

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

Jens Robben, Karim Barigou, Torsten
Kleinow

LIDAM Discussion Paper ISBA
2025 / 06

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>

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

Jens Robben^{1,2,*}, Karim Barigou³, and Torsten Kleinow^{1,2}

¹*Faculty of Economics and Business, University of Amsterdam, The Netherlands.*

²*RCLR, Research Centre for Longevity Risk, University of Amsterdam, The Netherlands.*

³*Institute of Statistics, Biostatistics and Actuarial Science (ISBA), Louvain Institute of Data Analysis and Modeling (LIDAM), UCLouvain, Louvain-la-Neuve, Belgium*

*Corresponding author: j.robbe@uva.nl

March 10, 2025

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) an environmental shock state for heat waves, and (3) a respiratory shock state that addresses mortality deviations caused by 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 healthcare providers and hospitals in managing risks and planning resources for potential future shocks.

Keywords: granular mortality modeling; regime-switching; environmental shocks; respiratory 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). 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 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 the 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 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.

Statistical methods used to analyze mortality rates are manifold. Examples include Lee and Carter (1992) who model age-specific death rates as a linear function of an unobserved period-specific intensity index, with parameters depending on age. Berger et al. (2023) use a novel survival modeling approach to analyze postoperative mortality in older patients and incorporate risk factors measured at hospital admission. Berger and Guglielmi (2001) use Bayesian random-effects logit models to model hospital mortality rates for a heart attack or acute myocardial infarction.

Our paper contributes to the statistical literature in three ways. First, we gather daily gridded temperature data from the Copernicus Climate Data Store and outline how to con-

vert 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. While previous studies have primarily focused 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 the death counts under various weather and respiratory scenarios.

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 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 $\mathcal{T} = \{1, 2, \dots, T\}$. 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. Figure 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 heightened vulnerability to environmental factors and illnesses such as influenza. Hereto, we consider 5-year age groups starting from the age of 65, i.e., $\mathcal{X} = \{65-69, 70-74, 75-79, 80-84, 85-89, 90+\}$. Figure 1 (right panel) presents the weekly death counts from 2013 to 2024 for individuals aged 90 and above 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. Furthermore, we observe a clear COVID-19 peak in Ile-de-France, where the death counts exceeded the usual level by more than three times. 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.




Figure 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.

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 (CDS, 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 et. al, 2018). This ensures that the temperature averages reflect areas where more people reside. Figure 2 illustrates a snapshot of the temperature and population data on a randomly selected date and year.




Figure 2: Left panel: the daily 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 highlight the boundaries of the French NUTS 2 regions in white.

We construct population-weighted daily temperature averages and compare them to region-specific high (95%) and low (5%) thresholds based on 2013–2024 data. We then average daily indicators for exceeding these thresholds weekly to create hot- and cold-week indices, see Figure 3 for an example. We also compute weekly averages of daily temperature anomalies by comparing observed temperatures to a region-specific daily baseline. While the hot-week index measures the frequency of hot days in a week, temperature anomalies reflect their severity. A similar interpretation applies to the cold-week index. We refer to Robben et al. (2024) for further details on this feature construction step.




Figure 3: Left panel: the weekly average of the hot-day indicator at the ISO week starting August 1, 2022. Right panel: the weekly average of the cold-day indicator at the ISO week starting January 14, 2013. We highlight the boundaries of the French NUTS 2 regions in grey.

Influenza data. We consult the French Sentinelles network, established in 1984 by Inserm and Sorbonne University, to collect near real-time epidemic data from $\approx 1\,300$ 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. Figure 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 norm.




Figure 4: We show the weekly number of new influenza-like incidences per 100 inhabitants (left) and the COVID-19 hospitalizations per 1 000 inhabitants (right) from 2013-2014 and 2020-2023, respectively, in Ile-de-France (FR10), Provence-Alpes-Côte d’Azur (FRL0), and Bretagne (FRH0).

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. Figure 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 discussed in this section.

Feature	Description
HI	Hot-week index: weekly average of daily indicators for unusually high temperatures.
CI	Cold-week index: weekly average of daily indicators for unusually low temperatures.
TA	Temperature anomalies: weekly average of daily deviations from baseline temperature levels.
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).

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

3 Modeling mortality with regime-switching dynamics

Figure 5 illustrates the proposed model. On the left, the red line visualizes the estimation of a robust seasonal baseline mortality trend, tailored to each considered age category and region 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 (environmental shock state): captures deviations caused by environmental shocks, such as elevated summer temperatures.
3. State 2 (respiratory shock state): addresses deviations driven by an increase in respiratory infections, including influenza and COVID-19.

The arrows in Figure 5 represent possible state transitions and their corresponding probabilities. The model assumes the system can only be in one of the three states at any time, and direct transitions between the environmental shock state (state 1) and respiratory shock state (state 2) are not allowed. 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 instance, if very high temperatures are observed at 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.

3.1 Weekly baseline mortality model

We design a region- and age-group-specific baseline to capture the overall seasonal mortality patterns by following the approach in Serfling (1963). Hereto, we incorporate seasonality




Figure 5: Schematic representation of the proposed regime-switching mortality model. The left panel shows a robust seasonal baseline trend (red line) for each age group and region, alongside observed death counts (grey). The right panel depicts the regime-switching structure, transitioning between three states: baseline mortality (state 0), excess mortality from environmental shocks (state 1), and excess mortality from respiratory infections (state 2). Arrows indicate state transitions, with probabilities driven by environmental and epidemic covariates.

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, $D_{x,t}^{(r)}$:

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

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

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

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 $\gamma_{x,p}^{(r)}$ ($p = 0, 1, \dots, 5$) 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:

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

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), environmental shock (state

1), or respiratory shock (state 2) state. For each state i , $D_{x,t}^{(r)}$ follows a Poisson distribution:

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

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

Each regime state corresponds to a distinct mortality pattern. For $S_t^{(r)} = 0$, the baseline state, we assume 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 environmental-shock state, only temperature covariates in $\mathbf{z}_t^{(r)}$ play a role, and we restrict $\boldsymbol{\alpha}_{1,x}$ to have non-zero values only for these variables. Lastly, for $S_t^{(r)} = 2$, the respiratory shock state, we assume only epidemic covariates affect mortality, with non-zero values in $\boldsymbol{\alpha}_{2,x}$.

To avoid overcomplicating the model, we assume 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 $\mathbf{P}_t^{(r)}$:

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

with $\boldsymbol{\beta}$ the parameter vector, and U_r a spatial effect to capture regional disparities. This spatial effect ensures that neighbouring 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):

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

where $R := |\mathcal{R}|$ is the region count, $\mathbf{D}$ a diagonal matrix with entries equal to the number of neighbours for each region, and $\mathbf{W}$ the adjacency matrix. The entry $W_{rr'}$ equals 1 if the regions r and r' are neighbors and 0 otherwise. The diagonal entries of $\mathbf{W}$ 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:

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

with $\boldsymbol{\beta} = (\boldsymbol{\beta}_{ij})_{ij}$. For transitions toward (or persistence in) the environmental shock state, only temperature-related covariates are included. For transitions toward (or persistence

in) the respiratory shock state, only 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):

$$\gamma^* = \arg \min_{\gamma} \left(-\ell_P(\gamma) + \sum_{p=0}^5 \lambda_p \left(\sum_{x \in \mathcal{X}} \gamma_{x,p}^\top (\mathbf{D} - \mathbf{W}) \gamma_{x,p} \right) \right),$$

where $\gamma = (\gamma_{x,p})_{x,p}$, with $\gamma_{x,p} = (\gamma_{x,p}^{(r)})_{r \in \mathcal{R}}$, for each age group $x \in \mathcal{X}$ and $p = 0, 1, \dots, 5$. Furthermore, $\ell_P(\gamma)$ is the Poisson log-likelihood, and the matrices $\mathbf{D}$ and $\mathbf{W}$ 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. 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 γ^* , 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 $\mathcal{T} = \{1, 2, \dots, T\}$, 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:

$$\boldsymbol{\theta} = \{\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2, \boldsymbol{\beta}_{01}, \boldsymbol{\beta}_{02}, \boldsymbol{\beta}_{11}, \boldsymbol{\beta}_{22}\}, \quad (4.1)$$

where $\boldsymbol{\alpha}_i = (\alpha_{i,x})_{x \in \mathcal{X}}$ 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 $\mathbf{U}$. To account for this, we integrate out $\mathbf{U}$:

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

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

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

where $\mathbf{Q}(\tau) := \tau[\mathbf{D} - \mathbf{W}]$ 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):

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

where $\mathbf{u}^*$ is the mode of the expected log joint probability distribution, and $\mathbf{S}^{-1}$ is the matrix of second-order derivatives evaluated at $\mathbf{u}^*$. 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:

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

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 $\mathbf{u}^*$ 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 $\mathbf{u}^*(\boldsymbol{\theta})$. 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 $\mathbf{u}_0^*$.
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, $\mathbf{u}_k^*$.
 - (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, $\mathbf{u}^*$, under the current optimal parameter vector, $\boldsymbol{\theta}_{k+1}^*$, and obtain $\mathbf{u}_{k+1}^*$.
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:

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

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):

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

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 $p_t^{ij}(\mathbf{z}_t^{(r)}; \boldsymbol{\beta}_\iota, u_{r,\iota})$, where $\mathbf{z}_t^{(r)}$ 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 $\mathbf{z}_t^{(r)}$.

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 $\mathcal{X} = \{65-69, 70-74, 75-79, 80-84, 85-89, 90+\}$, and weeks $\mathcal{T} = \{1, 2, \dots, 600\}$, where $t = 1$ corresponds to the first ISO week in 2013 and $t = 600$ to the 26th ISO week in 2024. We omit the weeks related to the first wave of the COVID-19 pandemic (weeks 12-16), as they heavily influence the overall performance of the baseline trend. Figure 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. Eastern French regions also saw a decline in mortality, while the Western regions experienced a slight increase (γ_1). The seasonal variation was most prominent in Northern(-Eastern) regions. Section A of the Online Appendix contains the parameter estimates for the five remaining age groups.




Figure 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.

6.1.2 Regime-switching mortality model

State-specific model specifications. We use the variables constructed from the temperature, influenza, and hospitalization data in Table 1 as features in the state-dependent Poisson models in (3.4). Since their impact on mortality differs by age, we use age-dependent parameters. We further assume shared parameters among two consecutive age groups, and focus, as such, on the reduced set of age groups: $\mathcal{X}_{\text{red}} = \{65-74, 75-84, 85+\}$.

