**Title:** Confronting the possibility of catastrophe: Identification, characterization, and simulation of mortality shocks via Hidden Markov Lee-Carter models

**Authors:** Jonas Schöley[1]

**Affiliations:**
[1] Max Planck Institute for Demographic Research; Rostock, Germany

**Corresponding authors:**
Jonas Schöley `schoeley@demogr.mpg.de`

**Keywords:** excess deaths; mortality shocks; Lee-Carter; forecasting; time-series; outlier detection

**Abstract:** I demonstrate a hidden Markov extension to the Lee-Carter model which allows for the analysis of past mortality shocks and their consideration in forecasts of future mortality. The model is capable of automatically and probabilistically identifying past periods of crisis mortality, and characterises these crisis via parameters of substantial demographic interest: the crisis age profile, the annual probability of entering or exiting a crisis, the historical distribution of crisis magnitudes, and excess mortality over the Lexis surface. Simultaneously, a normal mortality process, purged of crises, is estimated. The model naturally allows for stochastic forecasts of age specific mortality with the possibility of crises. Crises parameters are fully probabilistic and reflect estimation uncertainty. Alternatively, for scenario based forecasts, the parameters can be pre-specified.

# Introduction

Optimism is prevalent in the field of mortality forecasting as evident from the customary treatment of period shocks as outliers to be purged from the training data and ignored in the forecast. However, the COVID-19 pandemic and the resurgence of violent inter-state conflict has been a forceful reminder that the future is not “normal”, but holds the real potential for substantial deviations from stability. This paper concerns the modeling and stochastic forecasting of such rare and stark deviations – of mortality shocks – within the framework of a Hidden-MArkov-Lee-Carter (HIMALC) model.

The central idea of the HIMALC model is to separate the time evolution of death rates into crisis and non-crisis components. While the classical Lee-Carter drifted random walk remains the description for mortality dynamics during “normal” times, it is augmented by a separate process which produces mortality shocks during crises periods. A hidden Markov chain distinguishes between crisis and non-crisis periods.

As an analytical and predictive tool, a shock extension to the Lee-Carter model should ideally allow for the separation of the Lexis surface into regular mortality and excess mortality; allow for the automatic identification of past crisis periods wrt. timing, magnitude, length, significance, and crisis mortality age profile; and allow for stochastic forecasts of both future age specific mortality both with and without shocks. This work builds on past research on shock extensions to the Lee-Carter model, combining ideas to specify a model which meets the wish-list outlined above (see Table 1).

**Table 1:** Comparison of approaches for dealing with mortality shocks in Lee-Carter models.

| | | | Estimate shock | | | | | Forecast | |
|---|---|---|---|---|---|---|---|---|---|
| | Non-shock Lexis | Excess Lexis | Timing | Magnitude | Length | Probability | Age-profile | with shocks | without shocks |
| Shock removal | + | - | - | - | - | - | - | - | + |
| Shock indicator | + | + | - | + | - | + | - | + | + |
| Heavy tailed | - | - | - | - | - | + | - | + | - |
| Vanishing jumps | + | + | + | + | - | + | + | + | + |
| HIMALC | + | + | + | + | + | + | + | + | + |

The most basic and widely used way to deal with crisis periods in a Lee-Carter framework is to manually remove them from the fitting data, either year by year, or by only selecting time intervals assumed free of wars and pandemics. This, in effect, conditions the estimated Lee-Carter parameters to reflect mortality conditions during non-crisis periods allowing for forecasts of “normal” mortality. No inference about mortality shocks is done – as far as the fitting data is concerned, they don’t exist.

Similar in effect to exclusion, is the manual marking of crisis years with dummy coefficients [9]. This procedure still requires the manual identification of crisis periods, but the estimated coefficients allow inference about the shock magnitudes and significance. In principle, a separation of the fitted Lexis surface into normal and crisis mortality is possible as well. Forecasts with shocks can only be performed in an ad-hock fashion, e.g. by stochastic re-sampling of the shock in the forecast with the same annual probability as manually identified in the data [8].

Based on the realization that mortality shocks are non-normal and thus violate the assumption of the random-walk with drift for the mortality index, various proposals have been made to use heavy tailed distributions for the evolution of the mortality level over time [11, 5, 6]. While the heavy tailed approach allows for drastic jumps in the mortality index it misses two important features of mortality shocks: shocks have a different age pattern from normal mortality and shocks are serially correlated with returns to normal regimes.

