GRANULAR MORTALITY MODELING WITH TEMPERATURE AND EPIDEMIC SHOCKS: A THREE-STATE REGIME-SWITCHING APPROACH

Jens Robben, Karim Barigou, Torsten Kleinow

REPRINT | 2026 / 13

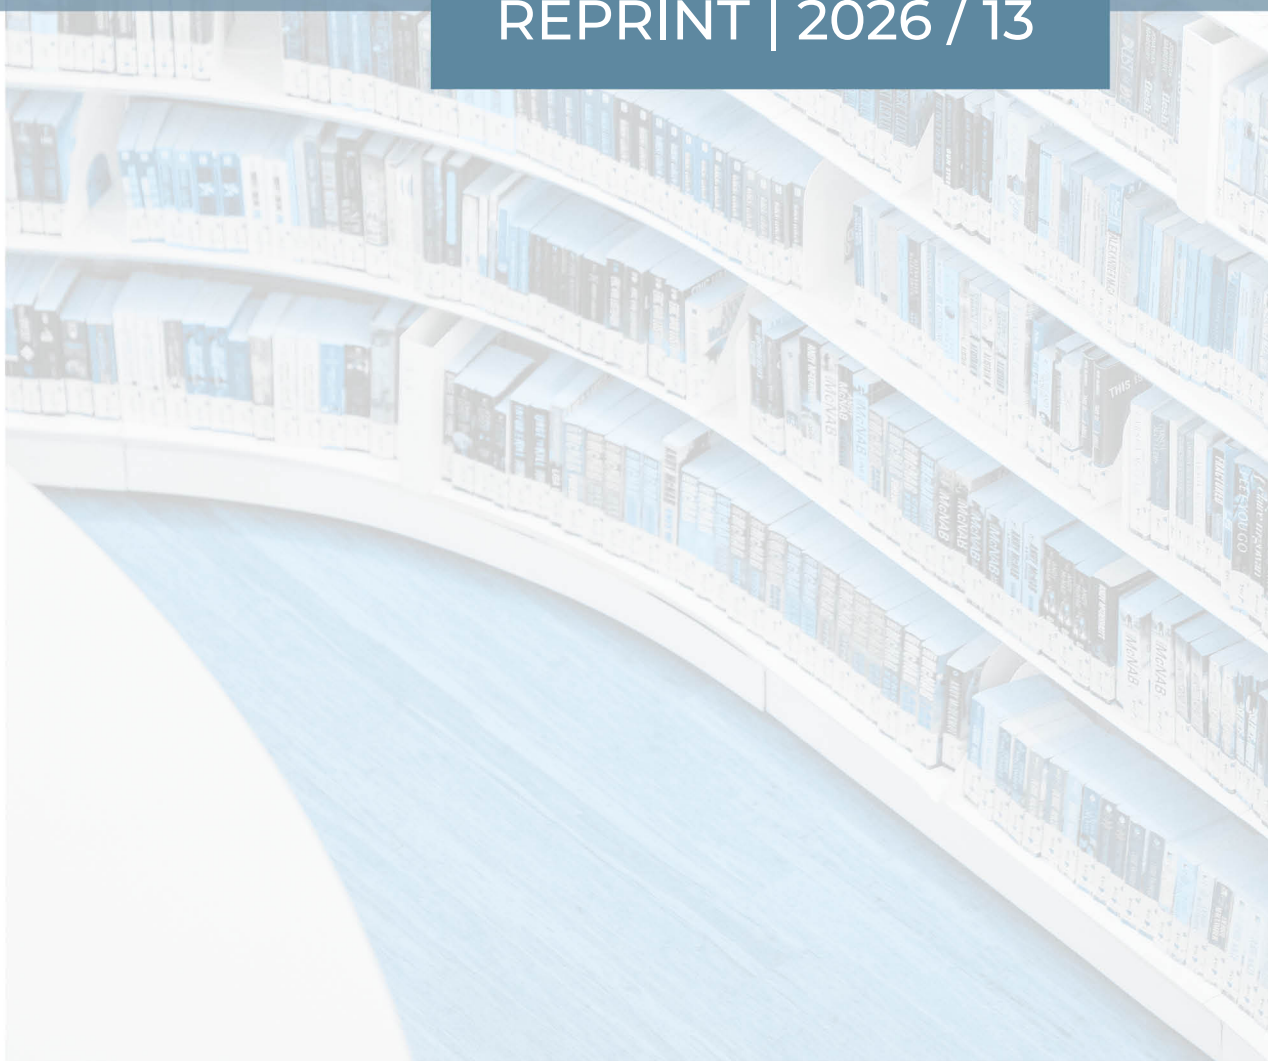

UCLouvain

ISBA | LIDAM
Louvain Institute of Data Analysis and Modeling in economics and statistics

ISBA

Voie du Roman Pays 20 - L1.04.01

B-1348 Louvain-la-Neuve

Email : lidam-library@uclouvain.be

https://uclouvain.be/en/research-institutes/lidam/isba/publication

Insurance Mathematics and Economics 128 (2026) 103250

Contents lists available at ScienceDirect

Insurance: Mathematics and Economics

journal homepage: www.elsevier.com/locate/ime

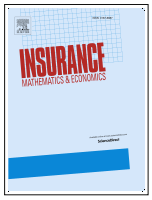

Granular mortality modeling with temperature and epidemic shocks: A three-state regime-switching approach

Jens Robben [a,b,*], Karim Barigou [c], Torsten Kleinow [a,b]

[a] Faculty of Economics and Business, University of Amsterdam, the Netherlands
[b] RCLR, Research Centre for Longevity Risk, University of Amsterdam, the Netherlands
[c] Institute of Statistics, Biostatistics and Actuarial Science (ISBA), Louvain Institute of Data Analysis and Modeling (LIDAM), UCLouvain, Louvain-la-Neuve, Belgium

ARTICLE INFO

Keywords:
Granular mortality
Regime-switching
Environmental shocks
Respiratory shocks

ABSTRACT

This paper develops a granular regime-switching framework to model mortality deviations from seasonal baseline trends driven by temperature and epidemic shocks. The framework features three states: (1) a baseline state that captures observed seasonal mortality patterns, (2) a summer shock state that captures heat waves and other high mortality events, and (3) a winter shock state that addresses mortality deviations caused by cold spells and strong outbreaks of respiratory diseases due to influenza and COVID-19. Transition probabilities between states are modeled using covariate-dependent multinomial logit functions. These functions incorporate, among others, lagged temperature and influenza incidence rates as predictors, allowing dynamic adjustments to evolving shocks. Calibrated on weekly mortality data across 21 French regions and six age groups, the regime-switching framework accounts for spatial and demographic heterogeneity. Under various projection scenarios for temperature and influenza, we quantify uncertainty in mortality forecasts through prediction intervals constructed using an extensive bootstrap approach. These projections can guide insurance companies, healthcare providers, and hospitals in managing risks and planning resources for potential future shocks.

1. Introduction

Mortality rates typically exhibit seasonal variation throughout the year, where the highest rates occur during winter (Marti-Soler et al., 2014; Gemmell et al., 2000; Mackenbach et al., 1992). External events, including increased respiratory infections, such as influenza outbreaks, and extreme weather conditions, such as heat waves, have historically led to substantial deviations from these expected seasonal trends, leading to what is known as excess mortality (Iuliano et al., 2018; Nielsen et al., 2011; Basu and Samet, 2002). A more granular understanding of these short-term mortality dynamics is important from an actuarial perspective as it supports the accurate pricing and underwriting of life contingent products and the assessment of short-term mortality risk. Moreover, since mortality is an important public health indicator, accurately estimating and forecasting mortality rates in the short term is essential for effective health care and public health actions (Hajat et al., 2010; Kok et al., 2010). In this paper, we analyze and aim to better understand the impact of factors such as increased temperatures or heightened respiratory infections on weekly, regional mortality rates using a novel regime-switching framework.

The proposed framework consists of two key components. The first component captures seasonal, age-specific mortality trends within each region, which we call the baseline model. The second component is a three-state regime-switching model, where mortality rates in the first regime follow the seasonal baseline trend. The second and third regimes amplify the baseline based on the current and past week's temperature and respiratory characteristics. This framework is particularly useful, as it enables the generation of future mortality scenarios under different weather and respiratory conditions. In addition, we provide guidance on the uncertainty surrounding these scenario projections. Such insights help to anticipate and respond to potential mortality spikes due to extreme weather events or respiratory outbreaks. Importantly, this paper focuses on short-term mortality projections. Long-term projections, which would require accounting for future improvements in heat resilience and public health interventions, are beyond the scope of this study.

Related work has been done in the EuroMOMO and FluMOMO projects, established to operate a coordinated and timely monitoring of mortality rates within 28 participating European countries. Similarly to our setting, they use the Serfling model as a baseline, incorporating

* Corresponding author.
E-mail address: j.robben@uva.nl (J. Robben).

https://doi.org/10.1016/j.insmatheco.2026.103250
Received 11 July 2025; Received in revised form 2 April 2026; Accepted 15 April 2026
Available online 17 April 2026

seasonality through Fourier terms (Serfling, 1963). Vestergaard et al. (2020) and Nørgaard et al. (2021) use the EuroMOMO framework to estimate excess mortality during the first and second waves of the COVID-19 pandemic. Nielsen et al. (2018) use the FluMoMo model to calculate influenza-attributable mortality during the winter seasons 2010/2011 - 2016/2017 in Denmark, whereas Nielsen et al. (2019) use it to analyze influenza-attributable mortality in Europe during the winter season 2017/2018.

The epidemiological and medical literature has extensively investigated the associations between temperature, epidemic variables, and mortality. Gasparrini et al. (2015) conduct a systematic assessment of the impact of temperature on mortality across multiple countries worldwide. Their findings reveal that, on average, 7.71% of the total daily death counts could be attributed to cold and heat. Song et al. (2017) provide an overview of systematic reviews to summarize evidence regarding the impact of ambient temperature on morbidity and mortality. Iuliano et al. (2018) estimate global influenza-associated respiratory deaths for 1999-2015. Their results reveal a higher annual mortality burden than previously reported and find significant variation across regions and age groups. Wang et al. (2022) conduct a systematic analysis to estimate COVID-19 excess mortality.

In the actuarial literature, the inclusion of temperature and epidemiological variables in mortality models has only recently received increasing attention. Regarding temperature, Seklecka et al. (2017) introduce a new temperature-related mortality (TRM) model that incorporates temperature as an exogenous factor into a stochastic mortality model. Their model results in improved annual mortality forecasts that are used for the pricing of life insurance products. In a subsequent study, Seklecka et al. (2019) investigate how the parameters of the TRM model influence its performance when projecting mortality rates. In addition, Li and Tang (2022) investigate the joint extremes of monthly temperature and mortality using bivariate extreme value theory, and find a strong dependence between extremes of cold weather and old-age death counts. Furthermore, Guibert et al. (2025) integrate temperature effects into a multi-population mortality model for projecting mortality rates, while Milhaud and Gouthon (2024) extend longevity models to incorporate the impacts of heat waves on mortality and predict future excess mortality under various climate change scenarios. Lastly, Bégin et al. (2025) propose a seasonal overlay for age-period-cohort mortality models, Barigou et al. (2025) incorporate climate information from the Actuaries Climate Index (ACI) into mortality models to enhance out-of-sample prediction accuracy, and Min et al. (2025) and Robben and Barigou (2025) use a combination of a seasonally adjusted stochastic mortality model and a distributed lag non-linear model to capture the delayed and non-linear effects of extreme temperature on mortality. In the context of epidemiological variables, Feng and Garrido (2011) combined epidemiological and actuarial frameworks to design insurance policies for infectious diseases, using case studies from historical epidemics such as the Great Plague and SARS. Additionally, Schnürch et al. (2022), Robben et al. (2022), Zhou and Li (2022), Goes et al. (2023), Robben and Antonio (2024), and van Berkum et al. (2025) propose various approaches to account for the impact of COVID-19 on stochastic mortality projection models.

Our paper contributes to the actuarial literature in three ways. First, we gather daily gridded temperature data from the Copernicus Climate Data Store and outline how to convert these data into weekly, region-specific population-weighted temperature levels. Furthermore, we use epidemic data from the French Sentinelles network and hospital admission records from Santé Publique France to capture influenza outbreaks and COVID-19 excess mortality. Second, we introduce a novel mortality modeling framework that integrates a granular age- and region-specific baseline with a regime-switching model that captures mortality shocks driven by environmental and epidemic factors. Using covariate-driven transition probabilities and spatial dependencies, this framework offers a comprehensive and interpretable approach to understanding how these shocks impact mortality patterns across various regions and age groups. Although previous studies have focused primarily on the effects of temperature or epidemic shocks on mortality, this paper is among the first to integrate both factors within a unified mortality modeling framework. Third, we quantify uncertainty in the model's estimates and forecasts and construct prediction intervals for death counts under various weather and respiratory scenarios. In doing so, we account for multiple sources of uncertainty, i.e., parameter, spatial, state, and Poisson uncertainty, using an extensive bootstrap approach. This enables us to isolate and compare the contributions of each component to overall uncertainty. We find that Poisson noise is the largest source of uncertainty when predicting actual death counts, but uncertainty about shock states also plays a substantial role, particularly for older ages.

The paper is organized as follows: Section 2 discusses the notation and the sources of mortality, environmental, and epidemic data considered. Section 3 specifies the model, consisting of the baseline model in Section 3.1 and the regime-switching model in Section 3.2. We outline the calibration procedure in Section 4, and the uncertainty quantification in- and out-of-sample predictions in Section 5. We outline the results of the practical application to 21 French regions and six older age groups in Section 6. Section 7 concludes.

2. Notations and data

Notations. Let $d_{x,t}^{(r)}$ denote the number of deaths in region r for age group x in week t, where t represents the number of ISO weeks elapsed since a selected start date. We adhere to the ISO 8601 standard, where each ISO week is 7 days long, and ISO week 1 starts on the Monday of the week that contains the first Thursday of the year. We define the set of regions as $\mathcal{R}$, the set of age groups as $\mathcal{X}$, and the range of weeks under consideration as $\mathcal{T} = \{1, 2, \ldots, T\}$. Although the methodologies presented in this paper apply to both men and women, our analysis uses combined data for both sexes. Furthermore, let $E_{x,t}^{(r)}$ be the exposure-to-risk in region $r \in \mathcal{R}$ at age group x in week t. The observed weekly death rate $m_{x,t}^{(r)}$ then equals $d_{x,t}^{(r)}/E_{x,t}^{(r)}$.

Mortality data. We retrieve weekly death counts in Metropolitan France from the first ISO week of 2013 to the 26th ISO week of 2024 ($T = 600$) from Eurostat. Deaths are categorized by sex, 5-year age groups, and NUTS 2 regions (Eurostat, 2024a). NUTS (Nomenclature of Territorial Units for Statistics) classifies European administrative areas into hierarchical levels (NUTS 1, 2, 3) for statistical purposes. The NUTS 2 regions in Metropolitan France correspond to the 22 former regions. We exclude Corsica due to the unavailability of weekly death counts. Fig. 1 (left panel) displays the considered French NUTS 2 regions for which Table A.1 in Suppl. Mat. A lists the corresponding names.

This paper focuses on the elderly population, given their increased vulnerability to environmental factors and diseases such as influenza. Hereto, we consider 5-year age groups starting from the age of 65, i.e., $\mathcal{X} = \{65\text{–}69, 70\text{–}74, 75\text{–}79, 80\text{–}84, 85\text{–}89, 90+\}$. Fig. 1 (right panel) presents the weekly death counts from 2013 to 2024 for individuals aged 90 and older in Ile-de-France (FR10), Provence-Alpes-Côte d'Azur (FRL0), and Bretagne (FRH0). The data reveal a pronounced seasonal variation, where death counts increase during winter and decrease during summer. We also observe a clear COVID-19 peak in Ile-de-France, where the death counts exceeded the usual level by more than three times. Furthermore, excess mortality during the winters of 2014/15 and 2016/17 can be largely attributed to severe influenza outbreaks, occasionally intensified by concurrent cold spells, such as the European cold wave of January 2017 (Mølbak et al., 2015; Vestergaard et al., 2017). Figures A.1 and A.2 in Suppl. Mat. A provide a more detailed overview of the structure of the death count data. In order to calculate weekly mortality rates, it is necessary to construct a weekly exposure measure. Suppl. Mat. A also details this construction.

