

A unified approach to mortality modelling using state-space framework: characterisation, identification, estimation and forecasting

Man Chung Fung*

Data61, CSIRO, Australia

Gareth W. Peters

Department of Statistical Science, University College London, 1–19 Torrington Place, WC1E 7HB, United Kingdom

Oxford Mann Institute, Oxford University

Systemic Risk Center, London School of Economics

Pavel V. Shevchenko

Department of Applied Finance and Actuarial Studies, Macquarie University, NSW 2109, Australia

Abstract

This paper explores and develops alternative statistical representations and estimation approaches for dynamic mortality models. The framework we adopt is to reinterpret popular mortality models such as the Lee–Carter class of models in a general state-space modelling methodology, which allows modelling, estimation and forecasting of mortality under a unified framework. We propose alternative model identification constraints which are more suited to statistical inference in filtering and parameter estimation. We then develop a class of Bayesian state-space models which incorporate a priori beliefs about the mortality model characteristics as well as for more flexible and appropriate assumptions relating to heteroscedasticity that present in observed mortality data. To study long-term mortality dynamics, we introduce stochastic volatility to the period effect. The estimation of the resulting stochastic volatility model of mortality is performed using a recent class of Monte Carlo procedure known as the class of particle Markov chain Monte Carlo methods. We illustrate the framework using Danish male mortality data, and show that incorporating heteroscedasticity and stochastic volatility markedly improves model fit despite an increase of model complexity. Forecasting properties of the enhanced models are examined with long-term and short-term calibration periods on the reconstruction of life tables.

Keywords

Mortality modelling; State-space model; Stochastic volatility; Bayesian inference; Particle Markov chain Monte Carlo

1. Introduction

An ageing population is a major challenge that many countries are facing today. The problem arises from the fact that fertility rates are declining while life expectancy has been increasing in the past

*Correspondence to: Man Chung Fung, Data61, CSIRO, Pembroke Road, Marsfield, NSW 2122, Australia. Tel: (612) 94905873. E-mail: simon.fung@csiro.au

several decades without any sign of slowing down. The adverse financial outcome of people living longer than expected, and hence the possibility of outliving their retirement savings, is known as longevity risk. This long-term demographic risk has significant implications for societies and manifests as a systematic risk for pension plans and annuity providers. Policymakers rely on mortality projection to determine appropriate pension benefits and to understand the costing of different economic assumptions and regulations regarding the age of retirement of a given population. For instance, in the United Kingdom and Australia defined-benefit pension plans prior to 2000s had limited exposure to effects of longevity risk since high equity returns on pension fund wealth management portfolios were masking the impact of longevity risk, however, post-2000 declining equity returns coupled with record low interest rate financial environments has demonstrated the significance of decades of longevity improvements, posing a very real problem for pension schemes. Furthermore, by regulation, insurers who offer retirement income products are required to hold additional reserving capital to cover longevity risk. A key input to address longevity risk is the development of advanced mortality modelling methodology, so that human longevity can be predicted with better accuracy and any uncertainties can be accounted for in mortality forecasting.

Since the introduction of the Lee–Carter (LC) model (Lee & Carter, 1992), a range of stochastic mortality models have been proposed in the literature. Renshaw & Haberman (2003) and Renshaw & Haberman (2006) introduce multiple period effects and cohort effect to capture the change of mortality with respect to year and year-of-birth, respectively, to the LC model. Cairns *et al.* (2006) proposed a two-factor period effect mortality model, known as the Cairns–Blake–Dowd (CBD) model, for pensioner ages. A cohort extension of the CBD model was studied in Cairns *et al.* (2009). Plat (2009) draws on the strengths of the LC model and the CBD model to produce an age–period–cohort model that includes a term to capture young mortality dynamics. In these well-known cases it is common practice in actuarial settings to estimate stochastic mortality models based on a singular value decomposition (SVD) approach (Lee & Carter, 1992; Koissi *et al.*, 2006) or via a maximum likelihood-based approach if a discrete Poisson regression setting is considered (Brouhns *et al.*, 2002; Cairns *et al.*, 2009).

A common feature of the estimation methods adopted in the frameworks mentioned above is that the dynamics of the period effect, the stochastic latent processes, are not directly incorporated into joint parameter and state estimation, and instead form a component of a second stage of estimation. Typically this involves specifying a model for the period effect for forecasting purpose only after an estimation is performed. Such approaches often suffer from a statistical lack of efficiency compared to methods that perform joint static model parameter estimation and latent process filtering. *Hence, the first argument we make is that recasting different classes of mortality models in a state-space formulation can better facilitate state-space-based inference under either frequentist or Bayesian estimation. This is especially true in the case that the inference is performed jointly on the latent process and static model parameters, rather than in a less statistically efficient two-stage procedure.*

We note that in the classical actuarial approach that utilises the two-stage procedure, a wide variety of model choices may be adopted to model the latent stochastic dynamics such as period effect, cohort effect, etc. These models may be time series models such as random walks with drift, seasonal autoregressive integrated moving average (ARIMA) model, long memory processes such as generalized autoregressive moving average (GARMA) model or stochastic volatility model such as generalized autoregressive conditional heteroscedasticity (GARCH) model, etc. It is upto the modeller to perform statistical estimation, model selection and model criticism stages such as via Box–Jenkins type procedures to select an appropriate model

choice. In the one-stage estimation framework of the state-space model, exactly the same flexible classes of model are also admissible for the state equation, so there is no loss in model generality when working with the state-space formulation. As we will discuss in linear Gaussian state-space choices such as under seasonal ARIMA (SARIMA) class of models there are optimal state-space estimators based on Kalman filters, but in more general classes of non-linear time series models one can use Sequential Monte Carlo (SMC) filtering as discussed below. Note also that one can always carry out model diagnostics in the state-space framework to decide whether the model proposed is suitable given the data. In addition, there are various formal statistical criteria designed for model ranking as well as model selection which are applicable for the one-stage procedure in state-space setting.

Typically the studies carried out in practice and in the literature have the feature that only mortality data from the past several decades is considered. For many countries, age-specific death rates are evolving rather smoothly except for some potential change of trend in the last 50 years or so in some developed countries. Besides ARIMA models, structural change model have been proposed to take into account the trend-changing behaviour of the period effect (Li *et al.*, 2011; van Berkum *et al.*, 2016). Despite this, the implication of including earlier periods that exhibit significant volatility of mortality, which can be attributed to some life-critical events such as wars and epidemics, is still not yet being investigated. The ability to incorporate such structural information into a mortality model is greatly facilitated when recasting the model in a state-space formulation. Furthermore, extensions to mortality models that can also be facilitated in a state-space formulation are increasingly able to be considered and may better explain the stochastic dynamics of such processes. These include features such as time-varying volatility; cross-sectional volatility between different age groups; extremal dependence features; cohort effects; structure breaks in regimes; long memory or persistence in mortality features in different age groups; cointegration and non-stationarity features; as well as regression-based structures that decompose mortality according to categorical features such as official death causes, regional categories, etc.

Moreover, additional stochastic factor models such as two- and three-factor models can be easily considered. This can be particularly relevant when modelling features such as trends in excess mortality in particular age groups resulting from disease epidemics (Zucs *et al.*, 2005; Dawood *et al.*, 2012), cold and heat waves (Fouillet *et al.*, 2006; Analitis *et al.*, 2008) and other effects such as medical impairments, occupational hazards, hazardous pursuits, geographical location of residence and ethnic origin, see Eloranta *et al.* (2012) and England & Haberman (1993). *Hence, the second argument we make is that all these different model structures can readily be encoded in state-space model structures. Furthermore, they can be consistently combined in joint estimation procedures in such state-space model structures in either frequentist and Bayesian formulations, whilst also admitting consistent joint forecasting models for predictive purposes.*

A variety of state-space model approaches exist in a range of different literatures, in this paper we propose to begin with the widely adopted frameworks typically introduced in state-space modelling settings in Harvey (1989) or West & Harrison (1997), which we develop to address some of the aforementioned issues. In contrast to the SVD and Poisson regression estimation approaches where the period effect is treated as parameter without any temporal structure in the first-stage estimation, period effect is regarded as a latent process with a Markovian structure under the state-space approach. In other words, a state-space formulation permits modelling, estimation and forecasting of mortality under a unified framework. Recent progress in sampling-based techniques has allowed statistical inference to be conducted on sophisticated state-space models that can incorporate multiple latent-driving factors which may exhibit non-linear and non-Gaussian stochastic dynamics.

We take advantage of this development and utilise realistic model to capture the long-term volatility structure of mortality time series.

Pedroza (2006) and Kogure & Kurachi (2010) consider Bayesian estimation of the LC model in state-space form. A maximum likelihood approach is studied in De Jong & Tickle (2006). Here, we extend such frameworks to show how to adopt a combination of filtering procedures with Rao-Blackwellisation to obtain gradient-based Fisher score equation recursions to accurately and efficiently perform optimal filtering of the latent state process, in the sense of mean square error minimisation, and recursive least squares estimation for the static model parameters jointly in a recursive manner. Furthermore, we extend such state-space models to incorporate non-linear and non-Gaussian features in the state-space structure that no longer admit simple Kalman filter forward-backward algorithm recursions, leading us to more cutting edge filtering techniques based on SMC methods. In this regard, we estimate and examine the LC model with heteroscedasticity using both gradient-based maximum likelihood and Bayesian analysis. Alternative models that have tried to include such features include, for example, the Poisson regression in Brouhns *et al.* (2002) and Czado *et al.* (2005), who aimed to replace the homogeneous additive error term in the LC model by a Poisson error structure. Also, we note a recently developed framework for modelling death counts with common latent risk factors via credit risk plus methodology with model estimation via Markov chain Monte Carlo (MCMC) in Hirz *et al.* (2017). However, we argue that the state-space formulation allows heteroscedasticity to be accounted for in a more straightforward manner.

Note that Pedroza (2006) also considers a LC model with heteroscedasticity structure. Our approach in this paper, however, is more general compared to Pedroza (2006) as we explain in details the machinery behind the gradient-based maximum likelihood estimation and sampling-based inference, not confining to the restricted class of linear and Gaussian state-space models as previous studies focus on. Such an effort is important especially for researchers and practitioners who are interested in more realistic and sophisticated models which would fall into the class of non-linear and non-Gaussian models, while not wanting to be handicapped by the complications arising from model estimations.

Through reformulation and extensions of the LC type mortality models in a state-space model structure, we investigate several key properties observed in mortality data. First, the cross-sectional variance-covariance matrix between age group structures is non-homogeneous. Second, examination of mortality data over a long period indicates that volatility of the evolution of death rates is not constant, that is, heteroscedasticity is present. We show that the incorporation of a second stochastic volatility latent factor will allow us to identify the periods in which mortality demonstrates heightened volatility. This will aid in interpretation and forecasting from such models. Specifically, we introduce a stochastic volatility model for the period effect, aiming to capture long-term mortality dynamics. The state-space framework provides a natural platform to analyse stochastic volatility models (Kim *et al.*, 1998; Chib *et al.*, 2002). In this paper we develop a particle Markov chain Monte Carlo (PMCMC) (Andrieu *et al.*, 2010). Bayesian model formulation in order to estimate the resulting stochastic volatility model of mortality jointly with the other latent stochastic factors and the static model parameters.

We introduce to mortality modelling the estimation framework based around the PMCMC algorithm which utilises SMC (Doucet *et al.*, 2001; Peters *et al.*, 2012) to obtain required quantities in Metropolis-Hastings algorithms that has found many applications in a variety of areas, for

example, finance (Peters *et al.*, 2013), economics (Flury & Shephard, 2011), non-life insurance (Peters *et al.*, 2010b), risk management (Targino *et al.*, 2015) and computational biology (Golightly & Wilkinson, 2011). We apply this powerful tool in mortality modelling and it allows us to develop efficient algorithms to estimate a stochastic volatility extension of the LC model.

Recently, there are growing interests in applying a state-space framework to the pricing and hedging of longevity risk. For example, Kogure & Kurachi (2010) considers the LC model in state-space form and show that a pricing risk premium for longevity instruments can be derived based on the maximum entropy principle. As another example, Liu & Li (2016a) investigates longevity hedging strategies accounting for population basis risk¹ by formulating multi-population mortality models in the state-space framework. Their approach relies crucially on the determination of the sensitivities of the hedge portfolio and the liability with respect to the hidden states of a state-space multi-population mortality model. In another paper, Liu & Li (2016b) addresses the issues of trend-changing behaviours of mortality rates by considering a locally linear CBD model in state-space representation. In addition, they consider a state-space hedging method including drift risk where the hidden states play an important role. We believe that our general presentation of the state-space framework in the paper, including the machinery such as PMCMC developed in the statistics literature and extended models we proposed, will further facilitate the potentials of utilising state-space methods in modelling human mortality and the risk management of longevity risk.

The paper is organised as follows. In section 2 we give an overview of the conventional mortality modelling and estimation methodology in the literature. A state-space approach for mortality modelling is formulated and discussed in section 3. Section 4 is devoted to state-space inference for stochastic mortality models in a frequentist approach. Section 5 focusses on Bayesian inference for dynamic mortality models in state-space framework. In section 6 we analyse Danish mortality data based on the enhanced models and methodologies proposed in the paper. Section 7 provides concluding remarks.

2. Classical Bayesian and Frequentist Approaches

In this section we first briefly recall some important definitions on mortality modelling. We then review stochastic mortality models that are commonly found in the literature. Standard estimation procedures under frequentist and Bayesian approaches are discussed.

2.1. Definitions and notation

Hereafter, we use the following standard definitions from actuarial literature on mortality modelling (Dickson *et al.*, 2009; Pitacco *et al.*, 2009). Let T_x be a random variable representing the remaining lifetime of a person aged x . The cumulative distribution function and survival function of T_x are written as ${}_xq_x = P(T_x \leq \tau)$ and ${}_x p_x = P(T_x > \tau)$, respectively. For a person aged x , the force of mortality at age $x + \tau$ is defined as

$$\mu_{x+\tau} := \lim_{h \rightarrow 0} \frac{1}{h} P(T_x < \tau + h | T_x > \tau) = - \frac{d}{d\tau} \ln {}_x p_x \quad (1)$$

¹ Population basis risk refers to the risk that the mortality experience of a portfolio being hedged is different to the mortality experience underlying an index-based longevity hedging instrument. This discrepancy causes the hedge to be less effective.

Let $f_x(t)$ be the density function of T_x , then from (1) we have ${}_tq_x = \int_0^t f_x(s)ds = \int_0^t {}_s p_x \mu_{x+s} ds$. The central death rate for a x -year-old, where $x \in \mathbb{N}$, is defined as

$$m_x := \frac{q_x}{\int_0^1 {}_s p_x ds} = \frac{\int_0^1 {}_s p_x \mu_{x+s} ds}{\int_0^1 {}_s p_x ds} \tag{2}$$

which is a weighted average of the force of mortality (here $q_x := {}_1q_x$). Under the so-called piecewise constant force of mortality assumption, that is $\mu_{x+s} = \mu_x$, where $0 \leq s < 1$ and $x \in \mathbb{N}$, we have, from (2), $m_x = \mu_x$. Moreover, if a Poisson assumption is made for the actual number of deaths, then the resulting maximum likelihood estimate of the force of mortality $\hat{\mu}_x$ (and hence \hat{m}_x) is given by $\hat{\mu}_x = D_x/E_x = \hat{m}_x$, where D_x is the number of deaths recorded at age x last birthday and the exposure-to-risk E_x the average number of people aged x last birthday, during the observation year. Note that E_x is approximated by an estimate of the population aged x last birthday in the middle of the observation year. We refer to \hat{m}_x as the crude death rate.

In the above setup it is assumed that the force of mortality μ is deterministic. The stochastic case can be handled by the intensity-based framework where death time is modelled as the first jump time of a doubly stochastic process (Biffis, 2005). Hereafter, we treat the force of mortality $\mu_{x+t}(t)$, the central death rate $m_{x,t}$ and the crude death rate $\hat{m}_{x,t}$ as stochastic processes. For a detailed discussion of the background of stochastic mortality modelling in discrete-time and continuous-time, see Cairns *et al.* (2008).

2.2. Stochastic mortality models

One of the most widely considered examples of stochastic factor model in the context of mortality modelling is the approach first presented in Lee & Carter (1992) who proposed a stochastic mortality model for the age-specific crude death rate $\hat{m}_{x,t}$, where $x = x_1, \dots, x_p$ and $t = 1, \dots, T$ represent age (or age group) and year (time), respectively. Under the model, the dynamics of the log crude death rates, $y_{x,t} = \ln \hat{m}_{x,t}$, is given by²

$$y_{x,t} = \alpha_x + \beta_x \kappa_t + \varepsilon_{x,t}, \quad \varepsilon_{x,t} \stackrel{iid}{\sim} N(0, \sigma_\varepsilon^2) \tag{3}$$

where $N(0, \sigma_\varepsilon^2)$ denotes a Gaussian distribution with zero mean and variance σ_ε^2 . The vector $\alpha = \alpha_{x_1:x_p} := [\alpha_{x_1}, \dots, \alpha_{x_p}]$ represents the age-profile of the log death rates and $\beta = \beta_{x_1:x_p}$ measures the sensitivity of death rates for different age group to a change of the time series κ_t . The period effect, κ_t , for forecasting purpose, is assumed to satisfy the equation

$$\kappa_t = \kappa_{t-1} + \theta + \omega_t, \quad \omega_t \stackrel{iid}{\sim} N(0, \sigma_\omega^2) \tag{4}$$

where $\varepsilon_{x,t}$ and ω_t are independent.

Under this specification, it is clear that the LC model is not identifiable, since (3) is invariant up to some linear transformations of the parameters:

$$y_t = \alpha + \beta \kappa_t + \varepsilon_t = \alpha + \beta c + \frac{\beta}{d} ((\kappa_t - c)d) + \varepsilon_t = \tilde{\alpha} + \tilde{\beta} \tilde{\kappa}_t + \varepsilon_t \tag{5}$$

where $\tilde{\alpha} = \alpha + \beta c$, $\tilde{\beta} = \beta/d$ and $\tilde{\kappa}_t = (\kappa_t - c)d$.

² Alternatively, one may treat the LC model as a model for the log central death rate $\ln m_{x,t} = \alpha_x + \beta_x \kappa_t$. The distinction of the crude and central death rate is of particular importance when one considers a Poisson regression setup of death counts (discussed in section 2.3.2) where the dynamics of the central death rate is being modelled (Cairns *et al.*, 2009; Dowd *et al.*, 2010).

Table 1. Several popular stochastic mortality models.

Models	Dynamics
Lee & Carter (1992)	$\ln(m_{x,t}) = \alpha_x + \beta_x \kappa_t$
Renshaw & Haberman (2003)	$\ln(m_{x,t}) = \alpha_x + \sum_{i=1}^k \beta_x^{(i)} \kappa_t^{(i)}$
Renshaw & Haberman (2006)	$\ln(m_{x,t}) = \alpha_x + \beta_x^{(1)} \kappa_t + \beta_x^{(2)} \zeta_{t-x}$
Currie (2009)	$\ln(m_{x,t}) = \alpha_x + \kappa_t + \zeta_{t-x}$
Cairns <i>et al.</i> (2006)	$\text{logit}(q_{x,t}) = \kappa_t^{(1)} + \kappa_t^{(2)}(x - \bar{x})$
Cairns <i>et al.</i> (2009)	$\text{logit}(q_{x,t}) = \kappa_t^{(1)} + \kappa_t^{(2)}(x - \bar{x}) + \zeta_{t-x}$
Plat (2009)	$\ln(m_{x,t}) = \alpha_x + \kappa_t^{(1)} + \kappa_t^{(2)}(\bar{x} - x) + \kappa_t^{(3)}(\bar{x} - x)^+ + \zeta_{t-x}$

To overcome this identification issue when estimating the LC model, one has to impose a non-unique choice of constraints to restrict the model to an identifiable class. It is standard practice in actuarial literature to consider the following two constraints:

$$\sum_{x=x_1}^{x_p} \beta_x = 1, \quad \sum_{t=1}^T \kappa_t = 0 \tag{6}$$

as suggested in Lee & Carter (1992) to remedy the identifiability issue. This choice of constraints is equivalent to fixing $c = (1/T) \sum_{t=1}^T \kappa_t$ and $d = \sum_{x=x_1}^{x_p} \beta_x$. Consequently we have $\sum_{t=1}^T \tilde{\kappa}_t = 0$ and $\sum_{x=x_1}^{x_p} \tilde{\beta}_x = 1$. The reason for these particular form of identification constraints relates to the fact that the constraint on the path space of the stochastic factor $\kappa_1, \dots, \kappa_T$ is intended to have the effect of centring the κ_t values over the range $t \in \{1, \dots, T\}$, such that the structure is designed to capture age-period effects with the α_x terms incorporating the main age effects, averaged over time and the bilinear terms $\beta_x \kappa_t$ incorporating the age-specific period trends (relative to the main age effects).

Since the introduction of the LC model it has found a widespread uptake of this class of factor model in both practice, where the LC model is now used as a benchmark methodology by the US Bureau of the Census, and in academia where a range of stochastic mortality model extensions have been proposed in the literature, see Table 1.

We note here that Renshaw & Haberman (2003) and Renshaw & Haberman (2006) introduces multi-period $(\sum_{i=1}^k \beta_x^{(i)} \kappa_t^{(i)})$ and cohort factor (ζ_{t-x}) , respectively, to the LC method. Currie (2009) considers a simplified version of the model in Renshaw & Haberman (2006). Cairns *et al.* (2006) propose to model $\text{logit}(q_{x,t}) := \ln(q_{x,t} / (1 - q_{x,t}))$ instead of log death rates and \bar{x} is the average age in the sample range. An addition of cohort factor is studied in Cairns *et al.* (2009). Plat (2009) introduces a model which combines the desirable features of the previous models and include a term $(\bar{x} - x)^+ := \max(\bar{x} - x, 0)$ to capture better young age mortality. The specification of identification constraints for the LC type models, that is for those where the log death rate is being modelled in Table 1, is discussed in Hunt & Villegas (2015).

2.3. Two-stage estimation approaches: frequentist view

Several “classical” approaches to LC model estimation have been proposed in the literature, though they typically involve a two-stage procedure looking first at the observation equation as a regression (ignoring the latent factor structure explicitly) and then in the second stage they fit time series models to the latent factor structures. A good overview of such methods is obtained in

Pitacco *et al.* (2009). This two-stage procedure is at odds with modern state-space modelling procedures which have been progressively moving towards joint parameter estimation and latent state estimation in frequentist and Bayesian formulations, which will be discussed in subsequent sections. This is reflected in the first attempt to improve the calibration approaches as reflected in the comment in Cairns *et al.* (2011) where they highlight that the “... *key element of the proposed framework is our single-stage approach to model fitting and process parameter estimation*”. Such sentiments, relating to consistent single-stage joint estimation are also echoed in the work of Czado *et al.* (2005).

2.3.1. Multi-factor LC SVD-based two-stage calibration

One of the most commonly adopted approaches to estimate stochastic mortality models is via SVD. We use the multi-period (*k*-factor) LC model (Renshaw & Haberman, 2003) with identification constraints given by

$$\sum_{t=1}^T \kappa_t^{(i)} = 0, \quad \sum_{x=x_1}^{x_p} \beta_x^{(i)} = 1 \tag{7}$$

where $i = 1, \dots, k$, as an example to illustrate the methodology below (Koissi *et al.*, 2006).