In state 1, we model heatwave-related shocks by incorporating weekly averages of daily temperature anomalies and extreme temperature indices with up to two weeks' lag, based on evidence from [Armstrong \(2006\)](#) and [Gasparrini and Armstrong \(2011\)](#). These features provide a good measure of the heat frequency and intensity within a given week. We model

the expected death counts in regime state 1 as:

$$\log \mathbb{E} \left[D_{x,t}^{(r)} \mid S_t^{(r)} = 1 \right] = \log \hat{b}_{x,t}^{(r)} + \sum_{a \in \mathcal{X}_{\text{red}}} \left(\alpha_{1,1}^{(a)} \text{TA}_t^{(r)} + \alpha_{1,2}^{(a)} \text{TA}_{t-1}^{(r)} + \alpha_{1,3}^{(a)} \text{TA}_{t-2}^{(r)} + \alpha_{1,4}^{(a)} \text{HI}_t^{(r)} + \alpha_{1,5}^{(a)} \text{HI}_{t-1}^{(r)} + \alpha_{1,6}^{(a)} \text{HI}_{t-2}^{(r)} \right) \mathbb{1} \{x \subset a\}, \quad (6.1)$$

where the $\alpha_{1,j}^{(a)}$'s are age-specific parameters, $\hat{b}_{x,t}^{(r)}$ refers to the baseline death count estimated as in Section 4.1, and $\mathbb{1}\{x \subset a\}$ is an indicator function which equals one whenever the smaller age group $x \in \mathcal{X}$ is a subset of the bigger age group $a \in \mathcal{X}_{\text{red}}$. We refer to Table 1 for the notation of the covariates, where, for instance, $\text{TA}_{t-1}^{(r)}$ refers to the weekly average of the daily temperature anomalies at week $t-1$ in region r (lag 1).

In state 2, we model mortality shocks from influenza activity and COVID-19 hospitalizations, and incorporate a cold indicator due to its additional impact on mortality. Since respiratory-related mortality effects are often delayed, we use a maximum lag of 3 weeks, based on evidence from Wong et al. (2004) and Lytras et al. (2019). Furthermore, we group the lags in short-term (weeks 0–1) and mid-term (weeks 2–3) intervals to reduce multicollinearity issues. Long-term impacts are not considered in this paper. We model the expected death counts in state 2 as:

$$\log \mathbb{E} \left[D_{x,t}^{(r)} \mid S_t^{(r)} = 2 \right] = \log \hat{b}_{x,t}^{(r)} + \sum_{a \in \mathcal{X}_{\text{red}}} \left(\alpha_{2,1}^{(a)} \text{IA}_{t,t-1}^{(r)} + \alpha_{2,2}^{(a)} \text{IA}_{t-2,t-3}^{(r)} + \alpha_{2,3}^{(a)} \text{CI}_{t,t-1}^{(r)} + \alpha_{2,4}^{(a)} \text{CI}_{t-2,t-3}^{(r)} + \alpha_{2,5}^{(a)} \text{HA}_{t,t-1}^{(r)} + \alpha_{2,6}^{(a)} \text{HA}_{t-2,t-3}^{(r)} \right) \mathbb{1} \{x \subset a\}, \quad (6.2)$$

where $\text{IA}_{t,t-1}^{(r)}$ and $\text{IA}_{t-2,t-3}^{(r)}$ refer to the influenza activity index at week t averaged over lags 0–1 and 2–3, respectively. Analogous notation is introduced for the cold indicator (CI) and the COVID-19 hospitalization index (HA). For the respiratory-related features, we use exceedances over their 75% quantile since extremer values, rather than moderate fluctuations, are more likely to contribute to significant mortality shocks.

Modeling regime transition probabilities. We construct multinomial logistic models for the transition probabilities, as defined in (3.6). Following Section 3.2, state transition probabilities are independent of age group x . To capture regional variations, we introduce a spatial effect U_r , which we model using an ICAR process. Specifically, we define:

$$\begin{aligned} \text{logit } p_t^{01} \left(\mathbf{z}_t^{(r)}; \boldsymbol{\beta}, u_r \right) &= \beta_{01,0} + \beta_{01,1} \text{HI}_t + U_r \\ \text{logit } p_t^{02} \left(\mathbf{z}_t^{(r)}; \boldsymbol{\beta}, u_r \right) &= \beta_{02,0} + \beta_{02,1} \text{IA}_{t,t-1}^{(r)} + \beta_{02,2} \text{HA}_{t,t-1}^{(r)} + U_r \\ \text{logit } p_t^{11} \left(\mathbf{z}_t^{(r)}; \boldsymbol{\beta}, u_r \right) &= \beta_{11,0} + \beta_{11,1} \text{HI}_t + \beta_{11,2} \text{HI}_{t-1} + \beta_{11,3} \text{HI}_{t-2} + U_r \\ \text{logit } p_t^{22} \left(\mathbf{z}_t^{(r)}; \boldsymbol{\beta}, u_r \right) &= \beta_{22,0} + \beta_{22,1} \text{IA}_{t,t-1}^{(r)} + \beta_{22,2} \text{IA}_{t-2,t-3}^{(r)} + \beta_{22,3} \text{HA}_{t,t-1}^{(r)} + \beta_{22,4} \text{HA}_{t-2,t-3}^{(r)} + U_r. \end{aligned} \quad (6.3)$$

Here, we focus solely on features capturing short-term mortality impacts to model the transition probabilities from the baseline to one of the shock states. This facilitates a direct transition to the shock state whenever we observe increased temperature levels or respiratory infections. Conversely, to determine whether to remain in the current state, we incorporate lagged features to account for state persistence and delayed effects.

Parameter estimates. We follow Section 4.2.3 and apply the EM-algorithm to estimate the regime-switching parameter vector θ . 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 $\tau^* = 10$.

Figure 7 (left panel) shows the parameter estimates for α_1 . Following Suppl. Mat. B.6, we construct 95% confidence intervals using the Fisher Information matrix based on 25 000 bootstrap samples. Elevated temperature anomalies in the current week are strongly associated to excess mortality, especially for those aged 85+. At lag 1, effects differ depending on the interaction between heat frequency (hot-week index) and severity (temperature anomalies). At lag 2, a harvesting effect is likely to occur across all age groups for severely hot weeks. However, the wide confidence intervals for the hot-week index, potentially due to collinearity with temperature anomalies, suggest caution in interpreting this effect.




Figure 7: Parameter estimates for α_1 (environmental shock state) and α_2 (respiratory shock state) in the Poisson models of the regime-switching model, along with their 95% confidence intervals based on 25 000 bootstrap samples.

The right panel of Figure 7 shows the estimates and 95% confidence intervals for the parameter vector α_2 , which captures the effects of increased influenza activity or COVID-19 hospitalizations on mortality. Results reveal significant excess mortality from high influenza activity in both the short- and mid-term, especially for ages 85+. Cold temperatures also significantly increase mortality in both periods and COVID-19 hospitalizations show a significant short-term impact on excess mortality, increasing with age. We observe a modest harvesting effect at lags 2–3 for the COVID-19 hospitalizations.

Figure 8 displays the estimates and 95% confidence intervals of the parameter vectors β_{01} , β_{02} , β_{11} , and β_{22} in the transition probabilities of (6.3). Transitions from the baseline to the environmental shock state are more likely with a high current-week hot-week index. Transitions to the respiratory shock state increase with higher short-term influenza activity and COVID-19 hospitalizations. Once in the environmental shock state, the probability of remaining in this state increases with the current-week hot-week index but decreases with the previous week's index. In the respiratory shock state, the likelihood of staying increases with higher short-term influenza activity or hospital admissions. The long-term effects (lags 2–3) are insignificant.

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




Figure 8: Estimates of the parameter vectors in the transitional probabilities and their 95% confidence intervals based on 25 000 bootstrap samples: β_{01} (top left), β_{02} (top right), β_{11} (bottom left), and β_{22} (bottom right). The label ‘INT’ refers to the intercept




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

6.2 Quantifying uncertainty in in-sample model predictions

6.2.1 Best-estimate in-sample fit

We use the estimated baseline mortality model and construct a 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:

$$s_{\text{BE},t}^{(r)} = \arg \max_{j \in \{0,1,2\}} \hat{P} \left(S_t^{(r)} = j \mid \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}^*, \mathbf{u}^* \right), \quad (6.4)$$

where $\boldsymbol{\theta}^*$ and $\mathbf{u}^*$ are the optimal regime-switching parameter vector and the mode of the spatial effect. The sets $\mathcal{D}_t^{(r)}$ and $\mathcal{Z}_t^{(r)}$ 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:

$$\hat{d}_{\text{BE},x,t}^{(r)} = \hat{\mathbb{E}} \left[D_{x,t}^{(r)} \mid S_t^{(r)} = s_{\text{BE},t}^{(r)} \right].$$

Figure 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 coloured 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 increased influenza activity during the winter seasons of 2014/2015 and 2016/2017, and those related to heatwaves in the summer seasons of 2015 and 2022.




Figure 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).

6.2.2 Uncertainty in 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. Figure 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 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 relatively limited parameter and spatial uncertainty is more pronounced in the younger old-age groups.




Figure 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).

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 $\tilde{T}$. 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 smaller calibration period $\tilde{T}$. For instance, we recompute the temperature anomalies based on a temperature baseline calibrated on $\tilde{T}$. Second, we recalibrate the baseline mortality model in Section 4.1, and the three-state regime-switching model in Section 3.2, on $\tilde{T}$. Third, we 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 temperature anomalies, 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.

To avoid constructing a point forecast for the exposure-to-risk, we focus on the death rates per 1 000 person-years. This approach facilitates making comparisons across different regions. Figure 12 illustrates the 95% prediction intervals for the weekly death rates in Pays de la Loire and Provence-Alpes-Côte d’Azur for the 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 Figure 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.




Figure 12: We present the 95% prediction intervals for the weekly death rates per 1 000 person-years 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, while the black line visualizes the observed death rates per 1 000 person-years.