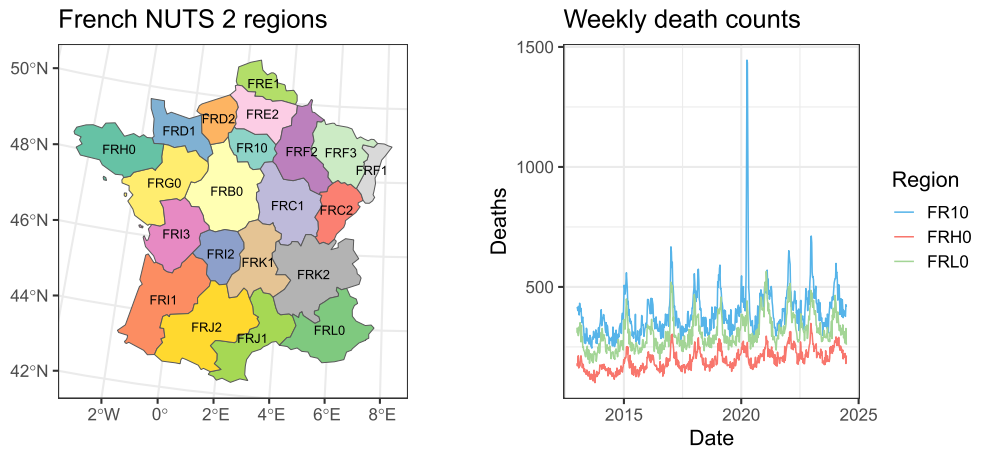

Fig. 1. Left panel: The French NUTS 2 regions. Right panel: Weekly death counts for the 90+ age group in Ile-de-France (FR10), Provence-Alpes-Côte d'Azur (FRL0), and Bretagne (FRH0) from 2013-2024.

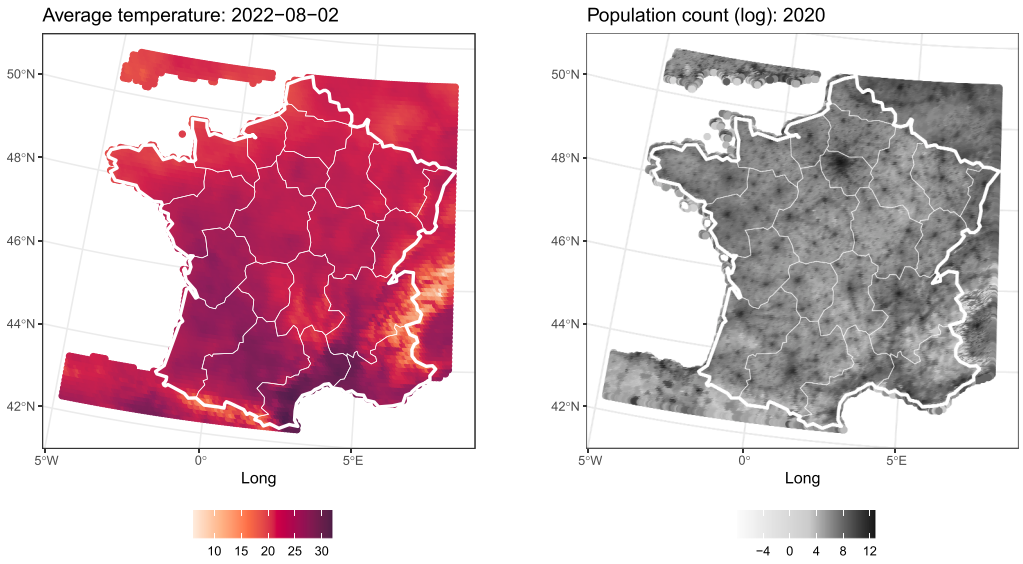

Fig. 2. Left panel: The average temperature in a spatial grid covering Metropolitan France on August 2, 2022. Right panel: Population count (log-scale) in 2020 on a spatial grid covering Metropolitan France. We show the boundaries of the NUTS 2 regions in white.

Temperature data. We use E-OBS gridded meteorological data from the Copernicus Climate Data Store (CDS) to extract daily average temperature levels in Metropolitan France at a 0.10° resolution (Copernicus Climate Change Service, Climate Data Store, 2020). To align with the NUTS 2-level mortality data, we aggregate daily temperatures into population-weighted averages within each NUTS 2 region using NASA's gridded population data (CIESIN, 2018), see Robben et al. (2024) for further details. This ensures that the temperature averages reflect areas where more people reside. Fig. 2 illustrates a snapshot of the temperature and population data on a randomly selected date and year.

For each day, we define region-specific high (95%) and low (5%) temperature thresholds based on the considered period 2013-2024. We then average daily exceedances above or below these thresholds over each week to create an Excess Heat Index (EHI) and an Excess Cold Index (ECI), see Armstrong (2006):

$$\begin{aligned} \mathrm{EHI}_t^{(r)} &= \frac{1}{7}\sum_{d=1}^{7}\left(\mathrm{Tavg}_{t,d}^{(r)} - q_{0.95,t,d}^{(r)}\right)_+ \\ \mathrm{ECI}_t^{(r)} &= \frac{1}{7}\sum_{d=1}^{7}\left(q_{0.05,t,d}^{(r)} - \mathrm{Tavg}_{t,d}^{(r)}\right)_+, \end{aligned} \tag{2.1}$$

for each week $t \in \mathcal{T}$, and $r \in \mathcal{R}$. Here, $\mathrm{Tavg}_{t,d}^{(r)}$ refers to the daily average temperature at day d of week t in region r, and $q_{0.95,t,d}^{(r)}$ and $q_{0.05,t,d}^{(r)}$ denote the 95% and 5% temperature thresholds of daily temperatures in the considered period 2013–2024. Fig. 3 shows the EHI and ECI at two randomly selected weeks.

Influenza data. We consult the French Sentinelles network, established in 1984 by Inserm and Sorbonne University, to collect near real-time epidemic data from ≈ 1 300 general practitioners across France (Valleron et al., 1986). For this study, we extract weekly influenza incidence rates per 100 inhabitants at the NUTS 2 level. These incidences are based on patient visits with fever, myalgia, and respiratory symptoms. Fig. 4 (left panel) shows influenza trends in Île-de-France (FR10), Provence-Alpes-Côte d'Azur (FRL0), and Bretagne (FRH0) from 2013 to 2024. The figure highlights peaks such as in the 2014/2015 winter season and a very low influenza activity in 2020/2021 due to the imposed COVID-19 lockdown measures. Since we aim to explain mortality deviations from a seasonal baseline, we work with so-called influenza anomalies. These anomalies indicate whether influenza activity in a given week was higher or lower than its seasonal average.

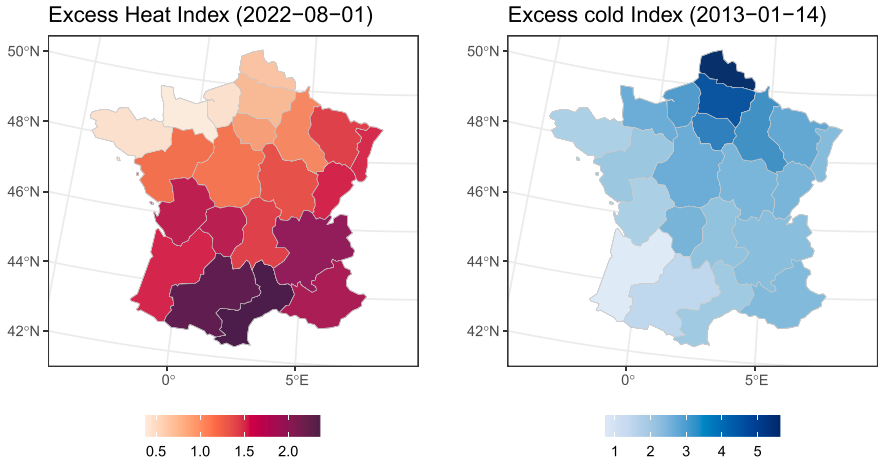

Fig. 3. Left panel: The excess heat index (EHI) at the ISO week starting August 1, 2022. Right panel: The excess cold index (ECI) at the ISO week starting January 14, 2013. We highlight the boundaries of the French NUTS 2 regions in grey.

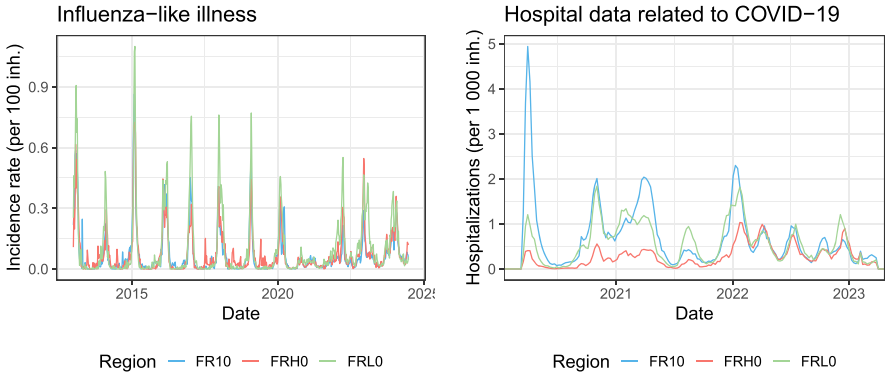

Fig. 4. We show the weekly influenza-like incidences per 100 inhabitants (2013–2024, left panel) and the COVID-19 hospitalizations per 1 000 inhabitants (2020–2023, right panel) in Ile-de-France (FR10), Provence-Alpes-Côte d'Azur (FRL0), and Bretagne (FRH0).

Table 1
Variables derived from the temperature, influenza, and COVID-19 hospitalization data.

Feature	Description
ECI	Excess Cold Index: weekly average of daily exceedances below a low temperature threshold (5% percentile).
EHI	Excess Heat Index: weekly average of daily exceedances above a high temperature threshold (95% percentile).
IA	Influenza activity index: influenza anomalies, deviations from expected seasonal influenza incidences (per 100 inhabitants).
HA	Hospital admission: weekly count of new hospitalizations for COVID-19 (per 1 000 inhabitants).

COVID-19 hospitalizations. We use daily data on COVID-19 hospitalizations per 1 000 inhabitants in French NUTS 2 regions, sourced from Santé Publique France, to capture the impact of COVID-19 on mortality. We aggregate this data, available from March 19, 2020, to March 31, 2023, weekly. Post-March 2023 hospitalization rates are set to zero. Fig. 4 shows weekly COVID-19 hospitalizations in Île-de-France (FR10), Provence-Alpes-Côte d'Azur (FRL0), and Bretagne (FRH0), and highlights clear regional variations in the timing and intensity of COVID-19 waves.

Feature overview. Table 1 summarizes the variables constructed from the environmental, influenza, and COVID-19 hospital admission data, as discussed in this section.

3. Modeling mortality with regime-switching dynamics

Fig. 5 illustrates the proposed model. On the left, the red line visualizes the estimation of a robust seasonal baseline mortality trend, tailored to each age category and region considered within France. However, the observed death counts (in gray) deviate from this baseline due to various potential factors. Hereto, we introduce a regime-switching model, illustrated on the right, which alternates between three states:

1. State 0 (baseline state): applies when the baseline accurately reflects the true underlying death counts.
2. State 1 (summer shock state): captures deviations caused by elevated summer temperatures, e.g., heat waves.
3. State 2 (winter shock state): addresses deviations driven by a cold spell or an increase in respiratory infections, including influenza and COVID-19.

The arrows in Fig. 5 represent state transitions and their corresponding probabilities. As described in Section 2, transition probabilities depend on environmental and respiratory-related covariates. Incorporating such covariates into the transition probabilities helps to determine the current state. For example, if very high temperatures are observed at the time point t, the model will facilitate a transition to state 1. As such, regime states are not entirely hidden but influenced by observable conditions.

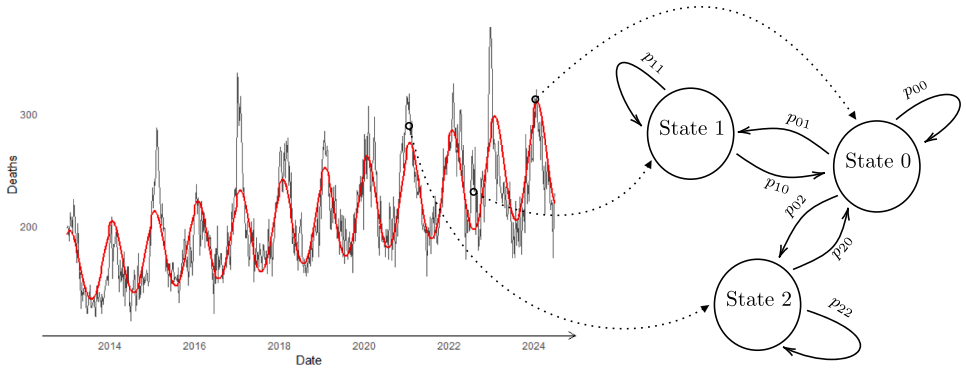

Fig. 5. Schematic representation of the regime-switching mortality model for the French NUTS 2 region Pays de la Loire (FRG0) and age group 90+. The red line shows the seasonal baseline trend for each age group and region, alongside observed death counts (grey). The right panel depicts the regime-switching model, transitioning between three states: baseline mortality (state 0), excess mortality from heat waves (state 1), and excess mortality due to cold spells and respiratory infections (state 2). Arrows indicate state transitions, with probabilities driven by temperature and epidemic covariates. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

In addition, the model assumes that the system can only be in one of the three states at any time, and direct transitions between the summer shock state (state 1) and the winter shock state (state 2) are not allowed. To support this latter assumption, we examined the empirical dependence between the drivers of these regimes. We find that the Pearson correlation between the excess heat and excess cold indices is very small (≈ –0.06), with similar values for the correlation between the excess heat index and the influenza activity index (≈ 0.01) and between the excess heat index and the hospital admission index (≈ –0.03). These findings indicate that heat-related and winter-related mortality drivers are largely non-overlapping and represent distinct phenomena.

3.1. Weekly baseline mortality model

We design a regional and age-group-specific baseline to capture the overall seasonal mortality patterns by following the approach in Serfling (1963). Hereto, we incorporate seasonality using Fourier terms and assume that the observed weekly death counts in region r and at age group x are realizations from a Poisson-distributed random variable, $D_{x,t}^{(r)}$:

$$D_{x,t}^{(r)} \sim \text{Poisson}\left(E_{x,t}^{(r)}\, \mu_{x,t}^{(r)}\right), \tag{3.1}$$

where the structure of $\mu_{x,t}^{(r)}$ is given by Serfling (1963):

$$\begin{aligned}\log \mu_{x,t}^{(r)} =& \gamma_{x,0}^{(r)} + \gamma_{x,1}^{(r)} t + \gamma_{x,2}^{(r)} \sin\left(\frac{2\pi w(t)}{52.18}\right) + \gamma_{x,3}^{(r)} \cos\left(\frac{2\pi w(t)}{52.18}\right) \\ &+ \gamma_{x,4}^{(r)} \sin\left(\frac{2\pi w(t)}{26.09}\right) + \gamma_{x,5}^{(r)} \cos\left(\frac{2\pi w(t)}{26.09}\right),\end{aligned} \tag{3.2}$$

with $w(t)$ the ISO-week number in the year $y(t)$ of week t, and 52.18 represents the average number of weeks in a year. This baseline structure is tailored to each region, age, and week. The parameters $\gamma_{x,p}^{(r)}$ ($p = 0, 1, \ldots, 5$) are region- and age-group-specific and account for the level, trend, and seasonal age-specific variations. We outline the estimation process of this baseline in Section 4.1 and denote the estimated expected weekly baseline death counts as:

$$\hat{b}_{x,t}^{(r)} := \hat{E}\left[D_{x,t}^{(r)}\right] = E_{x,t}^{(r)}\, \hat{\mu}_{x,t}^{(r)}, \tag{3.3}$$

where $\hat{\mu}_{x,t}^{(r)}$ represents the fitted baseline structure in (3.2).

3.2. Regime-switching mortality framework

Consider the estimated baseline deaths, $\hat{b}_{x,t}^{(r)}$, as fixed. We assume that the distribution of $D_{x,t}^{(r)}$ depends on the state of an underlying Markov chain $S_t^{(r)}$, which varies by region and time. At each time point t, $S_t^{(r)}$ can be in a baseline (state 0), summer shock (state 1), or winter shock (state 2) state. For each state i, $D_{x,t}^{(r)}$ follows a Poisson distribution:

$$D_{x,t}^{(r)} \mid S_t^{(r)} = i \sim \text{POI}\left(\hat{b}_{x,t}^{(r)} \cdot \exp\left[\left(\boldsymbol{z}_t^{(r)}\right)^\top \boldsymbol{\alpha}_{i,x}\right]\right), \tag{3.4}$$

where $\boldsymbol{\alpha}_{i,x}$ captures the effects of the covariate vector $\boldsymbol{z}_t^{(r)}$ (e.g., temperature, influenza, COVID-19) on deviations from baseline mortality for the age group x in regime state i. This age dependency allows the model to account for how covariates, such as extreme temperatures, influence excess mortality differently across the age groups considered.