Stage 1a – Observation Equation Estimation Stage: We first notice that the constraint $\sum_{t=1}^T \kappa_t^{(i)} = 0$ will lead to an estimator for the level α given by

$$\hat{\alpha}_x = \frac{1}{T} \sum_{y=1}^T y_{x,t} \tag{8}$$

Stage 1b – Observation Equation Estimation Stage: The next stage is to de-trend the observations $\{y_{1:T}\}$ by the level estimate $\hat{\alpha}$ and then to perform a SVD on the resulting ($p \times T$) matrix of residual observations to obtain the decomposition:

$$\text{SVD}[y_{1:T} - \hat{\alpha}] = \sum_{i=1}^b \rho_i \mu_i \nu_i^T \tag{9}$$

where T denotes transposition and ρ_i , for $i \in \{1, \dots, b\}$, are the descending singular values where b is the rank of the data matrix. Here μ_i and ν_i are the corresponding left and right singular vectors of the singular value ρ_i with dimension p and T , respectively. For a k -rank, where $k \leq b$, approximation of the matrix, we have

$$y_{1:T} - \hat{\alpha} = \sum_{i=1}^k \rho_i \mu_i \nu_i^T + \varphi_{1:T} \tag{10}$$

where $\varphi_{1:T} = \sum_{i=k+1}^b \rho_i \mu_i \nu_i^T$ is the k -rank residuals. We then identify $\tilde{\beta}^{(i)} = \mu_i$ and $\tilde{\kappa}^{(i)} = \rho_i \nu_i$, for $i = 1, \dots, k$. One then performs the transformation

$$\kappa_t^{(i)} = \tilde{\kappa}_t^{(i)} \sum_x \tilde{\beta}_x^{(i)}, \quad \beta_x^{(i)} = \frac{\tilde{\beta}_x^{(i)}}{\sum_x \tilde{\beta}_x^{(i)}} \tag{11}$$

to ensure the constraints $\sum_x \beta_x^{(i)} = 1$, for $i = 1, \dots, k$, are satisfied.

Stage 2 – Latent Process Factor Estimation Stage: At this stage³, the estimation of the latent factors can be performed by specifying a time series model structure such as ARIMA model for each of the factors:

$$\kappa_t^{(j)} = \theta^{(j)} + \sum_{r=1}^p \kappa_{t-r}^{(j)} + \sum_{s=1}^q \varepsilon_{t-s}^{(j)} + \varepsilon_t \tag{12}$$

or alternatively one could fit the equivalent vector autoregressive model structure not treating each factor as independent in the time series specification. One would typically perform this stage of estimation via the Yule–Walker equations, see for instance discussions in Tsay & Tiao (1984). Under such specifications, one then obtain closed form distributions and estimators for period effect latent factor forecasts that can be substituted into the observation model for forecasts of the mortality by age in future forecast horizons and used to construct life tables.

2.3.2. Regression-based approaches

It is important to note that the SVD approach assumes homoscedasticity in the error structure. Therefore, to account for heteroscedasticity in mortality data for different ages, Brouhns *et al.* (2002) propose to model death counts, instead of death rates, via Poisson regression where the addition error term in the LC approach is replaced by Poisson random variation. Specifically, the number of death $D_{x,t}$ is modelled as

$$D_{x,t} \sim \text{Poisson}(E_{x,t}m_{x,t}(\Phi)) \tag{13}$$

where $E_{x,t}$ is the death exposure, $m_{x,t}(\Phi)$ is a model of the central death rate and Φ is the parameter vector according to the model being used, including time dynamic factors such as period and cohort effect, see, for example, Table 1. The parameter vector is then estimated by maximising the log-likelihood function, which is given by

$$l(\Phi; D, E) = \sum_t \sum_x (D_{x,t} \ln(E_{x,t}m_{x,t}(\Phi)) - E_{x,t}m_{x,t}(\Phi) - \ln(D_{x,t}!)) \tag{14}$$

where $D_{x,t}!$ indicates the factorial of $D_{x,t}$. Times series models are then used to model the time dynamic factors forming a second-stage estimation procedure for forecasting purpose. Note that the CBD type models can be estimated under this approach since we have $q_{x,t} = 1 - \exp\{-m_{x,t}\}$ (Cairns *et al.*, 2009).

Remark 2.1 In all the discussed cases above, there is the general idea that the two-stage estimation approaches (SVD and regression) treat the unobserved factors corresponding to for instance a period effect κ_t and a cohort effect ζ_{t-x} as parameters. For forecasting purpose, these dynamics factors are then modelled as time series, typically under the ARIMA framework. In this paper we argue that a more consistent approach involves embedding the specification of the model formally within a state-space model structure and to perform the estimation via a joint combination of filtering and static parameter estimation, which can be achieved either in Bayesian (posterior-based) or frequentist (likelihood-based) settings. We will demonstrate both in this paper.

2.4. Estimation approaches: Bayesian view

From the Bayesian modelling perspective there are few papers that study stochastic mortality models, the main papers in this area involve the works of Czado *et al.* (2005), Pedroza (2006), Kogure *et al.* (2009) and Cairns *et al.* (2011). As observed in these studies, there are many possible advantages to

³ We omit here the refitting procedure for κ suggested in Lee & Carter (1992).

adopting a Bayesian approach for mortality modelling, especially in the context of small populations which may also have substantial quantities of missing data.

An important point to note is that all Bayesian model formulations to date in the mortality modelling literature, that we are aware of, have utilised what would, in modern statistical approaches be considered rudimentary sampling-based approaches to performing Bayesian estimation of the LC type models. The criticism here can be levelled in two ways:

1. The first relates to the fact that in these Bayesian formulations the latent dynamic process states are still treated in the MCMC sampling procedures as if they were a set of static model parameters. The issues with doing this have been mentioned in numerous places, see for example, Carter & Kohn (1994). Recently new approaches to such inference in Bayesian models have been developed to avoid having to make univariate conjugate Gibbs or Metropolis-within-Gibbs steps for the latent processes. The reason for this is that it is known in general to be very inefficient in performing inference and can be prone to misleading posterior inference results due to poor mixing performance of the Markov chain for a finite computational budget. Detailed discussions have been provided on such problems in Andrieu *et al.* (2010) and subsequently in work such as Chopin *et al.* (2013), and the specific case to population-based state-space models in ecology in Peters *et al.* (2010a). Pedroza (2006) also follows this more efficient sampling approach for mortality modelling based on Carter & Kohn (1994).
2. Second, all existing MCMC sampling-based approaches we are aware of for Bayesian inference in the mortality modelling literature tends to neglect the issue of model identification in the likelihood which can cause issues in the Bayesian formulation. In fact, some approaches implement identification constraint in the Bayesian model and develop an MCMC sampler that tries to impose the identification constraint in such a manner that the resulting Markov chain may not be consistent with preserving the correct invariant stationary distribution if one applies the constraints inappropriately. We investigate this issue in a separate paper (Peters *et al.*, 2017).

These two considerations need to be resolved to update the approaches to more efficient sampling approaches with enhanced specifications of the model formulation to deal with such issues directly. In particular, modern approaches to such model estimations are to treat the latent unobserved process not as static parameters but as a state-space model in which filtering-based methods (Kalman filter variants, SMC) can be utilised for the latent process estimations jointly with consistent estimation of the “static” model parameters. We will detail such estimation procedures which are also consistent with imposing specific identification constraints of relevance to the LC model formulations, that are developed to ensure the correct invariant Bayesian posterior model is preserved by the Markov chain sampler and filters developed.

Remark 2.2 (Likelihood Identification Issues and Bayesian Modelling): We note the fact that model parameters that are not identified in the likelihood pose no formal problem in a Bayesian analysis. Identification is a property of the likelihood function, whereas Bayesian inference simply uses the likelihood function to map through the data from prior beliefs to posterior beliefs. However, it is often the case that working with unidentified likelihood functions is usually unsatisfactory from a practical perspective as it may lead to partial identification issues in the posterior or problematic multimodality in the posterior. In general if one utilises a proper prior distribution it may act to provide a “near-identification” in the sense that one considers parameter restrictions as limiting forms of prior densities, then there is at least a functional equivalence between introducing prior information about parameters, and imposing identifying restrictions.

3. State-Space Formulations of Mortality Models

We are now in a position to present an alternative representations of stochastic mortality modelling based on state-space methodology (Harvey, 1989; West & Harrison, 1997). A key advantage of this approach is that the two-stage estimation and forecasting procedure under the SVD or Poisson regression maximum likelihood approaches can be combined in a single setting. The improved statistical consistency of a single-stage approach is recognised in Cairns *et al.* (2011). Another key advantage comes from the recent progress in sampling-based techniques in the estimation of state-space models. The advancement allows statistical inference to be conducted on sophisticated state-space models. We take advantage of this development and utilise realistic model aiming to capture long-term mortality dynamics.

A general state-space model consists of a state equation

$$\phi_t = a(\phi_{t-1}, \boldsymbol{\mu}_t) \quad (15)$$

and an observation equation

$$z_t = b(\phi_t, \boldsymbol{\nu}_t) \quad (16)$$

where the states ϕ_t form a hidden/latent Markov process with disturbance $\boldsymbol{\nu}_t$, and the observed time series data z_t depends only on ϕ_t and disturbance $\boldsymbol{\nu}_t$. Here $a(\cdot)$ and $b(\cdot)$ are possibly non-linear functions, and the states ϕ_t and observations z_t can be multi-dimensional.

It is clear that the models in Table 1 specify the observation equation of different state-space models that can be considered. For example, for the multi-period LC model (Renshaw & Haberman, 2003), the observed data is $z_{x,t} = \ln(\hat{m}_{x,t})$ for different age x and the latent states are the period effects $\phi_t = (\kappa_t^{(1)}, \dots, \kappa_t^{(k)})$. We also note that multi-population (i.e. multi-curve) structures can be incorporated in the following state-space models in a number of different ways and the approaches we will develop for estimation will accommodate such settings. In the following sub-sections we will discuss a few different classes of mortality models that are difficult to deal with in the approaches mentioned in section 2, but can be handled straightforwardly in state-space framework.

3.1. LC model with heteroscedasticity (LC-H model)

We present here a state-space formulation of the LC model with heteroscedasticity structure. In this context, the heteroscedasticity refers to a relaxation of the constant single degree of freedom diagonal covariance assumption typically made on the observation vector for each year t across the panel of age group stratifications $\mathbf{y}_t = (y_{x_1,t}, y_{x_2,t}, \dots, y_{x_p,t})$. Within this state-space model structure we propose an alternative identification constraint which is tailored for the estimation under the state-space approach.

The LC model with heteroscedasticity structure can be written in state-space form by combining the processes $\mathbf{y}_t = (y_{x_1,t}, \dots, y_{x_p,t})$ and κ_t into one dynamical system

$$\mathbf{y}_t = \boldsymbol{\alpha} + \boldsymbol{\beta}\kappa_t + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \stackrel{iid}{\sim} N(0, \boldsymbol{\Sigma}) \quad (17a)$$

$$\kappa_t = \kappa_{t-1} + \theta + \omega_t, \quad \omega_t \stackrel{iid}{\sim} N(0, \sigma_\omega^2) \quad (17b)$$

where $\alpha = \alpha_{x_1: x_p}$, $\beta = \beta_{x_1: x_p}$ and κ_t is the latent state of the resulting linear Gaussian state-space model. Here, Σ is a p by p diagonal matrix with $\sigma_{\varepsilon, x_1: x_p}^2$ on the diagonal. We refer to this model as LC-H model, and the special case with $\sigma_{\varepsilon, x_i}^2 = \sigma_\varepsilon^2$, $i \in \{1, \dots, p\}$, as LC model.

Instead of the identification constraint (6), we suggest an alternative constraint which is simpler and more readily applicable to Monte Carlo-based procedures such as MCMC and SMC. Our formulation of the identification constraint is given by setting

$$\alpha_{x_1} = \text{constant}, \quad \beta_{x_1} = \text{constant} \tag{18}$$

Such a choice is a valid identification constraint since if one of the elements of each α and β are known (here we have arbitrarily chosen α_{x_1} and β_{x_1} ; in general one can choose α_{x_i} and β_{x_j} , where $i, j \in \{1, \dots, p\}$), then a non-trivial linear transformation in (5) is not allowed anymore; that is, we must have $c = 0$ and $d = 1$. Note that implementing the proposed constraint is straightforward in both maximum likelihood and Bayesian setting compared to the constraint (6).

Remark 3.1 As discussed in section 2.4, one needs to be careful about the way how the constraint (6) is implemented in a MCMC-based sampling setting. An advantage of the proposed constraint (18) here is that it is error proof in this setup. Note that regardless of which constraint one chooses, a natural interpretation of α and β is that α represents the general pattern of the age-specific log death rates while β captures the sensitivity of log death rates with respect to a change of the period effect. Therefore, one can naturally set a value for α_{x_1} as the average of the log death rate at age x_1 . On the other hand, there is no restriction on what value β_{x_1} can take. However, we recommend it to be in the range of $[0.01, 1]$, as too small or too large value may lead to numerical issues such as overflow/underflow problems. We also emphasise here that the resulting fit and forecasting ability of the model will not be affected by the choice of constraint since an identification constraint serves only to identify the model in a unique way.

3.2. Two-factor LC model with age-based heteroscedasticity (LC2-H model)

A natural extension of the LC-H model is to include a second stochastic factor for the cohort effect. We denote this model by LC2-H model. The cohort effect (Renshaw & Haberman, 2006) can be modelled under the state-space framework as follows:

$$\begin{bmatrix} y_{x_1, t} \\ y_{x_2, t} \\ \vdots \\ y_{x_p, t} \end{bmatrix} = \begin{bmatrix} \alpha_{x_1} \\ \alpha_{x_2} \\ \vdots \\ \alpha_{x_p} \end{bmatrix} + \begin{bmatrix} \beta_{x_1}^{(1)} & \beta_{x_1}^{(2)} & 0 & \cdots & 0 \\ \beta_{x_2}^{(1)} & 0 & \beta_{x_2}^{(2)} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \beta_{x_p}^{(1)} & 0 & 0 & \cdots & \beta_{x_p}^{(2)} \end{bmatrix} \begin{bmatrix} \kappa_t \\ \zeta_t^{x_1} \\ \zeta_t^{x_2} \\ \vdots \\ \zeta_t^{x_p} \end{bmatrix} + \begin{bmatrix} \varepsilon_{x_1, t} \\ \varepsilon_{x_2, t} \\ \vdots \\ \varepsilon_{x_p, t} \end{bmatrix} \tag{19}$$

where $\zeta_t^x := \zeta_{t-x}$. The state equation can be expressed as

$$\begin{bmatrix} \kappa_t \\ \zeta_t^{x_1} \\ \zeta_t^{x_2} \\ \vdots \\ \zeta_t^{x_{p-1}} \\ \zeta_t^{x_p} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & \vartheta & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \end{bmatrix} \begin{bmatrix} \kappa_{t-1} \\ \zeta_{t-1}^{x_1} \\ \zeta_{t-1}^{x_2} \\ \vdots \\ \zeta_{t-1}^{x_{p-1}} \\ \zeta_{t-1}^{x_p} \end{bmatrix} + \begin{bmatrix} \theta \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} \omega_t^\kappa \\ \omega_t^{\zeta_1} \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix} \tag{20}$$

Here, we assume κ_t is a random walk with drift process and an autoregressive model of order 1 (AR(1)) process is assumed for the cohort effect, that is $\zeta_t^{x_1} = \vartheta \zeta_{t-1}^{x_1} + \omega_t^{\zeta_1}$, where $|\vartheta| < 1$. Note that, from (20), we have $\zeta_t^{x_i} = \zeta_{t-1}^{x_i}$ for $i = 2, \dots, p$, which is the defining property of the cohort effect and consequently we are only required to model the dynamics of $\zeta_t^{x_1}$. We can write the model (19)–(20) in the following form:

$$y_t = \alpha + B[\kappa_t, \zeta_t]^\top + \varepsilon_t, \quad \varepsilon_t \stackrel{iid}{\sim} N(0, \Sigma) \tag{21a}$$

$$\kappa_t = \kappa_{t-1} + \theta + \omega_t^\kappa, \quad \omega_t^\kappa \stackrel{iid}{\sim} N(0, \sigma_{\omega^\kappa}^2) \tag{21b}$$

$$\zeta_t = C\zeta_{t-1} + D\omega_t^\zeta, \quad \omega_t^{\zeta_1} \stackrel{iid}{\sim} N(0, \sigma_{\omega^\zeta}^2) \tag{21c}$$

where B is the p by $p + 1$ matrix in (19), C the corresponding p by p sub-matrix in (20) and D a zero p by p matrix except for the (1, 1) element with value 1.

For further details on the formulation, estimation and forecasting of cohort models in state-space framework, see Fung *et al.* (2017).

3.3. Three-factor LC model with dynamic time-based heteroscedasticity (LC3-H2 model)

We can further extend the LC2-H model by adding a third factor for the dynamic of volatility in the observation vector over time. Specifically, a time dependent stochastic factor $\sqrt{\gamma_t^y}$, which is independent to age (or age group), is incorporated into the observation noise term as follows:

$$y_t = \alpha + B[\kappa_t, \zeta_t]^\top + \sqrt{\gamma_t^y} \varepsilon_t, \quad \varepsilon_t \stackrel{iid}{\sim} N(0, \Sigma) \tag{22a}$$

$$\kappa_t = \kappa_{t-1} + \theta + \omega_t^\kappa, \quad \omega_t^\kappa \stackrel{iid}{\sim} N(0, \sigma_{\omega^\kappa}^2) \tag{22b}$$

$$\zeta_t = C\zeta_{t-1} + D\omega_t^\zeta, \quad \omega_t^{\zeta_1} \stackrel{iid}{\sim} N(0, \sigma_{\omega^\zeta}^2) \tag{22c}$$

$$\gamma_t^y = a(b - \gamma_{t-1}^y) + \gamma_{t-1}^y + \sigma \sqrt{\gamma_{t-1}^y} \varepsilon_t^y, \quad \varepsilon_t^y \stackrel{iid}{\sim} N(0, 1) \tag{22d}$$

where γ_t^y is a process obtained via an Euler discretisation of a square Bessel process corresponding to the Cox–Ingersoll–Ross process given by

$$d\gamma_t^y = a(b - \gamma_t^y)dt + \sigma \sqrt{\gamma_t^y} dW_t$$

where $2ab \geq \sigma^2$ to ensure γ_t^y is strictly positive. Such a dynamic volatility factor can be used to explain time-varying periods of heightened observation variance, which potentially occur in some populations over time. These may be attributed to disease, war, famine, environmental factors or shocks as well as changes in migration and immigration patterns that could influence the volatility of the observed death counts in different age groups. Note that this is one possible way to introduce stochastic volatility to mortality models where the stochastic volatility factor enters via the

observation noise term. Another possible choice is to consider stochastic volatility through the state equation described in section 3.4 which will be the focus of this paper.

3.4. Multi-factor model with stochastic volatility in the latent process (LCSV and LCSV-H models)

A common assumption in mortality modelling is that the period effect is derived from a discretisation of a random walk with drift process. Such a process may be sufficient for modelling simple dynamics, but can be insufficient if time-varying periods of volatility are present in the time series. In fact, much of the literature focusses mainly on capturing the trend of the period effect κ_t for the past several decades where mortality time series for many countries are reasonably smooth.

Here, we extend the LC framework to incorporate stochastic volatility in the latent process. As a result, the impact of epidemics, natural disasters, medical breakthrough or wars on the evolution of mortality can be taken into account. This will produce different structural effects on the calibration and importantly on the forecasting when compared to the previously developed model of LC3-H2. We refer (23a)–(23c) as LC stochastic volatility model which we denote as LCSV model:

$$y_t = \alpha + \beta\kappa_t + \epsilon_t, \quad \epsilon_t \stackrel{iid}{\sim} N(0, \sigma_\epsilon^2 \mathbf{1}_p) \tag{23a}$$

$$\kappa_t = \kappa_{t-1} + \theta + \omega_t, \quad \omega_t | \gamma_t \sim N(0, \exp\{\gamma_t\}) \tag{23b}$$

$$\gamma_t = \lambda_1 \gamma_{t-1} + \lambda_2 + \eta_t, \quad \eta_t \stackrel{iid}{\sim} N(0, \sigma_\gamma^2) \tag{23c}$$

The log-volatility process γ_t is introduced in the state equation for κ_t via the error term ω_t . The process γ_t is an AR(1) with $|\lambda_1| < 1$ and the mean reverting level is given by $\lambda_2/(1 - \lambda_1)$. A heteroscedasticity structure can be introduced in (23a) and will be referred to as LCSV-H model. Specifically, the LCSV-H model is given by

$$y_t = \alpha + \beta\kappa_t + \epsilon_t, \quad \epsilon_t \stackrel{iid}{\sim} N(0, \Sigma) \tag{24a}$$

$$\kappa_t = \kappa_{t-1} + \theta + \omega_t, \quad \omega_t | \gamma_t \sim N(0, \exp\{\gamma_t\}) \tag{24b}$$

$$\gamma_t = \lambda_1 \gamma_{t-1} + \lambda_2 + \eta_t, \quad \eta_t \stackrel{iid}{\sim} N(0, \sigma_\gamma^2) \tag{24c}$$

Cohort effect can also be incorporated into the LCSV model as follows:

$$y_t = \alpha + B[\kappa_t, \zeta_t]^\top + \epsilon_t, \quad \epsilon_t \stackrel{iid}{\sim} N(0, \sigma_\epsilon^2 \mathbf{1}_p) \tag{25a}$$

$$\kappa_t = \kappa_{t-1} + \theta + \omega_t, \quad \omega_t | \gamma_t \sim N(0, \exp\{\gamma_t\}) \tag{25b}$$

$$\gamma_t = \lambda_1 \gamma_{t-1} + \lambda_2 + \eta_t, \quad \eta_t \stackrel{iid}{\sim} N(0, \sigma_\gamma^2) \tag{25c}$$

$$\zeta_t = C\zeta_{t-1} + D\omega_t^\zeta, \quad \omega_t^{\zeta 1} \stackrel{iid}{\sim} N(0, \sigma_{\omega^\zeta}^2) \tag{25d}$$

Compared to the LC3-H2 model where stochastic volatility is included in the observation noise term, introducing stochastic volatility in the latent period process has the advantage, in terms of simplicity and ease of interpretation, that the variability of mortality data in the time dimension is captured purely by the latent process.

4. Frequentist State-Space Inference

Given these different state-space model structures, the next task to consider is the inference for the joint single-stage state and parameter estimation. In this section, we consider full likelihood-based joint inference procedures based on filtering and gradient estimation. To achieve this we must describe both filtering in linear Gaussian and non-linear/non-Gaussian filtering via SMC method (particle filters) and their application to gradient-based estimation in the marginal likelihood, having integrated out the latent state processes. We will do this in a general way and then present particular examples of relevance to this paper.