6.4 Short-term mortality forecasting with scenario-based approaches

To forecast death rates in the short-term, we use temperature scenarios based on the RCP 2.6, RCP 4.5, and RCP 8.5 pathways from the Copernicus Climate Change Service (Copernicus Climate Change Service, 2021). For influenza, we generate moderate, high, and severe incidence scenarios using a Bayesian stochastic SIRS 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 explored. We include all technical details in Suppl. Mat. C. Figure shows the results of the constructed scenarios for Pays de la Loire and Provenances-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 temperature anomalies, extreme temperature indicators, and the influenza activity index (anomalies) based on




Figure 13: Projections for the daily average temperature based on the RCP 2.6, 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, high, and severe incidence scenario for influenza activity using the point-wise 50%, 80%, and 95% quantiles of the posterior predictive samples, respectively. The top panels show the result for Pays de la Loire, while the bottom panels depict the results for Provinces-Alpes-Côte d’Azur.

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 three prediction intervals corresponding to three different scenarios: scenario 1 combines the moderate influenza scenario and the RCP 2.6 temperature projection, scenario 2 focuses on the high influenza scenario and the RCP 4.5 temperature projection, and, lastly, scenario 3 combines the severe influenza scenario with the RCP 8.5 temperature projection. Figure 14 shows the 95% prediction intervals for the death counts at age group 90+ in the NUTS 2 regions Centre-Val de Loire, Pays de la Loire, Languedoc-Roussillon, and Alpes Côte d’Azur where the intervals in green, blue, and red correspond to scenario 1, 2, and 3, respectively. The prediction intervals for scenarios 1 and 2 only differ slightly, which can be explained by the more severe winter temperatures in the RCP 2.6 temperature projection compared to the RCP 4.5 scenario. Such scenario-based mortality forecasts can offer insurers useful insights into how different climate and influenza scenarios impact future mortality risks. Section F of the Online Appendix contains the prediction intervals for all age groups and NUTS 2 regions considered in this paper.

We estimate the excess deaths relative to the baseline mortality model under the three considered scenarios. 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 counts with the baseline setting, expressed as a relative increase. We do this analysis for each region and age group.




Figure 14: Based on the moderate, high, and severe influenza scenario, we present 95% prediction intervals for the weekly death rates per 1 000 person-years in Centre-Val de Loire, Pays de la Loire, Languedoc-Roussillon, and Provence-Alpes-Côte d’Azur for the age group 90+. The predictions range from the 27th ISO week of 2024 to the 52nd ISO week of 2029. The intervals take into account parameter, spatial, state and Poisson uncertainty 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.

Figure 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. Based on the generated scenarios for influenza and temperature, we find that the southern NUTS 2 regions are generally more affected across all considered scenarios. Scenario 1 leads to moderate increases in deaths compared to the baseline, while Scenario 3 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.




Figure 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 the three presented scenarios.

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, an environmental shock state for periods of heightened mortality due to heatwaves, and a respiratory shock state to account for 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 increased influenza activity or hospital admissions. In the environmental 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 hot week. In the respiratory shock state, the amplification of baseline mortality is associated with increased influenza activity, cold temperatures, and COVID-19-related hospitalizations. The model also captures harvesting effects. 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 Jr 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.

Data and code availability statement

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>.

Funding statement

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.

Conflict of interest disclosure

The authors declare no conflict of interest.

References

- Armstrong, B. (2006). Models for the relationship between ambient temperature and daily mortality. *Epidemiology*, 17(6):624–631.
- Basu, R. and Samet, J. M. (2002). Relation between elevated ambient temperature and mortality: A review of the epidemiologic evidence. *Epidemiologic Reviews*, 24(2):190–202.
- Berger, J. O. and Guglielmi, A. (2001). Bayesian and conditional frequentist testing of a parametric model versus nonparametric alternatives. *Journal of the American Statistical Association*, 96(453):174–184.
- Berger, M., Kowark, A., Rossaint, R., Coburn, M., Schmid, M., and POSE-Study, G. (2023). Modeling postoperative mortality in older patients by boosting discrete-time competing risks models. *Journal of the American Statistical Association*, 118(544):2239–2249.
- Besag, J. and Kooperberg, C. (1995). On conditional and intrinsic autoregressions. *Biometrika*, 82(4):733–746.
- Bobb, J. F., Peng, R. D., Bell, M. L., and Dominici, F. (2014). Heat-related mortality and adaptation to heat in the United States. *Environmental Health Perspectives*, 122(8):811–816.
- Bucyibaruta, G., Dean, C., and Torabi, M. (2023). A discrete-time susceptible-infectious-recovered-susceptible model for the analysis of influenza data. *Infectious Disease Modelling*, 8(2):471–483.
- 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.
- Copernicus Climate Change Service (2021). Temperature and precipitation climate impact indicators from 1970 to 2100 derived from European climate projections. <https://cds.climate.copernicus.eu>. DOI: <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.
- Eurostat (2024a). 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.
- Eurostat (2024b). Population on 1 January by age, sex and NUTS 3 region. <https://doi.org/10.2908/DEMO.R.PJANAGGR3>. Accessed January 22, 2024.
- Ferguson, N. M., Cummings, D. A., Fraser, C., Cajka, J. C., Cooley, P. C., and Burke, D. S. (2006). Strategies for mitigating an influenza pandemic. *Nature*, 442(7101):448–452.
- Gasparrini, A. and Armstrong, B. (2011). The impact of heat waves on mortality. *Epidemiology*, 22(1):68.
- 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., and Armstrong, B. (2015). Mortality risk attributable to high and low ambient temperature: A multicountry observational study. *The Lancet*, 386(9991):369–375.
- Gemmell, I., McLoone, P., Boddy, F., Dickinson, G. J., and Watt, G. (2000). Seasonal variation in mortality in Scotland. *International Journal of Epidemiology*, 29(2):274–279.
- Hajat, S., O’Connor, M., and Kosatsky, T. (2010). Health effects of hot weather: from awareness of risk factors to effective health protection. *The Lancet*, 375(9717):856–863.
- 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. *The*

- Lancet*, 391(10127):1285–1300.
- Jdanov, D. A., Galarza, A. A., Shkolnikov, V. M., Jasilionis, D., Németh, L., Leon, D. A., Boe, C., and Barbieri, M. (2021). The short-term mortality fluctuation data series, monitoring mortality shocks across time and space. *Scientific Data*, 8(1):235.
- Kok, G., Jonkers, R., Gelissen, R., Meertens, R., Schaalma, H., and de Zwart, O. (2010). Behavioural intentions in response to an influenza pandemic. *BMC Public Health*, 10:1–19.
- Laplace, P. S. (1774). Mémoire sur la probabilité de causes par les événements. *Mémoire de l’académie royale des sciences*.
- Lee, R. D. and Carter, L. R. (1992). Modeling and forecasting us mortality. *Journal of the American statistical association*, 87(419):659–671.
- Li, L., Li, H., and Panagiotelis, A. (2024). Boosting domain-specific models with shrinkage: an application in mortality forecasting. *International Journal of Forecasting*.
- Louis, T. A. (1982). Finding the observed information matrix when using the EM algorithm. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 44(2):226–233.
- Lytras, T., Pantavou, K., Mouratidou, E., and 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.
- Mackenbach, J., Kunst, A., and Looman, C. (1992). Seasonal variation in mortality in the Netherlands. *Journal of Epidemiology & Community Health*, 46(3):261–265.
- 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.
- Meng, L. and 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)*, pages 1–6. IEEE.
- Meng, X.-L. and Rubin, D. B. (1991). Using EM to obtain asymptotic variance-covariance matrices: The SEM algorithm. *Journal of the American Statistical Association*, 86(416):899–909.
- Nielsen, J., Krause, T. G., and Mølbak, K. (2018). Influenza-associated mortality determined from all-cause mortality, Denmark 2010/11-2016/17: the FluMOMO model. *Influenza and Other Respiratory Viruses*, 12(5):591–604.
- Nielsen, J., Mazick, A., Glismann, S., and Mølbak, K. (2011). Excess mortality related to seasonal influenza and extreme temperatures in Denmark, 1994-2010. *BMC Infectious Diseases*, 11:1–13.
- Nielsen, J., Vestergaard, L. S., Richter, L., Schmid, D., Bustos, N., Asikainen, T., Trebbien, R., Denissov, G., Innos, K., Virtanen, M. J., ..., and 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? *Clinical Microbiology and Infection*, 25(10):1266–1276.
- 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.
- Oakes, D. (1999). Direct calculation of the information matrix via the EM. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(2):479–482.
- Robben, J., Antonio, K., and Kleinow, T. (2024). The short-term association between environmental variables and mortality: evidence from Europe. *arXiv preprint arXiv:2405.18020*.
- Serfling, R. E. (1963). Methods for current statistical analysis of excess pneumonia-influenza deaths. *Public Health Reports*, 78(6):494–506.

- Simonsen, L., Reichert, T. A., Viboud, C., Blackwelder, W. C., Taylor, R. J., and Miller, M. A. (2005). Impact of influenza vaccination on seasonal mortality in the US elderly population. *Archives of Internal Medicine*, 165(3):265–272.
- Song, X., Wang, S., Hu, Y., Yue, M., Zhang, T., Liu, Y., Tian, J., and Shang, K. (2017). Impact of ambient temperature on morbidity and mortality: an overview of reviews. *Science of the Total Environment*, 586:241–254.
- Spall, J. C. (2000). Adaptive stochastic approximation by the simultaneous perturbation method. *IEEE Transactions on Automatic Control*, 45(10):1839–1853.
- Stone Jr, B., Vargo, J., Liu, P., Habeeb, D., DeLucia, A., Trail, M., Hu, Y., and Russell, A. (2014). Avoided heat-related mortality through climate adaptation strategies in three US cities. *PLOS One*, 9(6):e100852.
- Tierney, L. and Kadane, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, 81(393):82–86.
- Valleron, A.-J., Bouvet, E., Garnerin, P., Ménares, J., Heard, I., Letrait, S., and Lefaucheux, J. (1986). A computer network for the surveillance of communicable diseases: the French experiment. *American Journal of Public Health*, 76(11):1289–1292.
- Vestergaard, L. S., Nielsen, J., Richter, L., Schmid, D., Bustos, N., Braeye, T., Denissov, G., Veideman, T., Luomala, O., Möttönen, T., and 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.
- 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. *The Lancet*, 399(10334):1513–1536.
- Wong, C.-M., Chan, K.-P., Hedley, A. J., and Peiris, J. M. (2004). Influenza-associated mortality in Hong Kong. *Clinical Infectious Diseases*, 39(11):1611–1617.

Supplementary material for “Granular mortality modeling with temperature and epidemic shocks: a three-state regime-switching approach”

A Data overview

NUTS 2 regions. Table A.1 lists the names of the NUTS 2 regions in Metropolitan France together with their corresponding NUTS 2 geocode.