Each regime state corresponds to a distinct mortality pattern. For $S_t^{(r)} = 0$, the baseline state, we assume that mortality rates follow the seasonal baseline trend $\hat{b}_{x,t}^{(r)}$, where no additional covariate effects influence mortality. We set $\boldsymbol{\alpha}_{0,x} = 0$. For $S_t^{(r)} = 1$, the summer shock state, only the excess heat covariate in $\boldsymbol{z}_t^{(r)}$ plays a role, and we restrict $\boldsymbol{\alpha}_{1,x}$ to have non-zero values only for this variable. Lastly, for $S_t^{(r)} = 2$, the winter shock state, we assume only the excess cold index and the epidemic covariates affect mortality, with non-zero values in $\boldsymbol{\alpha}_{2,x}$.

To avoid overcomplicating the model, we assume that the Markov chain $S_t^{(r)}$ is independent of the age group x. As a result, state transitions affect all ages simultaneously, but through age-dependent parameters in (3.4) we can still account for age-specific differences in the impact on mortality. The state transition probabilities are described by a time-varying, region-specific transition matrix $\boldsymbol{P}_t^{(r)}$:

$$\boldsymbol{P}_t^{(r)} = \begin{pmatrix} 1 - p_t^{01}\left(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}, U_r\right) & p_t^{01}\left(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}, U_r\right) & p_t^{02}\left(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}, U_r\right) \\ \quad - p_t^{02}\left(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}, U_r\right) & & \\ 1 - p_t^{11}\left(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}, U_r\right) & p_t^{11}\left(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}, U_r\right) & 0 \\ 1 - p_t^{22}\left(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}, U_r\right) & 0 & p_t^{22}\left(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}, U_r\right) \end{pmatrix},$$

with $\boldsymbol{\beta}$ the parameter vector, and U_r a spatial effect to capture regional disparities. This spatial effect ensures that neighboring regions are more likely to have similar transition probabilities. In this paper, we apply an Intrinsic Conditional Auto-Regressive (ICAR) model to the spatial effect vector (Besag and Kooperberg, 1995):

$$\boldsymbol{U} = (U_1, U_2, \ldots, U_R) \sim \mathcal{N}\left(\boldsymbol{0}, [\tau \cdot (\boldsymbol{D} - \boldsymbol{W})]^{-1}\right), \tag{3.5}$$

where $R := |\mathcal{R}|$ is the region count, $\boldsymbol{D}$ a diagonal matrix with entries equal to the number of neighbors for each region, and $\boldsymbol{W}$ the adjacency matrix. The entry $W_{rr'}$ equals 1 if the regions r and r' are neighbors and 0 otherwise. The diagonal entries of $\boldsymbol{W}$ are zero. The precision parameter τ controls the level of spatial autocorrelation among the regions.

We model the probability of transitioning from state i to state j using a multinomial logit specification. We therefore define $\mathcal{J}_i$ as the set of

possible states, excluding the baseline state, to which the system can transition starting from state i. Specifically, $\mathcal{J}_0 = \{1,2\}$, $\mathcal{J}_1 = \{1\}$, and $\mathcal{J}_2 = \{2\}$. For $ij \notin \{12, 21\}$, the probability of transitioning from i to j is:

$$p_t^{ij}\left(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}, u_r\right) = \begin{cases} \dfrac{\exp\left(\left(\boldsymbol{z}_t^{(r)}\right)^\top \boldsymbol{\beta}_{ij} + U_r\right)}{1 + \sum_{j' \in \mathcal{J}_i} \exp\left(\left(\boldsymbol{z}_t^{(r)}\right)^\top \boldsymbol{\beta}_{ij'} + U_r\right)} & j \neq 0 \\ \dfrac{1}{1 + \sum_{j' \in \mathcal{J}_i} \exp\left(\left(\boldsymbol{z}_t^{(r)}\right)^\top \boldsymbol{\beta}_{ij'} + U_r\right)} & j = 0, \end{cases} \tag{3.6}$$

with $\boldsymbol{\beta} = (\boldsymbol{\beta}_{ij})_{ij}$. For transitions toward (or persistence in) the summer shock state, only the excess heat index is included. For transitions toward (or persistence in) the winter shock state, only the excess cold index, the epidemic and COVID-19 hospitalization variables are considered. Other coefficients are set to zero.

4. Model calibration

4.1. Weekly baseline mortality model

We fit the weekly baseline mortality model in Section 3.1 using a Poisson generalized linear model (GLM) to capture region- and age-group-specific seasonal patterns in the death counts. To ensure smooth parameter estimates across neighboring regions, see (3.2), we add a penalty term to the objective function (Robben et al., 2024; Li et al., 2024):

$$\boldsymbol{\gamma}^* = \underset{\boldsymbol{\gamma}}{\operatorname{argmin}} \left(-\ell_P(\boldsymbol{\gamma}) + \sum_{p=0}^{5} \lambda_p \left(\sum_{x \in \mathcal{X}} \boldsymbol{\gamma}_{x,p}^\top (\boldsymbol{D} - \boldsymbol{W}) \boldsymbol{\gamma}_{x,p} \right) \right),$$

where $\boldsymbol{\gamma} = (\boldsymbol{\gamma}_{x,p})_{x,p}$, with $\boldsymbol{\gamma}_{x,p} = (\gamma_{x,p}^{(r)})_{r \in \mathcal{R}}$, for each age group $x \in \mathcal{X}$ and $p = 0, 1, \ldots, 5$. Furthermore, $\ell_P(\boldsymbol{\gamma})$ is the Poisson log-likelihood, and the matrices $\boldsymbol{D}$ and $\boldsymbol{W}$ are defined in (3.5). The smoothness parameters λ_p determine the amount of smoothing applied, with larger values imposing similar parameter estimates across regions. We select optimal λ_p-values using Un-Biased Risk Estimation scores (Craven and Wahba, 1978; Wood, 2004). We calculate the expected baseline death counts for region r, week period t, and age group x in (3.2) using the optimal parameter vector $\boldsymbol{\gamma}^*$, and denote the result as $\hat{b}_{x,t}^{(r)}$.

4.2. Regime-switching mortality framework

4.2.1. Notation

Let $\mathcal{D}_T$ denote the collection of age-specific death counts observed throughout the calibration period $\mathcal{T} = \{1, 2, \ldots, T\}$, i.e., $\mathcal{D}_T = \{D_{x,t}^{(r)} \mid r \in \mathcal{R}, x \in \mathcal{X}, t \in \mathcal{T}\}$. We introduce a similar notation for the collection of all states, $\mathcal{S}_T$, and covariate vectors, $\mathcal{Z}_T$. Lastly, we denote the parameter vector for the regime-switching process as:

$$\boldsymbol{\theta} = \left\{\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_{01}, \boldsymbol{\beta}_{02}, \boldsymbol{\beta}_{11}, \boldsymbol{\beta}_{22}\right\}, \tag{4.1}$$

where $\boldsymbol{\alpha}_i = (\boldsymbol{\alpha}_{i,x})_{x \in \mathcal{X}}$ for $i = 1, 2$. The vector $\boldsymbol{\theta}$ consists of the parameters in (3.4) and (3.6).

4.2.2. From the complete-data to the incomplete-data log-likelihood

The complete-data likelihood assumes that the observed death counts $\mathcal{D}_T$ and latent regime states $\mathcal{S}_T$ are fully known. However, this likelihood depends on the spatial random effect $\boldsymbol{U}$. To account for this, we integrate out $\boldsymbol{U}$:

$$\begin{aligned} \mathcal{L}\left(\boldsymbol{\theta}, \tau \mid \mathcal{D}_T, \mathcal{S}_T, \mathcal{Z}_T\right) &= \int_{\boldsymbol{u}} \mathbb{P}\left(\mathcal{D}_T, \mathcal{S}_T, \boldsymbol{u} \mid \mathcal{Z}_T; \boldsymbol{\theta}, \tau\right) d\boldsymbol{u} \\ &= \int_{\boldsymbol{u}} \mathbb{P}\left(\mathcal{D}_T, \mathcal{S}_T \mid \boldsymbol{u}, \mathcal{Z}_T; \boldsymbol{\theta}\right) f(\boldsymbol{u}; \tau) d\boldsymbol{u}, \end{aligned} \tag{4.2}$$

where $\mathbb{P}\left(\mathcal{D}_T, \mathcal{S}_T \mid \boldsymbol{u}, \mathcal{Z}_T; \boldsymbol{\theta}\right)$ is the complete-data likelihood conditionally on the spatial effect $\boldsymbol{u}$, and $f(\boldsymbol{u}; \tau)$ is the prior distribution of $\boldsymbol{u}$, which follows a multivariate normal distribution with precision parameter τ, see (3.5):

$$f(\boldsymbol{u}; \tau) = \frac{1}{(2\pi)^{R/2} (\det(\boldsymbol{Q}(\tau)^{-1}))^{1/2}} e^{-\frac{1}{2} \boldsymbol{u}^\top \boldsymbol{Q}(\tau) \boldsymbol{u}},$$

where $\boldsymbol{Q}(\tau) := \tau[\boldsymbol{D} - \boldsymbol{W}]$ is called the precision matrix that governs the degree of spatial autocorrelation across the regions.

It is computationally challenging for regime-switching models to evaluate the integral in (4.2). To simplify this, we apply Laplace's approximation method to avoid a direct calculation of the integral expression (Laplace, 1774; Tierney and Kadane, 1986):

$$\begin{aligned} &\log \mathcal{L}(\boldsymbol{\theta}, \tau \mid \mathcal{D}_T, \mathcal{S}_T, \mathcal{Z}_T) \\ &= \log \mathbb{P}\left(\mathcal{D}_T, \mathcal{S}_T \mid \boldsymbol{u}^*, \mathcal{Z}_T; \boldsymbol{\theta}\right) + \log f(\boldsymbol{u}^*; \tau) + \frac{R}{2} \log(2\pi) - \frac{1}{2} \log(\det(\boldsymbol{S}^{-1})), \end{aligned} \tag{4.3}$$

where $\boldsymbol{u}^*$ is the mode of the expected log joint probability distribution, and $\boldsymbol{S}^{-1}$ is the matrix of second-order derivatives evaluated at $\boldsymbol{u}^*$. Suppl. Mat. B.2 and B.3 provide details on the calculation of the Laplace-approximated log-likelihood and the involved Hessian.

We are not able to directly optimize the complete-data log-likelihood in (4.3) because the regime states $S_t^{(r)}$ are unobserved. Instead, we focus on maximizing the incomplete-data log-likelihood, which aggregates over all possible regime state configurations:

$$\log \tilde{\mathcal{L}}(\boldsymbol{\theta}, \tau \mid \mathcal{D}_T, \mathcal{Z}_T) = \sum_{\mathcal{S}_T \in \{0,1,..,N\}^{R \cdot T}} \log \mathcal{L}(\boldsymbol{\theta}, \tau \mid \mathcal{D}_T, \mathcal{S}_T, \mathcal{Z}_T),$$

where $N + 1 := 3$ represents the number of regime states. However, it is practically impossible to directly evaluate this expression, as it would require us to sum over a very large number of possible state combinations. Therefore, we apply the Expectation-Maximization (EM) algorithm to optimize the incomplete-data log-likelihood.

4.2.3. Optimizing the incomplete-data log-likelihood using the EM-algorithm

Since the values of the spatial effect $\boldsymbol{u}^*$ depend on the parameter vector $\boldsymbol{\theta}$, we incorporate Laplace's approximation of the complete-data log-likelihood into the EM-algorithm to iteratively update both the parameter vector $\boldsymbol{\theta}$ and the spatial effect $\boldsymbol{u}^*(\boldsymbol{\theta})$. Below is a summary and explanation of the key steps of which Suppl. Mat. B.4 provides the details:

1. Initialization: Start with initial estimates $\boldsymbol{\theta}_0^*$ and $\boldsymbol{u}_0^*$.
2. EM-algorithm (iteration $k + 1$):
   (a) E-step: Take the expected value of the complete-data log-likelihood with respect to the regime state process, given the observed data, covariate information, and the current best-estimates of the parameter vector, $\boldsymbol{\theta}_k^*$, and spatial effect, $\boldsymbol{u}_k^*$.
   (b) M-step: Update the parameter vector $\boldsymbol{\theta}$ by maximizing the expected complete-data log-likelihood and obtain $\boldsymbol{\theta}_{k+1}^*$.
   (c) Update spatial effect: Update the spatial effect, $\boldsymbol{u}^*$, under the current optimal parameter vector, $\boldsymbol{\theta}_{k+1}^*$, and obtain $\boldsymbol{u}_{k+1}^*$.
3. Convergence check: Repeat steps (a), (b), and (c) until we reach an absolute convergence of ϵ (a predefined small threshold) in the expected complete-data log-likelihood.

The final step is to estimate the precision parameter τ of the ICAR prior. We do this maximizing the profile log-likelihood over a grid of candidate values $\mathcal{G}$. For each $\tau \in \mathcal{G}$, we apply the EM algorithm to obtain the optimal parameter vector $\boldsymbol{\theta}^*(\tau)$. We then select the optimal τ^* as the value that maximizes the profile log-likelihood, see Suppl. Mat. B.5.

We can then assess the uncertainty in the estimated parameters of the regime-switching process by deriving the so-called Fisher-Information matrix (FIM). Existing methods, such as the ones proposed by Louis (1982), Meng and Rubin (1991) and Oakes (1999), often involve computational challenges. Therefore, we follow the approach of Meng and Spall (2017), who propose a Monte Carlo-based method using

simultaneous perturbation stochastic approximation to estimate the Hessian matrix of the incomplete data log-likelihood (Spall, 2000). Suppl. Mat. B.6 outlines the estimation procedure.

5. Multi-layered uncertainty in mortality predictions

When projecting weekly death counts with the proposed model, we account for parameter, spatial, regime-switching state, and Poisson uncertainty. We account for these uncertainties to improve the reliability of the mortality predictions. We do this by generating B bootstrap samples according to the below outlined prediction process. For each sample ι, we proceed as follows:

1. **Parameter uncertainty.** We account for parameter uncertainty in the model's predictions by drawing randomly from the following multivariate normal distributions:

$$\boldsymbol{\gamma}_\iota \sim \mathcal{N}\big(\boldsymbol{\gamma}^*, \boldsymbol{\Sigma}_1(\boldsymbol{\gamma}^*)\big), \qquad \boldsymbol{\theta}_\iota \sim \mathcal{N}\big(\boldsymbol{\theta}^*, \boldsymbol{\Sigma}_2(\boldsymbol{\theta}^*)\big), \tag{5.1}$$

   where $\boldsymbol{\gamma}^*$ and $\boldsymbol{\theta}^*$ are the optimal parameter vectors in the baseline and regime-switching model, respectively. We construct the variance-covariance matrices $\boldsymbol{\Sigma}_1(\boldsymbol{\gamma}^*)$ and $\boldsymbol{\Sigma}_2(\boldsymbol{\theta}^*)$ from the Fisher information of their corresponding log-likelihood functions. The construction of $\boldsymbol{\Sigma}_2(\boldsymbol{\theta}^*)$ is more intricate and detailed in Suppl. Mat. B.6.
2. **Spatial uncertainty.** We account for regional disparities in the transition probabilities by drawing a random sample from the ICAR model, specified in (3.5):

$$\boldsymbol{u}_\iota = \big(u_{1,\iota}, u_{2,\iota}, \dots, u_{R,\iota}\big) \sim \mathcal{N}\big(\mathbf{0}, [\tau \cdot (\boldsymbol{D} - \boldsymbol{W})]^{-1}\big).$$

3. **State uncertainty.** The state uncertainty reflects the uncertainty about the true values of the latent regime states when making predictions. To account for this, we generate regime state trajectories using the estimated parameters in the transition probabilities, see (3.6).

   We start by selecting an initial state i. This is either the most probable state at the end of the calibration period (for forecasting) or the most probable initial state (for in-sample prediction). The next state is then generated by sampling from a discrete distribution with possible values $j \in \{0, 1, 2\}$, weighted by transition probabilities $p_t^{ij}(\boldsymbol{z}_t^{(r)}; \boldsymbol{\beta}_\iota, u_{r,\iota})$, where $\boldsymbol{z}_t^{(r)}$ is the covariate vector, $u_{r,\iota}$ is the sampled spatial effect, and $\boldsymbol{\beta}_\iota$ is the sampled parameter vector contained in $\boldsymbol{\theta}_\iota$. We repeat this iterative procedure for all t in the prediction period $\mathcal{T}_{\text{pred}}$. We denote the resulting regime state trajectory as $s_{t,\iota}^{(r)}$ for $t \in \mathcal{T}_{\text{pred}}$.
4. **Poisson uncertainty.** Using the sampled states $s_{t,\iota}^{(r)}$, we sample death counts from the state-dependent Poisson distributions in (3.4) with the parameter vector $\boldsymbol{\theta}_\iota$. We calculate the baseline death counts using (3.1) and (3.2). This corresponds to either the calibrated baseline (for in-sample prediction) or a forecasted mean from the Poisson GLM (for out-of-sample forecasting). In the latter case, we need to use scenarios for the covariate vectors $\boldsymbol{z}_t^{(r)}$.