Under the classical maximum likelihood approach, parameters are estimated by maximising a model's log-likelihood function. In the case of state-space models in the form of (15)–(16), the likelihood is in two forms: the complete data likelihood, assuming ϕ_0 fixed, is given by

$$p_{\psi}(\phi_{1:T}, z_{1:T}) = \prod_{t=1}^T p_{\psi}(z_t | \phi_t) p_{\psi}(\phi_t | \phi_{t-1}) \tag{26}$$

and the marginal likelihood, typically used for the static model-based inference, is given by

$$p_{\psi}(z_{1:T}) = \int \prod_{t=1}^T p_{\psi}(z_t | \phi_t) p_{\psi}(\phi_t | \phi_{t-1}) d\phi_t \tag{27}$$

where ψ denotes the d -dimensional parameter vector of the model. Two challenges now arise. The first is that typically the integral in (27) cannot be evaluated in closed form, except for linear Gaussian state-space model systems. The second issue is that the gradient equations for such a state-space model marginal likelihood, even if they can be calculated in closed form, requires a non-linear multiple equation solver. For this reason it is common to adopt a solution based on a recursive estimation using gradient and Hessian information from the marginal likelihood. Under the gradient-based approach, the optimal parameter vector can be found by iterations, where the $(m + 1)$ th estimate is obtained by

$$\psi^{(m+1)} = \psi^{(m)} - \left[\nabla_{\psi}^2 \ell(\psi^{(m)}) \right]^{-1} \nabla_{\psi} \ell(\psi^{(m)}) \tag{28}$$

where $\ell(\psi)$, $\nabla_{\psi} \ell(\psi)$ and $-\nabla_{\psi}^2 \ell(\psi)$ denote the log-likelihood function, the gradient (or score) vector and the Hessian information matrix of the log-likelihood function, respectively, defined with respect to grad and Laplacian differential operators given by

$$\begin{aligned} [\nabla_{\psi}]_i &:= \frac{\partial}{\partial \psi_i}, \quad \forall i \in \{1, \dots, n\} \\ [\nabla_{\psi}^2]_{ij} &:= \frac{\partial^2}{\partial \psi_i \partial \psi_j}, \quad \forall i, j \in \{1, \dots, n\} \end{aligned} \tag{29}$$

The iterating scheme will stop once certain criterion is met, for example, when the magnitude of the score vector is small enough. This will be illustrated using the LC-H model as an example in section 4.2.

The results developed are based on the marginal likelihood of the state-space model, with generic static model parameters ψ for observations $z_{1:T} = z_1, \dots, z_T$ having integrated out latent states

ϕ_1, \dots, ϕ_T , denoted by $p_\psi(z_{1:T})$. We are then interested in forming recursive filtering to integrate the complete data likelihood to find the marginalised likelihood and then working with recursive gradient-based estimation to update static model parameters in Newton–Descent type algorithm, or for linear Gaussian systems a recursive least squares-based approach.

As observed in Poyiadjis *et al.* (2005) and Poyiadjis *et al.* (2011), it is useful to consider two classes of identities for the gradient and Hessian of the marginalised likelihood, given by the Fisher’s identity and the Louis’ identity, respectively, according to

$$\begin{aligned} \nabla_\psi p_\psi(z_{1:T}) &= \int \nabla_\psi \ln p_\psi(\phi_{1:T}, z_{1:T}) p_\psi(\phi_{1:T} | z_{1:T}) d\phi_{1:T} \\ -\nabla_\psi^2 p_\psi(z_{1:T}) &= \nabla_\psi \ln p_\psi(z_{1:T}) \nabla_\psi \ln p_\psi(z_{1:T})^\top - \frac{\nabla_\psi^2 \ln p_\psi(z_{1:T})}{p_\psi(z_{1:T})} \end{aligned} \tag{30}$$

where

$$\begin{aligned} \frac{\nabla_\psi^2 p_\psi(z_{1:T})}{p_\psi(z_{1:T})} &= \int \nabla_\psi \ln p_\psi(\phi_{1:T}, z_{1:T}) \nabla_\psi \ln p_\psi(\phi_{1:T}, z_{1:T})^\top p_\psi(\phi_{1:T} | z_{1:T}) d\phi_{1:T} \\ &\quad + \int \nabla_\psi^2 \ln p_\psi(\phi_{1:T}, z_{1:T}) p_\psi(\phi_{1:T} | z_{1:T}) d\phi_{1:T} \end{aligned} \tag{31}$$

An important point about these recursions is that the integrals for the gradient vector and Hessian matrix are expressed in terms of the path-space distribution $p_\psi(\phi_{1:T} | z_{1:T})$. In the case of the linear Gaussian dynamics this distribution can be obtained based on variations of the Kalman filter recursion, however, when the state-space model is non-linear or non-Gaussian this distribution must be estimated via Monte Carlo methods. The most efficient of these methods for state-space modelling purposes is known as the class of SMC methods (particle filters). In this case it will be more accurate from the perspective of the variance of the estimated gradient and Hessian matrices to utilise the filter distribution-based estimators in a recursive fashion based on the local estimates of the distributions $p_\psi(\phi_t | z_{1:t})$, for each $t \in \{1, \dots, T\}$ rather than the path-space estimator which is based on distribution $p_\psi(\phi_{1:T} | z_{1:T})$ at the final time T . The result explaining the difference in estimation precision for the gradient and Hessian, from the perspective of variance of the solution on the path-space distribution versus filter distributions is provided in theorem 1 of Poyiadjis *et al.* (2011). This motivates the need to work with the filter recursions.

To achieve this one can replace in the Fisher and Louis’ identities the path-space quantities $p_\psi(\phi_{1:T}, z_{1:T})$ and $p_\psi(\phi_{1:T} | z_{1:T})$ by the filter quantities given by $p_\psi(\phi_t, z_{1:t})$ and $p_\psi(\phi_t | z_{1:t})$. After this substitution, one may utilise the following recursive formulations to evaluate the gradient and Hessian, see Poyiadjis *et al.* (2005) and Poyiadjis *et al.* (2011). In this case the Fisher identity is recursively given by

$$\begin{aligned} \nabla_\psi \ln p_\psi(z_{1:t}) &= \int \nabla_\psi \ln p_\psi(\phi_t, z_{1:t}) p_\psi(\phi_t | z_{1:t}) d\phi_t \\ \nabla_\psi \ln p_\psi(\phi_t, z_{1:t}) &= \frac{p_\psi(z_{1:t-1}) p_\psi(z_t | \phi_t)}{p_\psi(\phi_t, z_{1:t})} \int p_\psi(\phi_t | \phi_{t-1}) p_\psi(\phi_{t-1} | z_{1:t-1}) \\ &\quad \times [\nabla_\psi \ln p_\psi(z_t | \phi_t) + \nabla_\psi \ln p_\psi(\phi_t | \phi_{t-1}) + \nabla_\psi \ln p_\psi(\phi_{t-1}, z_{1:t-1})] d\phi_{t-1} \\ p_\psi(\phi_t, z_{1:t}) &= p_\psi(z_{1:t-1}) p_\psi(z_t | \phi_t) \int p_\psi(\phi_t | \phi_{t-1}) p_\psi(\phi_{t-1} | z_{1:t-1}) d\phi_{t-1} \end{aligned}$$

The recursive form of Luis' identity is given by

$$\begin{aligned} \frac{\nabla_{\psi}^2 p_{\psi}(z_{1:t})}{p_{\psi}(\phi_t, z_{1:t})} &= \int \nabla_{\psi} \ln p_{\psi}(\phi_t, z_{1:t}) \nabla_{\psi} \ln p_{\psi}(\phi_t, z_{1:t})^{\top} p_{\psi}(\phi_t | z_{1:t}) d\phi_t \\ &\quad + \int \nabla_{\psi}^2 \ln p_{\psi}(\phi_t, z_{1:t}) p_{\psi}(\phi_t | z_{1:t}) d\phi_t \\ \nabla_{\psi}^2 \ln p_{\psi}(\phi_t, z_{1:t}) &= \frac{\nabla_{\psi}^2 p_{\psi}(\phi_t, z_{1:t})}{p_{\psi}(\phi_t, z_{1:t})} - \nabla_{\psi} \ln p_{\psi}(\phi_t, z_{1:t}) \nabla_{\psi} \ln p_{\psi}(\phi_t, z_{1:t})^{\top} \\ \nabla_{\psi}^2 p_{\psi}(\phi_t, z_{1:t}) &= p_{\psi}(z_{1:t-1}) p_{\psi}(z_t | \phi_t) \int p_{\psi}(\phi_t | \phi_{t-1}) p_{\psi}(\phi_{t-1} | z_{1:t-1}) \\ &\quad \times \left\{ \left[\nabla_{\psi} \ln p_{\psi}(z_t | \phi_t) + \nabla_{\psi} \ln p_{\psi}(\phi_t | \phi_{t-1}) + \nabla_{\psi} \ln p_{\psi}(\phi_{t-1}, z_{1:t-1}) \right] \right. \\ &\quad \times \left. \left[\nabla_{\psi} \ln p_{\psi}(z_t | \phi_t) + \nabla_{\psi} \ln p_{\psi}(\phi_t | \phi_{t-1}) + \nabla_{\psi} \ln p_{\psi}(\phi_{t-1}, z_{1:t-1}) \right]^{\top} \right. \\ &\quad \left. + \left[\nabla_{\psi}^2 \ln p_{\psi}(z_t | \phi_t) + \nabla_{\psi}^2 \ln p_{\psi}(\phi_t | \phi_{t-1}) + \nabla_{\psi}^2 \ln p_{\psi}(\phi_{t-1}, z_{1:t-1}) \right] \right\} d\phi_{t-1} \end{aligned}$$

In general, the solution to these recursions can be achieved via SMC as detailed in Poyiadjis *et al.* (2005) and Poyiadjis *et al.* (2011).

In the following sections, we will illustrate the use of these recursive identities for the special case of the LC-H model where the state-space takes a linear Gaussian form. In this case the integrals and recursive evaluation of the gradient and Hessian can be written in closed form. To proceed we first introduce the optimal filter recursion, with respect to minimisation of mean-squared error, in the case of linear Gaussian state-space models, known as the Kalman filter (Kalman, 1960; Harvey, 1989).

4.1. Closed form filter recursions for LC-H model

The aim of filtering is to obtain the distribution of the latest state given observations. For a general state-space model, (15)–(16), the filtering density $\pi(\phi_t | z_{1:t})$ at time t can be calculated sequentially by first assuming the filtering density $\pi(\phi_{t-1} | z_{1:t-1})$ at time $t-1$ is known. Then the one-step ahead predictive density for the state is given by

$$\pi(\phi_t | z_{1:t-1}) = \int \pi(\phi_t | \phi_{t-1}) \pi(\phi_{t-1} | z_{1:t-1}) d\phi_{t-1} \tag{32}$$

From Bayes' formula and the structure of the conditional dependency of the state-space model, one can obtain the filtering density as

$$\pi(\phi_t | z_{1:t}) = \frac{\pi(\phi_t | z_{1:t-1}) \pi(z_t | \phi_t)}{\int \pi(\phi_t | z_{1:t-1}) \pi(z_t | \phi_t) d\phi_t} \tag{33}$$

For non-linear and non-Gaussian state-space models, numerical techniques such as SMC methods are required to estimate the filtering density, see Doucet *et al.* (2001), Doucet *et al.* (2000) and Liu (2008).

In the case of the LC-H model, since it is a linear and Gaussian state-space model, the filtering distribution can be obtained analytically via Kalman filtering. In particular, we can find the conditional distributions of the key quantities in the filtering recursions are all Gaussian distributions as follows:

$$\kappa_{t-1} | \mathbf{y}_{1:t-1} \sim N(m_{t-1}, C_{t-1}) \tag{34a}$$

$$\kappa_t | \mathbf{y}_{1:t-1} \sim N(a_t, R_t) \tag{34b}$$

$$\mathbf{y}_t | \mathbf{y}_{1:t-1} \sim N(\mathbf{f}_t, \mathbf{Q}_t) \tag{34c}$$

$$\kappa_t | \mathbf{y}_{1:t} \sim N(m_t, C_t) \tag{34d}$$

where the recursive nature of these distributions arises from the recursions of the sufficient statistics:

$$a_t = m_{t-1} + \theta, \quad R_t = C_{t-1} + \sigma_\omega^2 \tag{35a}$$

$$\mathbf{f}_t = \boldsymbol{\alpha} + \boldsymbol{\beta}a_t, \quad \mathbf{Q}_t = \boldsymbol{\beta}\boldsymbol{\beta}^\top R_t + \boldsymbol{\Sigma} \tag{35b}$$

$$m_t = a_t + R_t \boldsymbol{\beta}^\top \mathbf{Q}_t^{-1} (\mathbf{y}_t - \mathbf{f}_t), \quad C_t = R_t - R_t \boldsymbol{\beta}^\top \mathbf{Q}_t^{-1} \boldsymbol{\beta} R_t \tag{35c}$$

That is, given the filtering distribution at $t - 1$, (34a), the filtering distribution at t is given by (34d) using (35a)–(35c).

4.2. Closed form gradient-based estimation via score and Hessian recursions for LC-H model

For the LC-H model, the log-likelihood function $\ell(\boldsymbol{\psi}) := \ln \pi(\mathbf{y}_{1:T} | \boldsymbol{\psi})$ is given by

$$\ell(\boldsymbol{\psi}) = \ln \left(\prod_{t=1}^T \pi(\mathbf{y}_t | \mathbf{y}_{1:t-1}, \boldsymbol{\psi}) \right) = -\frac{pT}{2} \ln 2\pi - \frac{1}{2} \sum_{t=1}^T (\ln |\mathbf{Q}_t| + \boldsymbol{\nu}_t^\top \mathbf{Q}_t^{-1} \boldsymbol{\nu}_t) \tag{36}$$

where $\boldsymbol{\nu}_t := \mathbf{y}_t - \mathbf{f}_t$, and $\boldsymbol{\psi} = (\alpha_{x_2:x_p}, \boldsymbol{\beta}_{x_2:x_p}, \theta, \sigma_{\varepsilon, x_1:x_p}^2, \sigma_\omega^2)$ is an n -dimensional parameter vector. The log-likelihood function (54) can be derived directly from (34c).

It can be shown that (Harvey, 1989) the elements of the score vector and the information matrix are given in closed form for the LC-H model according to the expressions:

$$\frac{\partial \ell}{\partial \psi_i} = \frac{1}{2} \sum_{t=1}^T \left\{ \text{tr} \left[\left(\mathbf{Q}_t^{-1} \frac{\partial \mathbf{Q}_t}{\partial \psi_i} \right) (\mathbf{1}_p - \mathbf{Q}_t^{-1} \boldsymbol{\nu}_t \boldsymbol{\nu}_t^\top) \right] + 2 \frac{\partial \boldsymbol{\nu}_t^\top}{\partial \psi_i} \mathbf{Q}_t^{-1} \boldsymbol{\nu}_t \right\}, \quad i = 1, \dots, n \tag{37}$$

where $\text{tr}[\cdot]$ denotes the trace operator and

$$-E \left[\frac{\partial^2 \ell}{\partial \psi_i \partial \psi_j} \right] = \frac{1}{2} \sum_{t=1}^T \left[\text{tr} \left(\mathbf{Q}_t^{-1} \frac{\partial \mathbf{Q}_t}{\partial \psi_i} \mathbf{Q}_t^{-1} \frac{\partial \mathbf{Q}_t}{\partial \psi_j} \right) \right] + E \left[\sum_{t=1}^T \frac{\partial \boldsymbol{\nu}_t^\top}{\partial \psi_i} \mathbf{Q}_t^{-1} \frac{\partial \boldsymbol{\nu}_t}{\partial \psi_j} \right], \quad i, j = 1, \dots, n \tag{38}$$

and the expectation operator $E[\cdot]$ on the second term in (38) can be dropped (since the expressions are asymptotically equivalent). In order to evaluate the score vector and the information matrix, we need

$$\frac{\partial \boldsymbol{\nu}_t}{\partial \psi_i} = -\frac{\partial \boldsymbol{\alpha}}{\partial \psi_i} - \frac{\partial \boldsymbol{\beta}}{\partial \psi_i} a_t - \boldsymbol{\beta} \frac{\partial a_t}{\partial \psi_i} \tag{39}$$

and

$$\frac{\partial \mathbf{Q}_t}{\partial \psi_i} = \frac{\partial \boldsymbol{\beta}}{\partial \psi_i} R_t \boldsymbol{\beta}^\top + \boldsymbol{\beta} \frac{\partial R_t}{\partial \psi_i} \boldsymbol{\beta}^\top + \boldsymbol{\beta} R_t \frac{\partial \boldsymbol{\beta}^\top}{\partial \psi_i} + \frac{\partial \boldsymbol{\Sigma}}{\partial \psi_i} \tag{40}$$

The expressions (39) and (40) require, for $t = 1, \dots, T$ and $i = 1, \dots, n$

$$\frac{\partial a_t}{\partial \psi_i} = \frac{\partial m_{t-1}}{\partial \psi_i} + \frac{\partial \theta}{\partial \psi_i} \quad \text{and} \quad \frac{\partial R_t}{\partial \psi_i} = \frac{\partial C_{t-1}}{\partial \psi_i} + \frac{\partial \sigma_\omega^2}{\partial \psi_i} \tag{41}$$

The expressions in (41) in turn require, for $t = 1, \dots, T-1$ and $i = 1, \dots, n$

$$\frac{\partial m_t}{\partial \psi_i} = \frac{\partial a_t}{\partial \psi_i} + \frac{\partial R_t}{\partial \psi_i} \beta^\top Q_t^{-1} \nu_t + R_t \frac{\partial \beta}{\partial \psi_i} Q_t^{-1} \nu_t - R_t \beta^\top Q_t^{-1} \frac{\partial Q_t}{\partial \psi_i} Q_t^{-1} \nu_t + R_t \beta^\top Q_t^{-1} \frac{\partial \nu_t}{\partial \psi_i} \tag{42}$$

and

$$\begin{aligned} \frac{\partial C_t}{\partial \psi_i} = & \frac{\partial R_t}{\partial \psi_i} - \frac{\partial R_t}{\partial \psi_i} \beta^\top Q_t^{-1} \beta R_t - R_t \frac{\partial \beta}{\partial \psi_i} Q_t^{-1} \beta R_t \\ & + R_t \beta^\top Q_t^{-1} \frac{\partial Q_t}{\partial \psi_i} Q_t^{-1} \beta R_t - R_t \beta^\top Q_t^{-1} \frac{\partial \beta}{\partial \psi_i} R_t - R_t \beta^\top Q_t^{-1} \beta \frac{\partial R_t}{\partial \psi_i} \end{aligned} \tag{43}$$

Note that $\frac{\partial m_0}{\partial \psi_i} = \frac{\partial C_0}{\partial \psi_i} = 0$ for $i = 1, \dots, n$ and the required differentiation matrices $\frac{\partial \alpha}{\partial \psi_i}, \frac{\partial \beta}{\partial \psi_i}, \frac{\partial \Sigma}{\partial \psi_i}, \frac{\partial \theta}{\partial \psi_i}$ and $\frac{\partial \sigma_\omega^2}{\partial \psi_i}$ are displayed in Appendix A. The gradient-based estimation for the LC-H model is described in Algorithm 1.

Algorithm 1 Gradient-based approach for estimating parameters ψ

- 1: Initialise $\psi = \psi^{(0)}$; specify m_0 and C_0 .
 - 2: **while** stopping criterion is not met **do**
 - 3: Count the number of iteration performed as m ;
 - 4: Run Kalman filter using $\psi^{(m)}$; obtain (39) and (40) for $t = 1, \dots, T$;
 - 5: Evaluate the score vector $\nabla_\psi \ell(\psi^{(m)})$ using (37);
 - 6: Evaluate the information matrix $-E[\nabla_\psi^2 \ell(\psi^{(m)})]$ given by (38);
 - 7: Set $\psi^{(m+1)} = \psi^{(m)} + [E[\nabla_\psi^2 \ell(\psi^{(m)})]]^{-1} \nabla_\psi \ell(\psi^{(m)})$.
 - 8: **end while**
-

5. Bayesian State-Space Inference

In contrast to the classical maximum likelihood approach where parameters are deterministic but unknown, in a Bayesian view the parameters are treated as random variables. In this way the Bayesian paradigm can take parameter uncertainty into account and incorporate in a consistent manner a priori beliefs on the important model parameters, as encoded through the prior.

In this section, we aim to develop modern approaches to Bayesian inference for state-space modelling that do not rely on potentially inefficient sampling approaches based on Gibbs or Metropolis-within-Gibbs for the latent state process. Instead we will introduce to stochastic mortality modelling state-space models of two classes of Bayesian inference:

- Linear Gaussian stochastic mortality models: a Rao-Blackwellised Gibbs sampler approach which is based on a combination of Metropolis-within-Gibbs and Gibbs sampling steps for the static model parameters, combined with a forward-backward Kalman filter recursion for the

state process. We assume the proposed constraints (18) throughout, avoiding the constraint issue when performing MCMC as discussed in section 2.4.

- Non-linear/non-Gaussian stochastic mortality models: in the case of non-linear and or non-Gaussian state-space model dynamics such as the stochastic volatility models of LC3-H2 and the LCSV models, the sampler we develop is based on novel developments of the particle Metropolis–Hastings samplers of Andrieu *et al.* (2010) adapted to the stochastic mortality models. In particular, we consider a combination of Rao-Blackwellised Kalman filter and particle filter for the latent state process full posterior conditionals, combined with a combination of Metropolis-within-Gibbs and Gibbs sampling steps for the static model parameters.

In general under all the Bayesian approaches we consider, we aim to obtain the joint posterior density:

$$\pi(\kappa_{0:T}, \boldsymbol{\psi} | \mathbf{y}_{1:T}) \tag{44}$$

of the states is $\kappa_{0:T}$ as well as the parameters, $\boldsymbol{\psi}$, given the observations $\mathbf{y}_{1:T}$. We begin with the first case of the linear Gaussian state-space stochastic mortality models and we use the LC-H model as an example where the parameter vector is $\boldsymbol{\psi} := (\alpha_{x_2:x_p}, \beta_{x_2:x_p}, \theta, \sigma_{\varepsilon, x_1:x_p}^2, \sigma_{\omega}^2)$ as we use the constraint proposed in (18).

5.1. Linear Gaussian state-space inference

We develop an efficient approach involving a combined Gibbs sampling conjugate model sampler for the marginal target distributions of the static model parameters along with a forward-backward Kalman filter sampler for the latent process $\kappa_{1:T}$. A sample of the targeted density is obtained via Gibbs sampling where M is the number of MCMC iterations (Algorithm 2).