NUTS 2	Region Name	NUTS 2	Region Name
FR10	Île-de-France	FRG0	Pays de la Loire
FRB0	Centre - Val de Loire	FRH0	Bretagne
FRC1	Bourgogne	FRI1	Aquitaine
FRC2	Franche-Comté	FRI2	Limousin
FRD1	Basse-Normandie	FRI3	Poitou-Charentes
FRD2	Haute-Normandie	FRJ1	Languedoc-Roussillon
FRE1	Nord-Pas de Calais	FRJ2	Midi-Pyrénées
FRE2	Picardie	FRK1	Auvergne
FRF1	Alsace	FRK2	Rhône-Alpes
FRF2	Champagne-Ardenne	FRL0	Provence-Alpes-Côte d’Azur
FRF3	Lorraine		

Table A.1: The NUTS 2 regions in Metropolitan France that constitute the set $\mathcal{R}$.

Population exposures. We obtain the population count of individuals aged 65 and older as of January 1 of each year, disaggregated by sex, age and NUTS 2 region, from Eurostat (Eurostat, 2024b). Let t represent any (ISO-)week within the specified time range $\mathcal{T}$, and let $y(t)$ denote the year corresponding to that week. Then, following the approach outlined by Jdanov et al. (2021), we estimate the weekly exposure as:

$$E_{x,t}^{(r)} = \frac{P_{x,y(t)}^{(r)} + P_{x,y(t)+1}^{(r)}}{2 \cdot 52.18}, \quad (\text{A.1})$$

where $P_{x,y(t)}^{(r)}$ represents the population count at age x in region $r \in \mathcal{R}$ on January 1 of year $y(t)$. The value 52.18 reflects the average number of weeks in a year. Whereas this approach assumes a constant weekly exposure every year, a more refined method would be to apply linear interpolation to ensure a smoother transition of weekly exposures between consecutive years. Figures A.1 and A.2 provide a detailed overview of the population count data.

Exploratory plots. Figure A.1 illustrates the stacked death counts and population exposures for the 21 French NUTS 2 regions over time, aggregated across age groups. Table A.1 provides the corresponding region names for each NUTS 2 geocode. The left panel shows a sharp spike in death counts at the beginning of 2020, which can be attributed to the COVID-19 pandemic. Additionally, we observe a spike during the winter seasons of 2014/2015 and 2016/2017 which coincides with severe influenza outbreaks. The right panel shows that population exposure remains constant within each year’s ISO weeks, see (A.1).

Figure A.2 displays the stacked death counts and population exposures for the 21 French NUTS 2 regions, categorized by age group and aggregated across all considered ISO weeks. The highest number of deaths occurs in the oldest age group (90+), despite having the




Figure A.1: Stacked death counts (left) and exposures (right) for the 21 French NUTS 2 regions from the first ISO week of 2013 till the 26th ISO week of 2024. The death counts per region and time point are aggregated among the different age groups under consideration.




Figure A.2: Stacked death counts (left) and exposures (right) for the 21 French NUTS 2 regions by age group. The death counts per region and age group are aggregated among the different ISO weeks under consideration.

smallest population exposure. Conversely, the youngest age group (65–69) has the largest population exposure but the lowest number of deaths.

B Details on the calibration procedure of the regime-switching model

B.1 Notation

In addition to the notation introduced in Section 4.2.1, we denote the collection of age-specific death counts observed up to time point $t \in \mathcal{T}$ as:

$$\mathcal{D}_t = \left\{ \mathcal{D}_{x,t}^{(r)} \mid r \in \mathcal{R}, x \in \mathcal{X} \right\},$$

where $\mathcal{D}_{x,t}^{(r)} = \{D_{x,t'}^{(r)} \mid t' \leq t\}$. We introduce similar notation for the collection of all states, $\mathcal{S}_t$, and covariate vectors, $\mathcal{Z}_t$, up to time point $t \in \mathcal{T}$:

$$\mathcal{S}_t = \left\{ \mathcal{S}_t^{(r)} \mid r \in \mathcal{R} \right\}, \quad \mathcal{Z}_t = \left\{ \mathcal{Z}_t^{(r)} \mid r \in \mathcal{R} \right\},$$

with $\mathcal{S}_t^{(r)} = \{S_{t'}^{(r)} \mid t' \leq t\}$ and $\mathcal{Z}_t^{(r)} = \{z_{t'}^{(r)} \mid t' \leq t\}$.

B.2 Complete data log-likelihood

We first derive the complete-data likelihood of all death counts and regime states, conditional on the spatial effect $\mathbf{u}^*$ and the covariate information $\mathcal{Z}_T$. We hereby assume that the temporal evolution of death counts and regime states is independent across regions, conditionally on the spatial effect $\mathbf{u}^*$ and region-specific covariates. Any spatial dependence across regions is assumed to be entirely captured through the spatial effect $\mathbf{u}^*$. We obtain:

$$\begin{aligned} \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T \mid \mathbf{u}^*, \mathcal{Z}_T; \boldsymbol{\theta}) &= \prod_{r \in \mathcal{R}} \mathbb{P} \left(\left(\mathcal{D}_{x,T}^{(r)} \right)_{x \in \mathcal{X}}, \mathcal{S}_T^{(r)} \mid u_r^*, \mathcal{Z}_T^{(r)}; \boldsymbol{\theta} \right) \\ &= \prod_{r \in \mathcal{R}} \left\{ \mathbb{P} \left(\left(D_{x,1}^{(r)} \right)_{x \in \mathcal{X}} = \left(d_{x,1}^{(r)} \right)_{x \in \mathcal{X}}, S_1^{(r)} = s_1^{(r)} \mid u_r^*, \mathcal{Z}_1^{(r)}; \boldsymbol{\theta} \right) \cdot \right. \\ &\quad \left. \prod_{t=2}^T \mathbb{P} \left(\left(D_{x,t}^{(r)} \right)_{x \in \mathcal{X}} = \left(d_{x,t}^{(r)} \right)_{x \in \mathcal{X}}, S_t^{(r)} = s_t^{(r)} \mid u_r^*, \mathcal{Z}_t^{(r)}, \left(\mathcal{D}_{x,t-1}^{(r)} \right)_{x \in \mathcal{X}}, \mathcal{S}_{t-1}^{(r)}; \boldsymbol{\theta} \right) \right\}, \end{aligned}$$

where we apply the law of conditional probability $T - 1$ times to obtain the last equality. Next, within each region r , we impose a temporal Markov property on the latent regime state processes, and rewrite the likelihood as:

$$\begin{aligned} \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T \mid \mathbf{u}^*, \mathcal{Z}_T; \boldsymbol{\theta}) &= \prod_{r \in \mathcal{R}} \left\{ \mathbb{P} \left(\left(d_{x,1}^{(r)} \right)_{x \in \mathcal{X}}, s_1^{(r)} \mid u_r^*, \mathcal{Z}_1^{(r)}; \boldsymbol{\theta} \right) \cdot \right. \\ &\quad \left. \prod_{t=2}^T \mathbb{P} \left(\left(d_{x,t}^{(r)} \right)_{x \in \mathcal{X}}, s_t^{(r)} \mid u_r^*, \mathcal{Z}_t^{(r)}, \left(\mathcal{D}_{x,t-1}^{(r)} \right)_{x \in \mathcal{X}}, \mathcal{S}_{t-1}^{(r)}; \boldsymbol{\theta} \right) \right\} \\ &= \prod_{r \in \mathcal{R}} \left\{ \mathbb{P} \left(\left(d_{x,1}^{(r)} \right)_{x \in \mathcal{X}} \mid s_1^{(r)}, \mathbf{z}_1^{(r)}; \boldsymbol{\alpha}_{s_1^{(r)}} \right) \cdot \mathbb{P} \left(s_1^{(r)} \mid u_r^*, \mathbf{z}_1^{(r)}; \boldsymbol{\beta} \right) \cdot \right. \\ &\quad \left. \prod_{t=2}^T \mathbb{P} \left(\left(d_{x,t}^{(r)} \right)_{x \in \mathcal{X}} \mid s_t^{(r)}, \mathbf{z}_t^{(r)}; \boldsymbol{\alpha}_{s_t^{(r)}} \right) \cdot \mathbb{P} \left(s_t^{(r)} \mid s_{t-1}^{(r)}, u_r^*, \mathbf{z}_t^{(r)}; \boldsymbol{\beta} \right) \right\}. \end{aligned}$$

Next, we assume that the death counts random variables $D_{x,t}^{(r)}$ are independent among the different age categories $x \in \mathcal{X}$, conditionally on the regime state $S_t^{(r)}$. Under this

assumption, we can rewrite the likelihood to:

$$\begin{aligned} \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T \mid \mathbf{u}^*, \mathcal{Z}_T; \boldsymbol{\theta}) \\ = \prod_{r \in \mathcal{R}} \left\{ \left[\prod_{x \in \mathcal{X}} \mathbb{P}(d_{x,1}^{(r)} \mid s_1^{(r)}, \mathbf{z}_1^{(r)}; \boldsymbol{\alpha}_{s_1^{(r)},x}) \right] \cdot \mathbb{P}(s_1^{(r)} \mid u_r^*, \mathbf{z}_1^{(r)}; \boldsymbol{\beta}) \cdot \right. \\ \left. \prod_{t=2}^T \left[\prod_{x \in \mathcal{X}} \mathbb{P}(d_{x,t}^{(r)} \mid s_t^{(r)}, \mathbf{z}_t^{(r)}; \boldsymbol{\alpha}_{s_t^{(r)},x}) \right] \cdot \mathbb{P}(s_t^{(r)} \mid s_{t-1}^{(r)}, u_r^*, \mathbf{z}_t^{(r)}; \boldsymbol{\beta}) \right\}. \end{aligned}$$