Jump processes allow for the time evolution of mortality to exhibit stochastic discontinuities [3, 10, 4]. If the jumps are assumed to be independent, reversals to normal mortality after a shock must be specified in a separate process, e.g. via a drift intensity which depends on the deviation from a target mortality level [3], via a jump effect which exponentially vanishes over time [4]. By specifying a two-state $\{\text{normal}, \text{crisis}\}$ hidden Markov model for the shocks, crisis entry and exit are embedded within the same process, with explicit probabilities for entry and exit, which allows for the straightforward simulation crises of varying length. I follow Liu and Li [10] in allowing for mortality shocks to have an age profile distinct from changes in mortality during normal periods.

The HIMALC model is closely related to the "Switching Regime" extension to the Lee-Carter model [7]. Hainaut [7] demonstrated a two state hidden Markov chain with age interaction but did not apply the model to mortality shocks, instead focusing on the identification of varying mortality improvement regimes.

## Model specification

I express the HIMALC model within a Bayesian hierarchical model framework where a age-period surface of death counts, $D_{x,t}$, over single ages $x$ and years $t$ is generated by

$$D_{x,t} \sim \text{NegativeBinomial}(\lambda_{x,t} E_{x,t}, \phi),$$

where $\lambda_{x,t}$ are age and time specific death rates and $E_{x,t}$ the corresponding person-years of exposure. The $\phi$ parameter controls the degree of overdispersion. In the classical Lee-Carter model death rates are expressed as $\lambda_{x,t} = \exp(\alpha_x + \beta_x \kappa_t)$, where the an overall age effect $\alpha_x$ is augmented by the age-time interaction $\beta_x \kappa_t$, with $\kappa_t$ being a time-index of mortality, and $\beta_x$ indicating the age specific sensitivity of mortality to period changes.

We augment the classical Lee-Carter specification by the shock-component $\gamma_x \eta_t$, yielding

$$\lambda_{x,t} = \exp(\alpha_x + \beta_x \kappa_t + \gamma_x \eta_t).$$

The *shock-index* $\eta_t$ allows for a proportional shift in the overall mortality profile at time $t$ by factor $\exp(\eta_t)$, distributed over ages via the shock-profile $\gamma_x$.

Crucially, $\eta_t$ is the realization of a Hidden Markov process

$$\eta_t = Z_t U_t,$$

with $Z_t$ a random variable over the binary state space $S = \{\text{crisis}(1), \text{noncrisis}(0)\}$ and $U$ a random variable indicating the magnitude of the shock. We specify

$$Z_t \sim \text{Bernoulli}(p_t),$$

with the crisis probability $p_t$ linked to the Markov-Chain transition probabilities via

$$p_t = z_{t-1} \theta_{1 \longrightarrow 1} + (1 - z_{t-1}) \theta_{0 \longrightarrow 1}.$$

Thus, HIMALC model estimates the probability of the next year being a crisis, given no current crisis, $\theta_{0 \longrightarrow 1}$, and conversely, the probability of remaining in an existing crisis, $\theta_{1 \longrightarrow 1}$.

I model the distribution of mortality shocks as

$$U_t \sim \text{Gamma}(a, b).$$

The requirement of every mortality shock increasing death rates is fulfilled by the strictly positive support of the Gamma distribution.

The Lee-Carter specification is completed by expressing the time-evolution of mortality under non-crisis conditions as a random walk with drift

$$\kappa_t = \delta + \kappa_{t-1} + \epsilon_t, \text{ and } \epsilon_t \sim \text{Normal}(0, \sigma^2).$$

## Model demonstration

I fit the HIMALC model to HMD data for England & Wales [2] using Markov Chain Monte Carlo via the PyMC framework for Bayesian inference [1]. A proper prior setup is crucial and will be explained in the full draft.

Figure 1 shows the model estimates. Panel A exhibits the Lexis surface of mortality rates as predicted from the model (not to be confused with the fitting data). The predicted surface clearly features the 1918/19 Spanish flu and the conflict related deaths of both world wars. These crises were not manually coded into the model – they have been identified by the hidden Markov process as crises periods. Panel B shows the surface of mortality shocks by period and age in terms of excess rate ratios. These are estimated as part of the model fitting via the product $\exp(\gamma_x \eta_t)$. During the intersection of the 1918 influenza pandemic with the fighting of World War 1, mortality rates among young men were up to 17 times elevated from the normal levels. The joint dynamics of the normal mortality index $k_t$ and the shock mortality index $\eta_t$ can be seen in panel C. Mortality shocks are visible in $\eta_t$ whereas the Lee-Carter random walk shows no signs of shocks and is on a downward trajectory throughout the training period. The distribution of shocks is plotted in panel D. This distribution allows for probabilistic statements on the magnitude and frequency of shocks. Panel E shows the estimated age profile of the mortality shocks, which is concentrated on ages below 50 peaking at age 25. This pattern is typical for excess deaths during both world wars. An extension of the model will alow for the age profile to vary for different types of crises.