Algorithm 2 Rao-Blackwellised forward-backward Kalman Filter and Gibbs sampling for $\pi(\kappa_{0:T}, \boldsymbol{\psi} | \mathbf{y}_{1:T})$

- 1: Initialise: $\boldsymbol{\psi} = \boldsymbol{\psi}^{(0)}$.
 - 2: **for** $i = 1, \dots, M$ **do**
 - 3: Sample $\kappa_{0:T}^{(i)}$ from $\pi(\kappa_{0:T} | \boldsymbol{\psi}^{(i-1)}, \mathbf{y}_{1:T})$ via FFBS (Section 5.1.1).
 - 4: **for** $h = 1, \dots, n$ **do**
 - 5: Sample $\psi_h^{(i)}$ from $\pi(\psi_h | \kappa_{0:T}^{(i)}, \boldsymbol{\psi}_{-h}^{(i)}, \mathbf{y}_{1:T})$,
 - 6: where $\boldsymbol{\psi}_{-h}^{(i)} = (\psi_1^{(i)}, \dots, \psi_{h-1}^{(i)}, \psi_{h+1}^{(i)}, \dots, \psi_n^{(i)})$.
 - 7: **end for**
 - 8: **end for**
-

The general block Gibbs sampling algorithm steps require to sample from the full conditional densities $\pi(\kappa_{0:T} | \boldsymbol{\psi}, \mathbf{y}_{1:T})$ and $\pi(\boldsymbol{\psi} | \kappa_{0:T}, \mathbf{y}_{1:T})$, which are shown in the following.

5.1.1. Sampling from the full conditional density $\pi(\kappa_{0:T} | \boldsymbol{\psi}, \mathbf{y}_{1:T})$

Samples of the full conditional density $\pi(\kappa_{0:T} | \boldsymbol{\psi}, \mathbf{y}_{1:T})$ can be obtained via the so-called forward-filtering-backward sampling (FFBS) procedure (Carter & Kohn, 1994). This sampling approach is also utilised in Pedroza (2006). We can write

$$\pi(\kappa_{0:T} | \boldsymbol{\psi}, \mathbf{y}_{1:T}) = \prod_{t=0}^T \pi(\kappa_t | \kappa_{t+1:T}, \boldsymbol{\psi}, \mathbf{y}_{1:T}) = \prod_{t=0}^T \pi(\kappa_t | \kappa_{t+1}, \boldsymbol{\psi}, \mathbf{y}_{1:t}) \tag{45}$$

where the last term in the product, $\pi(\kappa_T | \boldsymbol{\psi}, \mathbf{y}_{1:T})$, is distributed as $N(m_T, C_T)$ which is obtained from the last iteration of the Kalman filtering procedure.

Once we draw a sample κ_T from $N(m_T, C_T)$, then (45) suggests that we can draw recursively and backwardly κ_t from $\pi(\kappa_t | \kappa_{t+1}, \boldsymbol{\psi}, \mathbf{y}_{1:t})$, where $t = T-1, T-2, \dots, 1, 0$. Moreover, we have

$$\kappa_t | \kappa_{t+1}, \boldsymbol{\psi}, \mathbf{y}_{1:t} \sim N(h_t, H_t) \tag{46}$$

where

$$h_t = m_t + C_t R_{t+1}^{-1} (\kappa_{t+1} - a_{t+1}) \tag{47a}$$

$$H_t = C_t - C_t R_{t+1}^{-1} C_t \tag{47b}$$

which can be derived based on Kalman smoother (Petris *et al.*, 2009).

The FFBS procedure is displayed in Algorithm 3. Note that the prior distribution for $\boldsymbol{\kappa}_0$ can be set to be vague to run the Kalman filter; the output of the algorithm includes the posterior distribution of $\boldsymbol{\kappa}_0$.

Algorithm 3 FFBS Algorithm: forward filtering-backward sampling

- 1: Run Kalman filter to obtain m_T and C_T .
 - 2: Sample κ_T from $N(m_T, C_T)$.
 - 3: **for** $t = T-1, \dots, 0$ **do**
 - 4: Sample κ_t from $N(h_t, H_t)$ using the sample κ_{t+1} obtained in the previous step.
 - 5: **end for**
-
-

5.1.2. Sampling from the full conditional density $\pi(\boldsymbol{\psi} | \boldsymbol{\kappa}_{0:T}, \mathbf{y}_{1:T})$

The first thing to observe is that under the reparameterisation of the identification constraints (18), the following Gibbs sampling stages can be performed exactly.

We assume that the prior for $(\alpha_{x_2:x_p}, \beta_{x_2:x_p}, \theta, \sigma_{\varepsilon, x_1:x_p}^2, \sigma_\omega^2)$ are given by

$$\alpha_x \sim N(\tilde{\boldsymbol{\mu}}_\alpha, \tilde{\boldsymbol{\sigma}}_\alpha^2), \quad \beta_x \sim N(\tilde{\boldsymbol{\mu}}_\beta, \tilde{\boldsymbol{\sigma}}_\beta^2), \quad \theta \sim N(\tilde{\boldsymbol{\mu}}_\theta, \tilde{\boldsymbol{\sigma}}_\theta^2) \tag{48a}$$

$$\sigma_{\varepsilon, x}^2 \sim \text{IG}(\tilde{a}_\varepsilon, \tilde{b}_\varepsilon), \quad \sigma_\omega^2 \sim \text{IG}(\tilde{a}_\omega, \tilde{b}_\omega) \tag{48b}$$

where $\text{IG}(\tilde{a}_\omega, \tilde{b}_\omega)$ denotes an inverse- γ distribution with mean $\tilde{b}_\omega / (\tilde{a}_\omega - 1)^2$ and variance $\tilde{b}_\omega^2 / ((\tilde{a}_\omega - 1)^2 (\tilde{a}_\omega - 2))$ for $\tilde{a}_\omega > 2$. We assume that the priors for all parameters are independent. In this case the posterior densities of parameters are of the same type as the prior densities, a so-called conjugate prior. The posterior distribution for each parameter is given by (we write, for ease of notation, $\mathbf{y} = \mathbf{y}_{1:T}$, $\boldsymbol{\kappa} = \boldsymbol{\kappa}_{0:T}$ and family $\boldsymbol{\psi}_{-\lambda}$ means “parameter vector $\boldsymbol{\psi}$ without the parameter λ ”):

$$\alpha_x | \mathbf{y}_{1:T}, \boldsymbol{\kappa}, \boldsymbol{\psi}_{-\alpha_x} \sim N\left(\frac{\tilde{\boldsymbol{\mu}}_\alpha \sigma_{\varepsilon, x}^2 + \tilde{\boldsymbol{\sigma}}_\alpha^2 \sum_t (y_{xt} - \beta_x \kappa_t)}{\tilde{\boldsymbol{\sigma}}_\alpha^2 T + \sigma_{\varepsilon, x}^2}, \frac{\tilde{\boldsymbol{\sigma}}_\alpha^2 \sigma_{\varepsilon, x}^2}{\tilde{\boldsymbol{\sigma}}_\alpha^2 T + \sigma_{\varepsilon, x}^2}\right) \tag{49}$$

$$\beta_x | \mathbf{y}_{1:T}, \boldsymbol{\kappa}, \boldsymbol{\psi}_{-\beta_x} \sim N \left(\frac{\tilde{\sigma}_\beta^2 \sum_t (y_{xt} - \alpha_x) \kappa_t + \tilde{\mu}_\beta \sigma_{\varepsilon,x}^2}{\tilde{\sigma}_\beta^2 \sum_t \kappa_t^2 + \sigma_{\varepsilon,x}^2}, \frac{\tilde{\sigma}_\beta^2 \sigma_{\varepsilon,x}^2}{\tilde{\sigma}_\beta^2 \sum_t \kappa_t^2 + \sigma_{\varepsilon,x}^2} \right) \tag{50}$$

$$\theta | \mathbf{y}_{1:T}, \boldsymbol{\kappa}, \boldsymbol{\psi}_{-\theta} \sim N \left(\frac{\tilde{\sigma}_\theta^2 \sum_{t=1}^T (\kappa_t - \kappa_{t-1}) + \tilde{\mu}_\theta \sigma_\omega^2}{\tilde{\sigma}_\theta^2 T + \sigma_\omega^2}, \frac{\tilde{\sigma}_\theta^2 \sigma_\omega^2}{\tilde{\sigma}_\theta^2 T + \sigma_\omega^2} \right) \tag{51}$$

$$\sigma_{\varepsilon,x}^2 | \mathbf{y}_{1:T}, \boldsymbol{\kappa}, \boldsymbol{\psi}_{-\sigma_{\varepsilon,x}^2} \sim \text{IG} \left(\tilde{a}_\varepsilon + \frac{\tilde{p}T}{2}, \tilde{b}_\varepsilon + \frac{1}{2} \sum_{t=1}^T (y_{xt} - (\alpha_x + \beta_x \kappa_t))^2 \right) \tag{52}$$

$$\sigma_\omega^2 | \mathbf{y}_{1:T}, \boldsymbol{\kappa}, \boldsymbol{\psi}_{-\sigma_\omega^2} \sim \text{IG} \left(\tilde{a}_\omega + \frac{T}{2}, \tilde{b}_\omega + \frac{1}{2} \sum_{t=1}^T (\kappa_t - (\kappa_{t-1} + \theta))^2 \right) \tag{53}$$

5.2. Non-linear/non-Gaussian state-space inference

In the case of non-linear/non-Gaussian state-space model dynamics such as the stochastic volatility models of LC3-H2 and the LCSV models, the sampler we develop is based on novel developments of the particle Metropolis–Hastings samplers of Andrieu *et al.* (2010) adapted to the stochastic mortality models. In particular, we consider a combination of Rao-Blackwellised Kalman filter and particle filter for the latent state process full posterior conditionals, combined with a combination of Metropolis-within-Gibbs and Gibbs sampling steps for the static model parameters, both embedded within a PMCMC framework. We will illustrate the idea of this methodology for the LCSV model where a stochastic volatility dynamics is included in the latent process for the period effect.

5.2.1. Estimation for the LCSV mortality model

The static parameter vector is denoted as $\boldsymbol{\psi} = (\alpha_{x_2:x_p}, \beta_{x_2:x_p}, \theta, \sigma_\varepsilon^2, \sigma_\gamma^2, \lambda_1, \lambda_2, \gamma_0)$. Note that we treat γ_0 as a static parameter and our task is to obtain samples from the joint posterior distribution:

$$\pi(\boldsymbol{\kappa}_{0:T}, \gamma_{1:T}, \boldsymbol{\psi} | \mathbf{y}_{1:T}) \tag{54}$$

In this setting one can try a number of different approaches, the first would be to sample jointly from the full posterior distribution (54) via PMCMC methods to be described below. A second approach would be to combine PMCMC methods within a block-Gibbs-based sampler such as the following sampling scheme, where we apply Gibbs sampling to sample from the full conditional densities:

$$\pi(\boldsymbol{\kappa}_{0:T} | \boldsymbol{\psi}, \gamma_{1:T}, \mathbf{y}_{1:T}) \tag{55a}$$

$$\pi(\boldsymbol{\psi} | \boldsymbol{\kappa}_{0:T}, \gamma_{1:T}, \mathbf{y}_{1:T}) \tag{55b}$$

$$\pi(\gamma_{1:T} | \boldsymbol{\psi}, \boldsymbol{\kappa}_{0:T}, \mathbf{y}_{1:T}) \tag{55c}$$

Note that sampling from (55a) can be achieved by the FFBS procedure described in Algorithm 3, as one can apply Kalman filtering since $\gamma_{1:T}$ is assumed to be given. The only difference compared to section 5.1.1 is that the term σ_ω^2 is replaced by $\exp\{\gamma_t\}$ in Kalman filtering.

In the following, we provide details on how to sample from either the full posterior (54) or from full conditionals such as the density in (55c), via PMCMC method. Sampling from the posteriors of the static parameters (55b) is detailed in section 5.2.4.

5.2.2. PMCMC for mortality models

In this section, we explain the generic form of the PMCMC methodology that can be applied in a range of approaches for state-space stochastic mortality models. In general a PMCMC sampling method is a class of MCMC method where SMC algorithm is used as a proposal distribution within a MCMC algorithm. Though this seems trivial, it is actually based on a key observation that by using such a filter within the MCMC, the dimension of the acceptance probability in the Metropolis–Hastings acceptance–rejection stage is significantly reduced and can therefore facilitate much better mixing performance of the resulting Markov chain, reducing variance in estimation, see discussion in detail in Andrieu *et al.* (2010).

To bring out the essence of PMCMC, we first discuss a generic approach to sample from a target distribution:

$$\pi(\boldsymbol{\phi}_{1:T}, \boldsymbol{\psi} | \mathbf{z}_{1:T}) \tag{56}$$

where $\boldsymbol{\phi}_{1:T}$ and $\boldsymbol{\psi}$ are the latent state and static parameters of a general state-space model. Note, the state processes in this context are generally non-linear and potentially non-Gaussian.

From the perspective of obtaining the most efficiently mixing Markov chain to sample from this posterior, the ideal proposal distribution for constructing the Markov chain for $(\boldsymbol{\phi}'_{1:T}, \boldsymbol{\psi}')$ is easily seen to be given by

$$q(\boldsymbol{\psi}' | \boldsymbol{\psi}) p_{\boldsymbol{\psi}'}(\boldsymbol{\phi}'_{1:T} | \mathbf{z}_{1:T}) \tag{57}$$

where $q(\boldsymbol{\psi}' | \boldsymbol{\psi})$ is a proposal for the parameters and the proposal for the latent state, $p_{\boldsymbol{\psi}'}(\boldsymbol{\phi}'_{1:T} | \mathbf{z}_{1:T})$, is from the state equation (given $\boldsymbol{\psi}'$). Here $(\boldsymbol{\phi}_{1:T}, \boldsymbol{\psi})$ is the current state at MCMC iteration $j - 1$ and $(\boldsymbol{\phi}'_{1:T}, \boldsymbol{\psi}')$ is the proposed next move at MCMC iteration j .

In this ideal case, the acceptance probability of this ideal proposal is given by

$$\alpha((\boldsymbol{\phi}'_{1:T}, \boldsymbol{\psi}'), (\boldsymbol{\phi}_{1:T}, \boldsymbol{\psi})) = 1 \wedge \frac{p(\boldsymbol{\phi}'_{1:T}, \boldsymbol{\psi}' | \mathbf{z}_{1:T}) q(\boldsymbol{\psi} | \boldsymbol{\psi}') p_{\boldsymbol{\psi}}(\boldsymbol{\phi}_{1:T} | \mathbf{z}_{1:T})}{p(\boldsymbol{\phi}_{1:T}, \boldsymbol{\psi} | \mathbf{z}_{1:T}) q(\boldsymbol{\psi}' | \boldsymbol{\psi}) p_{\boldsymbol{\psi}'}(\boldsymbol{\phi}'_{1:T} | \mathbf{z}_{1:T})} \tag{58}$$

$$= 1 \wedge \frac{p_{\boldsymbol{\psi}'}(\boldsymbol{\phi}_{1:T} | \mathbf{z}_{1:T}) p(\boldsymbol{\psi}' | \mathbf{z}_{1:T}) q(\boldsymbol{\psi} | \boldsymbol{\psi}') p_{\boldsymbol{\psi}}(\boldsymbol{\phi}_{1:T} | \mathbf{z}_{1:T})}{p_{\boldsymbol{\psi}}(\boldsymbol{\phi}_{1:T} | \mathbf{z}_{1:T}) p(\boldsymbol{\psi} | \mathbf{z}_{1:T}) q(\boldsymbol{\psi}' | \boldsymbol{\psi}) p_{\boldsymbol{\psi}'}(\boldsymbol{\phi}'_{1:T} | \mathbf{z}_{1:T})} \tag{59}$$

$$= 1 \wedge \frac{p(\boldsymbol{\psi}' | \mathbf{z}_{1:T}) q(\boldsymbol{\psi} | \boldsymbol{\psi}')}{p(\boldsymbol{\psi} | \mathbf{z}_{1:T}) q(\boldsymbol{\psi}' | \boldsymbol{\psi})} \tag{60}$$

$$= 1 \wedge \frac{p_{\boldsymbol{\psi}'}(\mathbf{z}_{1:T}) p(\boldsymbol{\psi}') q(\boldsymbol{\psi} | \boldsymbol{\psi}')}{p_{\boldsymbol{\psi}}(\mathbf{z}_{1:T}) p(\boldsymbol{\psi}) q(\boldsymbol{\psi}' | \boldsymbol{\psi})} \tag{61}$$

where $r_1 \wedge r_2 := \min(r_1, r_2)$. A desirable property of the ideal proposal is that the acceptance probability depends only on the marginal likelihood, together with the prior and proposal for the static parameters. This is optimal in the sense that the dimension of the numerator and denominator

is reduced significantly to the static model parameter dimensions, and not including explicitly the path-space latent process dimensions, a reduction of $d \times T$ dimensions for a d -dimensional state vector ϕ_t . However, clearly one can never achieve this goal as it requires perfect knowledge of $p_\psi(\phi'_{1:T} | z_{1:T})$ as well as the ability to sample this distribution, both of which are unachievable except in the special case of the Linear Gaussian case explained in section 5.1.1.

To circumvent this problem, the particle marginal Metropolis–Hastings sampler (PMMH; Andrieu *et al.*, 2010) applies SMC method to obtain an approximate of the state transition density (which is also the state proposal):

$$\hat{p}_\psi(\phi_{1:T} | z_{1:T}) = \sum_{i=1}^N w_T^{(i)} \delta_{\phi_{1:T}^{(i)}}(\phi_{1:T}) \tag{62}$$

where $w_T^{(i)}$ is the importance weight, $\delta_x(X)$ a Dirac mass function centred at X and a proposed next move of the latent state is drawn from this discrete approximate distribution. Moreover, a by-product of a SMC algorithm is the marginal likelihood, $\hat{p}_\psi(z_{1:T})$, which has the following important property:

Lemma 5.1 A SMC proposal admits as a by-product an unbiased estimator of the marginal likelihood $p_\psi(z_{1:T})$ given by

$$\hat{p}_\psi(z_{1:T}) := \prod_{t=2}^T \hat{p}_\psi(z_t | z_{1:t-1}) \tag{63}$$

where a SMC approximation with N -particles produces, for all t

$$\hat{p}_\psi(z_t | z_{1:t-1}) = \frac{1}{N} \sum_{i=1}^N w_t^{(i)} \tag{64}$$

which is an unbiased particle estimate of $p_\psi(z_t | z_{1:t-1})$. This non-trivial unbiasedness was first presented in Del Moral (2004) and has since been utilised to great advantage as explained in Chopin *et al.* (2013). In addition, the variance of this estimator typically only grows linearly with T .

The unbiased approximate marginal likelihoods are then used in the acceptance probability (61):

$$\alpha((\phi'_{1:T}, \psi'), (\phi_{1:T}, \psi)) = 1 \wedge \frac{\hat{p}_{\psi'}(z_{1:T}) p(\psi') q(\psi | \psi')}{\hat{p}_\psi(z_{1:T}) p(\psi) q(\psi' | \psi)} \tag{65}$$

Due to the unbiasedness of the estimated marginal likelihood, Andrieu *et al.* (2010) show that, even though only SMC approximates are used (with finite number of particles N), the invariant distribution of PMMH is the target distribution $\pi(\phi_{1:T}, \psi | z_{1:T})$.

To apply PMCMC for an efficient estimation of the LCSV model, we first notice that we can obtain explicitly the posteriors of static parameters via conjugate priors. As a result we are only required to sample from the density $\pi(\gamma_{1:T} | \psi, \kappa_{0:m}, y_{1:T})$, instead of the joint density $\pi(\gamma_{1:T}, \psi | \kappa_{0:m}, y_{1:T})$. It turns out that there is a class of PMCMC algorithm, called particle-independent Metropolis–Hastings sampler (PIMH), which provide a mechanism to sample exactly from $\pi(\gamma_{1:T} | \psi, \kappa_{0:m}, y_{1:T})$.

Our approach to sampling from the joint posterior distribution, $\pi(k_{0:T}, \gamma_{1:T}, \boldsymbol{\psi} | \mathbf{y}_{1:T})$, of the LCSV model is summarised in Algorithm 4.

Algorithm 4 Sampling from $\pi(k_{0:T}, \gamma_{1:T}, \boldsymbol{\psi} | \mathbf{y}_{1:T})$

- 1: Initialise: $\boldsymbol{\psi} = \boldsymbol{\psi}^{(0)}, \gamma_{1:T} = \gamma_{1:T}^{(0)}$.
 - 2: **for** $i = 1, \dots, M$ **do**
 - 3: Sample $\kappa_{0:T}^{(i)}$ from $\pi(\kappa_{0:T} | \gamma_{1:T}^{(i-1)}, \boldsymbol{\psi}^{(i-1)}, \mathbf{y}_{1:T})$ via FFBS;
 - 4: Sample $\gamma_{1:T}^{(i)}$ from $\pi(\gamma_{1:T} | \kappa_{0:T}^{(i)}, \boldsymbol{\psi}^{(i-1)}, \mathbf{y}_{1:T})$ via PIMH (Section 5.2.3);
 - 5: **for** $h = 1, \dots, n$ **do**
 - 6: Sample $\psi_b^{(i)}$ from $\pi(\psi_b | \kappa_{0:T}^{(i)}, \gamma_{1:T}^{(i)}, \boldsymbol{\psi}_{-b}^{(i)}, \mathbf{y}_{1:T})$,
 - 7: where $\boldsymbol{\psi}_{-b}^{(i)} = (\psi_1^{(i)}, \dots, \psi_{b-1}^{(i)}, \psi_{b+1}^{(i)}, \dots, \psi_n^{(i-1)})$ via conjugate prior.
 - 8: **end for**
 - 9: **end for**
-

5.2.3. PIMH: sampling from $\pi(\gamma_{1:T} | \boldsymbol{\psi}, \kappa_{0:n}, \mathbf{y}_{1:T})$

We first note that $\pi(\gamma_{1:T} | \boldsymbol{\psi}, \kappa_{0:n}, \mathbf{y}_{1:T}) = \pi_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:n})$ given the structure of the LCSV model. Using an independent proposal density, $q_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:n})$, in the Metropolis–Hastings algorithm, the acceptance probability is given by

$$\alpha(\gamma'_{1:T}, \gamma_{1:T}) = 1 \wedge \frac{\pi_{\boldsymbol{\psi}}(\gamma'_{1:T} | \kappa_{0:n}) q_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:n})}{\pi_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:n}) q_{\boldsymbol{\psi}}(\gamma'_{1:T} | \kappa_{0:n})} \tag{66}$$

Ideally, one may take $q_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:n}) = \pi_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:n})$. However, in most cases such an ideal choice is impossible to sample from and to evaluate. The PIMH sampler proposes instead to use the SMC approximation $\hat{\pi}_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:n})$ as the proposal density and calculate the acceptance probability as

$$\alpha(\gamma'_{1:T}, \gamma_{1:T}) = 1 \wedge \frac{\hat{\pi}_{\boldsymbol{\psi}}(\kappa_{0:n})'}{\hat{\pi}_{\boldsymbol{\psi}}(\kappa_{0:n})[j-1]} \tag{67}$$

where $\hat{\pi}_{\boldsymbol{\psi}}(\kappa_{0:n})'$ and $\hat{\pi}_{\boldsymbol{\psi}}(\kappa_{0:n})[j-1]$ are unbiased marginal likelihoods estimated by SMC (see Lemma 5.1) in the current MCMC iteration j and the previous iteration $j-1$, respectively. It can be shown that the invariant distribution of the PIMH sampler is the target distribution $\pi_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:n})$ (Andrieu *et al.*, 2010).