Next, we construct confidence intervals for the predicted death counts by calculating the 2.5%, 50%, and 97.5% percentiles of the B bootstrap samples. As a result, we obtain a range of outcomes in which the predicted death counts may fall with a 95% confidence rather than a point estimate. Lastly, note that we can switch any source of uncertainty on or off, independently of each other. This allows us to control for the factors contributing to the uncertainty and to detect the sources of uncertainty in the predicted death counts.

6. Case study on 21 French NUTS 2 regions

6.1. Baseline and regime-switching model: Specification and calibration

6.1.1. Baseline mortality model

We calibrate the baseline parameters $\gamma_{x,p}^{(r)}$, using the strategy in Section 4.1, on the set $\mathcal{R}$ of French NUTS 2 regions, age groups $\mathcal{X} = \{65\text{–}69, 70\text{–}74, 75\text{–}79, 80\text{–}84, 85\text{–}89, 90+\}$, and weeks $\mathcal{T} = \{1, 2, ., 600\}$, where $t = 1$ corresponds to the first ISO week in 2013 and $t = 600$ to the 26th ISO week in 2024. To limit the impact of extreme mortality peaks on the estimation of the baseline model, such as those observed during the COVID-19 pandemic, we identify and smooth outliers in each age- and region-specific death count series $d_{x,t}^{(r)}$. Hereto, we decompose the series $d_{x,t}^{(r)}$ into a trend, seasonal, and remainder component based on the Seasonal-Trend decomposition using Loess (STL) algorithm (Cleveland et al., 1990). Outliers in the remainder series are detected using Tukey's interquartile range criterion and replaced by linearly interpolated values (Tukey, 1977).

Fig. 6 shows the estimated parameters for the age group 90+. The top left and middle panels show that Île-de-France had the lowest overall death rate in 2013 (γ_0) and experienced the largest mortality decline. The eastern French regions also saw a decline in mortality, while the Western regions experienced a slight increase (γ_1). The seasonal variation is most prominent in the northern and north-eastern regions. Section A of the Online Appendix contains the parameter estimates for the five remaining age groups.

6.1.2. Regime-switching mortality model

Model selection. We first select the variables to be included in the state-specific specifications of the Poisson model in (3.4) and in the multinomial logistic models describing the transition probabilities in (3.6). The variables considered are listed in Table 1: the Excess Cold and Heat Index (ECI, EHI), the influenza activity index (IA), and the weekly number of new hospital admissions due to COVID-19 (HA). In the Poisson model of state 1, the summer shock state, we restrict the candidate variables to the EHI and its lags. The same candidate variables are used for modeling the transition probabilities toward or persistence in state 1. In the Poisson model of state 2, the winter shock state, we restrict the candidate variables to the ECI, IA, and HA, along with their lags. The same candidate variables are used for modeling the transition probabilities toward or persistence in state 2. For respiratory-related variables, we use exceedances over their 75% quantile since extreme values, rather than moderate fluctuations, are more likely to contribute to significant mortality shocks.

For each candidate variable, we consider the present-week value (lag 0) to capture immediate effects, the one-week lag (lag 1) to capture short-term effects, and the average of lags 2-3 to represent medium-term effects. We do not focus on long-term effects in this paper. Although heat effects on mortality are typically immediate and short-lived, cold and influenza-related impacts tend to manifest over longer periods, sometimes up to 21 days after exposure (see, e.g., Braga et al., 2001; Pattenden et al., 2003; Wong et al., 2004; Lytras et al., 2019). Therefore, we allow the model to consider lags for up to three weeks, but rely on the model-selection procedure to exclude lags that are not relevant in either the Poisson or multinomial logistic models.

As discussed in Section 3.2, the parameters in the Poisson model specifications vary with age $x \in \mathcal{X} = \{65\text{–}69, 70\text{–}74, 75\text{–}79, 80\text{–}84, 85\text{–}89, 90+\}$, whereas the parameters of the transition probability models are assumed to be age-invariant. This implies that regime transitions affect all age groups simultaneously, while the influence of the explanatory variables on mortality may still differ across ages. Furthermore, we only allow transitions toward the summer shock state whenever the excess heat index is non-zero. Similarly, transitions to the winter shock state are only possible whenever the excess cold index, influenza activity index, or COVID-19 hospitalizations are non-zero.

We then select the variables using a stepwise evaluation of several Poisson model specifications for the summer and winter shock states and the multinomial logistic models describing transitions between or persistence in these two states. Starting from a baseline specification that solely includes lag 0 age-specific effects in the Poisson models and intercept-only models for the transition probabilities, additional lagged and interaction terms are sequentially tested based on their impact on

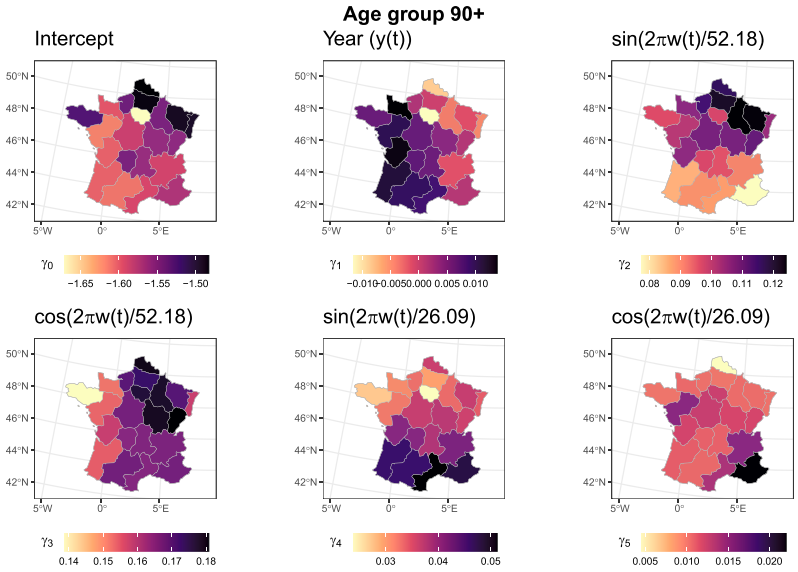

Fig. 6. The estimated parameters in the baseline mortality model, i.e., $\gamma_{x,p}^{(r)}$, for the age group 90+ in the French NUTS 2 regions.

the Bayesian Information Criterion (BIC). Variables or lags that reduce the BIC are added to the model, while those increasing BIC are excluded. We provide further details of the variable selection process and corresponding BIC values for each considered model specification in Suppl. Mat. C. In addition, Suppl. Mat. D.1 compares the model fit and parameter estimates when using different quantile threshold definitions in (2.1), and shows that conclusions remain largely unchanged.

State-specific Poisson model specifications. Based on the BIC-based variable selection procedure from Suppl. Mat. C, we model the expected death counts in the summer shock state as:

$$\log \mathbb{E}\left[D_{x,t}^{(r)} \mid S_t^{(r)} = 1\right] = \log \hat{b}_{x,t}^{(r)} + \alpha_{1,0} + \alpha_{1,1}^{(x)} \mathrm{EHI}_t^{(r)} + \alpha_{1,2}^{(x)} \mathrm{EHI}_{t-1}^{(r)}, \tag{6.1}$$

where $\alpha_{1,0}$ refers to an overall intercept, the $\alpha_{1,j}^{(x)}$'s are age-specific parameters for $j = 1, 2$, and the offset $\hat{b}_{x,t}^{(r)}$ refers to the baseline death count estimated as in Section 4.1. Here, $\mathrm{EHI}_{t-1}^{(r)}$ refers to the excess heat index in region r at week $t-1$ (lag 1). Note that the medium-term effect (average over lags 2-3) was not selected, as including it increased the overall BIC (see Suppl. Mat. C).

For the winter shock state (state 2), the optimally selected model for the expected death counts is:

$$\begin{aligned}\log \mathbb{E}\left[D_{x,t}^{(r)} \mid S_t^{(r)} = 2\right] = \log \hat{b}_{x,t}^{(r)} &+ \alpha_{2,0} + \alpha_{2,1}^{(x)} \mathrm{ECI}_t^{(r)} + \alpha_{2,2}^{(x)} \mathrm{ECI}_{t-1}^{(r)} + \alpha_{2,3}^{(x)} \mathrm{ECI}_{t-2,t-3}^{(r)} \\ &+ \alpha_{2,4}^{(x)} \mathrm{IA}_t^{(r)} + \alpha_{2,5}^{(x)} \mathrm{IA}_{t-1}^{(r)} + \alpha_{2,6}^{(x)} \mathrm{IA}_{t-2,t-3}^{(r)} \\ &+ \alpha_{2,7}^{(x)} \mathrm{HA}_t^{(r)} + \alpha_{2,8}^{(x)} \mathrm{HA}_{t-1}^{(r)} \\ &+ \alpha_{2,9}^{(x)} \left(\mathrm{ECI}_t^{(r)} \cdot \mathrm{HA}_t^{(r)}\right) + \alpha_{2,10}^{(x)} \left(\mathrm{ECI}_{t-1}^{(r)} \cdot \mathrm{HA}_{t-1}^{(r)}\right),\end{aligned} \tag{6.2}$$

where $\mathrm{ECI}_{t-2,t-3}^{(r)}$ and $\mathrm{IA}_{t-2,t-3}^{(r)}$ refer to the excess cold and influenza activity indices averaged over lags 2-3, respectively. The selected model structure includes immediate, short-, and medium-term effects of excess cold and influenza activity on mortality, as well as immediate and short-term effects for COVID-19 hospital admissions. Interactions between the cold index and hospital admissions at lags 0 and 1 were also selected by BIC. One plausible interpretation is that COVID-19 transmission tends to increase in colder conditions and decrease in warmer ones (Kang et al., 2021; Notari, 2021).

Modeling regime transition probabilities. We construct multinomial logistic models for the transition probabilities, as defined in (3.6). We select the model structures based on BIC, as detailed in Suppl. Mat. C. We obtain:

$$\begin{aligned}\operatorname{logit} p_t^{01}\left(\boldsymbol{z}_t^{(r)};\, \boldsymbol{\beta}, u_r\right) &= \beta_{01,0} + \beta_{01,1} \mathrm{EHI}_t^{(r)} + U_r \\ \operatorname{logit} p_t^{02}\left(\boldsymbol{z}_t^{(r)};\, \boldsymbol{\beta}, u_r\right) &= \beta_{02,0} + \beta_{02,1} \mathrm{ECI}_t^{(r)} + U_r \\ \operatorname{logit} p_t^{11}\left(\boldsymbol{z}_t^{(r)};\, \boldsymbol{\beta}, u_r\right) &= \beta_{11,0} + \beta_{11,1} \mathrm{EHI}_t^{(r)} + U_r \\ \operatorname{logit} p_t^{22}\left(\boldsymbol{z}_t^{(r)};\, \boldsymbol{\beta}, u_r\right) &= \beta_{22,0} + \beta_{22,1} \mathrm{ECI}_{t-2,t-3}^{(r)} + U_r,\end{aligned} \tag{6.3}$$

where U_r represents the spatial effect introduced in (3.5) and (3.6). Following Section 3.2, the parameters governing the transition probabilities are independent of the age group x. Based on BIC, we select much simpler structures compared to the ones in (6.1) and (6.2). Specifically, we only retain variables related to the excess cold and heat indices.

Parameter estimates. We follow Section 4.2.3 and apply the EM-algorithm to estimate the regime-switching parameter vector $\boldsymbol{\theta}$. Following Suppl. Mat. B.5, we apply profile likelihood maximization to select the precision parameter τ. We use the grid $\mathcal{G} = \{0.001, 0.01, 0.1, 1, 10, 100, 1000\}$ as candidate values and select $\tau^* = 1$.

Fig. 7 (top panel) shows the parameter estimates for $\boldsymbol{\alpha}_1$. Following Suppl. Mat. B.6, we construct 95% confidence intervals using the Fisher Information matrix based on 25 000 bootstrap samples. Excess heat in the current week is strongly associated with excess mortality, with the effect increasing with age. At a lag of one week, the effects remain positive but decrease with age. Notably, when the current week shows no excess heat, the overall effect can become positive for large values of last week's excess heat indicator or negative for moderate values. This latter situation reflects the so-called harvesting or mortality displacement effect (Hajat and Kosatky, 2010), which disappears in case of extreme heat (Pascal et al., 2018).

The bottom panel of Fig. 7 shows the estimates and 95% confidence intervals for the parameter vector $\boldsymbol{\alpha}_2$, which captures the effects of cold spells, increased influenza activity, and COVID-19 hospitalizations on mortality. First, we find that cold spells are not associated with excess mortality in the current week (lag 0), except for the oldest age group (90+). At lags 1 and 2-3, this effect becomes more pronounced and also increases with age. This indicates that cold has a more delayed and longer-lasting impact on excess mortality, see e.g. (Pascal et al., 2018). Second, a higher influenza activity index in the current week is associated with higher excess mortality. These associations are stronger among older age groups. The effect of lag 1 is less pronounced but increases again at lags 2-3. The larger estimated effect at lags 2-3 has also been found in Wong et al. (2004). Third, the effect of high COVID-19 hospital admissions appears to be more short-lived, being strongest in

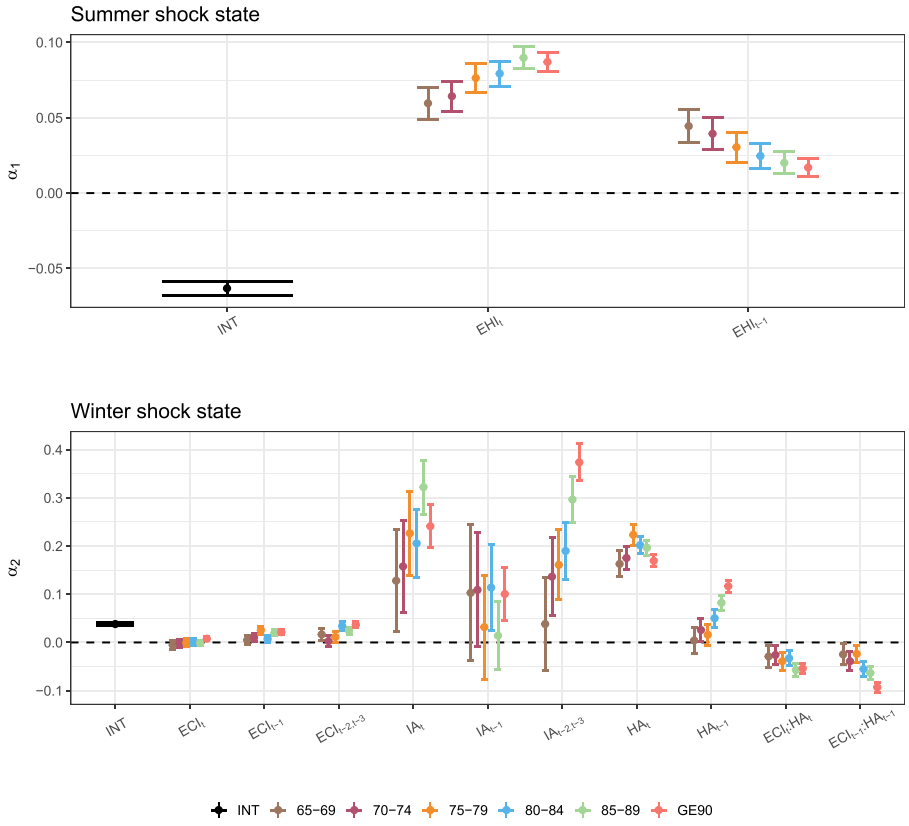

Fig. 7. Parameter estimates for $\boldsymbol{\alpha}_1$ (state 1, summer shock state) and $\boldsymbol{\alpha}_2$ (state 2, winter shock state) in the Poisson models of the regime-switching model, along with their 95% confidence intervals based on 25 000 bootstrap samples. The label 'INT' refers to the intercept.

the current week and only significant at lag 1 among older individuals (80+). Fourth, regarding the interaction effect, we find that a simultaneous cold spell and high COVID-19 hospitalization rates do not result in an additional increase in deaths. A possible explanation is that deaths attributable to cold spells and those attributable to respiratory diseases such as COVID-19 largely involve the same vulnerable population (primarily older and frail individuals). As such, these risk factors affect overlapping rather than distinct groups, suggesting a competing-risk dynamic, leading to a negative parameter value for the interaction effect (Andersen et al., 2012). Another contributing factor may be that during a pandemic outbreak, many vulnerable individuals are hospitalized, limiting their exposure to cold spells. Consequently, those remaining exposed are generally more resilient, which may lead to lower observed mortality during such periods.