We will treat $\rho_{s_1^{(r)}} := \mathbb{P}(s_1^{(r)} \mid u_r^*, \mathbf{z}_1^{(r)}; \boldsymbol{\beta})$ as additional parameters in the likelihood and obtain:

$$\begin{aligned} \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T \mid \mathbf{u}^*, \mathcal{Z}_T; \boldsymbol{\theta}) \\ = \prod_{r \in \mathcal{R}} \left\{ \left[\sum_{j=0}^N \mathbb{1}\{S_1^{(r)} = j\} \cdot \left(\prod_{x \in \mathcal{X}} \mathbb{P}(d_{x,1}^{(r)} \mid S_1^{(r)} = j, \mathbf{z}_1^{(r)}; \boldsymbol{\alpha}_{j,x}) \right) \cdot \rho_j \right] \cdot \right. \\ \left. \prod_{t=2}^T \left[\sum_{j=0}^N \sum_{i=0}^N \mathbb{1}\{S_{t-1}^{(r)} = i, S_t^{(r)} = j\} \cdot \left(\prod_{x \in \mathcal{X}} \mathbb{P}(d_{x,t}^{(r)} \mid S_t^{(r)} = j, \mathbf{z}_t^{(r)}; \boldsymbol{\alpha}_{j,x}) \right) \cdot \right. \right. \\ \left. \left. p_t^{ij}(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta}) \right] \right\}. \end{aligned}$$

Taking the logarithm leads to:

$$\begin{aligned} \log \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T \mid \mathbf{u}^*, \mathcal{Z}_T; \boldsymbol{\theta}) \\ = \sum_{r \in \mathcal{R}} \left\{ \sum_{x \in \mathcal{X}} \sum_{j=0}^N \mathbb{1}\{S_1^{(r)} = j\} \cdot \log \mathbb{P}(d_{x,1}^{(r)} \mid S_1^{(r)} = j, \mathbf{z}_1^{(r)}; \boldsymbol{\alpha}_{j,x}) + \right. \\ \sum_{j=0}^N \mathbb{1}\{S_1^{(r)} = j\} \cdot \log \rho_j + \\ \sum_{x \in \mathcal{X}} \sum_{t=2}^T \sum_{j=0}^N \mathbb{1}\{S_t^{(r)} = j\} \cdot \log \mathbb{P}(d_{x,t}^{(r)} \mid S_t^{(r)} = j, \mathbf{z}_t^{(r)}; \boldsymbol{\alpha}_{j,x}) + \\ \left. \sum_{t=2}^T \sum_{j=0}^N \sum_{i=0}^N \mathbb{1}\{S_{t-1}^{(r)} = i, S_t^{(r)} = j\} \cdot \log p_t^{ij}(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta}) \right\}, \end{aligned} \tag{B.1}$$

where (3.6) specifies the transition probabilities $p_t^{ij}(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta})$, and $\mathbb{P}(d_{x,t}^{(r)} \mid S_t^{(r)} = j, \mathbf{z}_t^{(r)}; \boldsymbol{\alpha}_{j,x})$ represents the probability mass function of the state-dependent Poisson distributions given in (3.4).

Following (4.3) in the main paper, we can then calculate the (unconditional) expected

complete-data log-likelihood as:

$$\begin{aligned}
& \log \mathcal{L}(\boldsymbol{\theta}, \tau \mid \mathcal{D}_T, \mathcal{S}_T, \mathcal{Z}_T) \\
&= \sum_{r \in \mathcal{R}} \left\{ \sum_{x \in \mathcal{X}} \sum_{j=0}^N \mathbb{1} \{S_1^{(r)} = j\} \cdot \log \mathbb{P} \left(D_{x,1}^{(r)} = d_{x,1}^{(r)} \mid S_1^{(r)} = j, \mathbf{z}_1^{(r)}; \boldsymbol{\alpha}_{j,x} \right) + \right. \\
&\quad \sum_{j=0}^N \mathbb{1} \{S_1^{(r)} = j\} \cdot \log \rho_j + \\
&\quad \sum_{x \in \mathcal{X}} \sum_{t=2}^T \sum_{j=0}^N \mathbb{1} \{S_t^{(r)} = j\} \cdot \log \mathbb{P} \left(D_{x,t}^{(r)} = d_{x,t}^{(r)} \mid S_t^{(r)} = j, \mathbf{z}_t^{(r)}; \boldsymbol{\alpha}_{j,x} \right) + \\
&\quad \left. \sum_{t=2}^T \sum_{j=0}^N \sum_{i=0}^N \mathbb{1} \{S_{t-1}^{(r)} = i, S_t^{(r)} = j\} \cdot \log p_t^{ij} \left(\mathbf{z}_t^{(r)}; \mathbf{u}_r^*, \boldsymbol{\beta} \right) \right\} - \\
&\quad \frac{1}{2} \log(\det(\mathbf{Q}(\tau)^{-1})) - \frac{1}{2} (\mathbf{u}^*(\boldsymbol{\theta}))^\top \mathbf{Q}(\tau) \mathbf{u}^*(\boldsymbol{\theta}) - \frac{1}{2} \log(\det(\mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*(\boldsymbol{\theta})))),
\end{aligned} \tag{B.2}$$

where

$$\begin{aligned}
& \mathbf{u}^* := \mathbf{u}^*(\boldsymbol{\theta}) = \arg \max_{\mathbf{u}} \left\{ \mathbb{E} [\log \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T, \mathbf{u} \mid \mathcal{Z}_T; \boldsymbol{\theta}, \tau)] \right\} \\
& \mathbf{S}^{-1} := \mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*(\boldsymbol{\theta})) = - \frac{\partial^2 \mathbb{E} [\log \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T, \mathbf{u} \mid \mathcal{Z}_T; \boldsymbol{\theta}, \tau)]}{\partial \mathbf{u} \partial \mathbf{u}^\top} \Big|_{\mathbf{u}=\mathbf{u}^*(\boldsymbol{\theta})}.
\end{aligned} \tag{B.3}$$

Note that we consider $\rho_j := \mathbb{P}(S_1^{(r)} = j \mid \mathbf{u}_r^*, \mathbf{z}_1^{(r)}; \boldsymbol{\beta})$ in (B.2) as additional region-independent parameters in the model. Appendix B.3 details the calculation procedure of the Hessian $\mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*(\boldsymbol{\theta}))$.

B.3 Hessian of the expected complete-data log-likelihood

Following (B.3), we first take the expectation of the (unconditional) complete-data log-likelihood with respect to the regime state process:

$$\begin{aligned}
& \mathbb{E} [\log \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T, \mathbf{u} \mid \mathcal{Z}_T; \boldsymbol{\theta}, \tau)] \\
&= \mathbb{E} [\log \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T \mid \mathbf{u}, \mathcal{Z}_T; \boldsymbol{\theta})] + \log f(\mathbf{u}; \tau) \\
&= \mathbb{E} [\log \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T \mid \mathbf{u}, \mathcal{Z}_T; \boldsymbol{\theta})] - \frac{R}{2} \log(2\pi) - \frac{1}{2} \log(\det(\mathbf{Q}(\tau)^{-1})) - \frac{1}{2} \mathbf{u}^\top \mathbf{Q}(\tau) \mathbf{u},
\end{aligned}$$

with $\mathbb{P}(\mathcal{D}_T, \mathcal{S}_T \mid \mathbf{u}, \mathcal{Z}_T; \boldsymbol{\theta})$ the complete-data log-likelihood, given the spatial effect $\mathbf{u}$ (see (B.1)), and $\mathbf{Q}(\tau)$ the precision matrix that depends on the precision parameter τ .

We then calculate the expectation of the expression in (B.1), conditionally on the spatial effect $\mathbf{u}$, the death counts $\mathcal{D}_T$, and the covariate information $\mathcal{Z}_T$. We obtain:

$$\begin{aligned}
& \mathbb{E} [\log \mathbb{P}(\mathcal{D}_T, \mathcal{S}_T \mid \mathbf{u}, \mathcal{Z}_T; \boldsymbol{\theta})] \\
&= \sum_{r \in \mathcal{R}} \left\{ \sum_{x \in \mathcal{X}} \sum_{j=0}^N \mathbb{P} \left(S_1^{(r)} = j \mid \mathbf{u}_r, \mathbf{z}_1^{(r)}; \boldsymbol{\theta} \right) \cdot \log \mathbb{P} \left(d_{x,1}^{(r)} \mid S_1^{(r)} = j, \mathbf{z}_1^{(r)}; \boldsymbol{\alpha}_{j,x} \right) + \right. \\
&\quad \sum_{j=0}^N \mathbb{P} \left(S_1^{(r)} = j \mid \mathbf{u}_r, \mathbf{z}_1^{(r)}; \boldsymbol{\theta} \right) \cdot \log \rho_j + \\
&\quad \left. \sum_{x \in \mathcal{X}} \sum_{t=2}^T \sum_{j=0}^N \mathbb{P} \left(S_t^{(r)} = j \mid \mathbf{u}_r, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta} \right) \cdot \log \mathbb{P} \left(d_{x,t}^{(r)} \mid S_t^{(r)} = j, \mathbf{z}_t^{(r)}; \boldsymbol{\alpha}_{j,x} \right) + \right.
\end{aligned} \tag{B.4}$$

$$\sum_{t=2}^T \sum_{j=0}^N \sum_{i=0}^N \mathbb{P} \left(S_{t-1}^{(r)} = i, S_t^{(r)} = j \mid u_r, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta} \right) \cdot \log p_t^{ij} \left(\mathbf{z}_t^{(r)}; u_r, \boldsymbol{\beta} \right) \Big\},$$

where $\mathcal{D}_t^{(r)} := \{\mathcal{D}_{x,t}^{(r)} \mid x \in \mathcal{X}\}$ is the set of death counts for region r at all ages $x \in \mathcal{X}$ up to time t .