While panels A to E show the potential of the HIMALC model for analysis of mortality crises, panel F shows stochastic life expectancy forecasts, which reproduce the magnitude, age pattern and relative frequency of mortality shocks identified in the fitting data.

**Figure 1:** Hidden Markov Lee-Carter Model estimates for male mortality in England and Wales (Data: HMD).

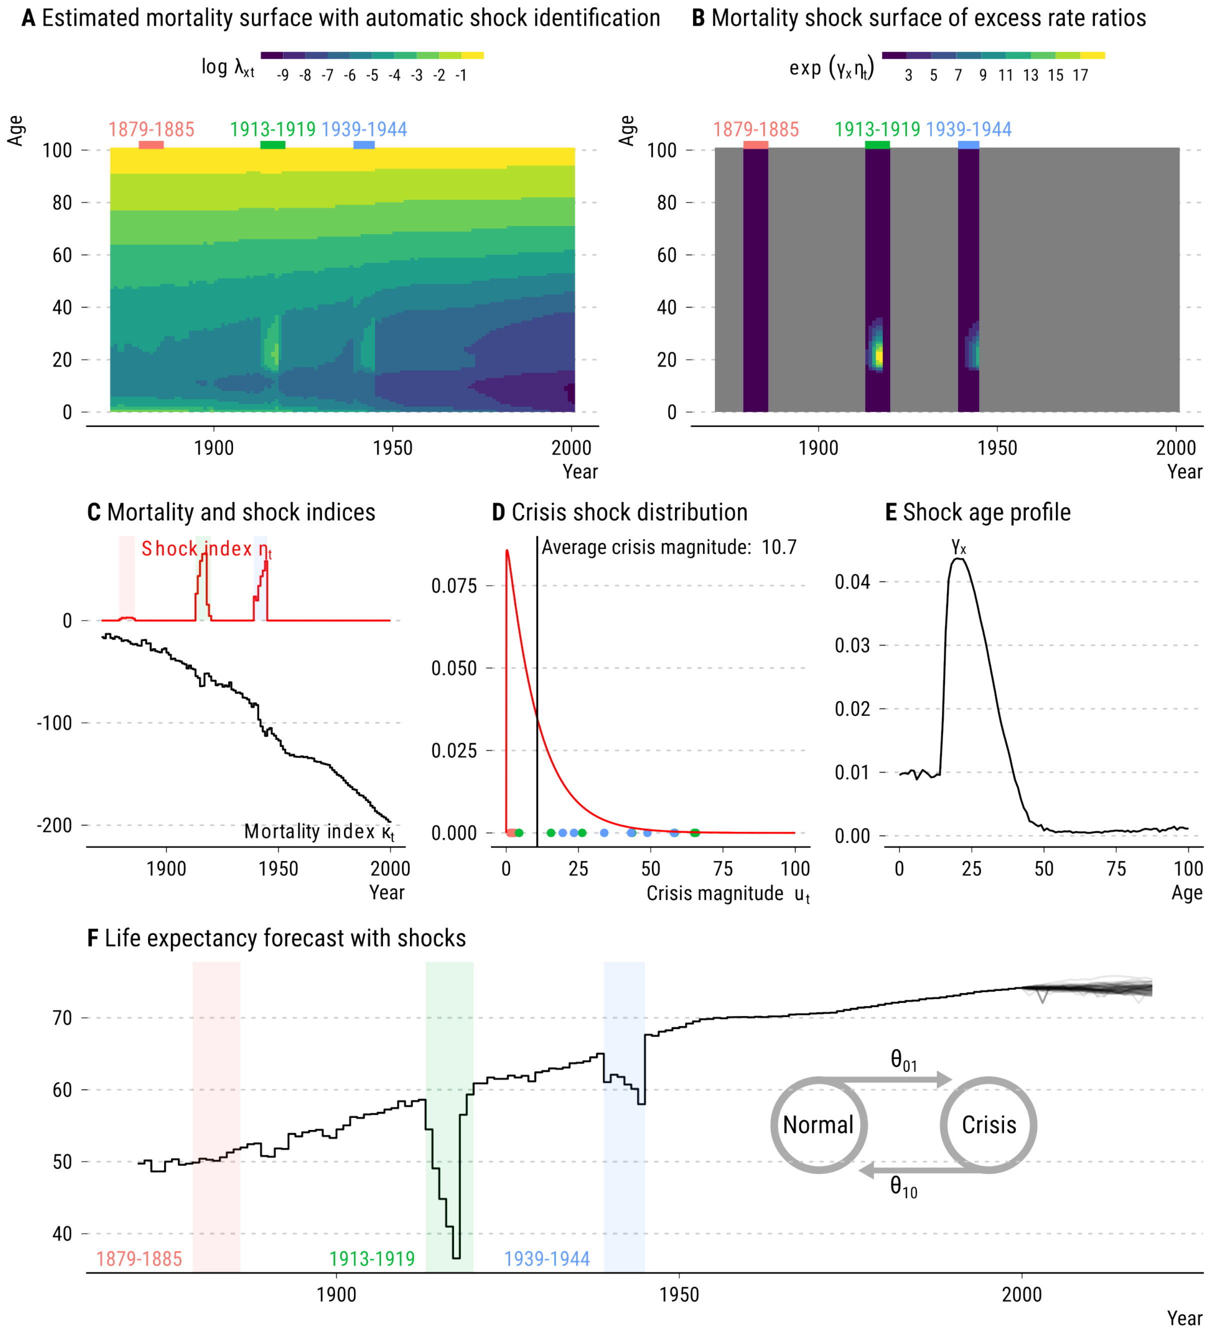

## Outlook

For proof-of-concept and model development purposes I focus on male mortality in England & Wales 1871–2000. The final model will be fit to data until the present, covering the COVID-19 period, and distinguish between different crisis age profiles.

# References

[1] O. Abril-Pla, V. Andreani, C. Carroll, L. Dong, C. J. Fonnesbeck, M. Kochurov, R. Kumar, J. Lao, C. C. Luhmann, O. A. Martin, M. Osthege, R. Vieira, T. Wiecki, and R. Zinkov. PyMC: A modern, and comprehensive probabilistic programming framework in Python. *PeerJ Computer Science*, 9:e1516, Sept. 2023. ISSN 2376-5992. doi:10.7717/peerj-cs.1516.

[2] M. Barbieri, J. R. Wilmoth, V. M. Shkolnikov, D. Glei, D. Jasilionis, D. Jdanov, C. Boe, T. Riffe, P. Grigoriev, and C. Winant. Data resource profile: The human mortality database (HMD). *International Journal of Epidemiology*, 44(5):1549–1556, June 2015. ISSN 0300-5771. doi:10.1093/ije/dyv105.

[3] E. Biffis. Affine processes for dynamic mortality and actuarial valuations. *Insurance: Mathematics and Economics*, 37(3):443–468, Dec. 2005. ISSN 01676687. doi:10.1016/j.insmatheco.2005.05.003.

[4] J. Goes, K. Barigou, and A. Leucht. Bayesian mortality modelling with pandemics: A vanishing jump approach. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 74(4):1150–1182, Aug. 2025. ISSN 0035-9254, 1467-9876. doi:10.1093/jrsssc/qlaf018.