Fig. 8 displays the estimates and 95% confidence intervals of the parameter vectors $\boldsymbol{\beta}_{01}$, $\boldsymbol{\beta}_{02}$, $\boldsymbol{\beta}_{11}$, and $\boldsymbol{\beta}_{22}$ in the transition probabilities of (6.3). The results indicate that transitions from the baseline state to the summer shock state become more likely when the current week's excess heat index is high. Similarly, transitions to the winter shock state are more probable when the current week's excess cold index is high. Once the system enters the summer shock state, the probability of remaining in this state decreases as the current week's excess heat index rises. In the winter shock regime, the probability of state persistence decreases with the value of the excess cold index at lags 2-3, although this effect is marginally insignificant. These findings suggest that the duration of the summer shock state tends to be shorter than that of the winter shock state.

Spatial effect. Fig. 9 visualizes the estimated mode of the expected log joint density, i.e., the spatial effect $\boldsymbol{u}^*$ (see Suppl. Mat. B.4). The mode is the largest in North/North-East France, implying that state transitions occur more frequently in these regions.

6.2. Quantifying uncertainty in-sample model predictions

6.2.1. Best-estimate in-sample fit

We use the estimated baseline mortality model and construct a Best-Estimate (BE) in-sample fit for the regime-switching framework. We do this by using the estimated marginal state probabilities. For any $t \in \mathcal{T}$, we define the BE regime-switching trajectory as:

$$s^{(r)}_{\text{BE},t} = \underset{j\in\{0,1,2\}}{\text{argmax}}\, \hat{P}\left(S^{(r)}_t = j \mid \mathcal{D}^{(r)}_t, \mathcal{Z}^{(r)}_t; \boldsymbol{\theta}^*, \boldsymbol{u}^*\right), \tag{6.4}$$

where $\boldsymbol{\theta}^*$ and $\boldsymbol{u}^*$ are the optimal regime-switching parameter vector and the mode of the spatial effect. The sets $\mathcal{D}^{(r)}_t$ and $\mathcal{Z}^{(r)}_t$ denote the collection of death counts and covariate vectors up to time t in region r. Suppl. Mat. B.4 outlines the calculation of the marginal state probabilities. We define the BE prediction for the weekly death counts, see (3.4), as:

$$\hat{d}^{(r)}_{\text{BE},x,t} = \hat{\mathbb{E}}\left[D^{(r)}_{x,t} \mid S^{(r)}_t = s^{(r)}_{\text{BE},t}\right]. \tag{6.5}$$

Fig. 10 presents the BE in-sample predicted weekly death counts in Pays de la Loire and Provence-Alpes-Côte d'Azur for the age groups 75-79 and 90+. Section B of the Online Appendix contains the figures for all French NUTS 2 regions and age groups. The colored bars at the bottom of each panel indicate the BE regime-switching trajectory. The figure shows that our model captures well the observed spikes in the weekly mortality patterns. The inclusion of the (lagged) hospital admissions index is useful to account for the impact of the COVID-19 pandemic on mortality. Furthermore, the model successfully captures excess deaths linked with cold snaps or increased influenza activity during the winter seasons of 2014/2015 and 2016/2017, and those related to heatwaves in the summers of 2017 and 2022.

To quantify the model performance, we calculate the Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) between the

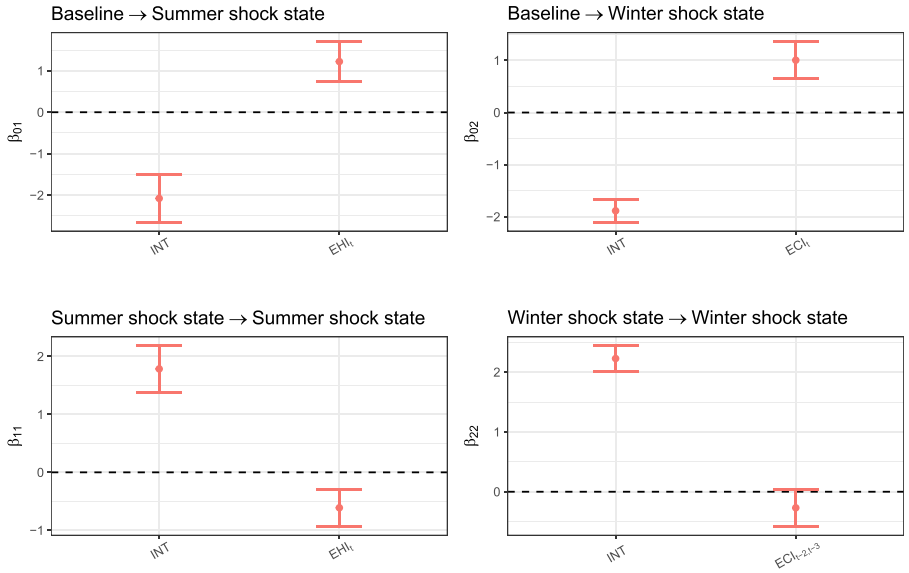

Fig. 8. Estimates of the parameter vectors in the transitional probabilities and their 95% confidence intervals based on 25 000 bootstrap samples: $\boldsymbol{\beta}_{01}$ (top left), $\boldsymbol{\beta}_{02}$ (top right), $\boldsymbol{\beta}_{11}$ (bottom left), and $\boldsymbol{\beta}_{22}$ (bottom right). The label 'INT' refers to the intercept.

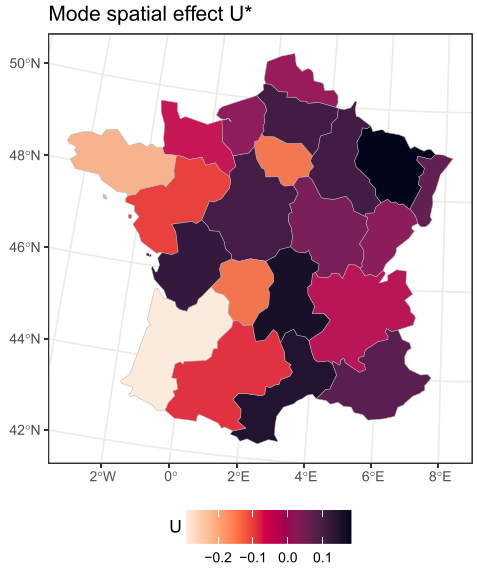

Fig. 9. The mode of the spatial effect $\boldsymbol{u}^*$ in the expected log joint density.

observed and estimated weekly death counts from both the baseline mortality model (see (3.3)) and from the BE prediction of the regime-switching model (see (6.5)). For each age group $x \in \mathcal{X}$, we compute the RMSE and MAE as:

$$\text{RMSE}_x = \sqrt{\frac{1}{n_x} \sum_{r\in\mathcal{R}} \sum_{t\in\mathcal{T}} \left(d_{x,t}^{(r)} - \hat{d}_{x,t}^{(r)}\right)^2},$$

$$\text{MAE}_x = \frac{1}{n_x} \sum_{r\in\mathcal{R}} \sum_{t\in\mathcal{T}} \left|d_{x,t}^{(r)} - \hat{d}_{x,t}^{(r)}\right|,$$

where n_x denotes the total number of observations for age group x, and $d_{x,t}^{(r)}$ and $\hat{d}_{x,t}^{(r)}$ denote the observed and estimated death counts, respectively. Similarly, we compute the RMSE and MAE per region (RMSE_r and MAE_r) by replacing the summation over regions with a summation over the age groups $x \in \mathcal{X}$.

Table 2 presents the results by age group, while Table 3 reports the results across the 21 French NUTS 2 regions. Overall, the regime-switching model achieves clearly lower RMSE and MAE values compared to the weekly baseline mortality model, with the improvement being more pronounced in older age groups. At the regional level, we observe the largest performance gains in Île-de-France (FR10), followed by Rhône-Alpes (FRK2) and Provence-Alpes-Côte d'Azur (FRL0), whereas the smallest improvements occur in Limousin (FRI2) and Basse-Normandie (FRD1). These findings suggest that, in the age groups and regions showing the largest improvement, the inclusion of environmental and respiratory covariates is most beneficial.

Table 2
Model performance evaluated using the root mean square error (RMSE) and mean absolute error (MAE) between observed and estimated weekly death counts for the baseline and regime-switching model (best-estimate prediction). Metrics are computed across six age groups and overall. Relative improvement (Rel. Imp.) reflects the proportional reduction in the error metric of the regime-switching model compared with the baseline.

	RMSE			MAE		
Age group	Baseline	Regime Switch	Rel. Imp.	Baseline	Regime Switch	Rel. Imp.
65–69	6.90	6.32	0.08	4.96	4.77	0.04
70–74	8.13	6.96	0.14	5.56	5.22	0.06
75–79	9.34	7.57	0.19	6.33	5.72	0.10
80–84	12.37	9.32	0.25	7.86	6.98	0.11
85–89	16.94	11.11	0.34	10.07	8.31	0.18
90+	27.51	16.22	0.41	15.07	11.55	0.23
Overall	15.26	10.16	0.33	8.31	7.09	0.15

6.2.2. *Uncertainty in-sample predictions*

We follow the methodology in Section 5 and quantify the uncertainty in the in-sample predicted weekly death counts of the regime-switching modeling framework. Using a bootstrap sampling approach with 25 000 samples, we construct 95% prediction intervals. Fig. 11 provides an example of how state, parameter, spatial, and Poisson uncertainty contribute to the intervals for weekly death counts in the age groups 75–79 and 90+ within the Pays de la Loire and Provence-Alpes-Côte d'Azur NUTS 2 region. Section C of the Online Appendix contains the in-sample prediction intervals for all age groups and NUTS 2 regions considered in this paper. First, state uncertainty effectively captures spikes in weekly mortality patterns whenever they correspond to external factors such as heatwaves, heightened influenza activity, or increases in COVID-19 hospitalization rates. Second, Poisson uncertainty is especially

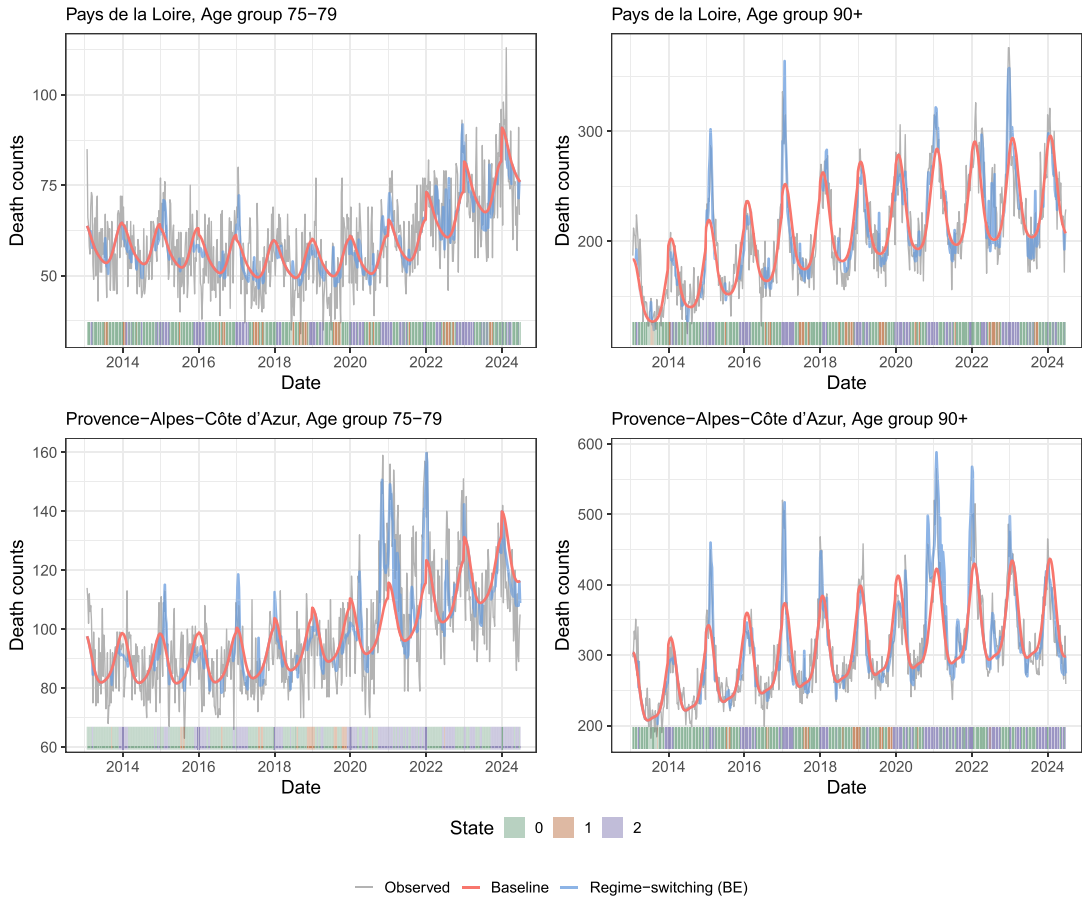

Fig. 10. The grey line represents the observed death counts from the first ISO week of 2013 to the 26th ISO week of 2024. The red line visualizes the estimated baseline number of deaths, while the blue line visualizes the best-estimate (BE) prediction of the death counts resulting from the regime-switching model. At the bottom, we show the BE regime-switching trajectory. We show the results for the age groups 75–79 (left) and 90+ (right) in Pays de la Loire (top) and Provence–Alpes-Côte d'Azur (bottom). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

prominent in the younger old-age groups. The underlying reason is the relatively small number of weekly death counts, and, consequently the more random fluctuations. In contrast, older age groups report higher weekly death counts, leading to substantially less Poisson uncertainty. Third, the parameter and spatial uncertainties are relatively minor.

Tables 4 and 5 show the contributions of the different sources of uncertainty to the total uncertainty across, respectively, the six considered age groups and the 21 NUTS 2 regions. We focus on the contributions of the state uncertainty, the combined parameter and spatial uncertainty, and the Poisson uncertainty. For instance, we calculate the contribution of the state uncertainty as the proportion of the red area (in Fig. 11) to the total area, corresponding to the 95% uncertainty bounds. We follow a similar strategy for the parameter and spatial, and the Poisson uncertainty. The summary tables indicate that the contribution of the state uncertainty increases with age, and the contribution of the Poisson uncertainty declines with age, but remains the most important source of uncertainty. The contributions of the parameter and spatial uncertainties remain relatively minor and constant across both age groups and regions. By region, state uncertainty is highest in the most populated area, Île-de-France (FR10), followed by the South-Eastern NUTS 2 regions of Rhône-Alpes (FRK2) and Provence-Alpes-Côte d'Azur (FRL0). This indicates that the summer and winter shocks have the most significant impact on older people and in high population areas.

6.3. *Out-of-sample back-testing*

We conduct out-of-sample back-testing to evaluate the model's performance on unseen data. Specifically, we limit the calibration period to the first ISO week of 2013 until the 26th ISO week of 2022, denoted as $\tilde{\mathcal{T}}$. The out-of-sample period ranges from the 27th week of 2022 to the 26th week of 2024, covering two years, and we refer to it as $\mathcal{T}_{\text{pred}}$.

We first reconstruct the features using the shorter calibration period $\tilde{\mathcal{T}}$. For instance, we recompute the excess heat and cold indices based on the 5% and 95% temperature quantiles. Second, we recalibrate the baseline mortality model in Section 4.1, and the three-state regime-switching model in Section 3.2, on $\tilde{\mathcal{T}}$. We represent the corresponding parameter estimates in Suppl. Mat. D.2. Importantly, these estimates closely align with those obtained from the original calibration period (see Figs. 7 and 8), which highlights the robustness and stability of the parameters across a different (shorter) calibration period.

We then generate projections for the death counts in $\mathcal{T}_{\text{pred}}$ while accounting for parameter, spatial, state, and Poisson uncertainty, as described in Section 5. In the back-testing, we explore a hypothetical scenario that uses observed values for the excess heat and cold index, influenza activity, and COVID-19 hospitalizations for the two-year prediction period. This allows us to gain insights into how well the model captures excess deaths conditionally on the observed evolution of these covariates.