In the EM algorithm, we take the expectation of the complete-data log-likelihood conditional on the current best-estimate of the spatial effect $\mathbf{u}$ and regime-switching parameter vector $\boldsymbol{\theta}$, see Appendix B.4. As such, the marginal and joint state probabilities in (B.4) are considered to be constant when computing the negative Hessian $\mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*(\boldsymbol{\theta}))$ from (B.3). The diagonal entries (r, r) of the negative Hessian equal:

$$- \sum_{t=2}^T \sum_{j=0}^N \sum_{i=0}^N \mathbb{P} \left\{ S_{t-1}^{(r)} = i, S_t^{(r)} = j \mid u_r, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta} \right\} \cdot \frac{\partial^2 \log p_t^{ij} \left(\mathbf{z}_t^{(r)}; u_r, \boldsymbol{\beta} \right)}{\partial u_r^2} + Q_{rr}(\tau),$$

where $Q_{rr}(\tau)$ is entry (r, r) of the precision matrix $\mathbf{Q}(\tau)$. Using the expression for the regime state transition probabilities in (3.6), we obtain:

$$\frac{\partial^2 \log p_t^{ij} \left(\mathbf{z}_t^{(r)}; u_r, \boldsymbol{\beta} \right)}{\partial u_r^2} = -p_t^{i0} \left(\mathbf{z}_t^{(r)}; u_r, \boldsymbol{\beta} \right) \cdot \left(1 - p_t^{i0} \left(\mathbf{z}_t^{(r)}; u_r, \boldsymbol{\beta} \right) \right),$$

for $i, j \in \{0, 1, \dots, N\}$ with $(ij) \notin \{(12), (21)\}$. Entry (r, s) of the Hessian is then equal to:

$$\begin{aligned} & \left(\sum_{t=2}^T \sum_{j=0}^N \sum_{i=0}^N \mathbb{P} \left\{ S_{t-1}^{(r)} = i, S_t^{(r)} = j \mid u_r, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta} \right\} \cdot p_t^{i0} \left(\mathbf{z}_t^{(r)}; u_r, \boldsymbol{\beta} \right) \cdot \right. \\ & \quad \left. \left(1 - p_t^{i0} \left(\mathbf{z}_t^{(r)}; u_r, \boldsymbol{\beta} \right) \right) \right) \cdot \mathbb{1}\{r = s\} + Q_{rs}(\tau). \end{aligned} \tag{B.5}$$

We conclude that the negative Hessian $\mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*(\boldsymbol{\theta}))$ is the sum of the precision matrix $\mathbf{Q}(\tau)$ and a diagonal matrix with as entries the second order partial derivatives of the log transition probabilities, summed over all the states and periods under consideration.

B.4 Expectation-maximization algorithm

We provide details on the proposed EM algorithm used to optimize the incomplete-data log-likelihood in Section 4.2.3. For iteration $k + 1$, we apply:

E-step. We compute the conditional expectation of the complete-data log-likelihood with respect to the regime-state process, given the current best estimates of the regime-switching parameter vector $\boldsymbol{\theta}_k^*$ and the spatial effect $\mathbf{u}_k^*$:

$$\begin{aligned} & \mathbb{E} [\log \mathcal{L}(\boldsymbol{\theta}, \tau \mid \mathcal{D}_T, \mathcal{S}_T, \mathcal{Z}_T) \mid \mathbf{u}_k^*, \mathcal{D}_T, \mathcal{Z}_T; \boldsymbol{\theta}_k^*] \\ &= \sum_{r \in \mathcal{R}} \left\{ \sum_{x \in \mathcal{X}} \sum_{j=0}^N \mathbb{P} \left(S_1^{(r)} = j \mid u_{k,r}^*, \mathbf{z}_1^{(r)}; \boldsymbol{\theta}_k^* \right) \cdot \log \mathbb{P} \left(d_{x,1}^{(r)} \mid S_1^{(r)} = j, \mathbf{z}_1^{(r)}; \boldsymbol{\alpha}_{j,x} \right) + \right. \\ & \quad \sum_{j=0}^N \mathbb{P} \left(S_1^{(r)} = j \mid u_{k,r}^*, \mathbf{z}_1^{(r)}; \boldsymbol{\theta}_k^* \right) \cdot \log \rho_j + \\ & \quad \left. \sum_{x \in \mathcal{X}} \sum_{t=2}^T \sum_{j=0}^N \mathbb{P} \left(S_t^{(r)} = j \mid u_{k,r}^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^* \right) \cdot \log \mathbb{P} \left(d_{x,t}^{(r)} \mid S_t^{(r)} = j, \mathbf{z}_t^{(r)}; \boldsymbol{\alpha}_{j,x} \right) + \right. \end{aligned} \tag{B.6}$$

$$\left. \sum_{t=2}^T \sum_{j=0}^N \sum_{i=0}^N \mathbb{P} \left(S_{t-1}^{(r)} = i, S_t^{(r)} = j \mid u_{k,r}^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^* \right) \cdot \log p_t^{ij} \left(\mathbf{z}_t^{(r)}; u_r, \boldsymbol{\beta} \right) \right\} - \frac{1}{2} \log(\det(\mathbf{Q}(\tau)^{-1})) - \frac{1}{2} (\mathbf{u}^*(\boldsymbol{\theta}))^\top \mathbf{Q}(\tau) \mathbf{u}^*(\boldsymbol{\theta}) - \frac{1}{2} \log(\det(\mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*(\boldsymbol{\theta})))),$$

where $u_{k,r}^*$ is the r -th component of the spatial effect $\mathbf{u}_k^*$, and $\boldsymbol{\beta}_k^*$ is a subset of the parameter vector $\boldsymbol{\theta}_k^*$, related to the parameters in the transition probabilities. We refer to the involved state probabilities, i.e., $\mathbb{P}(S_t^{(r)} = j \mid u_{k,r}^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^*)$ and $\mathbb{P}(S_{t-1}^{(r)} = i, S_t^{(r)} = j \mid u_{k,r}^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^*)$, as the filtered marginal and joint state probabilities, respectively. The term filtered indicates that these probabilities are conditioned only on information available up to and including time t . The computation of these marginal and joint state probabilities follows the iterative procedure described below.

1. Calculate the state-dependent Poisson death count and transition probabilities from (3.4) and (3.6), using the current best estimates of the regime-switching parameter vector $\boldsymbol{\theta}_k^*$ and the spatial effect $\mathbf{u}_k^*$.
2. For every region $r \in \mathcal{R}$, we calculate the filtered joint state probabilities by iterating over the following steps for $t = 2, 3, \dots, T$:

(a) For $t = 2$, we compute:

$$\begin{aligned} \mathbb{P} \left(S_1^{(r)} = i, S_2^{(r)} = j, \left(D_{x,2}^{(r)} \right)_{x \in \mathcal{X}} = \left(d_{x,2}^{(r)} \right)_{x \in \mathcal{X}} \mid u_{k,r}^*, \mathcal{D}_1^{(r)}, \mathcal{Z}_2^{(r)}; \boldsymbol{\theta}_k^* \right) \\ = \rho_{i,k}^* \cdot p_2^{ij} \left(\mathbf{z}_2^{(r)}; u_{k,r}^*, \boldsymbol{\beta}_k^* \right) \cdot \prod_{x \in \mathcal{X}} \mathbb{P} \left(D_{x,2}^{(r)} = d_{x,2}^{(r)} \mid S_2^{(r)} = j, \mathbf{z}_2^{(r)}; \boldsymbol{\alpha}_{j,x,k}^* \right), \end{aligned}$$

with $i, j \in \{0, 1, \dots, N\}$. For $t \geq 2$, we compute the probability:

$$\begin{aligned} \mathbb{P} \left(S_{t-1}^{(r)} = i, S_t^{(r)} = j, \left(D_{x,t}^{(r)} \right)_{x \in \mathcal{X}} = \left(d_{x,t}^{(r)} \right)_{x \in \mathcal{X}} \mid u_{k,r}^*, \mathcal{D}_{t-1}^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^* \right) \\ = \sum_{h=0}^N \mathbb{P} \left(S_{t-2}^{(r)} = h, S_{t-1}^{(r)} = i \mid u_{k,r}^*, \mathcal{D}_{t-1}^{(r)}, \mathcal{Z}_{t-1}^{(r)}; \boldsymbol{\theta}_k^* \right) \cdot p_t^{ij} \left(\mathbf{z}_t^{(r)}; u_{k,r}^*, \boldsymbol{\beta}_k^* \right) \cdot \\ \prod_{x \in \mathcal{X}} \mathbb{P} \left(D_{x,t}^{(r)} = d_{x,t}^{(r)} \mid S_t^{(r)} = j, \mathbf{z}_t^{(r)}; \boldsymbol{\alpha}_{j,x,k}^* \right), \end{aligned}$$

with $i, j \in \{0, 1, \dots, N\}$. In the above calculation, we relied on the assumption that the death counts are conditionally independent, given the regime state the system is in. The second and third terms are the transition probabilities and the state-dependent Poisson probabilities from step 1 and the first term is obtained from step 2.c of the previous iteration at time $t - 1$.

(b) Using step 2.a, we calculate:

$$\begin{aligned} \mathbb{P} \left(\left(D_{x,t}^{(r)} \right)_{x \in \mathcal{X}} = \left(d_{x,t}^{(r)} \right)_{x \in \mathcal{X}} \mid u_{k,r}^*, \mathcal{D}_{t-1}^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^* \right) \\ = \sum_{i=0}^N \sum_{j=0}^N \mathbb{P} \left(S_{t-1}^{(r)} = i, S_t^{(r)} = j, \left(D_{x,t}^{(r)} \right)_{x \in \mathcal{X}} = \left(d_{x,t}^{(r)} \right)_{x \in \mathcal{X}} \mid u_{k,r}^*, \mathcal{D}_{t-1}^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^* \right). \end{aligned}$$

(c) From step 2.a and 2.b, we obtain the time t filtered joint state probabilities for $i, j \in \{0, 1, \dots, N\}$:

$$\mathbb{P} \left(S_{t-1}^{(r)} = i, S_t^{(r)} = j \mid u_{k,r}^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^* \right)$$

$$= \frac{\mathbb{P}\left(S_{t-1}^{(r)} = i, S_t^{(r)} = j, \left(D_{x,t}^{(r)}\right)_{x \in \mathcal{X}} = \left(d_{x,t}^{(r)}\right)_{x \in \mathcal{X}} \mid u_{k,r}^*, \mathcal{D}_{t-1}^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^*\right)}{\mathbb{P}\left(\left(D_{x,t}^{(r)}\right)_{x \in \mathcal{X}} = \left(d_{x,t}^{(r)}\right)_{x \in \mathcal{X}} \mid u_{k,r}^*, \mathcal{D}_{t-1}^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^*\right)}.$$

After iterating over these steps for $t = 2, 3, \dots, T$, we obtain all filtered joint state probabilities. We obtain the filtered marginal state probabilities as:

$$\mathbb{P}\left(S_t^{(r)} = j \mid u_{k,r}^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^*\right) = \sum_{i=0}^N \mathbb{P}\left(S_{t-1}^{(r)} = i, S_t^{(r)} = j \mid u_{k,r}^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}_k^*\right),$$

for $j = 0, 1, \dots, N$ and $t = 2, 3, \dots, T$.