It remains to specify an SMC approximation $\hat{\pi}_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:n})$ (Appendix B). We use the so-called bootstrap filter, that is, the proposal distribution in the SMC algorithm to draw γ_t is given by the state equation (23c):

$$g_t(\gamma_t | \gamma_{1:t-1}, \kappa_{0:t}) := \pi_{\boldsymbol{\psi}}(\gamma_t | \gamma_{t-1}) \tag{68}$$

Consequently, the importance weight is evaluated as

$$\tilde{w}_t \propto \tilde{w}_{t-1} \pi_{\boldsymbol{\psi}}(\kappa_t | \gamma_t, \kappa_{t-1}) \tag{69}$$

where $\pi(\kappa_t | \gamma_t, \kappa_{t-1})$ is the incremental importance weight. Our approach for sampling from $\pi_{\boldsymbol{\psi}}(\gamma_{1:T} | \kappa_{0:T})$ is summarised in Algorithm 5 (together with Algorithm 6).

Algorithm 5 PIMH: sampling from $\pi_{\psi}(\gamma_{1:T}|\kappa_{0:T})$

- 1: Iteration $j = 0$: obtain an SMC approximation $\hat{\pi}_{\psi}(\gamma_{1:T}|\kappa_{0:T})$ via Algorithm 6. Draw $\gamma_{1:T}[0] \sim \hat{\pi}_{\psi}(\gamma_{1:T}|\kappa_{0:T})$ and obtain the corresponding marginal likelihood estimate $\hat{\pi}_{\psi}(\kappa_{0:n})[0]$.
- 2: **for** $j = 1, \dots, N_{PIMH}$ **do**
- 3: Obtain an SMC approximation $\hat{\pi}_{\psi}(\gamma_{1:T}|\kappa_{0:T})$ via Algorithm 6. Draw $\gamma'_{1:T} \sim \hat{\pi}_{\psi}(\gamma_{1:T}|\kappa_{0:T})$ and obtain the corresponding marginal likelihood estimate $\hat{\pi}_{\psi}(\kappa_{0:n})'$.
- 4: Draw $u \sim U(0, 1)$. If

$$u < \frac{\hat{\pi}_{\psi}(\kappa_{0:n})'}{\hat{\pi}_{\psi}(\kappa_{0:n})[j-1]}, \tag{70}$$
 set $\gamma_{1:T}[j] = \gamma'_{1:T}$ and $\hat{\pi}_{\psi}(\kappa_{0:n})[j] = \hat{\pi}_{\psi}(\kappa_{0:n})'$; otherwise set $\gamma_{1:T}[j] = \gamma_{1:T}[j-1]$ and $\hat{\pi}_{\psi}(\kappa_{0:n})[j] = \hat{\pi}_{\psi}(\kappa_{0:n})[j-1]$.
- 5: **end for**
- 6: Obtain $\gamma_{1:T}[N_{PIMH}]$ as a sample of $\pi_{\psi}(\gamma_{1:T}|\kappa_{0:T})$.

Algorithm 6 Bootstrap filter of $\pi_{\psi}(\gamma_{1:T}|\kappa_{0:T})$; see Appendix B

- 1: At $t = 1$: draw $\gamma_1^{(i)}$ from $\pi_{\psi}(\gamma_1|\gamma_0)$. Set $\tilde{w}_1^{(i)} = \pi_{\psi}(\kappa_1|\gamma_1, \kappa_0)$ and $w_1^{(i)} = \tilde{w}_1^{(i)} / \sum_{j=1}^N \tilde{w}_1^{(j)}$.
- 2: **for** $t = 2, \dots, T$ **do**
- 3: Draw $\gamma_t^{(i)}$ from $\pi_{\psi}(\gamma_t|\gamma_{t-1}^{(i)})$ and set

$$\gamma_{1:t}^{(i)} = (\gamma_{1:t-1}^{(i)}, \gamma_t^{(i)}) \tag{71}$$
- 4: Evaluate

$$\tilde{w}_t^{(i)} = \tilde{w}_{t-1}^{(i)} \cdot \pi_{\psi}(\kappa_t|\gamma_t^{(i)}, \kappa_{t-1}) \tag{72}$$
- 5: Normalise:

$$w_t^{(i)} = \frac{\tilde{w}_t^{(i)}}{\sum_{j=1}^N \tilde{w}_t^{(j)}} \tag{73}$$
- 6: Evaluate

$$N_{eff} = \left(\sum_{i=1}^N (w_t^{(i)})^2\right)^{-1} \tag{74}$$
- 7: If $N_{eff} < 0.8N$, resample $\gamma_{1:t}^{(i)}$ from $(w_t^{(j)}, \gamma_{1:t}^{(j)})_{j=1}^N$ and set $w_t^{(i)} = \frac{1}{N}$.
- 8: **end for**
- 9: Obtain $\hat{\pi}_{\psi}(\gamma_{1:T}|\kappa_{0:T}) = \sum_{i=1}^N w_T^{(i)} \delta_{\gamma_{1:T}^{(i)}}(\gamma_{1:T})$.

5.2.4. Sampling from $\pi(\psi|\kappa_{0:T}, \gamma_{1:T}, \mathbf{y}_{1:T})$

We assume the prior for $(\alpha_{x_2:x_p}, \beta_{x_2:x_p}, \theta, \sigma_{\epsilon}^2, \sigma_{\alpha}^2, \lambda_1, \lambda_2, \gamma_0)$ is given by

$$\alpha_x \sim N(\tilde{\mu}_{\alpha}, \tilde{\sigma}_{\alpha}^2), \quad \beta_x \sim N(\tilde{\mu}_{\beta}, \tilde{\sigma}_{\beta}^2), \quad \theta \sim N(\tilde{\mu}_{\theta}, \tilde{\sigma}_{\theta}^2), \quad \sigma_{\epsilon}^2 \sim \text{IG}(\tilde{a}_{\epsilon}, \tilde{b}_{\epsilon}) \tag{75a}$$

$$\sigma_{\gamma}^2 \sim \text{IG}(\tilde{a}_{\gamma}, \tilde{b}_{\gamma}), \quad \lambda_1 \sim N_{[-1, 1]}(\tilde{\mu}_{\lambda_1}, \tilde{\sigma}_{\lambda_1}^2), \quad \lambda_2 \sim N(\tilde{\mu}_{\lambda_2}, \tilde{\sigma}_{\lambda_2}^2), \quad \gamma_0 \sim N(\tilde{\mu}_{\gamma_0}, \tilde{\sigma}_{\gamma_0}^2) \tag{75b}$$

where $x = x_2, \dots, x_p$ and $N_{[-1, 1]}$ denotes a truncated Gaussian with support $[-1, 1]$. It is assumed that the priors for all parameters are independent.

Samples from the density $\pi(\psi|\kappa_{0:T}, \gamma_{1:T}, \mathbf{Y}_{1:T})$ are obtained by sampling from the following posteriors:

$$\alpha_x | \mathbf{y}_{1:T}, \kappa, \gamma, \psi_{-\alpha_x} \sim N\left(\frac{\tilde{\mu}_{\alpha} \sigma_{\epsilon}^2 + \tilde{\sigma}_{\alpha}^2 \sum_t (y_{xt} - \beta_x \kappa_t)}{T \tilde{\sigma}_{\alpha}^2 + \sigma_{\epsilon}^2}, \frac{\tilde{\sigma}_{\alpha}^2 \sigma_{\epsilon}^2}{T \tilde{\sigma}_{\alpha}^2 + \sigma_{\epsilon}^2}\right) \tag{76}$$

$$\beta_x | \mathbf{y}_{1:T}, \kappa, \gamma, \boldsymbol{\Psi}_{-\beta_x} \sim N \left(\frac{\tilde{\sigma}_\beta^2 \sum_t (y_{xt} - \alpha_x) \kappa_t + \tilde{\mu}_\beta \sigma_\varepsilon^2}{\tilde{\sigma}_\beta^2 \sum_t \kappa_t^2 + \sigma_\varepsilon^2}, \frac{\tilde{\sigma}_\beta^2 \sigma_\varepsilon^2}{\tilde{\sigma}_\beta^2 \sum_t \kappa_t^2 + \sigma_\varepsilon^2} \right) \quad (77)$$

$$\theta | \mathbf{y}_{1:T}, \kappa, \gamma, \boldsymbol{\Psi}_{-\theta} \sim N \left(\frac{\tilde{\mu}_\theta / \tilde{\sigma}_\theta^2 + \sum_t (\kappa_t - \kappa_{t-1}) / e^{\gamma t}}{1 / \tilde{\sigma}_\theta^2 + \sum_t 1 / e^{\gamma t}}, \frac{1}{1 / \tilde{\sigma}_\theta^2 + \sum_t 1 / e^{\gamma t}} \right) \quad (78)$$

$$\sigma_\varepsilon^2 | \mathbf{y}_{1:T}, \kappa, \gamma, \boldsymbol{\Psi}_{-\sigma_\varepsilon^2} \sim \text{IG} \left(\tilde{a}_\varepsilon + \frac{pT}{2}, \tilde{b}_\varepsilon + \frac{1}{2} \sum_{t=1}^T \sum_{x=1}^p (y_{xt} - (\alpha_x + \beta_x \kappa_t))^2 \right) \quad (79)$$

$$\sigma_\gamma^2 | \mathbf{y}_{1:T}, \kappa, \gamma, \boldsymbol{\Psi}_{-\sigma_\gamma^2} \sim \text{IG} \left(\tilde{a}_\gamma + \frac{T}{2}, \tilde{b}_\gamma + \frac{1}{2} \sum_{t=1}^T (\gamma_t - \lambda \gamma_{t-1})^2 \right) \quad (80)$$

$$\lambda_1 | \mathbf{y}_{1:T}, \kappa, \gamma, \boldsymbol{\Psi}_{-\lambda_1} \sim N_{[-1,1]} \left(\frac{\sigma_\gamma^2 \tilde{\mu}_{\lambda_1} + \tilde{\sigma}_{\lambda_1}^2 \sum_t \gamma_{t-1} \gamma_t}{\sigma_\gamma^2 + \tilde{\sigma}_{\lambda_1}^2 \sum_t \gamma_{t-1}^2}, \frac{\tilde{\sigma}_{\lambda_1}^2 \sigma_\gamma^2}{\sigma_\gamma^2 + \tilde{\sigma}_{\lambda_1}^2 \sum_t \gamma_{t-1}^2} \right) \quad (81)$$

$$\lambda_2 | \mathbf{y}_{1:T}, \kappa, \gamma, \boldsymbol{\Psi}_{-\lambda_2} \sim N \left(\frac{\sigma_\gamma^2 \tilde{\mu}_{\lambda_2} + \tilde{\sigma}_{\lambda_2}^2 \sum_t (\gamma_t - \lambda_1 \gamma_{t-1})}{\sigma_\gamma^2 + T \tilde{\sigma}_{\lambda_2}^2}, \frac{\tilde{\sigma}_{\lambda_2}^2 \sigma_\gamma^2}{\sigma_\gamma^2 + T \tilde{\sigma}_{\lambda_2}^2} \right) \quad (82)$$

$$\gamma_0 | \mathbf{y}_{1:T}, \kappa, \gamma, \boldsymbol{\Psi}_{-\gamma_0} \sim N \left(\frac{\sigma_\gamma^2 \tilde{\mu}_{\gamma_0} + \tilde{\sigma}_{\gamma_0}^2 \lambda \gamma_1}{\sigma_\gamma^2 + \tilde{\sigma}_{\gamma_0}^2 \lambda^2}, \frac{\tilde{\sigma}_{\gamma_0}^2 \sigma_\gamma^2}{\sigma_\gamma^2 + \tilde{\sigma}_{\gamma_0}^2 \lambda^2} \right) \quad (83)$$

where the posterior distributions are obtained similarly as in section 5.1.2.

6. Empirical Analysis: Danish Male Population

In this section, a comprehensive empirical study⁴ is conducted on Danish mortality data using the models summarised in Table 2. The LC, LC-H, LCSV and LCSV-H models are described in section 3. While the LC-H model addresses heteroscedasticity in the observation equation, the LCSV model attempts to incorporate stochastic volatility in the state dynamics. The LCSV-H model includes both features of the LC-H model and the LCSV model, thus allowing for a full consideration of variability in long-term mortality dynamics.

Note that we have omitted the LC2-H model (a cohort model) in this analysis as a detailed account of the formulation, Bayesian estimation and forecasting of state-space cohort models is considered in a separate paper, see Fung *et al.* (2017). We are also aware that there are growing interests in the literature to consider more realistic models for the latent factor; for example, Li *et al.* (2011), van Berkum *et al.* (2014) and Liu & Li (2016b) all study trend-changing stochastic behaviour for the latent period effects. For this reason we will focus on the LCSV and LCSV-H models, where stochastic volatility is introduced in the latent dynamics, instead of the LC3-H2 model where stochastic volatility enters via the observation equation, in our empirical studies.

⁴ We have also performed numerous simulation studies using synthetic data to confirm the effectiveness of our estimation approaches but these are omitted here for space considerations. They are available upon request.

Table 2. A summary of state-space mortality models considered in our empirical study.

Models	Name	Dynamics
Lee–Carter (LC) model	LC	(3)–(4)
LC model with heteroscedasticity	LC-H	(21)
LC stochastic volatility (SV) model	LCSV	(23)
LC SV model with heteroscedasticity	LCSV-H	(24)

The Human Mortality Database⁵ provides a particularly long time series of mortality data from year 1835 to 2011 for the Danish population, supplemented with a detailed document analysing the data (Andreev, 2002). The provision of a long time series is important to our analysis concerning stochastic volatility. In the past several decades, mortality trend for developed countries generally exhibit a rather smooth pattern. The inclusion of periods that involve wars, epidemics or other life-critical events are crucial factors in witnessing significant volatility in mortality time series. In the following, we analyse the population mortality from Denmark based on the models in Table 2 and Bayesian methodologies studied in this paper. We then examine the models in terms of the forecasting properties of death rates and life expectancies. We also comment on the linear trend assumption and jump-off bias in mortality forecasting.

6.1. Data description

The data set consists of Danish male population death rates for 21 age groups (0, 1–4, 5–9, ..., 95–99) from year 1835 to 2010 where we fix year 2010 as the end year. Figure 1 displays some of the time series of the log death rates for the Danish male population. It is clear that the multi-dimensional time series exhibit different volatility for different age groups, which justify the introduction of heteroscedasticity into the observation equation as discussed in section 3.1. We also observe that, mainly before 1950, there are periods that the volatility of death rates for some age groups are markedly different. Such a change of volatility in the temporal dimension suggests that stochastic volatility may be present in the underlying time period effect.

6.2. Estimation results

In our empirical study we focus on Bayesian inference and forecasting. We assume vague priors so that all inferences are mainly based on the data and the impact of the prior is not material. Taking the LCSV model as an example, we assume $\kappa_0 \sim N(0, 10)$, $\alpha_x \sim N(0, 10)$, $\beta_x \sim N(0, 10)$, $\theta \sim N(0, 10)$, $\sigma_\varepsilon^2 \sim \text{IG}(2.001, 0.001)$, $\sigma_\gamma^2 \sim \text{IG}(2.001, 0.001)$, $\lambda_1 \sim N(0, 10)$, $\lambda_2 \sim N(0, 10)$ and $\gamma_0 \sim N(0, 10)$, where $x \in \{x_2, \dots, x_p\}$. The number of iterations of the Markov chain is 15,000 with 5,000 burn-in. We fix $\alpha_{x_1} = \frac{1}{T} \sum_{t=1}^T y_{x_1,t}$ and $\beta_{x_1} = 0.2$ as an identification constraint; see Remark 3.1 for a discussion for our choice of values here.

Estimated values of the static parameters (except α and β) for the Danish mortality data (1835–2010) are shown in Table 3. The rest of the estimated parameters and states are displayed in Figure 2. Here, we only show the plots for the LCSV-H model since the corresponding

⁵ <http://www.mortality.org/> (accessed on September 2015).

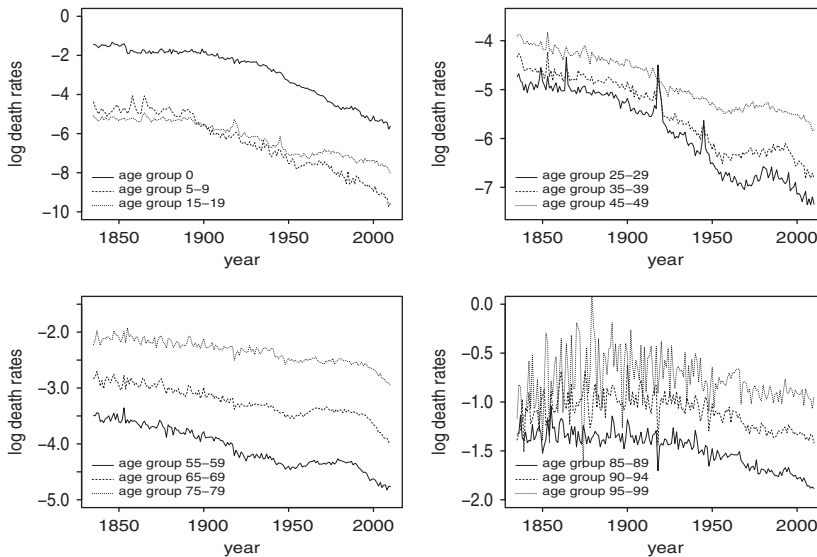


Figure 1. Time series of log death rates for Danish male population from year 1835 to 2010.

Table 3. Estimated values of the static parameters (except α and β) for the Danish male mortality data (1835–2010).

	LC	LC-H	LCSV	LCSV-H
θ	-0.11 (-0.17, -0.06)	-0.11 (-0.17, -0.06)	-0.11 (-0.15, -0.07)	-0.09 (-0.14, -0.04)
σ_{ϵ}^2	0.023 (0.022, 0.024)	Similar to Figure 2	0.023 (0.022, 0.024)	Figure 2
σ_{ω}^2	0.13 (0.09, 0.18)	0.15 (0.10, 0.21)	N.A.	N.A.
λ_1	N.A.	N.A.	0.989 (0.962, 0.999)	0.984 (0.949, 0.999)
λ_2	N.A.	N.A.	-0.025 (-0.11, 0.042)	-0.03 (-0.15, 0.05)
σ_{τ}^2	N.A.	N.A.	0.15 (0.03, 0.48)	0.25 (0.06, 0.67)
γ_0	N.A.	N.A.	-2.09 (-4.52, 0.23)	-2.11 (-5.04, 0.47)

The range in (,) represents 95% credible interval.

LC, Lee-Carter; LC-H, LC model with heteroscedasticity; LCSV, LC stochastic volatility; LCSV-H, LCSV model with heteroscedasticity.

figures obtained from the LC, LC-H and LCSV model are visually similar to the case of the LCSV-H model.

It is evident from Figure 2 that there are periods, namely 1850–1870, 1910–1920, 1930–1950, that the time effect κ exhibits higher volatility compared with other periods. We also observe that κ accelerates markedly downward after 1990 and is relatively smooth in the recent period 1950–2010. The filtering of the log-volatility process $\gamma_{1835:2010}$ (Figure 2) quantifies the volatility level (e^{γ_t}) of the time effect and gives further evidence on the stochastic volatility nature of mortality. To see more clearly the phenomenon of changing volatility, we plot the first difference $\Delta \bar{\kappa}_t = \bar{\kappa}_t - \bar{\kappa}_{t-1}$ in Figure 2 for the LCSV-H model, where $\bar{\kappa}_t$ denotes the posterior mean of κ_t , $t = 1836, \dots, 2010$. It shows evidently the change of volatility level in the latent process κ_t . The patterns of the estimated log-volatility $\gamma_{1835:2010}$ and the first difference $\Delta \bar{\kappa}_t$ clearly suggest that it is not appropriate to assume constant volatility (σ_{ω}^2) for the time effect.

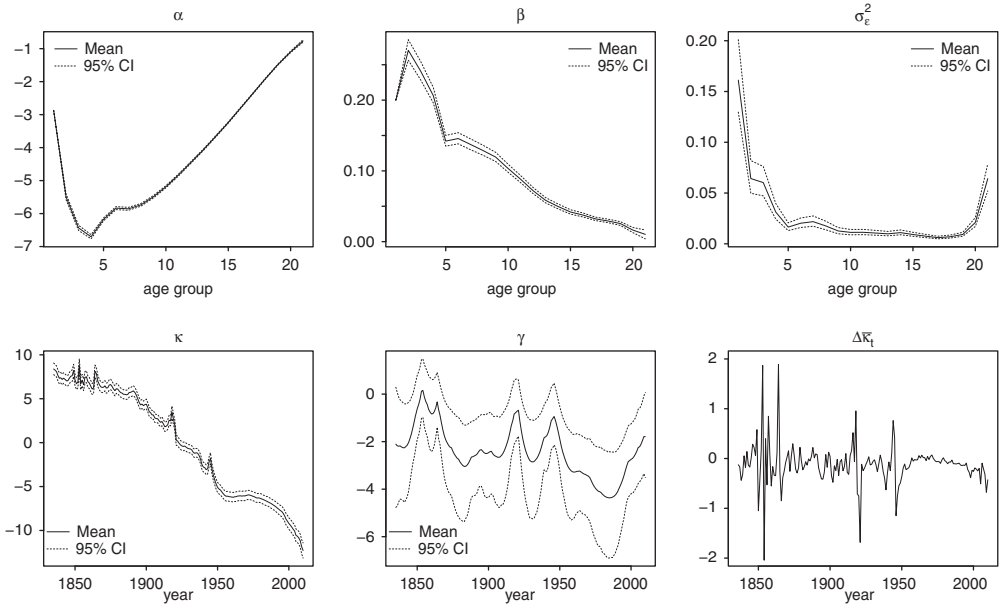


Figure 2. Estimation of (upper panels) α , β and σ_{ϵ}^2 ; (lower panels) time effect $\kappa_{1834:2010}$, log-volatility $\gamma_{1835:2010}$ and first difference $\Delta\bar{\kappa}_t$, for Danish male mortality data (1835–2010) using the LCSV-H model. LCSV-H, Lee–Carter stochastic volatility model with heteroscedasticity; CI, confidence interval.

The state-space modelling approach is able to uncover the age-specific heteroscedasticity structure hidden in the Danish mortality time series. Figure 2 reveals that variability is particularly high for the very young and very old age group. Implications of the heteroscedastic structure on forecasting will be discussed in section 6.4.

To investigate the forecasting properties of the stochastic volatility model, we also estimate the models based on calibration periods 1835–1990 and 1950–1990. Figures 3 and 4 show the estimated parameters and states for the LCSV-H model in those periods.