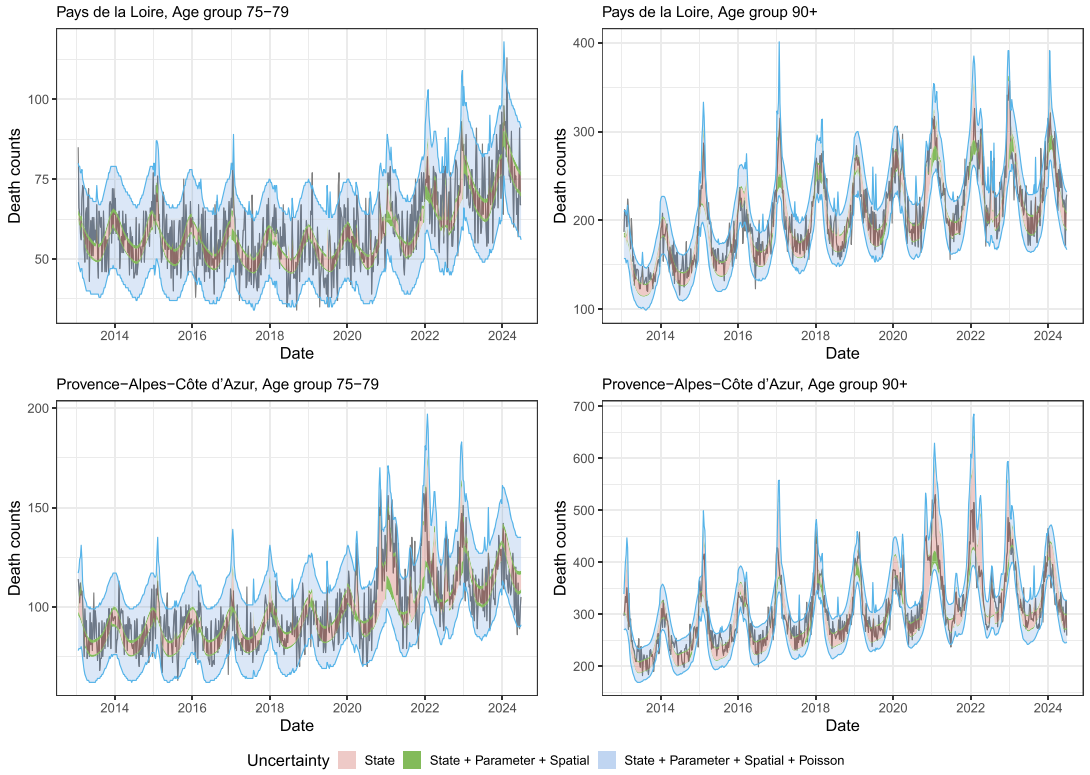

Fig. 11. We present 95% confidence intervals for the weekly death counts in Pays de la Loire (top) and Provence-Alpes-Côte d'Azur (bottom) for the age groups 75-79 (left), and 90+ (right) from the first ISO week of 2013 till the 26th ISO week of 2024. The intervals represent different sources of uncertainty: state uncertainty only (red), state, parameter, and spatial uncertainty (green), and state, parameter, spatial, and Poisson uncertainty (blue). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 3
Model performance evaluated using the root mean square error (RMSE) and mean absolute error (MAE) between observed and estimated weekly death counts for the baseline and regime-switching model (best-estimate prediction). Metrics are computed across the 21 French NUTS 2 regions and overall. Relative improvement (Rel. Imp.) reflects the proportional reduction in the error metric of the regime-switching model compared with the baseline.

	RMSE			MAE		
Region	Baseline	Regime Switch	Rel. Imp.	Baseline	Regime Switch	Rel. Imp.
FR10	43.08	19.30	0.55	18.77	13.74	0.27
FRB0	10.93	9.00	0.18	7.71	6.79	0.12
FRC1	9.65	7.97	0.17	6.69	5.89	0.12
FRC2	7.21	5.95	0.17	5.01	4.49	0.10
FRD1	7.75	6.87	0.11	5.79	5.30	0.08
FRD2	8.16	7.09	0.13	5.93	5.34	0.10
FRE1	13.68	11.19	0.18	9.92	8.45	0.15
FRE2	9.03	7.46	0.17	6.35	5.64	0.11
FRF1	10.71	7.35	0.31	6.32	5.40	0.15
FRF2	7.32	6.21	0.15	5.33	4.77	0.10
FRF3	12.30	9.12	0.26	8.15	6.89	0.15
FRG0	13.16	10.65	0.19	9.33	8.02	0.14
FRH0	12.18	10.49	0.14	8.84	7.93	0.10
FRI1	12.74	11.07	0.13	9.06	8.24	0.09
FRI2	5.91	5.45	0.08	4.33	4.03	0.07
FRI3	9.39	8.18	0.13	6.74	6.08	0.10
FRJ1	11.82	9.73	0.18	8.42	7.30	0.13
FRJ2	12.03	10.15	0.16	8.39	7.54	0.10
FRK1	8.62	7.24	0.16	5.92	5.33	0.10
FRK2	26.55	15.77	0.41	14.23	11.06	0.22
FRL0	19.65	14.79	0.25	13.25	10.66	0.20
Overall	15.26	10.16	0.33	8.31	7.09	0.15

Table 4
Contribution of the state uncertainty (State), parameter & spatial uncertainty (Par. & Spat.), and Poisson uncertainty (Poisson), to the total uncertainty, across the six considered age groups and based on 95% uncertainty bounds.

Age group	State	Par. & Spat.	Poisson
65–69	0.09	0.07	0.84
70–74	0.12	0.07	0.81
75–79	0.15	0.06	0.79
80–84	0.18	0.07	0.75
85–89	0.23	0.07	0.70
90+	0.29	0.07	0.64

Table 5
Contribution of the state uncertainty (State), parameter & spatial uncertainty (Par. & Spat.), and Poisson uncertainty (Poisson), to the total uncertainty, across the 21 French NUTS 2 regions and based on 95% uncertainty bounds.

Region	State	Par. & Spat.	Poisson	Region	State	Par. & Spat.	Poisson
FR10	0.31	0.08	0.61	FRG0	0.19	0.07	0.74
FRB0	0.19	0.05	0.76	FRH0	0.18	0.09	0.73
FRC1	0.16	0.06	0.78	FRI1	0.18	0.09	0.73
FRC2	0.13	0.06	0.81	FRI2	0.12	0.06	0.82
FRD1	0.14	0.06	0.80	FRI3	0.15	0.06	0.79
FRD2	0.15	0.06	0.79	FRJ1	0.18	0.06	0.76
FRE1	0.23	0.08	0.69	FRJ2	0.19	0.07	0.74
FRE2	0.16	0.06	0.78	FRK1	0.15	0.05	0.80
FRF1	0.16	0.07	0.77	FRK2	0.27	0.07	0.66
FRF2	0.14	0.05	0.81	FRL0	0.28	0.07	0.65
FRF3	0.19	0.06	0.75				

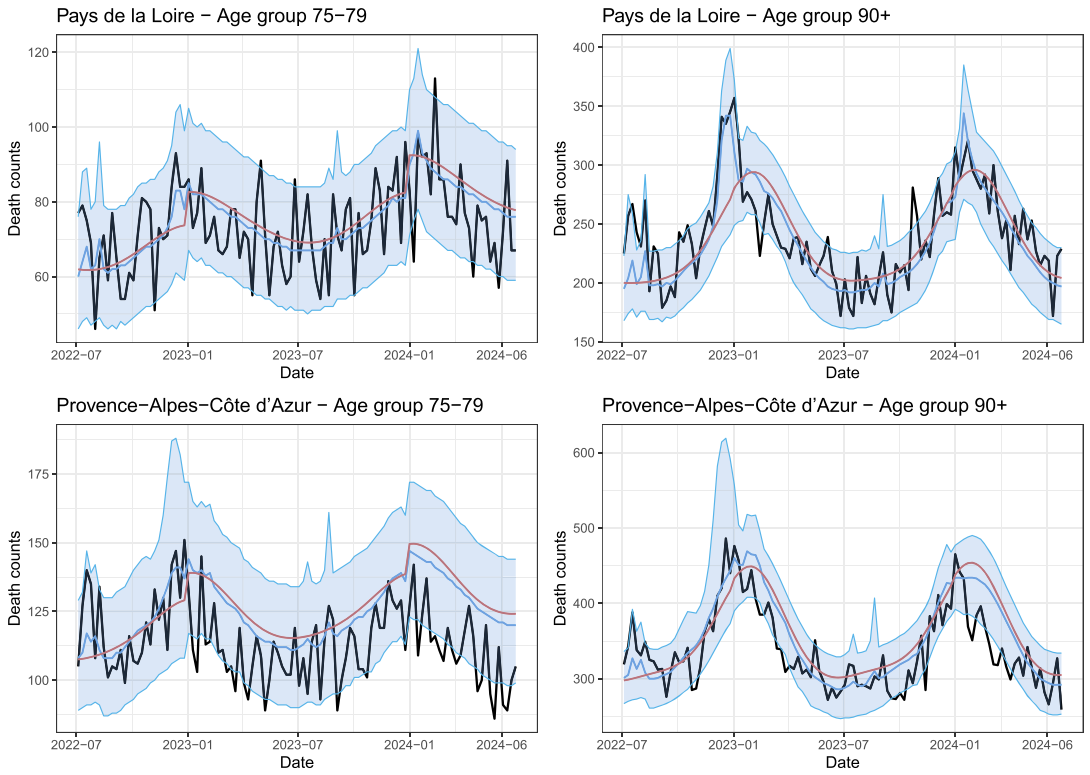

Fig. 12. We present the 95% prediction intervals for the weekly death counts in Pays de la Loire (top) and Provence-Alpes-Côte d'Azur (bottom) for the age groups 75–79 (left) and 90+ (right), from the 27th ISO week of 2022 to the 26th ISO week of 2024. The intervals are based on 25 000 bootstrap samples and consider parameter, spatial, state, and Poisson uncertainty. The blue lines show the 2.5%, 50%, and 97.5% quantile, the red line visualizes the estimated baseline number of deaths, and the black line shows the observed death counts. We also consider two other prediction intervals, see Suppl. Mat. E. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Fig. 12 illustrates the 95% prediction intervals for the weekly death counts in Pays de la Loire and Provence-Alpes-Côte d'Azur for age groups 75–79 and 90+. The remaining illustrations can be found in Section D of the Online Appendix. These intervals incorporate parameter, spatial, state, and Poisson uncertainty, as described in Section 5. As Fig. 12 indicates, the prediction intervals effectively capture the uncertainty inherent in the observed death counts. The model captures the excess deaths observed in the 2022/2023 winter season for the age group 90+. Furthermore, the prediction bounds are narrower for the winter of 2023/2024, aligning with the observed death rates.

As a quantitative comparison of the prediction performance between the weekly baseline and regime-switching model on the out-of-sample period $\mathcal{T}_{\text{pred}}$, we compute the RMSE and MAE between the observed and estimated death counts resulting from both models. Specifically, for the regime-switching model, we take the estimated death counts as resulting from the 50% quantile of the simulated regime state trajectories (blue line in Fig. 12). Tables 6 and 7 summarize the results across age groups and French NUTS 2 regions, respectively. The improvement in the predictions of the regime-switching model relative to the ones of the baseline model are more pronounced for older age groups and in Île-de-France (FR10), and the South-Eastern NUTS 2 regions Rhône-Alpes (FRK2) and Provence-Alpes-Côte d'Azur (FRL0).

Lastly, in Suppl. Mat. E, we perform an expanding-window backtest where the model is repeatedly re-estimated with increasing data and used to forecast 1-3 years ahead. We show that the prediction intervals capture well the observed variability and excess mortality periods. The best-estimate projection of the regime-switching model also improves accuracy by about 8% at the 1-year horizon, while gains at the 3-year horizon are a bit less pronounced (≈ 2% improvement).

Table 6
Evaluation of the prediction performance on the out-of-sample period using the Root mean square error (RMSE) and mean absolute error (MAE) between observed and predicted weekly death counts for the baseline and regime-switching model (50% quantile). Metrics are computed across six age groups and overall. Relative improvement (Rel. Imp.) reflects the proportional reduction in the error metric of the regime-switching model compared with the baseline.

	RMSE			MAE		
Age group	Baseline	Regime Switch	Rel. Imp.	Baseline	Regime Switch	Rel. Imp.
65–69	6.82	6.64	0.03	5.10	4.96	0.03
70–74	8.94	8.48	0.05	6.73	6.42	0.05
75–79	11.04	10.06	0.09	8.17	7.48	0.08
80–84	10.22	9.52	0.07	7.50	7.08	0.06
85–89	13.96	12.22	0.12	10.13	9.05	0.11
90+	24.46	21.46	0.12	17.03	15.18	0.11
Overall	13.82	12.37	0.11	9.11	8.36	0.08

6.4. *Short-term mortality forecasting with scenario-based approaches*

To avoid constructing a point forecast for the exposure-to-risk, we focus on forecasting the death rates per 1 000 person-years in the short-term. Hereto, we use temperature scenarios based on the RCP 4.5, and RCP 8.5 pathways from the Copernicus Climate Change Service (Copernicus Climate Change Service, 2021). For influenza, we generate moderate and severe incidence scenarios using a Bayesian stochastic SIRS

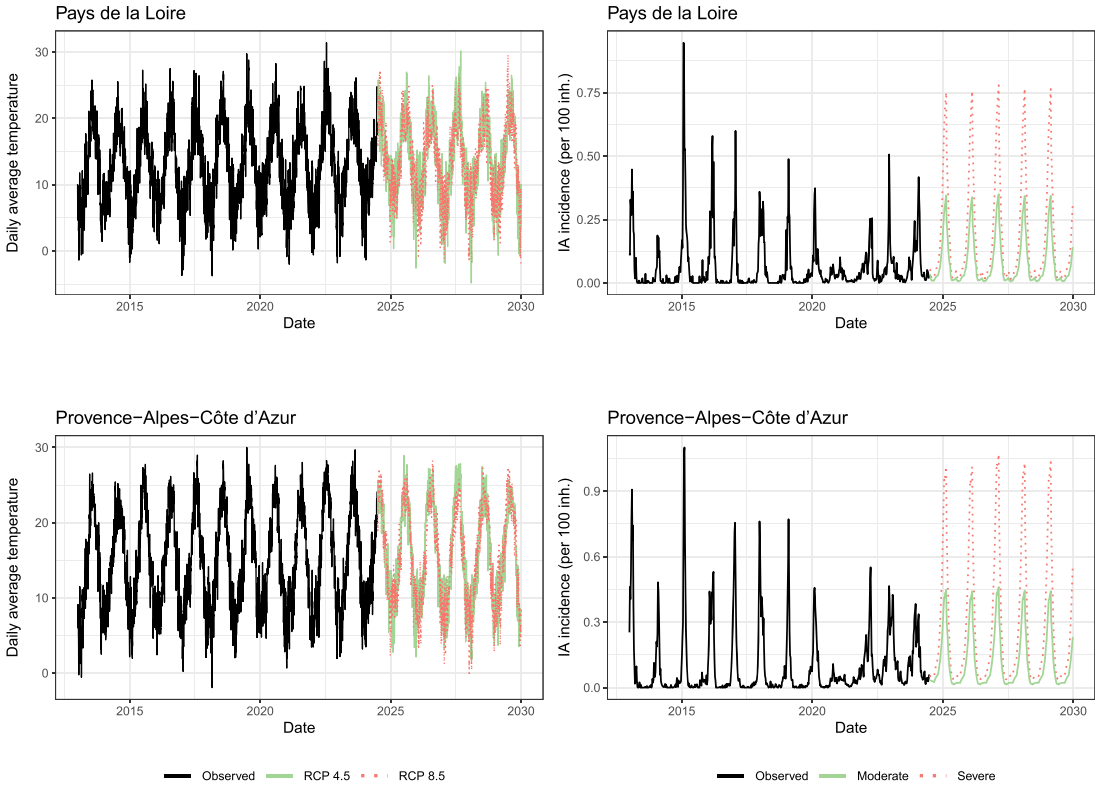

Fig. 13. Projections for the daily average temperature based on the RCP 4.5, and RCP 8.5 scenarios (left panel), and for the weekly influenza activity per 100 inhabitants based on the proposed Bayesian stochastic SIRS model (right panel). We provide a moderate and severe incidence scenario for influenza activity using the point-wise 50%, and 95% quantiles of the posterior predictive samples, respectively. The top panels show the results for Pays de la Loire, while the bottom panels depict the results for Provences-Alpes-Côte d'Azur.

Table 7
Evaluation of the prediction performance on the out-of-sample period using the Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) between observed and predicted weekly death counts for the baseline and regime-switching model (50% quantile). Metrics are computed across the 21 French NUTS 2 regions and overall. Relative improvement (Rel. Imp.) reflects the proportional reduction in the error metric of the regime-switching model compared with the baseline.