M-step. We update the regime-switching parameter vector $\boldsymbol{\theta}$ by solving:

$$\boldsymbol{\theta}_{k+1}^* = \arg \max_{\boldsymbol{\theta}} \mathbb{E} [\log \mathcal{L}(\boldsymbol{\theta}, \tau \mid \mathcal{D}_T, \mathcal{S}_T, \mathcal{Z}_T) \mid \mathbf{u}_k^*, \mathcal{D}_T, \mathcal{Z}_T; \boldsymbol{\theta}_k^*].$$

For computational convenience, we can split this expected complete-data log-likelihood into a part that only depends on the parameter vectors $\boldsymbol{\alpha}_j := (\boldsymbol{\alpha}_{j,x})_{x \in \mathcal{X}}$, for each $j \in \{1, 2\}$, and a part that only depends on the parameters $\beta_{01}, \beta_{02}, \beta_{11}$, and β_{22} . As such, we split the maximization problem into three separate maximization problems.

Update spatial effect. We update the mode $\mathbf{u}^*$ from (B.3) under the current optimal parameter vector $\boldsymbol{\theta}_{k+1}^*$:

$$\mathbf{u}_{k+1}^* := \mathbf{u}^*(\boldsymbol{\theta}_{k+1}^*) = \arg \max_{\mathbf{u}} \mathbb{E} [\log \mathcal{L}(\boldsymbol{\theta}_{k+1}^*, \tau \mid \mathcal{D}_T, \mathcal{S}_T, \mathcal{Z}_T) \mid \mathbf{u}_k^*, \mathcal{D}_T, \mathcal{Z}_T; \boldsymbol{\theta}_k^*].$$

B.5 Optimizing the precision parameter

The precision parameter τ appears in the ICAR prior distribution applied to the spatial effect $\mathbf{U}$ and controls the degree of spatial autocorrelation in the transition probabilities of the different regions (see (3.5) in the paper). We estimate τ by maximizing the so-called profile log-likelihood. Let $\mathcal{G}$ be a grid of candidate values for τ . For each $\tau \in \mathcal{G}$, we obtain the optimal parameter vector $\boldsymbol{\theta}^*(\tau)$ in the three-state regime-switching process by maximizing the incomplete-data log-likelihood using the EM-algorithm summarized in Section 4.2.3 of the paper:

$$\boldsymbol{\theta}^*(\tau) = \arg \max_{\boldsymbol{\theta}} \log \tilde{\mathcal{L}}(\boldsymbol{\theta}, \tau \mid \mathcal{D}_T, \mathcal{Z}_T).$$

The optimal precision parameter τ is then the value that maximizes the profile log-likelihood:

$$\tau^* = \arg \max_{\tau \in \mathcal{G}} \log \tilde{\mathcal{L}}(\boldsymbol{\theta}^*(\tau), \tau \mid \mathcal{D}_T, \mathcal{Z}_T).$$

B.6 Parameter uncertainty in the regime-switching model

We describe the procedure to calculate the Fisher Information Matrix (FIM) of the unknown parameter vector $\boldsymbol{\theta}$ in the regime-switching model according to the approach outlined in Meng and Spall (2017).

1. We apply the EM-algorithm to maximize the incomplete-data log-likelihood using the observed death counts dataset, $\mathcal{D}_T$, and the optimal precision parameter τ^* , as determined in Section B.5. This maximization results in the optimal parameter vector $\boldsymbol{\theta}^*(\tau^*)$ after reaching an absolute convergence in the EM-algorithm, see Section 4.2.3.

2. We define the expectation of the complete-data log-likelihood with respect to the conditional distribution of the regime state process $\mathcal{S}_T$, given the observed death counts dataset $\mathcal{D}_T$, the covariate information $\mathcal{Z}_T$, and the optimal parameter vector $\boldsymbol{\theta}^*$ and spatial effect $\mathbf{u}^*$, obtained from the EM-algorithm:

$$Q(\boldsymbol{\theta} \mid \mathcal{D}_T, \mathcal{Z}_T, \boldsymbol{\theta}^*) = \mathbb{E}[\log \mathcal{L}(\boldsymbol{\theta}, \tau^* \mid \mathcal{D}_T, \mathcal{S}_T, \mathcal{Z}_T) \mid \mathbf{u}^*, \mathcal{D}_T, \mathcal{Z}_T; \boldsymbol{\theta}^*],$$

where τ^* is the value of the precision parameter τ that maximizes the profile log-likelihood, see Section B.5. Subsequently, we calculate the derivative of the expected complete-data log-likelihood with respect to the parameter vector $\boldsymbol{\theta}$, evaluated in $\boldsymbol{\theta}^*$. We define:

$$\mathbf{G}(\boldsymbol{\theta}^* \mid \mathcal{D}_T, \mathcal{Z}_T) = \left. \frac{\partial Q(\boldsymbol{\theta} \mid \mathcal{D}_T, \mathcal{Z}_T, \boldsymbol{\theta}^*)}{\partial \boldsymbol{\theta}} \right|_{\boldsymbol{\theta}=\boldsymbol{\theta}^*}. \quad (\text{B.7})$$

Since the complete-data log-likelihood can be expressed as a sum of separate terms that depend on the parameter groups $\boldsymbol{\alpha}_1 := (\boldsymbol{\alpha}_{1,x})_{x \in \mathcal{X}}$, $\boldsymbol{\alpha}_2 := (\boldsymbol{\alpha}_{2,x})_{x \in \mathcal{X}}$, and $\boldsymbol{\beta} := (\boldsymbol{\beta}_0, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2)$ individually, we calculate the derivatives with respect to each parameter group separately.

Derivative w.r.t. $\boldsymbol{\alpha}_1$. We obtain:

$$\mathbf{G}(\boldsymbol{\alpha}_{1,x} \mid \mathcal{D}_T, \mathcal{Z}_T) = \sum_{r \in \mathcal{R}} \sum_{t=1}^T P\left(S_t^{(r)} = 1 \mid u_r^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}^*\right) \left(d_{x,t}^{(r)} - \hat{b}_{x,t}^{(r)} e^{(\mathbf{z}_t^{(r)})^\top \boldsymbol{\alpha}_{1,x}}\right) \mathbf{z}_t^{(1,r)}, \quad (\text{B.8})$$

for every $x \in \mathcal{X}$, where $\mathbf{z}_t^{(1,r)}$ is the subset of the covariate vector $\mathbf{z}_t^{(r)}$ for which the parameter coefficients in $\boldsymbol{\alpha}_{1,x}$ are not set to zero.

Derivative w.r.t. $\boldsymbol{\alpha}_2$. In a similar way, we obtain:

$$\mathbf{G}(\boldsymbol{\alpha}_{2,x} \mid \mathcal{D}_T, \mathcal{Z}_T) = \sum_{r \in \mathcal{R}} \sum_{t=1}^T P\left(S_t^{(r)} = 2 \mid u_r^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}^*\right) \left(d_{x,t}^{(r)} - \hat{b}_{x,t}^{(r)} e^{(\mathbf{z}_t^{(r)})^\top \boldsymbol{\alpha}_{2,x}}\right) \mathbf{z}_t^{(2,r)}, \quad (\text{B.9})$$

for every $x \in \mathcal{X}$, where $\mathbf{z}_t^{(2,r)}$ is the subset of the covariate vector $\mathbf{z}_t^{(r)}$ for which the parameter coefficients in $\boldsymbol{\alpha}_{2,x}$ are not set to zero.

Derivative w.r.t. $\boldsymbol{\beta}$. We calculate the derivative of the expected complete-data log-likelihood with respect to $\boldsymbol{\beta}_{ij}$, where $ij \in \{01, 02, 11, 22\}$ (see (3.6)):

$$\begin{aligned} \mathbf{G}(\boldsymbol{\beta}_{ij} \mid \mathcal{D}_T, \mathcal{Z}_T) &= \sum_{r \in \mathcal{R}} \sum_{t=2}^T \sum_{j'=0}^N \mathbb{P}\left(S_{t-1}^{(r)} = i, S_t^{(r)} = j' \mid u_r^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}^*\right) \frac{\partial \log p_t^{ij'}(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta})}{\partial \boldsymbol{\beta}_{ij}} \\ &\quad - \frac{1}{2} \frac{\partial \log(\det(\mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*)))}{\boldsymbol{\beta}_{ij}}. \end{aligned} \quad (\text{B.10})$$

We first calculate the derivative of the log transition probabilities with respect to $\boldsymbol{\beta}_{ij}$. A straightforward calculation shows that:

$$\frac{\partial \log p_t^{ij'}(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta})}{\partial \boldsymbol{\beta}_{ij}} = \begin{cases} -p_t^{ij}(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta}) \mathbf{z}_t^{(ij,r)} & j \neq j' \\ \left(1 - p_t^{ij}(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta})\right) \mathbf{z}_t^{(ij,r)} & j = j', \end{cases} \quad (\text{B.11})$$

where $\mathbf{z}_t^{(ij,r)}$ is the subset of the covariate vector $\mathbf{z}_t^{(r)}$ for which the parameter coefficients in β_{ij} are not set to zero.

Next, we calculate the derivative of the logarithm of the determinant of the Hessian $\mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*)$, see (B.5), with respect to the l -th entry of the parameter vector β_{ij} , i.e., β_{ijl} . We apply Jacobi's formula and obtain:

$$\frac{\partial \log(\det \mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*))}{\beta_{ijl}} = \text{tr} \left(\mathbf{S}(\boldsymbol{\theta}, \mathbf{u}^*) \frac{\partial \mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*)}{\partial \beta_{ijl}} \right).$$

Note that the derivative of the Hessian with respect to β_{ijl} is a diagonal matrix with entries:

$$\begin{aligned} \left(\frac{\mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*)}{\partial \beta_{ijl}} \right)_{rr} &= \sum_{t=2}^T \sum_{j'=0}^N \mathbb{P} \left(S_{t-1}^{(r)} = i, S_t^{(r)} = j' \mid u_r^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}^* \right) \cdot \\ &\quad p_t^{i0} \left(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta} \right) \cdot \left(2p_t^{i0} \left(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta} \right) - 1 \right) \cdot p_t^{ij} \left(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta} \right) \cdot \mathbf{z}_{t,l}^{(ij,r)}, \end{aligned}$$

for any $r \in \mathcal{R}$. This then leads to:

$$\begin{aligned} &\frac{\partial \log(\det \mathbf{S}^{-1}(\boldsymbol{\theta}, \mathbf{u}^*))}{\beta_{ij}} \\ &= \sum_{r \in \mathcal{R}} (\mathbf{S}(\boldsymbol{\theta}, \mathbf{u}^*))_{rr} \left[\sum_{t=2}^T \sum_{j'=0}^N \mathbb{P} \left(S_{t-1}^{(r)} = i, S_t^{(r)} = j' \mid u_r^*, \mathcal{D}_t^{(r)}, \mathcal{Z}_t^{(r)}; \boldsymbol{\theta}^* \right) \cdot \right. \\ &\quad \left. p_t^{i0} \left(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta} \right) \cdot \left(2p_t^{i0} \left(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta} \right) - 1 \right) \cdot p_t^{ij} \left(\mathbf{z}_t^{(r)}; u_r^*, \boldsymbol{\beta} \right) \cdot \mathbf{z}_{t,l}^{(ij,r)} \right]. \end{aligned} \tag{B.12}$$

We substitute the results from (B.11) and (B.12) into (B.10) to derive the gradient of the expected complete-data log-likelihood with respect to β_{ij} , for $ij \in \{01, 02, 11, 22\}$.