6.3. Model assessment

To compare the fit of the models to the data, we apply deviance information criterion (DIC) as a Bayesian measures of model complexity and fit (Spiegelhalter *et al.*, 2002). It is common to assess and compare models with latent variables using conditional DIC (Berg *et al.*, 2004; Celeux *et al.*, 2006). Specifically, we use the so-called conditional log-likelihood which is calculated as

$$\ln f(y_{1:T} | \Psi, \kappa_{1:T}) = \sum_{x=x_1}^{x_p} \sum_{t=1}^T \left(-\frac{1}{2} \ln 2\pi - \ln \sigma_{\epsilon,x} - \frac{1}{2} \left(\frac{y_{x,t} - (\alpha_x + \beta_x \kappa_t)}{\sigma_{\epsilon,x}} \right)^2 \right) \quad (84)$$

Note that the likelihood is conditional on parameters that include both static parameters and the latent process κ . Using the conditional log-likelihood function, the deviance is defined as

$$D(\Psi) = -2 \ln f(y_{1:T} | \Psi) + 2 \ln h(y_{1:T}) \quad (85)$$

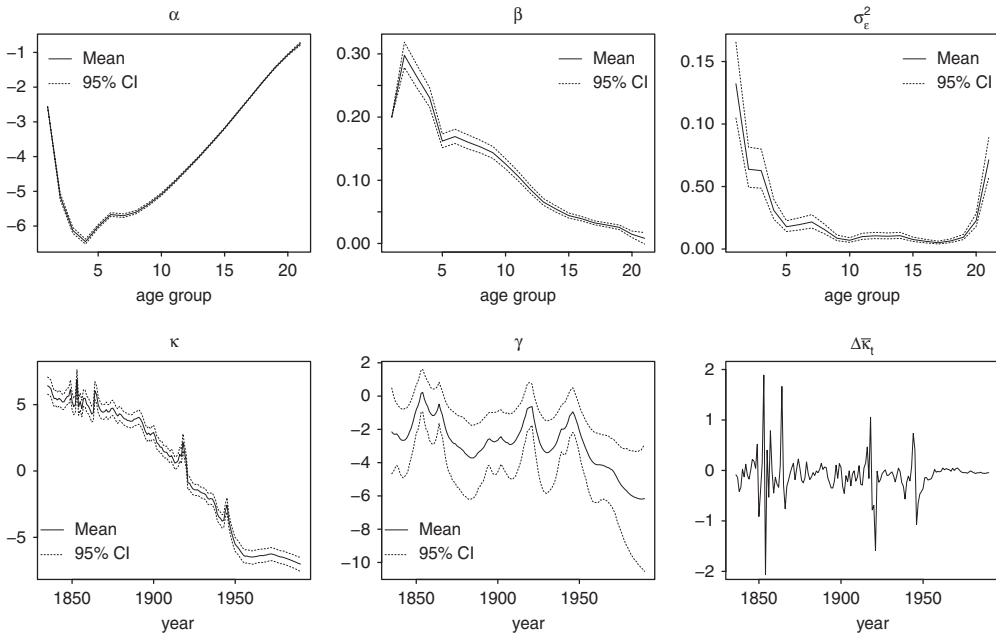


Figure 3. Estimation of (upper panels) α , β and $\sigma_{x_1:x_{21},\epsilon}^2$; (lower panels) time effect $\kappa_{1834:1900}$, log-volatility $\gamma_{1835:1990}$ and first difference $\Delta \bar{\kappa}_t$, for Danish male mortality data (1835–1990) using the LCSV-H model. LCSV-H, Lee–Carter stochastic volatility model with heteroscedasticity; CI, confidence interval.

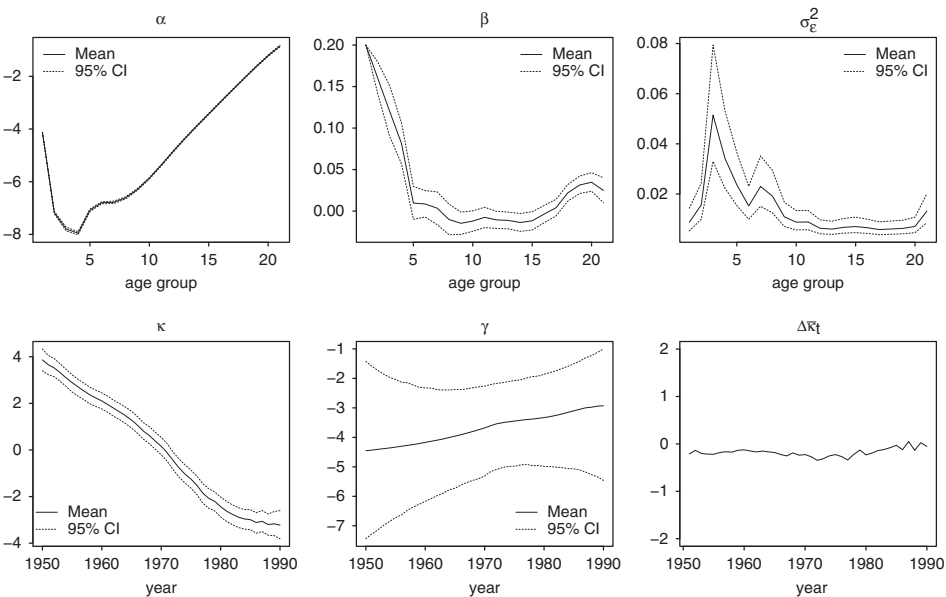


Figure 4. Estimation of (upper panels) α , β and $\sigma_{x_1:x_{21},\epsilon}^2$; (lower panels) time effect $\kappa_{1949:1990}$, log-volatility $\gamma_{1950:1990}$ and first difference $\Delta \bar{\kappa}_t$, for Danish male mortality data (1950–1990) using the LCSV-H model. LCSV-H, Lee–Carter stochastic volatility model with heteroscedasticity; CI, confidence interval.

Table 4. Deviance information criterion of models with different calibration periods.

Calibration period	1835–2010	1835–1990	1950–1990
LC	–3,218.6	–3,087.5	–1,567.3
LC-H	–4,469.1	–4,269.7	–1,793.6
LCSV	–3,250.8	–3,109.7	–1,559.7
LCSV-H	–4,518.3	–4,326.8	–1,794.1

LC, Lee–Carter; LC-H, LC model with heteroscedasticity; LCSV, LC stochastic volatility; LCSV-H, LCSV model with heteroscedasticity.

where $\Psi = (\boldsymbol{\psi}, \kappa_{1:T})$ and we assume $h(\mathbf{y}_{1:T}) = 1$ since in the models we consider it plays the role of a constant which is the same for competing models. The effective dimension, p_D , is evaluated as

$$p_D = \overline{D}(\Psi) - D(\overline{\Psi}) \tag{86}$$

where $\overline{D}(\Psi)$ and $\overline{\Psi}$ denote, respectively, the mean of $D(\Psi)$ and the mean of the posterior distribution of Ψ . The conditional DIC is then given by

$$\text{DIC} := \overline{D}(\Psi) + p_D = 2\overline{D}(\Psi) - D(\overline{\Psi}) \tag{87}$$

which can be evaluated straightforwardly using MCMC samples.

The DIC values for the models with different calibration periods are shown in Table 4,⁶ The LCSV-H model clearly outperforms other models for all calibration periods considered. The LC-H is the second best; in fact it even outperforms the LCSV model for long calibration periods 1835–2010 and 1835–1990. These results suggest that one is strongly encouraged to include heteroscedasticity structure when modelling Danish mortality data. Note that for long calibration periods, the incorporation of stochastic volatility can markedly improve model fit as can be seen by comparing the LCSV-H model and the LC-H model, as well as the comparison between LCSV model and the LC model.

6.4. Forecasting

In this section, we investigate the forecasting properties of the mortality models summarised in Table 2 where heteroscedasticity as well as stochastic volatility structures are incorporated. Our analysis is based on the forecasting distributions of (log) death rates and life expectancy. The Bayesian state-space framework allows us to obtain the forecasting distributions using MCMC samples which is shown below.

6.4.1. Death rates

For the LC (LC-H) model, the k -step ahead forecasting distribution of y_{T+k} , given $\mathbf{y}_{1:T}$, is given by

$$\pi(\mathbf{y}_{T+k} | \mathbf{y}_{1:T}) = \int \pi(\mathbf{y}_{T+k} | \boldsymbol{\kappa}_{T+k}, \boldsymbol{\psi}) \pi(\boldsymbol{\kappa}_{T+k} | \boldsymbol{\kappa}_{T+k-1}, \boldsymbol{\psi}) \dots \pi(\boldsymbol{\kappa}_T, \boldsymbol{\psi} | \mathbf{y}_{1:T}) d\boldsymbol{\psi} d\boldsymbol{\kappa}_{T:T+k} \tag{88}$$

⁶ The lower the DIC value, the better the model in terms of a trade-off of fit and complexity.

where $\boldsymbol{\psi}$ is the parameter vector for the LC (LC-H) model; (88) suggests that we can sample recursively to obtain the forecasting distribution, for $k \geq 1$, as follows:

$$\kappa_{T+k}^{(\ell)} \sim N\left(\kappa_{T+k-1}^{(\ell)} + \theta^{(\ell)}, (\sigma_{\omega}^2)^{(\ell)}\right) \tag{89a}$$

$$\mathbf{y}_{T+k}^{(\ell)} \sim N\left(\boldsymbol{\alpha}^{(\ell)} + \boldsymbol{\beta}^{(\ell)} \kappa_{T+k}^{(\ell)}, \boldsymbol{\Sigma}^{(\ell)}\right) \tag{89b}$$

where $\ell = 1, \dots, L$ and L is the number of MCMC iterations after burn-in. Here $\boldsymbol{\Sigma}$ is a diagonal matrix with $\sigma_{\epsilon,x}^2$ on the diagonal for the LC-H model and σ_{ϵ}^2 for the LC model. This procedure generates an estimate of the forecasting distribution.

Similarly, the forecasting distribution of \mathbf{y}_{T+k} , given \mathbf{y}_T , for the LCSV (LCSV-H) model is given by

$$\begin{aligned} \pi(\mathbf{y}_{T+k} | \mathbf{y}_{1:T}) &= \int \pi(\mathbf{y}_{T+k} | \kappa_{T+k}, \boldsymbol{\psi}) \pi(\kappa_{T+k} | \kappa_{T+k-1}, \gamma_{T+k}, \boldsymbol{\psi}) \dots \\ &\quad \pi(\gamma_{T+1} | \gamma_T, \boldsymbol{\psi}) \pi(\kappa_T, \gamma_T, \boldsymbol{\psi} | \mathbf{y}_{1:T}) d\boldsymbol{\psi} d\kappa_{T:T+k} d\gamma_{T:T+k} \end{aligned} \tag{90}$$

For $k \geq 1$, the forecasting distribution can be obtained by sampling recursively

$$\gamma_{T+k}^{(\ell)} \sim N\left(\lambda_1^{(\ell)} \gamma_{T+k-1}^{(\ell)} + \lambda_2^{(\ell)}, (\sigma_{\gamma}^2)^{(\ell)}\right) \tag{91a}$$

$$\kappa_{T+k}^{(\ell)} \sim N\left(\kappa_{T+k-1}^{(\ell)} + \theta^{(\ell)}, \exp\left\{\gamma_{T+k}^{(\ell)}\right\}\right) \tag{91b}$$

$$\mathbf{y}_{T+k}^{(\ell)} \sim N\left(\boldsymbol{\alpha}^{(\ell)} + \boldsymbol{\beta}^{(\ell)} \kappa_{T+k}^{(\ell)}, \boldsymbol{\Sigma}^{(\ell)}\right) \tag{91c}$$

where $\ell = 1, \dots, L$, and $\boldsymbol{\Sigma}$ is a diagonal matrix with $\sigma_{\epsilon,x}^2$ on the diagonal for the LCSV-H model and σ_{ϵ}^2 for the LCSV model.

Figure 5 shows the forecasted log death rates based on the LC-H, LCSV and LCSV-H model, using the LC model as a benchmark. We show age groups 5–9, 35–39, 65–69 and 95–99 as representatives of young, adult, old and very old age. The models are estimated using data for the period 1835–2010 and forecast for 30 years.

The heteroscedasticity structure, from the LC-H model, gives rise to materially larger forecasting intervals for the young and very old age group, while the forecasting interval for the age group 35–39 is narrower than predicted by the LC model. The LCSV model, on the other hand, produces a wider forecasting interval compared to the LC model except for the very old age group. The observed wider forecasting interval is due to the fact that the volatility level is increasing in the last estimation periods and is larger than σ_{ω}^2 estimated in the LC model. Moreover, as the estimated β_x is close to zero at older ages (Figure 2), the impact of the forecasted κ on the prediction of death rates diminished significantly as older ages are considered. The LCSV-H model exhibits similar features of the LC-H and the LCSV model. It is interesting to note that the forecasted means obtained from the different models are very similar and their differences mainly lie in the forecasting interval.

To illustrate further the forecasting property of the LCSV model, we estimate the models for the period 1835–1990 and plot 20-year out-of-sample forecasted log death rates in Figure 6. It turns out the forecasting intervals predicted by the LCSV model tends to be narrower than the LC model, as

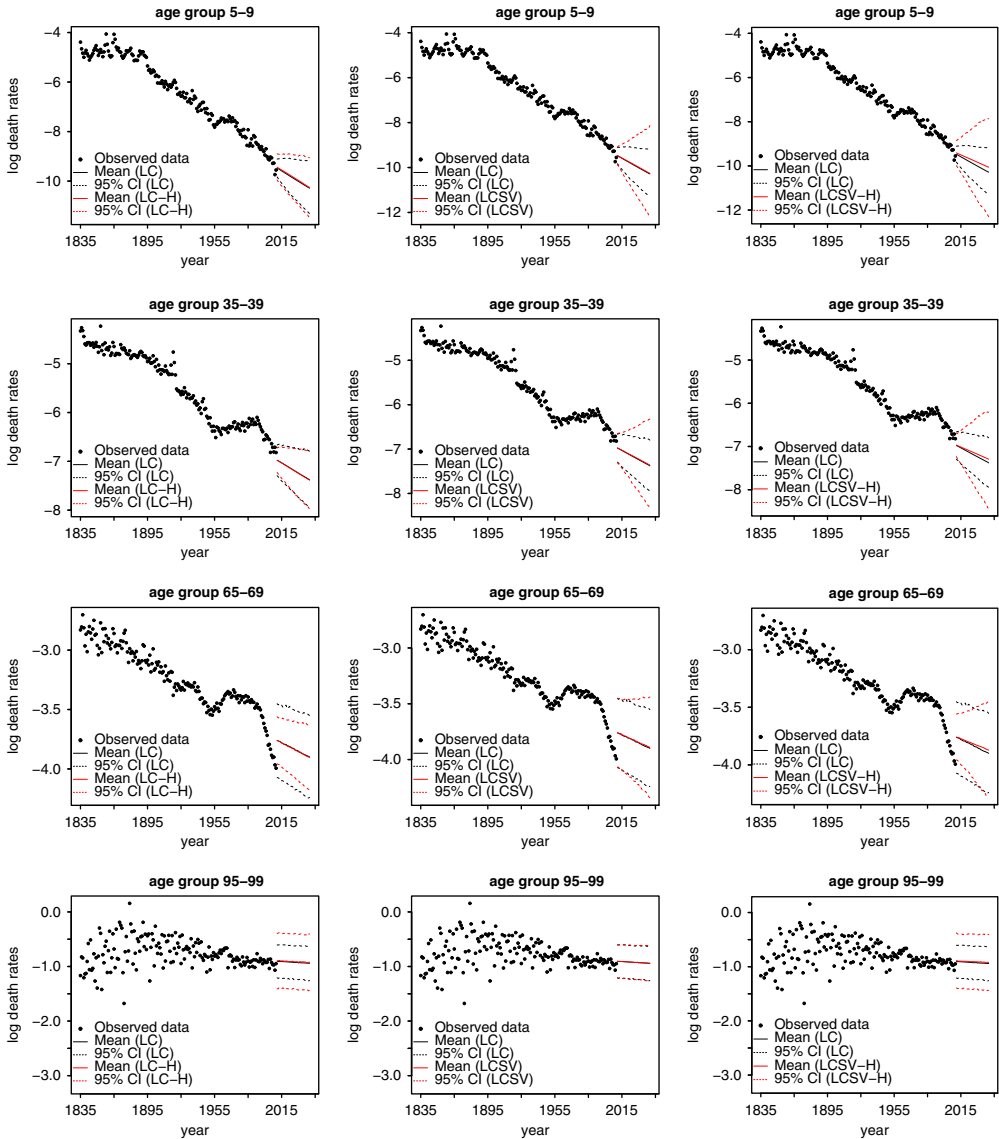


Figure 5. (Colour online) 30-year forecasted log death rates (2011–2041) for Danish male population under (left column) LC-H model, (middle column) LCSV model and (right column) LCSV-H model in comparison with LC model. Calibration period: 1835–2010. LC, Lee–Carter; LC-H, LC model with heteroscedasticity; LCSV, LC stochastic volatility; LCSV-H, LCSV model with heteroscedasticity; CI, confidence interval.

the estimated σ_{ω}^2 in the LC model is larger than the volatility level at the last estimation period for the LCSV model in this case. Note that the forecasted distributions produced by the LC-H model are biased compared to the benchmark LC model since the fitted rates at the last estimation period, that is year 1990, are different for the LC and LC-H model. This feature is known as jump-off error (Lee & Miller, 2001). One may remove this jump-off bias by forcing the forecasted death rates to start at

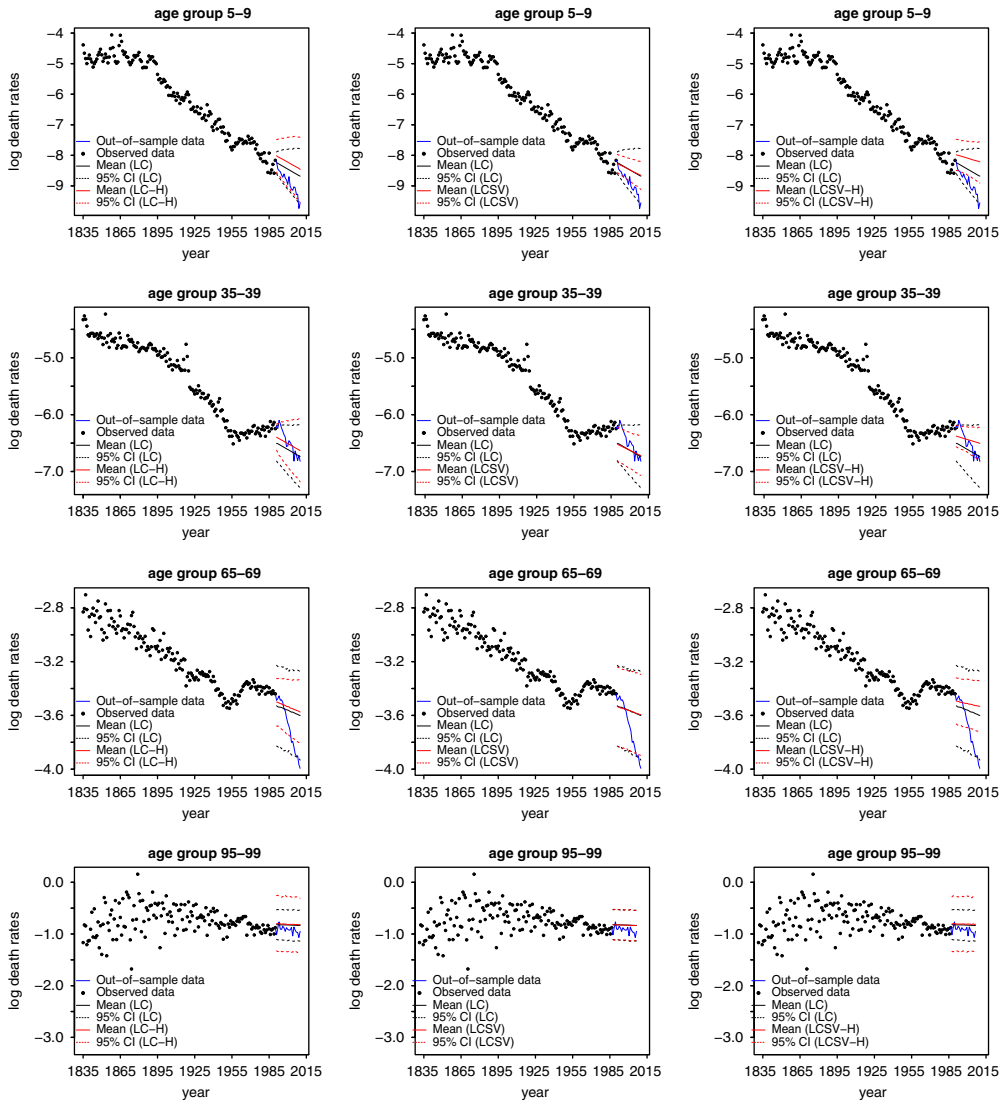


Figure 6. (Colour online) 20-year out-of-sample forecasted log death rates for Danish male population under (left column) LC-H model, (middle column) LCSV model and (right column) LCSV-H model in comparison with LC model. Calibration period: 1835–1990. LC, Lee–Carter; LC-H, LC model with heteroscedasticity; LCSV, LC stochastic volatility; LCSV-H, LCSV model with heteroscedasticity; CI, confidence interval.

the actual rates instead of the fitted rates (Bell, 1997; Shang *et al.*, 2011). In this paper we do not perform this procedure, however.

