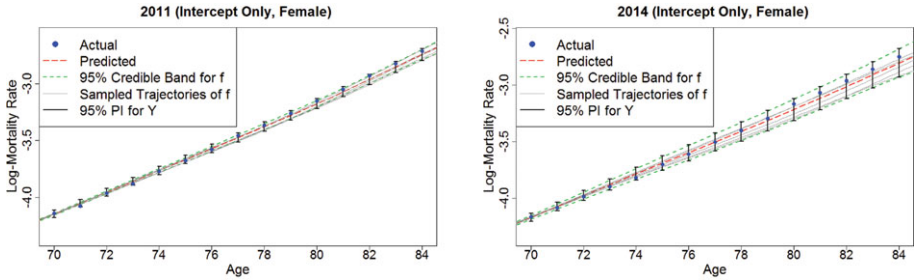


FIGURE B5: Mortality rate prediction for years 2011 and 2014 and Ages 71–84. Model is fit on Subset II US Female data with intercept-only mean function and squared-exponential kernel. (Color online)



TABLE B2  
GP MODELS FITTED BY AGE GROUPS WITH US FEMALE DATA.

$$\text{Quadratic } m(x) = \beta_0 + \beta_1^{ag} x_{ag} + \beta_1^{yr} x_{yr} + \beta_2^{ag} x_{ag}^2$$

Ages Fit	$\beta_0$	$\beta_1^{ag}$	$\beta_2^{ag}$	$\beta_1^{yr}$	$\eta^2$	$\sigma^2$	$\theta_{ag}$	$\theta_{yr}$
Extended [1, 84]	-25.224	-0.008	8.721e-04	8.678e-03	2.170e-01	1.187e-03	4.095	13.040
Younger [50, 69]	1.128	0.080	3.912e-05	-5.471e-03	4.311e-03	2.907e-04	5.695	5.487
Older [70, 84]	17.272	-0.038	9.071e-04	-1.151e-02	2.543e-03	1.334e-04	44.118	6.856
All [50, 84]	7.473	0.035	4.186e-04	-7.980e-03	2.814e-03	2.236e-04	5.574	5.249

All models used squared-exponential kernel and years 1999–2014.

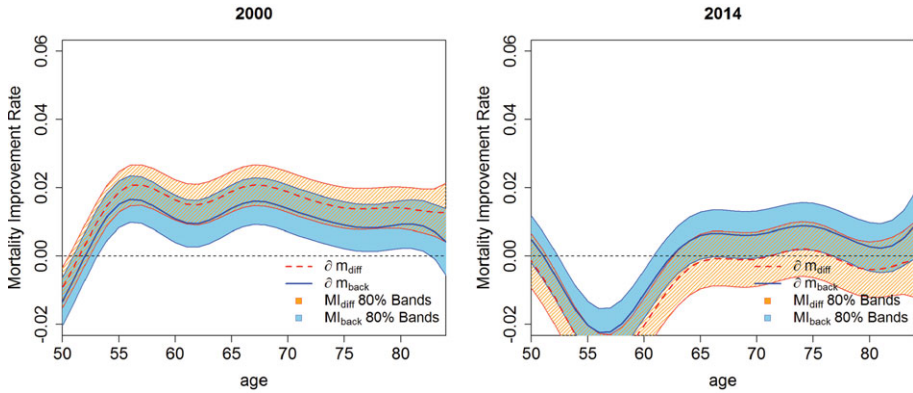


FIGURE B6: Estimated US Female mortality improvement using the differential GP model (instantaneous improvement) and the YoY improvement from the original GP model. We show the means and 80% uncertainty bands for  $M_{diff}^{GP}$  and  $M_{back}^{GP}$  for Ages 50–84 and years 2000 and 2014. Models used are fit to All Data. (Color online)