[5] G. Gungah and J. Narsoo. A novel EVT-modified Lee-Carter model for mortality forecasting : An application to extreme mortality events. *Journal of Statistics and Management Systems*, 25 (1):211–243, Jan. 2022. ISSN 0972-0510, 2169-0014. doi:10.1080/09720510.2021.1878629.

[6] Y. Guo and J. S.-H. Li. Robust parameter estimation for the Lee-Carter family: A probabilistic principal component approach. *Insurance: Mathematics and Economics*, 125:103164, Nov. 2025. ISSN 01676687. doi:10.1016/j.insmatheco.2025.103164.

[7] D. Hainaut. Multidimensional Lee–Carter model with switching mortality processes. *Insurance: Mathematics and Economics*, 50(2):236–246, Mar. 2012. ISSN 01676687. doi:10.1016/j.insmatheco.2011.11.003.

[8] R. Lee. The Lee-Carter Method for Forecasting Mortality, with Various Extensions and Applications. *North American Actuarial Journal*, 4(1):80–91, Jan. 2000. ISSN 1092-0277, 2325-0453. doi:10.1080/10920277.2000.10595882.

[9] R. D. Lee and L. R. Carter. Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, 87(419):659–671, Sept. 1992. doi:10.1080/01621459.1992.10475265.

[10] Y. Liu and J. S.-H. Li. The age pattern of transitory mortality jumps and its impact on the pricing of catastrophic mortality bonds. *Insurance: Mathematics and Economics*, 64: 135–150, Sept. 2015. ISSN 01676687. doi:10.1016/j.insmatheco.2015.05.005.

[11] C.-W. Wang, H.-C. Huang, and I.-C. Liu. A Quantitative Comparison of the Lee-Carter Model under Different Types of Non-Gaussian Innovations. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 36(4):675–696, Oct. 2011. ISSN 1018-5895, 1468-0440. doi:10.1057/gpp.2011.20.