Figure 7 shows the forecasting distributions of log death rates where we assume a shorter calibration period from 1950 to 1990. For all the models, the estimated β_x for all age groups, except for age groups 0, 1–4 and 5–9, are very close to zero. It is in fact expected since there is no clear downward

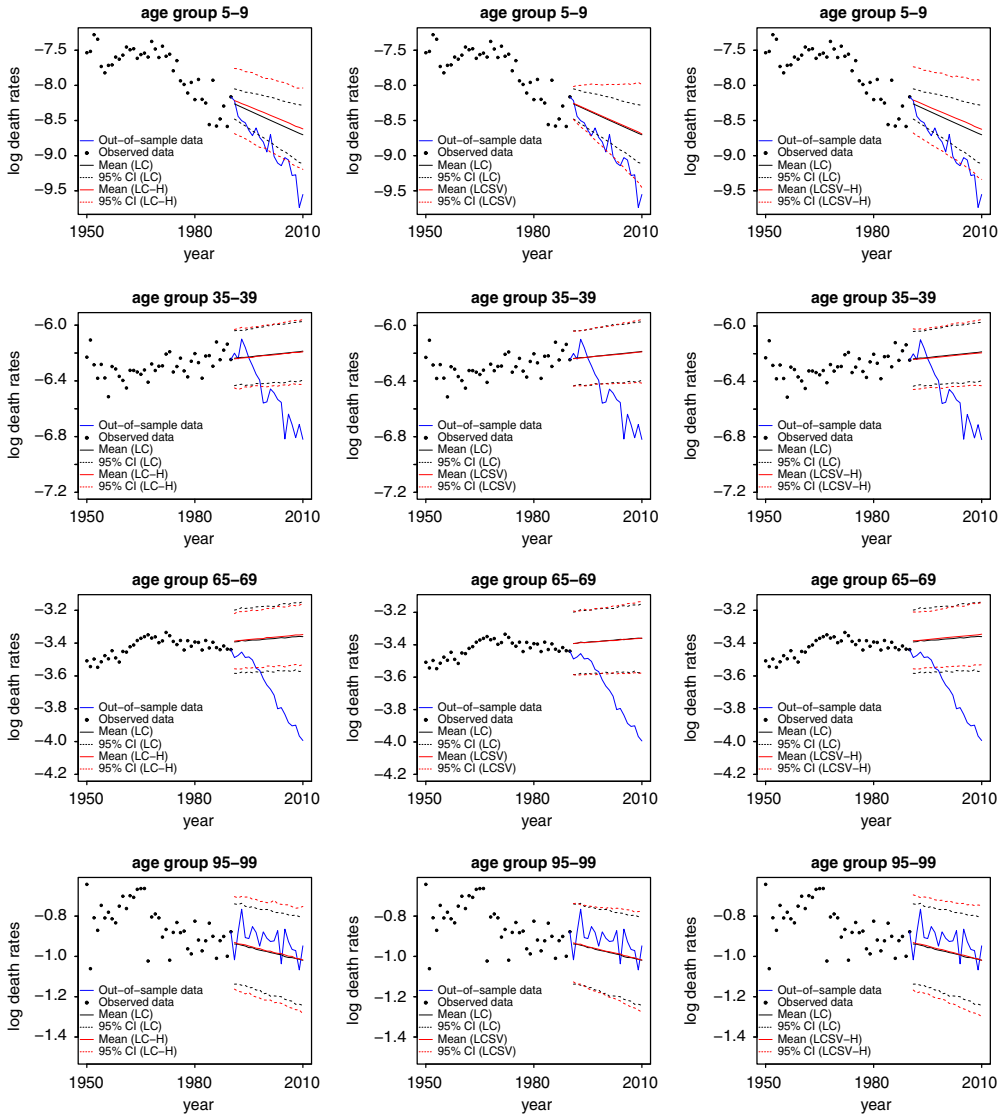


Figure 7. (Colour online) 20-year out-of-sample forecasted log death rates for Danish male population under (left column) LC-H model, (middle column) LCSV model and (right column) LCSV-H model in comparison with LC model. Calibration period: 1950–1990. LC, Lee–Carter; LC-H, LC model with heteroscedasticity; LCSV, LC stochastic volatility; LCSV-H, LCSV model with heteroscedasticity; CI, confidence interval.

trend in the observed mortality data besides the first few age groups, during the period 1950–1990. Therefore, there is only small difference between the forecasting distributions produced by the LC model and the LCSV model, except for young age groups. Note that there is a clear change of downward trend for some of the middle age groups for the Danish male mortality data as shown in Figure 7. It results in the out-of-sample data falling out of the lower bound of the 95% credible intervals and its consequence for the forecasting of life expectancy will be discussed in section 6.4.2 and generally in section 6.4.3.

By comparing the forecast performance using in-sample data from 1835 to 1990 and from 1950 to 1990 displayed in Figures 6 and 7, we expose the influence that leaving out important historical events, that may affect the mortality rates markedly in a population, can have on the ability to accurately model trend and volatility structures in population dynamics. In particular, we observe that one must be cautious as forecast performance can degrade markedly when important historical events are excluded from the sample as the forecast using data from 1835 to 1990 has clearly outperformed the forecast using only shorter calibration data from 1950 to 1990.

6.4.2. Life expectancy

Using the samples of the forecasted log death rates $y_{x,t}^{(\ell)} = \ln \hat{m}_{x,t}^{(\ell)}$, where $\ell = 1, \dots, L$ and L is the number of MCMC samples, we can obtain the so-called period life expectancy at different ages by constructing an abridged life table, since we use age group data, as follows (Koissi *et al.*, 2006; Yusuf *et al.*, 2014). We consider age group $x \in \{0, 1-4, 5-9, \dots, 95-99\}$ and \tilde{x} is defined as the initial age of age group x , that is $\tilde{x} \in \{0, 1, 5, \dots, 90, 95\}$. Define $n_{\tilde{x}}$ as the length of the interval of age group x (corresponds to \tilde{x}) and hence we have $n_0 = 1, n_1 = 4, n_5 = 5, \dots, n_{95} = 5$. We then calculate the (crude) death probability⁷ that a person aged \tilde{x} in year t will die in the next $n_{\tilde{x}}$ years as

$${}_{n_{\tilde{x}}}\hat{q}_{\tilde{x},t}^{(\ell)} = \frac{n_{\tilde{x}}\hat{m}_{x,t}^{(\ell)}}{1 + n_{\tilde{x}}(1 - a(\tilde{x}, n_{\tilde{x}}))\hat{m}_{x,t}^{(\ell)}} \tag{92}$$

where $a(\tilde{x}, n_{\tilde{x}})$ is the average fraction of the $n_{\tilde{x}}$ years lived by the people who is initially aged \tilde{x} in that interval. Using the assumption that deaths are distributed uniformly in the interval, we set $a(\tilde{x}, n_{\tilde{x}}) = 0.5$ for every \tilde{x} ⁸. The hypothetical number of people alive at age $\tilde{x} + n_{\tilde{x}}$, $l_{\tilde{x}+n_{\tilde{x},t}}^{(\ell)}$, is determined by $l_{\tilde{x}+n_{\tilde{x},t}}^{(\ell)} = l_{\tilde{x},t}^{(\ell)}(1 - n_{\tilde{x}}q_{\tilde{x},t}^{(\ell)})$, where $l_{0,t}^{(\ell)}$ is assumed to be 100,000. We can then calculate the number of deaths $n_{\tilde{x}}d_{\tilde{x},t}^{(\ell)} = l_{\tilde{x},t}^{(\ell)} - l_{\tilde{x}+n_{\tilde{x},t}}^{(\ell)}$ and the person-years lived $n_{\tilde{x}}L_{\tilde{x},t}^{(\ell)} = n_{\tilde{x}}(l_{\tilde{x}+n_{\tilde{x},t}}^{(\ell)} + a(\tilde{x}, n_{\tilde{x}}) \times n_{\tilde{x}}d_{\tilde{x},t}^{(\ell)})$. The total future lifetime of the $l_{\tilde{x},t}^{(\ell)}$ persons who attain age \tilde{x} is $T_{\tilde{x},t}^{(\ell)} = \sum_{i \geq \tilde{x}} n_{\tilde{x}}L_{i,t}^{(\ell)}$, where $i \in \{0, 1, 5, \dots, 90, 95\}$. Finally, a sample of the period life expectancy at age \tilde{x} is obtained as

$$e_{\tilde{x},t}^{(\ell)} = T_{\tilde{x},t}^{(\ell)} / l_{\tilde{x},t}^{(\ell)} \tag{93}$$

and the distributions are obtained in different forecasting year $t = T + k$, where $k \geq 1$.

Remark 6.1 (Period and cohort life expectancy): Period life expectancy assumes there is no trend for future death rates (it is evaluated based on the age-specific death rates in a fixed year t) while cohort life expectancy assumes death rates following the lifetime of a cohort and hence it takes mortality trend into account. For example, to evaluate period life expectancy at age 65 in year t , one needs $\{ {}_5q_{65,t}, {}_5q_{70,t}, \dots, {}_5q_{95,t} \}$ while for cohort life expectancy, $\{ {}_5q_{65,t+5}, {}_5q_{70,t+5}, \dots, {}_5q_{95,t+30} \}$ are used instead. However, the cohort life expectancy for people born in recent years cannot be evaluated

⁷ The death probability is “crude” in the sense that the crude death rate is used for the calculation. For a discussion of crude and true death probabilities, see Dowd *et al.* (2010).

⁸ The Human Mortality Database provides abridged life tables with specific values for $a(\tilde{x}, n_{\tilde{x}})$ in different period. For ease of comparison we use the typical assumption that $a(\tilde{x}, n_{\tilde{x}}) = 0.5$. For the special case when age, instead of age group is considered (i.e. $\tilde{x} = x$), the 1-year death probability $q_{x,t}$ in year t is defined as the ratio of death counts and the population at the beginning of the year. Assuming half of the deaths occurred during the first half of the year, then we have $q_{x,t} := D_{x,t} / (E_{x,t} + 0.5D_{x,t}) = \hat{m}_{x,t} / (1 + 0.5\hat{m}_{x,t})$ where $E_{x,t}$ is the population at the middle of the year. The $n_{\tilde{x}}$ -year death probability (92) for the general case when age group is considered can be derived similarly.

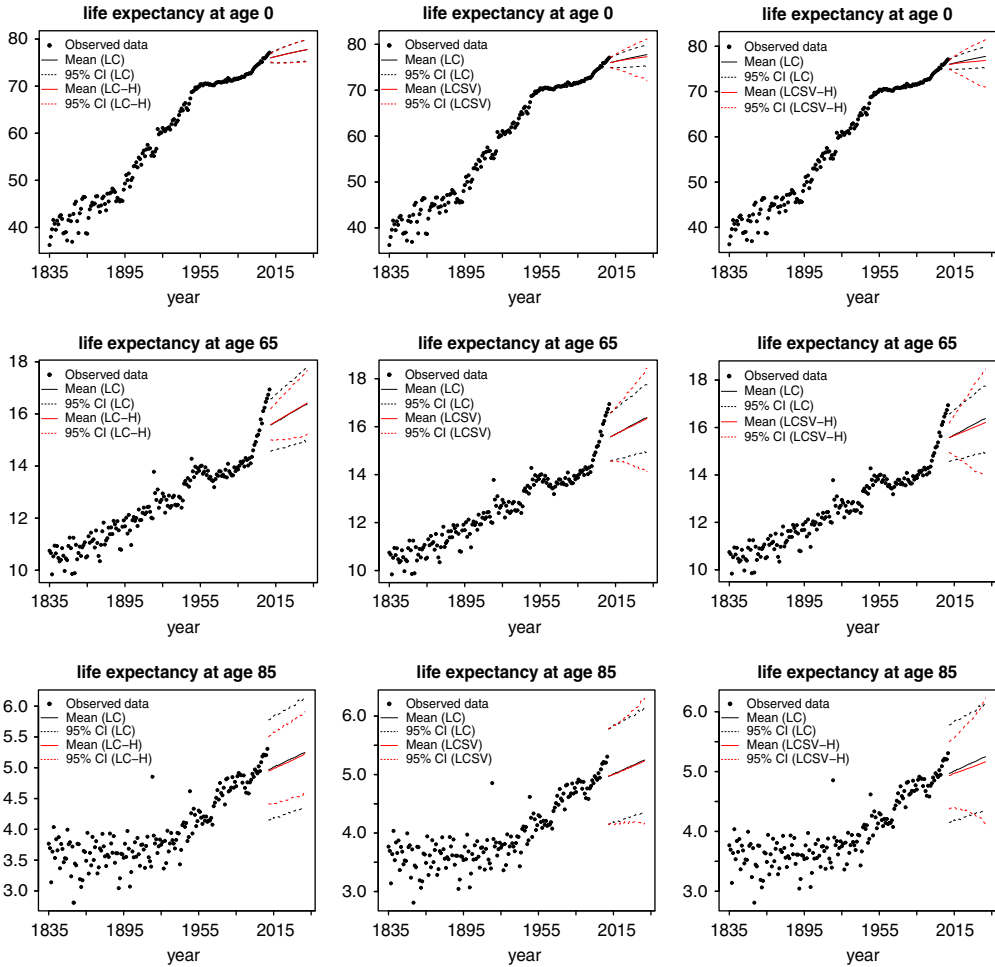


Figure 8. (Colour online) 30-year forecasted life expectancy (2011–2041) at birth, ages 65 and 85 for Danish male population under (left column) LC-H model, (middle column) LCSV model and (right column) LCSV-H model in comparison with LC model. Calibration period: 1835–2010. LC, Lee–Carter; LC-H, LC model with heteroscedasticity; LCSV, LC stochastic volatility; LCSV-H, LCSV model with heteroscedasticity; CI, confidence interval.

using data alone since some of the death rates data are yet to be observed. As a result we focus on period life expectancy so that our forecasts can be compared with the observed data.

Figure 8 shows the forecasted life expectancy at birth, age 65 and age 85 for all the models estimated using data from 1835 to 2010. Interestingly, the forecasted life expectancy at birth is very similar for the LC model and LC-H model for the calibration period 1835–2010. It reflects the fact that forecasting intervals of death rates produced by the LC-H model are wider for some age groups and narrower for other, compared to the LC model, as can be seen from Figure 5. It turns out that these effects almost cancel each other out for this case as death rates are aggregated for all age groups to form the life expectancy at birth distribution. This explanation does not apply to life expectancy at ages 65 and 85, however, since only forecasted death rates for age groups larger

than 65 and 85, respectively, are used to obtain the corresponding life expectancy distribution. As the forecasting intervals of death rates generated by the LC-H model tend to be narrower for old age groups compared with the LC model, the interval for the forecasted life expectancy at ages 65 and 85 distribution produced by the LC-H model is observably narrower than the LC model. Note that we only show four different age groups in Figure 5; projected death rates for other age groups are also relevant but are not displayed due to limited space. In contrast to the LC-H model, we can observe from Figure 5 that LCSV model produces generally wider forecasting intervals for death rates at different age groups compared to the LC model. Consequently, the forecasting intervals for life expectancy at various ages generated by the LCSV model is wider than the LC model, as shown in Figure 8. Similarly to the case of death rates forecasting, the LCSV-H model has both the features of the LC-H and LCSV model in terms of life expectancy prediction.

Results for forecasted life expectancy using data from 1835 to 1990 are shown in Figure 9. It is clear that the forecasting intervals for life expectancy at different ages produced by the LCSV model are narrower compared to the LC model, which can be traced back to the fact that the LCSV model predicts narrower forecasting intervals for death rates at various age groups than the prediction by the LC model, see Figure 6. Apparently the jump-off bias are quite different for the LC-H (LCSV-H) model and the LC model. Note also that the “cancel out” effect of aggregated death rates (Figure 6) for the life expectancy at birth for the LC-H model is less prominent for the calibration period 1835–1990 than the period 1835–2010.

As we use the fitted death rates instead of the observed death rates in the jump-off year (i.e. year 2010), there is a jump-off bias in the forecasted death rates. The forecasted life expectancy at age 65 is particularly sensitive to this jump-off bias. It comes from a sudden decline of death rates for age groups larger than 65 beginning in year 1990, hence a significant increase of life expectancy at age 65 is observed. The jump-off bias is significantly smaller when the calibration period 1835–1990 is considered, see Figure 9. As expected, the forecasted distributions of life expectancy at birth and age 65 are similar for all the models estimated using mortality data from year 1950–1990 (Figure 10). Note that the 95% credible intervals capture poorly the out-of-sample data in this case except for the life expectancy at age 85. It is a consequence of the sudden change of significant downward trend for the death rates observed in the middle age groups of the Danish mortality data starting from around 1990, see Figure 7, as well as the jump-off bias. We discuss about the linear trend assumption and jump-off bias in the next section.

6.4.3. Linear trend assumption and jump-off bias

In performing the forecasting of death rates and life expectancies, we use the models summarised in Table 2 where the LC-H, LCSV and LCSV-H models are variants of the LC model in which a linear trend of the period effect is assumed. However, for the Danish male mortality data that we used, the overall trend is reasonably linear for the whole period 1835–2010, but the same may not be said for the shorter period 1950–2010 as Figures 6 and 7 indicate.

In particular, we observe in Figure 7 that there is a clear change of trend for the death rates of middle age groups. Such a change of trend is difficult, if not impossible, to predict in terms of timing and magnitude. We also perform the analysis on French male mortality data and found similar patterns. For forecasting purpose, one may therefore argue that expert opinion will be an important factor in predicting mortality. Even though any change of mortality trend in the short run cannot be predicted

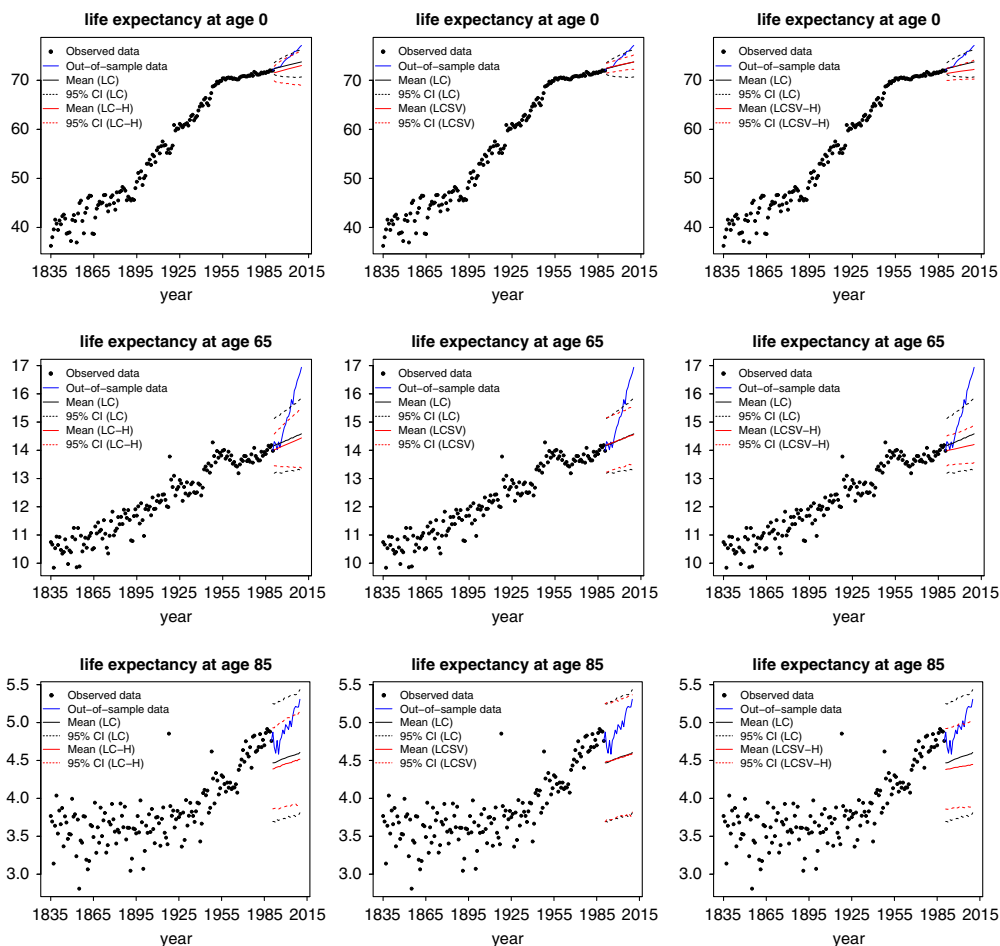


Figure 9. (Colour online) 20-year out-of-sample forecasted life expectancy (1991–2010) at birth, ages 65 and 85 for Danish male population under (left column) LC-H model, (middle column) LCSV model and (right column) LCSV-H model in comparison with LC model. Calibration period: 1835–1990. LC, Lee–Carter; LC-H, LC model with heteroscedasticity; LCSV, LC stochastic volatility; LCSV-H, LCSV model with heteroscedasticity; CI, confidence interval.

with reasonable accuracy using data alone, it can be *detected* if the instantaneous volatility of mortality is quantified. For example, using the LCSV-H model, the log-volatility γ is quantified and we observe from Figure 2 that γ started to increase around 1990 after several decades of declining. The change of volatility level not only affect the prediction intervals as discussed in previous sections, but also indicates that the change of mortality is heightened and one should be cautious whether a change of trend is taking place.

We also find that jump-off bias, discussed in sections 6.4.1–6.4.2, is an important factor in predicting deaths rates and life expectancies. We note that it is straight forward to remove the jump-off bias by adjusting (89)–(91) so that the actual death rates, instead of the fitted death rates, are used in the beginning of the forecasting period. We do not provide the corresponding plots in the paper but one can envisage the results simply by shifting the forecasted distribution so that the forecasted mean is

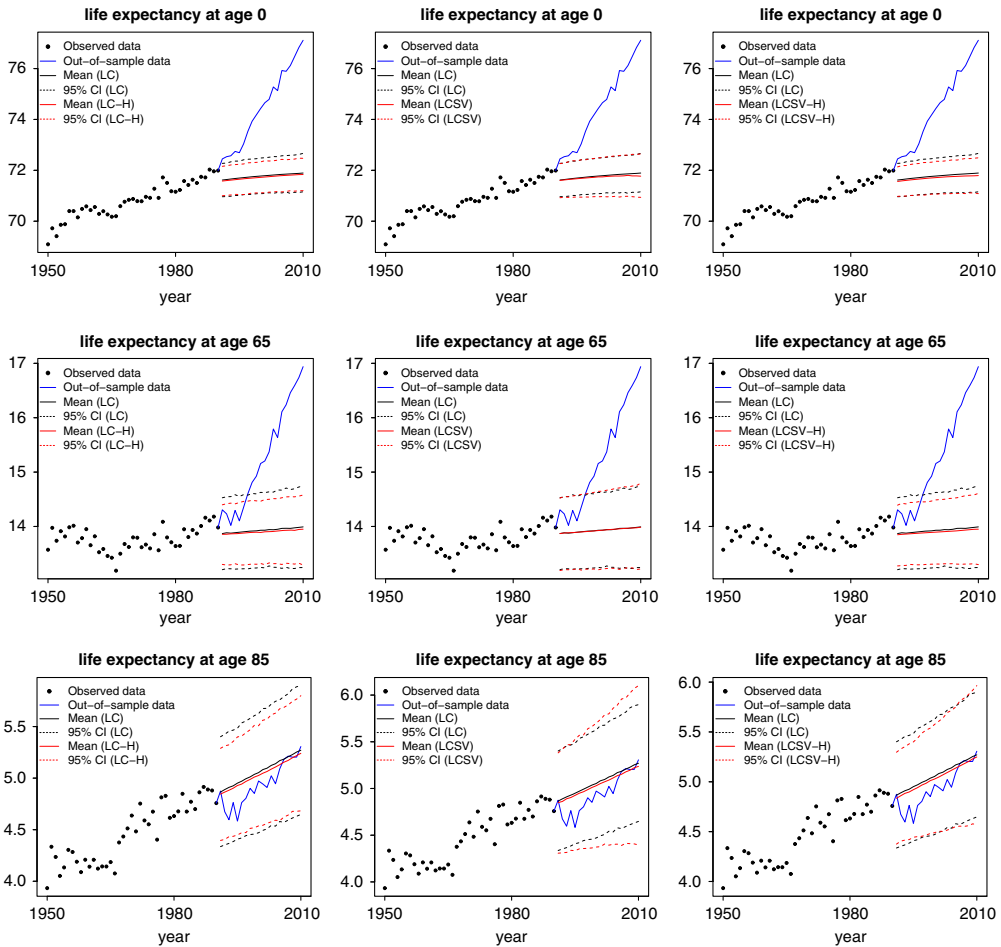


Figure 10. (Colour online) 20-year out-of-sample forecasted life expectancy (1991–2010) at birth, ages 65 and 85 for Danish male population under (left column) LC-H model, (middle column) LCSV model and (right column) LCSV-H model in comparison with LC model. Calibration period: 1950–1990. LC, Lee–Carter; LC-H, LC model with heteroscedasticity; LCSV, LC stochastic volatility; LCSV-H, LCSV model with heteroscedasticity; CI, confidence interval.

attached to the (in-sample) data at the end of the estimation year. Removing the jump-off bias will have a significant impact on the accuracy of mortality forecasting especially when data exhibit clear trending.