	RMSE			MAE		
Region	Baseline	Regime Switch	Rel. Imp.	Baseline	Regime Switch	Rel. Imp.
FR10	29.99	26.02	0.13	21.08	18.26	0.13
FRB0	11.06	9.74	0.12	8.06	7.35	0.09
FRC1	8.81	8.60	0.02	6.54	6.33	0.03
FRC2	6.32	6.15	0.03	4.71	4.63	0.02
FRD1	8.92	8.17	0.08	6.47	6.05	0.06
FRD2	9.25	8.59	0.07	6.78	6.48	0.04
FRE1	15.17	13.97	0.08	11.23	10.28	0.08
FRE2	9.57	8.80	0.08	7.02	6.64	0.06
FRF1	9.01	8.17	0.09	6.57	6.08	0.07
FRF2	8.00	7.36	0.08	6.06	5.63	0.07
FRF3	10.51	9.75	0.07	8.01	7.48	0.07
FRG0	15.91	13.63	0.14	11.22	10.02	0.11
FRH0	14.43	12.87	0.11	10.17	9.45	0.07
FRI1	13.98	13.24	0.05	9.93	9.48	0.05
FRI2	5.94	5.70	0.04	4.43	4.20	0.05
FRI3	10.61	9.72	0.08	7.46	7.02	0.06
FRJ1	12.05	10.62	0.12	8.82	8.09	0.08
FRJ2	12.50	11.73	0.06	8.96	8.57	0.04
FRK1	8.02	7.77	0.03	6.05	5.83	0.04
FRK2	21.11	18.45	0.13	15.40	13.59	0.12
FRL0	22.68	19.85	0.12	16.35	14.17	0.13
Overall	13.82	12.37	0.11	9.11	8.36	0.08

model (Bucyibaruta et al., 2023), which is able to generate forecasts for weekly new influenza cases based on posterior predictive samples. COVID-19 hospitalizations are set to zero, but alternative scenarios can be used. We include all technical details in Suppl. Mat. F. Fig. 13 shows the results of the constructed scenarios for Pays de la Loire and Provences-Alpes-Côte d'Azur. Section E of the Online Appendix contains the results for the other regions.

We forecast death rates from the 27th ISO-week of 2024 till the 52nd ISO week of 2029 for all NUTS 2 regions in France. Specifically, we construct the excess heat and cold index (see (2.1)), and the influenza activity index based on these temperature and influenza scenarios during the forecasting period. We then apply the forecasting strategy discussed in Section 5 to incorporate parameter, spatial, state, and Poisson uncertainty in the predicted mortality outcomes. We construct four prediction intervals corresponding to four different scenarios:

1. moderate influenza and RCP 4.5 scenario,
2. severe influenza and RCP 4.5 scenario,
3. moderate influenza and RCP 8.5 scenario,
4. severe influenza and RCP 8.5 scenario.

Fig. 14 shows the 95% prediction intervals for the death rates at age group 90+ in Pays de la Loire and Alpes Côte d'Azur. The left panels display the blue and orange intervals associated with scenario 1 (moderate influenza, RCP 4.5) and scenario 2 (severe influenza, RCP 4.5), while the right panels show the green and red intervals corresponding to scenario 3 (moderate influenza, RCP 8.5) and scenario 4 (severe influenza, RCP 8.5). Prediction intervals for scenarios 1 versus 2 and for scenarios 3 versus 4 differ substantially in winter because higher

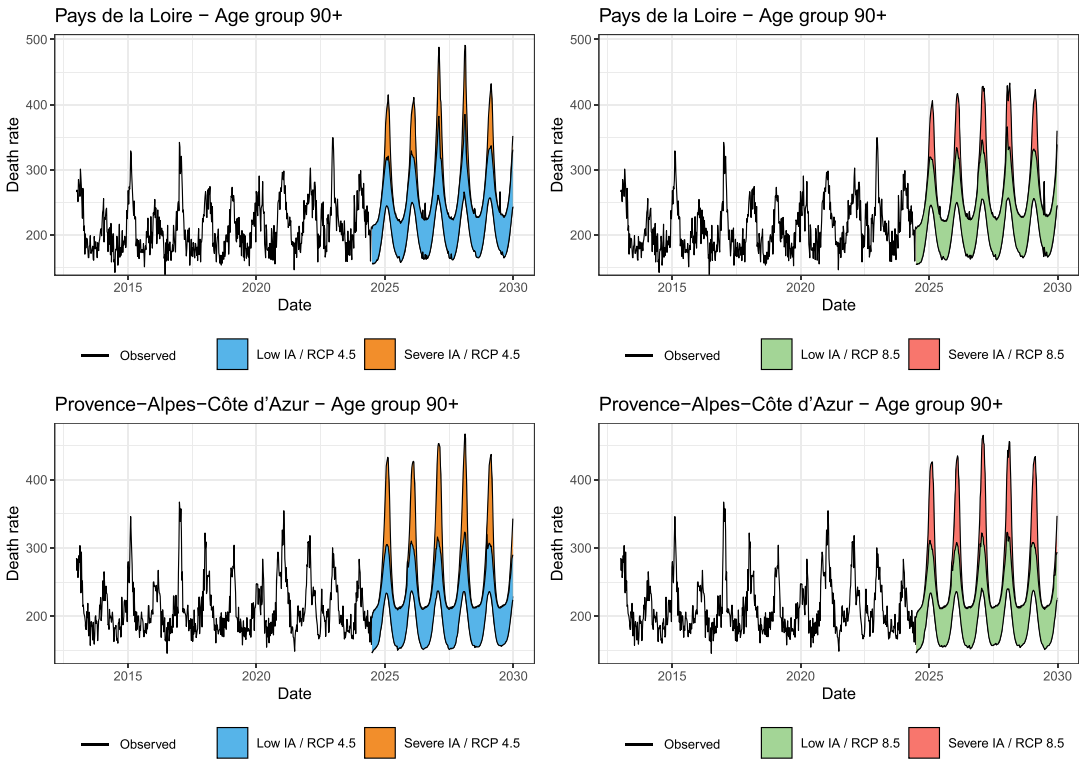

Fig. 14. We present 95% prediction intervals for the weekly death rates per 1 000 person-years in Pays de la Loire (top row) and Provence-Alpes-Côte d'Azur (bottom row) for the age group 90+, from the 27th ISO week of 2024 to the 52nd ISO week of 2029. The left panels show the results for the moderate (blue) and severe (orange) influenza scenario combined with the RCP 4.5 temperature projection, whereas the right panels display the results for the moderate (green) and severe (red) influenza scenario combined with the RCP 8.5 temperature projection. The intervals take into account parameter, spatial, state and Poisson uncertainty and are based on 25 000 bootstrap samples. The black line depicts the observed death rates per 1 000 person-years, from the first ISO week of 2013 to the 26th ISO week of 2024. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

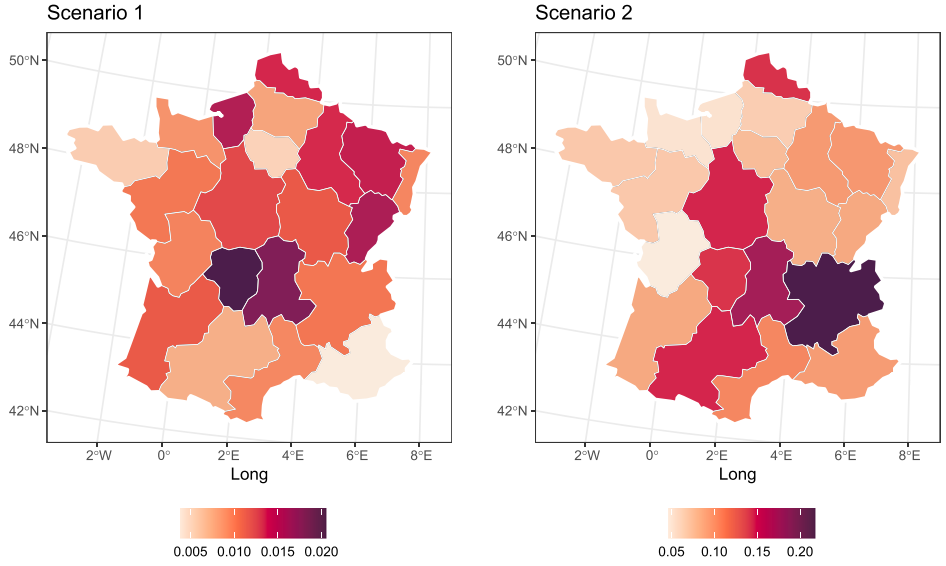

Fig. 15. We present the relative excess deaths compared to the baseline mortality model for the age group 90+ in the French NUTS 2 regions under scenario 1 (moderate influenza, RCP 4.5) and scenario 2 (severe influenza, RCP 4.5).

mortality is expected under the severe influenza scenario. For Pays de la Loire, the upper bounds of the intervals for scenarios 1 and 2 during the winters of 2026/27 and 2027/28 exceed those of scenarios 3 and 4. This occurs because, under RCP 4.5, the winters are projected to be colder than under RCP 8.5 (see Fig. 12). This pattern does not appear in Provence-Alpes-Côte d'Azur. Finally, during the summer weeks, the prediction intervals under RCP 4.5 differ little from those under RCP 8.5. As shown in Fig. 12, summer temperatures in both scenarios remain relatively moderate and similar, resulting in comparable bounds during the weeks. The scenario-based mortality projections can provide insurers with valuable insights into how differing climate and influenza conditions may shape future mortality risks. Section F of the Online

Appendix reports the prediction intervals of scenarios 1 and 2 for all age groups and NUTS-2 regions considered.

Next, we estimate the excess deaths relative to the baseline mortality model under the different scenarios. We put focus on scenarios 1 and 2 because of the minor differences between the RCP 8.5 and RCP 4.5 scenarios in the short-term. Hereto, we project the baseline mortality model on the forecasting period, from the 27th ISO week of 2024 until the 52nd ISO week of 2029, and compute the estimated forecasted number of deaths per 1 000 person-years per week. We sum up the results over the entire forecasting period. Next, we calculated the number of deaths (per 1 000 person-years per week) for the median quantile of each scenario. We then determine the relative excess deaths by comparing the scenario-based death rates with the baseline setting, expressed as a relative increase. We do this analysis for each region and age group. Fig. 15 shows the results for the 90+ age group, while Section G of the Online Appendix displays the results for the age groups 65–69, 70–74, 75–79, 80–84, and 85–89. The results highlight clear regional differences. Scenario 1 leads to moderate increases in deaths compared to the baseline, while Scenario 2 shows much sharper rises. Although the scenarios for the involved covariates (influenza, temperature) are purely hypothetical, they provide useful insights into how different conditions can impact regions in France very differently.

While our short-term mortality forecasts from the regime-switching model account for parameter, state, spatial, and Poisson uncertainty, the considered RCP scenarios provide only deterministic temperature paths and ignore climate uncertainty. To capture this, Avanzi et al. (2025) propose to generate multiple future climate trajectories to account for both aleatoric and climate model uncertainties. Specifically, they combine ensembles of climate models to capture model uncertainty, then apply bias correction against historical data, and, finally, add stochastic residuals to represent aleatoric uncertainty. Translating this to our setting, we could generate multiple stochastic temperature trajectories for each RCP scenario (RCP 4.5, 8.5) and feed these into the regime-switching mortality model. Each trajectory would produce a corresponding mortality projection, thereby incorporating climate projection uncertainty in addition to the existing parameter, state, spatial, and Poisson uncertainty described in Section 5. While this extension would allow for a more complete assessment of uncertainty, the scope of the current study puts more focus on illustrating how mortality evolves under a given temperature scenario rather than providing a full probabilistic climate-mortality forecast. As such, we leave it as an avenue for future research.

7. Conclusion

This paper introduces an intuitive framework for incorporating environmental and epidemic factors from detailed, open-source data into a weekly mortality modeling framework. Specifically, we develop a three-state regime-switching model, which includes a baseline state reflecting seasonal mortality trends, a summer shock state for periods of heightened mortality due to heatwaves, and a winter shock state to account for cold spells, and increased mortality linked to influenza or COVID-19 hospitalizations.

Based on a case study on 21 regions and six older age groups in France, we detect regime periods of higher mortality associated with heat waves and cold spells, increased influenza activity or hospital admissions. In the summer shock state, we estimate that the factor amplifying baseline mortality is highest for the oldest age groups. Furthermore, we find evidence of harvesting effects, reducing mortality in the weeks following a moderate hot week. In the winter shock state, the amplification of baseline mortality is associated with increased influenza activity, cold temperatures, and COVID-19-related hospitalizations. Through a back-test, we show that the model provides reliable short-term predictions of weekly mortality, aligning well with observed data.

We emphasize that the main purpose of this paper is to make short-term mortality predictions, and that the model in its current form is not intended for long-term predictions. To do this we must consider future improvements in heat resilience or the mitigating impact of public health adjustments such as increased vaccination rates or expanded cooling centers (Bobb et al., 2014; Ferguson et al., 2006). These interventions reduce heat- and respiratory-related mortality (Stone et al., 2014; Simonsen et al., 2005). The large-scale implementation of such measures usually takes several years and is therefore especially relevant to consider in long-term projections. However, this is beyond the scope of our short-term view in this paper.

While our model currently focuses on temperature, influenza, and COVID-19, we acknowledge that other drivers - including short-term air pollution spikes (Di et al., 2017), natural disasters such as flash floods or wildfires (Vinet et al., 2016; Tarín-Carrasco et al., 2021), economic crises (Falagas et al., 2009), and wars or terrorist attacks (Vandentorren et al., 2016) - can also lead to short-term deviations from baseline mortality. These additional fluctuations are now absorbed within the Poisson uncertainty of the model.

CRediT authorship contribution statement

Jens Robben: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Karim Barigou:** Writing – review & editing, Validation, Supervision, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization; **Torsten Kleinow:** Writing – review & editing, Supervision, Project administration, Investigation, Funding acquisition.

Data availability

The datasets used in this paper are publicly available. We consult the deaths by week, sex, 5-year age group and NUTS 2 region on Eurostat (https://ec.europa.eu/eurostat/web/products-datasets/-/demo_r_mwk2_05). We retrieve the daily average temperature from the E-OBS dataset on the Copernicus Climate Data Store (https://cds.climate.copernicus.eu/datasets/insitu-gridded-observations-europe), the incidences rates for influenza activity from the French Sentinelles network (https://www.sentiweb.fr/france/en/?page=table&maladie=25), and the hospital data related to the COVID-19 epidemic from the French government's open data platform (https://www.data.gouv.fr/fr/datasets/donnees-hospitalieres-relatives-a-lepidemie-de-covid-19/). The `R` code used to implement and analyze the case study presented in this paper is accessible in the GitHub repository: https://github.com/jensrobben/GranMoMoRS. The Online Appendix is available on the following GitHub link: https://jensrobben.github.io/articles/GranMoMoRS/Online_Appendix.pdf.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This study is part of the research programme at the Research Centre for Longevity Risk, a joint initiative of NN Group and the University of Amsterdam, with additional funding from the Dutch government's Public-Private Partnership (PPP) programme. Karim Barigou acknowledges the financial support from the AMF research funds in integrated risk management of financial institutions (Fonds AMF-GIRIF) from Université Laval.

Appendix A. Supplementary materials

The Supplementary Material (Suppl. Mat.) and Online Appendix associated with this article can be found in the online version at 10.1016/j.insmatheco.2026.103250.

References

Andersen, P.K., Geskus, R.B., de Witte, T., Putter, H., 2012. Competing risks in epidemiology: possibilities and pitfalls. Int. J. Epidemiol. 41 (3), 861–870.

Armstrong, B., 2006. Models for the relationship between ambient temperature and daily mortality. Epidemiology 17 (6), 624–631. https://doi.org/10.1097/01.ede.0000239732.50999.8f

Avanzi, B., Li, Y., Taylor, G., Wong, B., 2025. Dynamic financial analysis (DFA) of general insurers under climate change. arXiv preprint arXiv:2508.16444.

Barigou, K., Patten, M., Zhou, K.Q., 2025. Mortality modeling and forecasting with the actuaries climate index. arXiv preprint arXiv:2510.16266.

Basu, R., Samet, J.M., 2002. Relation between elevated ambient temperature and mortality: a review of the epidemiologic evidence. Epidemiol. Rev. 24 (2), 190–202. https://doi.org/10.1093/epirev/mxf007

Bégin, J.-F., Boudreault, M., Landry, T., 2025. Modelling seasonal mortality: an age–period–cohort approach. Insur. Math. Econ. 125, 103162.