7. Concluding Remarks

We developed and presented a comprehensive state-space framework for stochastic mortality modelling. The state-space approach has two key advantages. First, it puts modelling, estimation and forecasting of mortality in a unified framework in contrast to common practice in this area. Second, the methodology permits realistic and sophisticated mortality models to be estimated and forecasted, which could be difficult to handled using other approaches.

We show that many of the popular mortality models exist in the literature can be cast in state-space form. We then suggest several classes of mortality models that can be classified as linear Gaussian state-space models and non-linear/non-Gaussian state-space models. Our proposals are not exhaustive but aim to illustrate the flexibility of the methodology. In particular, we incorporate heteroscedasticity and stochastic volatility in mortality modelling, as an examination of mortality data suggests that volatility of death rates is not constant in the age and time dimension over a long time period. Moreover, we propose an alternative identification constraint for the LC type modelling, which is tailored for the state-space approach.

Frequestist state-space inference for stochastic mortality models is carried out and explained based on the gradient and Hessian of the marginalised likelihood developed recently in statistics literature. We also utilise a modern approach to Bayesian inference for state-space modelling hinged on the PMCMC framework. In particular, we develop a sampler using a combination of Rao-Blackwellised Kalman filter and particle filter for the latent state process full posterior conditionals, combined with Gibbs sampling steps for the static model parameters to estimate a stochastic volatility model for mortality proposed in this paper.

Using mortality data of Danish male population, we assess the extended models based on deviance conditional criterion. It is found that incorporating heteroscedasticity is a crucial improvement factor in model fitting, while model complexity is accounted for. The incorporation of stochastic volatility clearly enhances model performance for fitting of long-term mortality time series. Estimation results for long calibration period support the assumption of stochastic volatility. We show that forecasting can be carried out straightforwardly in state-space framework under a Bayesian setting. We examine the forecasting properties of the models using different calibration periods. The inclusion of heteroscedasticity and stochastic volatility substantially affects prediction intervals of death rate and life expectancy distributions. The linear trend assumption commonly found in mortality modelling and jump-off bias are discussed in light of the Danish mortality data.

State-space framework provides attractive features that are of importance to mortality modelling. The methods and results developed and shown in the paper will have significant implications for longevity risk management in actuarial applications. In particular, we anticipate the models and methods introduced in this paper can be employed to longevity risk applications such as the state-space longevity hedging methods proposed in Liu & Li (2016a) and Liu & Li (2016b) which is a topic of future research.

Acknowledgements

This research was supported by the CSIRO-Monash Superannuation Research Cluster, a collaboration among CSIRO, Monash University, Griffith University, the University of Western Australia, the University of Warwick and stakeholders of the retirement system in the interest of better outcomes for all. This research was also partially supported under the Australian Research Council's Discovery Projects funding scheme (project number: DP160103489). The authors also acknowledge project "Research on Urban Intelligence" from Research Organization of Information and Systems in Japan.

References

- Analitis, A., Katsouyanni, K., Biggeri, A., Baccini, M., Forsberg, B., Bisanti, L., Kirchmayer, U., Ballester, F., Cadum, E., Goodman, P., Hojs, A., Sunyer, J., Tittanen, P. & Michelozzi, P. (2008). Effects of cold weather on mortality: results from 15 European cities within the PHEWE project. *American Journal of Epidemiology*, **168**(12), 1397–1408.

- Andreev, K.F. (2002). *Evolution of the Danish Population from 1835 to 2000*. Odense University Press, Denmark.
- Andrieu, C., Doucet, A. & Holenstein, R. (2010). Particle Markov chain Monte Carlo methods. *Journal of Royal Statistical Society Series B*, 72, 269–342.
- Bell, W.R. (1997). Comparing and assessing time series methods for forecasting age-specific fertility and mortality rates. *Journal of Official Statistics*, 13(3), 279–303.
- Berg, A., Meyer, R. & Yu, J. (2004). Deviance information criterion for comparing stochastic volatility models. *Journal of Business and Economic Statistics*, 22(1), 107–120.
- Biffis, E. (2005). Affine processes for dynamic mortality and actuarial valuations. *Insurance: Mathematics and Economics*, 37(3), 443–468.
- Brouhns, N., Denuit, M. & Vermunt, J.K. (2002). A Poisson log-bilinear regression approach to the construction of projected lifetables. *Insurance: Mathematics and Economics*, 31, 373–393.
- Cairns, A., Blake, D. & 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., Blake, D. & Dowd, K. (2008). Modelling and management of mortality risk: a review. *Scandinavian Actuarial Journal*, 2(3), 79–113.
- Cairns, A., Blake, D., Dowd, K., Coughlan, G., Epstein, D., Ong, A. & 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.
- Cairns, A., Blake, D., Dowd, K., Coughlan, G. & Khalaf-Allah, M. (2011). Bayesian stochastic mortality modeling for two populations. *ASTIN Bulletin*, 41(1), 29–59.
- Carter, C.K. & Kohn, R. (1994). On Gibbs sampling for state-space models. *Biometrika*, 81(3), 541–553.
- Celeux, C., Forbes, F., Robert, C.P. & Titterton, D.M. (2006). Deviance information criteria for missing data models. *Bayesian Analysis*, 1(4), 651–674.
- Chib, S., Nardari, F. & Shephard, N. (2002). Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics*, 108, 281–316.
- Chopin, N., Jacob, P.E. & Papaspiliopoulos, O. (2013). SMC²: an efficient algorithm for sequential analysis of state-space models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(3), 397–426.
- Currie, I.D. (2009). Smoothing and forecasting mortality rates with P-splines. Available online at the address <http://www.ma.hw.ac.uk/~iain/research.talks.html> (accessed on 20th August 2015).
- Czado, C., Delwarde, A. & Denuit, M. (2005). Bayesian Poisson log-bilinear mortality projections. *Insurance: Mathematics and Economics*, 36, 260–284.
- Dawood, F.S., Iuliano, A.D., Reed, C., Meltzer, M.I., Shay, D.K., Cheng, P.-Y., Bandaranayake, D., Breiman, R.F., Brooks, W.A., Buchy, P., Feikin, D.R., Fowler, K.B., Gordon, A., Hien, N.T., Horby, P., Huang, Q.S., Katz, M.A., Krishnan, A., Lal, R., Montgomery, J.M., Molbak, K., Pebody, R., Presanis, A.M., Razuri, H., Steens, A., Tinoco, Y.O., Wallinga, J., Yu, H., Vong, S., Bresee, K. & Widdowson, M.A. (2012). Estimated global mortality associated with the first 12 months of 2009 pandemic influenza A H1N1 virus circulation: a modelling study. *The Lancet Infectious Diseases*, 12(9), 687–695.
- De Jong, P. & Tickle, L. (2006). Extending the Lee-Carter mortality forecasting. *Mathematical Population Studies*, 13, 1–18.
- Del Moral, P. (2004). *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer-Verlag, New York.
- Dickson, D., Hardy, M. & Waters, H.R. (2009). *Actuarial Mathematics for Life Contingent Risks*. Cambridge University Press.
- Doucet, A., De Freitas, J.F.G. & Gordon, N.J. (Eds.) (2001). *Sequential Monte Carlo Methods in Practice*. Springer, New York.

- Doucet, A., Godsill, S. & Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, **10**(3), 197–208.
- Dowd, K., Cairns, A., Blake, D., Coughlan, G., Epstein, D. & Khalaf-Allah, M. (2010). Evaluating the goodness of fit of stochastic mortality models. *Insurance: Mathematics and Economics*, **47**, 255–265.
- Eloranta, S., Lambert, P.C., Andersson, T.M., Czene, K., Hall, P., Björkholm, M. & Dickman, P.W. (2012). Partitioning of excess mortality in population-based cancer patient survival studies using flexible parametric survival models. *BMC Medical Research Methodology*, **12**(1), 1.
- England, P.D. & Haberman, S. (1993). A new approach to modeling excess mortality. *Journal of Actuarial Practice*, **1**, 85–117.
- Flury, T. & Shephard, N. (2011). Bayesian inference based only on simulated likelihood: particle filter analysis of dynamic economic models. *Econometric Theory*, **27**, 933–956.
- Fouillet, A., Rey, G., Laurent, F., Pavillon, G., Bellec, S., Guihenneuc-Jouyau, C., Clavel, J., Jougl, E. & Hémon, D. (2006). Excess mortality related to the August 2003 heat wave in France. *International Archives of Occupational and Environmental Health*, **80**(1), 16–24.
- Fung, M.C., Peters, G.W. & Shevchenko, P.V. (2017). Cohort effects in mortality modelling: a Bayesian state-space analysis. Preprint, available at SSRN: 2907868.
- Golightly, A. & Wilkinson, D.J. (2011). Bayesian parameter inference for stochastic biochemical network models using particle MCMC. *Interface Focus*, **1**(6), 807–820.
- Harvey, A.C. (1989). *Forecasting: Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Hirz, J., Schmock, U. & Shevchenko, P.V. (2017). Actuarial applications and estimation of extended CreditRisk+. *Risks* **5**, 1–23.
- Hunt, A. & Villegas, A.M. (2015). Robustness and convergence in the Lee-Carter model with cohort effects. *Insurance: Mathematics and Economics*, **64**, 186–202.
- Kalman, R.E. (1960). A new approach to linear filtering and prediction problems. *Transactions of the ASME – Journal of Basic Engineering*, **82**(Series D), 35–45.
- Kim, S., Shephard, N. & Chib, S. (1998). Stochastic volatility: likelihood inference and comparison of ARCH models. *Review of Economic Studies*, **65**, 361–393.
- Kogure, A., Kitsukawa, K. & Kurachi, Y. (2009). A Bayesian comparison of models for changing mortalities toward evaluating longevity risk in Japan. *Asia-Pacific Journal of Risk and Insurance*, **3**(2), <https://doi.org/10.2202/2153-3792.1036>.
- Kogure, A. & Kurachi, Y. (2010). A Bayesian approach to pricing longevity risk based on risk-neutral predictive distributions. *Insurance: Mathematics and Economics*, **46**, 162–172.
- Koissi, M., Shapiro, A.F. & Hognas, Goran (2006). Evaluating and extending Lee-Carter model for mortality forecasting: bootstrap confidence interval. *Insurance: Mathematics and Economics*, **38**, 1–20.
- Lee, R.D. & Carter, L.R. (1992). Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, **87**, 659–675.
- Lee, R.D. & Miller, T. (2001). Evaluating the performance of the Lee-Carter method for forecasting mortality. *Demography*, **38**(4), 537–549.
- Li, J., Chan, W. & Cheung, S. (2011). Structural changes in the Lee-Carter mortality indexes. *North American Actuarial Journal*, **15**(1), 13–31.
- Liu, J.S. (2008). *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag, New York.
- Liu, Y. & Li, J.S.-H. (2016a). It's all in the hidden states: a longevity hedging strategy with an explicit measure of population basis risk. *Insurance: Mathematics and Economics*, **70**, 301–319.
- Liu, Y. & Li, J.S.-H. (2016b). The locally linear Cairns-Blake-Dowd model: a note on delta-nuga hedging of longevity risk. *ASTIN Bulletin*, **47**(1), 79–151, 1–73.

- Pedroza, C. (2006). A Bayesian forecasting model: predicting U.S. male mortality. *Biostatistics*, 7(4), 530–550.
- Peters, G.W., Briers, M., Shevchenko, P.V. & Doucet, A. (2013). Calibration and filtering for multi factor commodity models with seasonality: incorporating panel data from futures contracts. *Methodology and Computing in Applied Probability*, 15(4), 841–874.
- Peters, G.W., Fan, Y. & Sisson, S.A. (2012). On sequential Monte Carlo, partial rejection control and approximate Bayesian computation. *Statistics and Computing*, 22(6), 1209–1222.
- Peters, G.W., Fung, M.C. & Shevchenko, P.V. (2017). Addressing identification problems in MCMC estimation of stochastic mortality models. Working paper.
- Peters, G.W., Hosack, G.R. & Hayes, K.R. (2010a). Ecological non-linear state-space model selection via adaptive particle Markov chain Monte Carlo (AdPMCMC). Available at arXiv preprint, available at arXiv:1005.2238.
- Peters, G.W., Wüthrich, M.V. & Shevchenko, P.V. (2010b). Chain ladder method: Bayesian bootstrap versus classical bootstrap. *Insurance: Mathematics and Economics*, 47(1), 36–51.
- Petris, G., Petrone, S. & Campagnoli, P. (2009). *Dynamic Linear Models with R*. Springer-Verlag, New York.
- Pitacco, E., Denuit, M., Haberman, S. & Olivieri, A. (2009). *Modelling Longevity Dynamics for Pensions and Annuity Business*. Oxford University Press.
- Plat, R. (2009). On stochastic mortality modeling. *Insurance: Mathematics and Economics*, 45, 393–404.
- Poyiadjis, G., Doucet, A. & Singh, S.S. (2005). Maximum likelihood parameter estimation in general state-space models using particle methods. In Proceedings of the American Statistical Association, Minneapolis, USA. August.
- Poyiadjis, G., Doucet, A. & Singh, S.S. (2011). Particle approximations of the score and observed information matrix in state-space models with application to parameter estimation. *Biometrika*, 98(1), 65–80.
- Renshaw, A. & Haberman, S. (2003). Lee-Carter mortality forecasting with age-specific enhancement. *Insurance: Mathematics and Economics*, 33, 255–272.
- Renshaw, A. & Haberman, S. (2006). A cohort-based extension to the Lee-Carter model for mortality reduction factors. *Insurance: Mathematics and Economics*, 38, 556–570.
- Shang, H.L., Booth, H. & Hyndman, R.J. (2011). Point and interval forecasts of mortality rates and life expectancy: a comparison of ten principal component methods. *Demography*, 25(5), 173–214.
- Spiegelhalter, D.J., Best, N.G., Carlin, B.P. & van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society, Series B*, 64, 583–639.
- Targino, R.S., Peters, G.W. & Shevchenko, P.V. (2015). Sequential Monte Carlo samplers for capital allocation under copula-dependent risk models. *Insurance: Mathematics and Economics*, 61, 206–226.
- Tsay, R.S. & Tiao, G.C. (1984). Consistent estimates of autoregressive parameters and extended sample autocorrelation function for stationary and nonstationary ARMA models. *Journal of the American Statistical Association*, 79(385), 84–96.
- van Berkum, F., Antonio, K. & Vellekoop, M. (2016). The impact of multiple structural changes on mortality predictions. *Scandinavian Actuarial Journal*, 2016(7), 581–603.
- West, M. & Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models*. Springer Series in Statistics.
- Yusuf, F., Martins, J.M. & Swanson, D.A. (2014). *Methods of Demographic Analysis*. Springer, Netherlands.
- Zucs, P., Buchholz, U., Haas, W. & Uphoff, H. (2005). Influenza associated excess mortality in Germany, 1985–2001. *Emerging Themes in Epidemiology*, 2(1), 1.

Appendix A: Differentiation Matrices in Gradient-Based Estimation

For the LC-H model, the parameter vector is denoted by $\boldsymbol{\psi} = (\boldsymbol{\alpha}_{x_2:x_p}, \boldsymbol{\beta}_{x_2:x_p}, \sigma_{\varepsilon, x_1:x_p}^2, \theta, \sigma_{\omega}^2)$ with dimension $n = 3p$, where p is the number of age group considered. We are required to evaluate $\frac{\partial \boldsymbol{\alpha}}{\partial \boldsymbol{\psi}_i}$, $\frac{\partial \boldsymbol{\beta}}{\partial \boldsymbol{\psi}_i}$, $\frac{\partial \Sigma}{\partial \boldsymbol{\psi}_i}$ and $\frac{\partial \sigma_{\omega}^2}{\partial \boldsymbol{\psi}_i}$ in the gradient-based estimation (section 4.2). Define

$$\boldsymbol{\psi}_{\alpha} := \boldsymbol{\psi}_{1:p-1} = \boldsymbol{\alpha}_{x_2:x_p} \quad (:= \boldsymbol{\alpha}_{-x_1}), \quad \boldsymbol{\psi}_{\beta} := \boldsymbol{\psi}_{p:2p-2} = \boldsymbol{\beta}_{x_2:x_p} \quad (:= \boldsymbol{\beta}_{-x_1})$$

$$\boldsymbol{\psi}_{\sigma_{\varepsilon}^2} := \boldsymbol{\psi}_{2p-1:3p-2} = \sigma_{\varepsilon, x_1:x_p}^2, \quad \boldsymbol{\psi}_{\theta} := \boldsymbol{\psi}_{3p-1} = \theta, \quad \boldsymbol{\psi}_{\sigma_{\omega}^2} := \boldsymbol{\psi}_{3p} = \sigma_{\omega}^2$$

Then we have

$$\frac{\partial (\boldsymbol{\alpha}_{-x_1})_j}{\partial (\boldsymbol{\psi}_{\alpha})_i} = \delta_{ij} = \frac{\partial (\boldsymbol{\beta}_{-x_1})_j}{\partial (\boldsymbol{\psi}_{\beta})_i}, \quad i, j = 1, \dots, p-1$$

$$\frac{\partial (\Sigma)_{jj}}{\partial (\boldsymbol{\psi}_{\sigma_{\varepsilon}^2})_i} = \delta_{ij}, \quad i, j = 1, \dots, p$$

$$\frac{\partial \theta}{\partial \boldsymbol{\psi}_{\theta}} = 1 = \frac{\partial \sigma_{\omega}^2}{\partial \boldsymbol{\psi}_{\sigma_{\omega}^2}}$$

where $\delta_{ij} = 1$ if $j = i$ and zero otherwise; Σ is a diagonal matrix with diagonal $\sigma_{\varepsilon, x_1:x_p}^2$. Note that $\frac{\partial (\boldsymbol{\alpha})_1}{\partial (\boldsymbol{\psi}_{\alpha})_i} = \frac{\partial (\boldsymbol{\beta})_1}{\partial (\boldsymbol{\psi}_{\beta})_i} = 0$ for $i = 1, \dots, p-1$, where $\boldsymbol{\alpha} = \boldsymbol{\alpha}_{x_1:x_p}$ and $\boldsymbol{\beta} = \boldsymbol{\beta}_{x_1:x_p}$.

Appendix B: A Review of SMC Method

SMC, also known as particle filtering, can be viewed as a generalisation of Kalman filtering in state-space modelling context. The method is based on importance sampling and it has become an essential sampling-based tool in many domains (Doucet *et al.*, 2001). In the following, we give a brief review of the method using the LCSV model, (23b)–(23c), as an example to derive a basic particle filtering algorithm. Our target density is the joint posterior distribution of the states for stochastic volatility:

$$\pi(\gamma_{1:t} | \boldsymbol{\kappa}_{0:t}) \tag{94}$$

where the parameters of the model are assumed to be known and is suppressed here for ease of notation. To apply importance sampling, we first calculate

$$\begin{aligned} \pi(\gamma_{1:t} | \boldsymbol{\kappa}_{0:t}) &= \frac{\pi(\boldsymbol{\kappa}_t | \gamma_{1:t}, \boldsymbol{\kappa}_{0:t-1}) \pi(\gamma_{1:t} | \boldsymbol{\kappa}_{0:t-1})}{\pi(\boldsymbol{\kappa}_t | \boldsymbol{\kappa}_{0:t-1})} \\ &= \frac{\pi(\boldsymbol{\kappa}_t | \gamma_{1:t}, \boldsymbol{\kappa}_{0:t-1}) \pi(\gamma_t | \gamma_{1:t-1}, \boldsymbol{\kappa}_{0:t-1})}{\pi(\boldsymbol{\kappa}_t | \boldsymbol{\kappa}_{0:t-1})} \pi(\gamma_{1:t-1} | \boldsymbol{\kappa}_{0:t-1}) \\ &= \frac{\pi(\boldsymbol{\kappa}_t | \gamma_t, \boldsymbol{\kappa}_{t-1}) \pi(\gamma_t | \gamma_{t-1})}{\pi(\boldsymbol{\kappa}_t | \boldsymbol{\kappa}_{0:t-1})} \pi(\gamma_{1:t-1} | \boldsymbol{\kappa}_{0:t-1}) \end{aligned} \tag{95}$$

The importance density is assumed to satisfy

$$g_{1:t}(\gamma_{1:t} | \boldsymbol{\kappa}_{0:t}) := g_t(\gamma_t | \gamma_{1:t-1}, \boldsymbol{\kappa}_{0:t}) g_{1:t-1}(\gamma_{1:t-1} | \boldsymbol{\kappa}_{0:t-1}) \tag{96}$$

and the importance weight is given by

$$\begin{aligned} \tilde{w}_t &= \frac{\pi(\kappa_t | \gamma_t, \kappa_{t-1}) \pi(\gamma_t | \gamma_{t-1})}{\pi(\kappa_t | \kappa_{t-1}) g_t(\gamma_t | \gamma_{1:t-1}, \kappa_{0:t})} \frac{\pi(\gamma_{1:t-1} | \kappa_{0:t-1})}{g_{0:t-1}(\gamma_{1:t-1} | \kappa_{0:t-1})} \\ &\propto \frac{\pi(\kappa_t | \gamma_t, \kappa_{t-1}) \pi(\gamma_t | \gamma_{t-1})}{g_t(\gamma_t | \gamma_{1:t-1}, \kappa_{0:t})} \tilde{w}_{t-1} \\ &:= \hat{w}_t \tilde{w}_{t-1} \end{aligned} \tag{97}$$

where \hat{w}_t is called the incremental importance weight. The normalised importance weights are then obtained as $w_t^{(i)} := \tilde{w}_t^{(i)} / \sum_{j=1}^N \tilde{w}_t^{(j)}$. To summarise, suppose we have N particle paths $(\gamma_{1:t-1}^{(i)}, w_{t-1}^{(i)})_{i=1}^N$ to approximate the density $\pi(\gamma_{1:t-1} | \kappa_{0:t-1})$ at time $t-1$. Then, from (96), the i th particle path at time t is given by $\gamma_{1:t}^{(i)} = (\gamma_{1:t-1}^{(i)}, \gamma_t^{(i)})$, where $\gamma_t^{(i)}$ is sampled from $g_t(\gamma_t | \gamma_{1:t-1}^{(i)}, \kappa_{0:t})$. The target density $\pi(\gamma_{1:t} | \kappa_{0:t})$ is approximated by $(\gamma_{1:t}^{(i)}, w_t^{(i)})_{i=1}^N$, where the normalised weight $w_t^{(i)}$ is obtained from (97) and normalisation is carried out.

The problem of degeneracy, that is a majority of the particle paths may have negligible weight, can be handled by resampling. Specifically, we define the so-called effective sample size

$$N_{eff} := \left(\sum_{i=1}^N (w_t^{(i)})^2 \right)^{-1} \tag{98}$$

At each time t , if N_{eff} is smaller than some threshold (e.g. 80% of N) then we draw N samples (denoted by $N(i)$, $i = 1, \dots, N$) from a multinomial distribution with probability weights $w_t^{(i)}$, $i = 1, \dots, N$, and replace the particle paths $\gamma_{1:t}^{(i)}$ by $\gamma_{1:t}^{(N(i))}$, and set $w_t^{(i)} = 1/N$. The resampling step allows to keep the particle paths in proportion to their weights and tend to discard those that have negligible weights.