Besag, J., Kooperberg, C., 1995. On conditional and intrinsic autoregressions. Biometrika 82 (4), 733–746. https://doi.org/10.1093/biomet/82.4.733

Bobb, J.F., Peng, R.D., Bell, M.L., Dominici, F., 2014. Heat-related mortality and adaptation to heat in the United States. Environ. Health Perspect. 122 (8), 811–816. https://doi.org/10.1289/ehp.1307392

Braga, A. L.F., Zanobetti, A., Schwartz, J., 2001. The time course of weather-related deaths. Epidemiology 12 (6), 662–667. https://doi.org/10.1097/00001648-200111000-00014

Bucyibaruta, G., Dean, C.B., Torabi, M., 2023. A discrete-time susceptible-infectious-recovered-susceptible model for the analysis of influenza data. Infect. Dis. Model. 8 (2), 471–483. https://doi.org/10.1016/j.idm.2023.04.008

Center for International Earth Science Information Network – CIESIN – Columbia University, 2018. Gridded population of the world, version 4 (GPWv4): population count, revision 11. Accessed 01-08-2023. https://doi.org/10.7927/H4JW8BX5

Cleveland, R.B., Cleveland, W.S., McRae, J.E., Terpenning, I., et al., 1990. STL: a seasonal-trend decomposition. J. Off. Stat 6 (1), 3–73.

Copernicus Climate Change Service, 2021. Temperature and precipitation climate impact indicators from 1970 to 2100 derived from European climate projections. https://doi.org/10.24381/cds.9eed87d5

Copernicus Climate Change Service, Climate Data Store, 2020. E-OBS daily gridded meteorological data for Europe from 1950 to present derived from in-situ observations. Accessed 01-09-2023. https://doi.org/10.24381/cds.151d3ec6

Craven, P., Wahba, G., 1978. Smoothing noisy data with spline functions: estimating the correct degree of smoothing by the method of generalized cross-validation. Numer. Math. 31 (4), 377–403.

Di, Q., Dai, L., Wang, Y., Zanobetti, A., Choirat, C., Schwartz, J.D., Dominici, F., 2017. Association of short-term exposure to air pollution with mortality in older adults. JAMA 318 (24), 2446–2456.

Eurostat, 2024b. Deaths by week, sex, 5-year age group and NUTS 2 region. https://doi.org/10.2908/DEMO_R_MWK2_05. Accessed October 8, 2024.

Falagas, M.E., Vouloumanou, E.K., Mavros, M.N., Karageorgopoulos, D.E., 2009. Economic crises and mortality: a review of the literature. Int. J. Clin. Pract. 63 (8), 1128–1135.

Feng, R., Garrido, J., 2011. Actuarial applications of epidemiological models. N. Am. Actuarial J. 15 (1), 112–136. https://doi.org/10.1080/10920277.2011.10597612

Ferguson, N.M., Cummings, D. A.T., Fraser, C., Cajka, J.C., Cooley, P.C., Burke, D.S., 2006. Strategies for mitigating an influenza pandemic. Nature 442 (7101), 448–452. https://doi.org/10.1038/nature04795

Gasparrini, A., Guo, Y., Hashizume, M., Lavigne, E., Zanobetti, A., Schwartz, J., Tobias, A., Tong, S., Rocklöv, J., Forsberg, B., Leone, M., De Sario, M., Bell, M.L., Guo, Y.-L.L., Wu, C.-f., Kan, H., Yi, S.-M., de Sousa Zanotti Stagliorio Coelho, M., Saldiva, P. H.N., Honda, Y., Kim, H., Armstrong, B., 2015. Mortality risk attributable to high and low ambient temperature: a multicountry observational study. Lancet 386 (9991), 369–375. https://doi.org/10.1016/S0140-6736(14)62114-0

Gemmell, I., McLoone, P., Boddy, F.A., Dickinson, G.J., Watt, G., 2000. Seasonal variation in mortality in Scotland. Int. J. Epidemiol. 29 (2), 274–279. https://doi.org/10.1093/ije/29.2.274

Goes, J., Barigou, K., Leucht, A., 2025. Bayesian mortality modelling with pandemics: a vanishing jump approach. J. R. Stat. Soc. Ser. C Appl. Stat. 74 (4), 1150–1182. https://doi.org/10.1093/jrsssc/qlaf018

Guibert, Q., Pincemin, G., Planchet, F., 2026. Impact of climate change on mortality: an extrapolation of temperature effects based on time series data in France. Int. J. Forecast.42 (2), 359–413.

Hajat, S., Kosatky, T., 2010. Heat-related mortality: a review and exploration of heterogeneity. J. Epidemiol. Community Health 64 (9), 753–760.

Hajat, S., O'Connor, M., Kosatsky, T., 2010. Health effects of hot weather: from awareness of risk factors to effective health protection. Lancet 375 (9717), 856–863. https://doi.org/10.1016/S0140-6736(09)61711-6

Iuliano, A.D., Roguski, K.M., Chang, H.H., Muscatello, D.J., Palekar, R., Tempia, S., Cohen, C., Gran, J.M., Schanzer, D., Cowling, B.J., et al., 2018. Estimates of global seasonal influenza-associated respiratory mortality: a modelling study. Lancet 391 (10127), 1285–1300. https://doi.org/10.1016/S0140-6736(17)33293-2

Kang, D., Ellgen, C., Kulstad, E., 2021. Possible effects of air temperature on COVID-19 disease severity and transmission rates. J. Med. Virol. 93 (9), 5358–5366.

Kok, G., Jonkers, R., Gelissen, R., Meertens, R., Schaalma, H., de Zwart, O., 2010. Behavioural intentions in response to an influenza pandemic. BMC Public Health 10, 1–19. https://doi.org/10.1186/1471-2458-10-174

Laplace, P.S., 1774. Mémoire sur la probabilité de causes par les évenements. Mémoire L'académie R. Sci. 6, 621–656.

Li, H., Tang, Q., 2022. Joint extremes in temperature and mortality: a bivariate POT approach. N. Am. Actuarial J. 26 (1), 43–63. https://doi.org/10.1080/10920277.2020.1823236

Li, L., Li, H., Panagiotelis, A., 2024. Boosting domain-specific models with shrinkage: an application in mortality forecasting. Int. J. Forecast.. https://doi.org/10.1016/j.ijforecast.2024.05.001

Louis, T.A., 1982. Finding the observed information matrix when using the EM algorithm. J. R. Stat. Soc. Ser. B Stat. Methodol. 44 (2), 226–233. https://doi.org/10.1111/j.2517-6161.1982.tb01203.x

Lytras, T., Pantavou, K., Mouratidou, E., Tsiodras, S., 2019. Mortality attributable to seasonal influenza in Greece, 2013 to 2017: variation by type/subtype and age, and a possible harvesting effect. Eurosurveillance 24 (14), 1800118. https://doi.org/10.2807/1560-7917.ES.2019.24.14.1800118

Mackenbach, J.P., Kunst, A.E., Looman, C.W., 1992. Seasonal variation in mortality in the Netherlands. J. Epidemiol. Community Health 46 (3), 261–265. https://doi.org/10.1136/jech.46.3.261

Marti-Soler, H., Gonseth, S., Gubelmann, C., Stringhini, S., Bovet, P., Chen, P.-C., Wojtyniak, B., Paccaud, F., Tsai, D.-H., Zdrojewski, T., et al., 2014. Seasonal variation of overall and cardiovascular mortality: a study in 19 countries from different geographic locations. PLoS ONE 9 (11), e113500. https://doi.org/10.1371/journal.pone.\protect\penalty-\@M0113500

Meng, L., Spall, J.C., 2017. Efficient computation of the Fisher information matrix in the EM algorithm. In: 2017 51st Annual Conference on Information Sciences and Systems (CISS). IEEE, pp. 1–6. https://doi.org/10.1109/CISS.2017.7926126

Meng, X.-L., Rubin, D.B., 1991. Using EM to obtain asymptotic variance-covariance matrices: the SEM algorithm. J. Am. Stat. Assoc. 86 (416), 899–909. https://doi.org/10.1080/01621459.1991.10475130

Milhaud, X., Gouthon, J.-L., 2024. Impact of heat waves on mortality: extension of longevity models to account for global warming. hal-04831333.

Min, J., Li, H., Nagler, T., Li, S., 2025. Assessing climate-driven mortality risk: a stochastic approach with distributed lag non-linear models. arXiv preprint arXiv:2506.00561.

Mølbak, K., Espenhain, L., Nielsen, J., Tersago, K., Bossuyt, N., Denissov, G., Baburin, A., Virtanen, M., Fouillet, A., Sideroglou, T., et al., 2015. Excess mortality among the elderly in European countries, December 2014 to February 2015. Eurosurveillance 20 (11), 21065.

Nielsen, J., Krause, T.G., Mølbak, K., 2018. Influenza-associated mortality determined from all-cause mortality, Denmark 2010/11-2016/17:tthe FluMOMO model. Influenza Other Respir. Viruses 12 (5), 591–604. https://doi.org/10.1111/irv.12564

Nielsen, J., Mazick, A., Glismann, S., Mølbak, K., 2011. Excess mortality related to seasonal influenza and extreme temperatures in Denmark, 1994–2010. BMC Infect. Dis. 11, 1–13. https://doi.org/10.1186/1471-2334-11-350

Nielsen, J., Vestergaard, L.S., Richter, L., Schmid, D., Bustos, N., Asikainen, T., Trebbien, R., Denissov, G., Innos, K., Virtanen, M.J., …, & Krause, T.G., 2019. European all-cause excess and influenza-attributable mortality in the 2017/18 season: should the burden of influenza B be reconsidered? Clin. Microbiol. Infect. 25 (10), 1266–1276. https://doi.org/10.1016/j.cmi.2019.02.011

Nørgaard, S.K., Vestergaard, L.S., Nielsen, J., Richter, L., Schmid, D., Bustos, N., Braye, T., Athanasiadou, M., Lytras, T., Denissov, G., et al., 2021. Real-time monitoring shows substantial excess all-cause mortality during second wave of COVID-19 in Europe, October to December 2020. Eurosurveillance 26 (2), 2002023. https://doi.org/10.2807/1560-7917.ES.2021.26.1.2002023

Notari, A., 2021. Temperature dependence of COVID-19 transmission. Sci. Total Environ. 763, 144390.

Oakes, D., 1999. Direct calculation of the information matrix via the EM. J. R. Stat. Soc. Ser. B (Stat. Methodol.) 61 (2), 479–482. https://doi.org/10.1111/1467-9868.00188

Pascal, M., Wagner, V., Corso, M., Laaidi, K., Ung, A., Beaudeau, P., 2018. Heat and cold related-mortality in 18 french cities. Environ. Int. 121, 189–198.

Pattenden, S., Nikiforov, B., Armstrong, B., 2003. Mortality and temperature in Sofia and London. J. Epidemiol. Community Health 57 (8), 628–633. https://doi.org/10.1136/jech.57.8.628

Robben, J., Antonio, K., 2024. Catastrophe risk in a stochastic multi-population mortality model. J. Risk Insur.. https://doi.org/10.1111/jori.12470

Robben, J., Antonio, K., Devriendt, S., 2022. Assessing the impact of the COVID-19 shock on a stochastic multi-population mortality model. Risks 10 (2), 26. https://doi.org/10.3390/risks10020026

Robben, J., Antonio, K., Kleinow, T., 2025. The short-term association between environmental variables and mortality: evidence from Europe. J. R. Stat. Soc. Ser. A Stat. Soc. 189 (2), 1131–1153. https://doi.org/10.1093/jrsssa/qnaf052

Robben, J., Barigou, K., 2025. A penalized distributed lag non-linear Lee-Carter framework for regional weekly mortality forecasting. arXiv preprint arXiv:2509.24087.

Schnürch, S., Kleinow, T., Korn, R., Wagner, A., 2022. The impact of mortality shocks on modelling and insurance valuation as exemplified by COVID-19. Ann. Actuarial Sci. 16 (3), 498–526. https://doi.org/10.1017/S1748499522000045

Seklecka, M., Pantelous, A.A., O'Hare, C., 2017. Mortality effects of temperature changes in the United Kingdom. J. Forecast. 36 (7), 824–841. https://doi.org/10.1002/for.2473

Seklecka, M., Pantelous, A.A., O'Hare, C., 2019. The impact of parameter uncertainty in insurance pricing and reserve with the temperature-related mortality model. J. Forecast. 38 (4), 327–345.

Serfling, R.E., 1963. Methods for current statistical analysis of excess pneumonia-influenza deaths. Public Health Rep. 78 (6), 494–506. https://doi.org/10.2307/4591848

Simonsen, L., Reichert, T.A., Viboud, C., Blackwelder, W.C., Taylor, R.J., Miller, M.A., 2005. Impact of influenza vaccination on seasonal mortality in the US elderly population. Arch. Intern. Med. 165 (3), 265–272. https://doi.org/10.1001/archinte.165.3.265

Song, X., Wang, S., Hu, Y., Yue, M., Zhang, T., Liu, Y., Tian, J., Shang, K., 2017. Impact of ambient temperature on morbidity and mortality: an overview of reviews. Sci. Total Environ. 586, 241–254. https://doi.org/10.1016/j.scitotenv.2017.01.212

Spall, J.C., 2000. Adaptive stochastic approximation by the simultaneous perturbation method. IEEE Trans. Autom. Control 45 (10), 1839–1853. https://doi.org/10.1109/TAC.2000.880982

Stone Jr, , B., Vargo, J., Liu, P., Habeeb, D., DeLucia, A., Trail, M., Hu, Y., Russell, A., 2014. Avoided heat-related mortality through climate adaptation strategies in three US cities. PLoS ONE 9 (6), e100852. https://doi.org/10.1371/journal.pone.\protect\penalty-\@M0100852

Tarín-Carrasco, P., Augusto, S., Palacios-Peña, L., Ratola, N., Jiménez-Guerrero, P., 2021. Impact of large wildfires on PM 10 levels and human mortality in portugal. Nat. Hazards Earth Syst. Sci. 21 (9), 2867–2880.

Tierney, L., Kadane, J.B., 1986. Accurate approximations for posterior moments and marginal densities. J. Am. Stat. Assoc. 81 (393), 82–86. https://doi.org/10.2307/2287970

Tukey, J.W., 1977. Exploratory data analysis. Reading/Addison-Wesley.

Valleron, A.-J., Bouvet, E., Garnerin, P., Ménares, J., Heard, I., Letrait, S., Lefaucheux, J., 1986. A computer network for the surveillance of communicable diseases: the French experiment. Am. J. Public Health 76 (11), 1289–1292. https://doi.org/10.2105/ajph.76.11.1289

van Berkum, F., Melenberg, B., Vellekoop, M., 2025. Estimating the impact of COVID-19 on mortality using granular data. Insur. Math. Econ. 121, 144–156. https://doi.org/10.1016/j.insmatheco.2025.01.001

Vandentorren, S., Paty, A.-C., Baffert, E., Chansard, P., Caserio-Schönemann, C., 2016. Syndromic surveillance during the paris terrorist attacks. Lancet 387 (10021), 846–847.

Vestergaard, L.S., Nielsen, J., Krause, T.G., Espenhain, L., Tersago, K., Sierra, N.B., Denissov, G., Innos, K., Virtanen, M.J., Fouillet, A., et al., 2017. Excess all-cause and influenza-attributable mortality in Europe, December 2016 to February 2017. Eurosurveillance 22 (14), 30506.

Vestergaard, L.S., Nielsen, J., Richter, L., Schmid, D., Bustos, N., Braeye, T., Denissov, G., Veideman, T., Luomala, O., Möttönen, T., Mølbak, K., 2020. Excess all-cause mortality during the COVID-19 pandemic in Europe – preliminary pooled estimates from the EuroMOMO network, March to April 2020. Eurosurveillance 25 (26), 2001214. https://doi.org/10.2807/1560-7917.ES.2020.25.26.2001214

Vinet, F., Boissier, L., Saint-Martin, C., 2016. Flashflood-related mortality in southern France: first results from a new database. In: 3rd European Conference on Flood Risk Management (FLOODrisk 2016). Vol. 7, p. 06001.

Wang, H., Paulson, K.R., Pease, S.A., Watson, S., Comfort, H., Zheng, P., Aravkin, A.Y., Bisignano, C., Barber, R.M., Alam, T., et al., 2022. Estimating excess mortality due to the COVID-19 pandemic: a systematic analysis of COVID-19-related mortality, 2020–21. Lancet 399 (10334), 1513–1536. https://doi.org/10.1016/S0140-6736(21)02796-3

Wong, C.-M., Chan, K.-P., Hedley, A.J., Peiris, J.S.M., 2004. Influenza-associated mortality in Hong Kong. Clin. Infect. Dis. 39 (11), 1611–1617. https://doi.org/10.1086/425315

Wood, S.N., 2004. Stable and efficient multiple smoothing parameter estimation for generalized additive models. J. Am. Stat. Assoc. 99 (467), 673–686.

Zhou, R., Li, J. S.-H., 2022. A multi-parameter-level model for simulating future mortality scenarios with COVID-alike effects. Ann. Actuarial Sci. 16 (3), 453–477. https://doi.org/10.1017/S1748499522000033