Using the derivatives from (B.8), (B.9), and (B.10), the derivative of the expected complete-data log-likelihood in (B.7) results in:

$$\mathbf{G}(\boldsymbol{\theta}^* \mid \mathcal{D}_T, \mathcal{Z}_T) = \begin{pmatrix} (\mathbf{G}(\boldsymbol{\alpha}_{1,x}^* \mid \mathcal{D}_T, \mathcal{Z}_T))_{x \in \mathcal{X}} \\ (\mathbf{G}(\boldsymbol{\alpha}_{2,x}^* \mid \mathcal{D}_T, \mathcal{Z}_T))_{x \in \mathcal{X}} \\ \mathbf{G}(\boldsymbol{\beta}_{01}^* \mid \mathcal{D}_T, \mathcal{Z}_T) \\ \mathbf{G}(\boldsymbol{\beta}_{02}^* \mid \mathcal{D}_T, \mathcal{Z}_T) \\ \mathbf{G}(\boldsymbol{\beta}_{11}^* \mid \mathcal{D}_T, \mathcal{Z}_T) \\ \mathbf{G}(\boldsymbol{\beta}_{22}^* \mid \mathcal{D}_T, \mathcal{Z}_T) \end{pmatrix} \in \mathbb{R}^q.$$

3. We generate M synthetic datasets of death counts by simulating from the state-dependent distribution function in (3.4), using the optimal parameter vector $\boldsymbol{\theta}^*(\tau^*)$. This involves two steps: (i) we first generate M sample paths for the state process $\mathcal{S}_T$ of the underlying Markov chain, based on the estimated transition probabilities in (3.6) using $\boldsymbol{\theta}^*(\tau^*)$, and (ii) we subsequently sample from the Poisson distribution conditioned on the simulated states. During this sampling process, we keep the covariates $\mathbf{z}_t^{(r)}$, for $j = 1, 2$, fixed. We denote the generated M synthetic death counts datasets as $\mathcal{D}'_{T,1}, \mathcal{D}'_{T,2}, \dots, \mathcal{D}'_{T,M}$ and the corresponding underlying simulated Markov chains as $\mathcal{S}'_{T,1}, \mathcal{S}'_{T,2}, \dots, \mathcal{S}'_{T,M}$.
4. We generate M perturbation vectors $\boldsymbol{\Delta}_k \in \mathbb{R}^q$ for $k = 1, 2, \dots, M$, where each $\boldsymbol{\Delta}_{kj}$ is independently drawn across different k and j . Each $\boldsymbol{\Delta}_{kj}$ follows a scaled Rademacher distribution with a scale factor c , meaning it takes the value $+c$ with a probability of

0.5 and $-c$ with a probability of 0.5.

5. For each $k = 1, \dots, M$, we define the first-order differences:

$$\delta \mathbf{G}_k = \mathbf{G}(\boldsymbol{\theta}^* + \boldsymbol{\Delta}_k | \mathcal{D}'_{T,k}, \mathcal{Z}_t) - \mathbf{G}(\boldsymbol{\theta}^* - \boldsymbol{\Delta}_k | \mathcal{D}'_{T,k}, \mathcal{Z}_t).$$

According to [Meng and Spall \(2017\)](#), we can then approximate the Hessian of the incomplete-data log-likelihood for death counts dataset $\mathcal{D}'_{T,k}$ as:

$$\hat{\mathbf{H}}_k(\boldsymbol{\theta}^* | \mathcal{D}'_{T,k}, \mathcal{Z}_t) = 0.5 \cdot \left(\frac{\delta \mathbf{G}_k}{2} [\Delta_{k1}^{-1}, \Delta_{k2}^{-1}, \dots, \Delta_{kq}^{-1}] + \left(\frac{\delta \mathbf{G}_k}{2} [\Delta_{k1}^{-1}, \Delta_{k2}^{-1}, \dots, \Delta_{kq}^{-1}] \right)^\top \right) \quad (\text{B.13})$$

Lastly, we approximate the Fisher information matrix by taking the negative average of all $\hat{\mathbf{H}}_k(\boldsymbol{\theta}^* | \mathcal{D}'_{T,k}, \mathcal{Z}_t)$. The diagonal entries of the inverse Fisher information matrix serve as measures of parameter uncertainty, enabling the construction of confidence intervals around the estimated parameter vector $\boldsymbol{\theta}^*$.

C Constructing temperature and influenza scenarios for short-term mortality forecasting

C.1 Temperature scenario

We use temperature projections based on the RCP 2.6, RCP 4.5, and RCP 8.5 scenarios from the Copernicus Climate Change Service ([Copernicus Climate Change Service, 2021](#)). These scenarios represent a low-emission scenario aiming to limit warming (RCP 2.6), a medium stabilization scenario (RCP 4.5), and a high-emission, worst-case scenario that leads to significant global warming (RCP 8.5). Since we focus on short-term mortality forecasting (up to 2030), differences between the RCP scenarios remain small because the climate system only responds slowly to changes in emissions. Over this short-term period, natural variations, such as ocean cycles or weather patterns, cause temperatures under the RCP 4.5 or RCP 2.6 scenario to occasionally exceed those under the RCP 8.5 scenario. However, our goal is to provide plausible temperature projections for the near future, rather than assessing long-term differences.

We extract bias-adjusted daily average temperature projections with a 5 km resolution from the 27th ISO week of 2024 to January 1, 2030. The bias adjustment reduces bias in climate model projections by comparing them with observed reference data. Using the population-weighted approach, we transform the daily temperature projections to regional-level (NUTS 2) projections for each RCP scenario. We then compute daily anomalies and extreme indicators from these regional temperature levels and average them weekly (see Section 2 of the paper). Figure C.1 shows the scenario projections for the daily average temperature in Pays de la Loire and Provence-Alpes-Côte d'Azur. Section E of the Online Appendix contains the results for the other regions.

C.2 Influenza scenario

To generate plausible scenarios for the weekly number of new influenza incidences, we apply the Bayesian stochastic SIRS model based on the framework introduced by [Bucyibaruta et al. \(2023\)](#). Unlike traditional SIRS models that track the total number of currently infected individuals, our approach focuses explicitly on modeling the weekly number of new infections. Let I_t denote the number of new influenza cases per 1 000 inhabitants in week t , as reported by the French Sentinelles network. Additionally, let S_t represent the number of susceptible individuals per 1 000 inhabitants at week t , and let R_t denote the individuals




Figure C.1: Projections for the daily average temperature based on the RCP 2.6, RCP 4.5, and RCP 8.5 scenarios (left panel), and for the weekly influenza activity per 1 000 inhabitants based on the proposed Bayesian stochastic SIRS model (right panel). We provide a moderate, high, and severe incidence scenario for influenza activity using the point-wise 50%, 80%, and 95% quantiles of the posterior predictive samples, respectively. The top panels show the result for Pays de la Loire, while the bottom panels depict the results for Provenances-Alpes-Côte d’Azur.

who are not susceptible to influenza infection in week t . This latter group includes those who are either still infected or have acquired immunity from a prior infection.

Given the number of susceptible and newly infected individuals at week $t - 1$, we model the number of new infections at week t as:

$$I_t \sim \text{POI}(\lambda_t),$$

$$\lambda_t = \frac{S_{t-1}}{N} \cdot (I_{t-1} + I_{t-2})^\kappa \cdot \exp(r_t + \epsilon_t),$$

where λ_t is the expected number of new infections, which depends on the susceptible population at week $t - 1$, the number of infections in the two preceding weeks, and the rate r_t incorporates seasonal variation. The parameter κ allows for non-linear transmission effects, and the error terms $\epsilon_t \sim \mathcal{U}[0, \zeta]$ allow to capture unexplained random variations that are not accounted for by the deterministic model structure. We update the susceptible and non-susceptible compartments recursively:

$$S_t = S_{t-1} - \frac{S_{t-1}}{N} \cdot (I_{t-1} + I_{t-2})^\kappa \cdot \exp(r_t + \epsilon_t) + \psi R_{t-1},$$

$$R_t = (1 - \psi) R_{t-1} + I_{t-1}.$$

Here, ψ represents the rate at which individuals transition back to the susceptible compartment. The number of non-susceptible individuals at week t , i.e., R_t , consists of those who were infected in the previous week as well as a fraction of those who were not susceptible a

week before. Lastly, we update the susceptible population S_t based on the previous week's susceptible population, the expected infections at week t , and the fraction of individuals re-entering susceptibility.

We impose the following structure for r_t to capture the seasonal dependence in the influenza infections (Bucyibaruta et al., 2023):

$$r_t = \varphi_{q(t)}$$

$$q(t) = \begin{cases} \min\left(\text{floor}\left(\frac{w(t)}{4}\right) + 1, 13\right) & \text{if } w(t) \bmod 4 \neq 0 \\ \text{floor}\left(\frac{w(t)}{4}\right) & \text{if } w(t) \bmod 4 = 0, \end{cases}$$

where $w(t)$ is the ISO-week number in the year $y(t)$ of time point t . Since some years contain 53 ISO-weeks, we restrict the maximum number of period effects to 13.

We apply Bayesian inference through Markov Chain Monte Carlo (MCMC) methods. Hereto, we use the JAGS (Just Another Gibbs Sampler) implementation to estimate the posterior distribution of the weekly new infections I_t . Furthermore, we assign independent normal priors with zero mean and unit variance to the parameters ϕ_j , for $j = 1, 2, \dots, 13$, and uniform priors to the parameters ψ , κ , and ζ . The model structure supports forecasting by treating future infections as missing data. More specifically, we extend the observed time series of infections by appending NA values for the weeks we aim to forecast. JAGS interprets these NA entries as parameters to be sampled, which enables the MCMC algorithm to produce posterior predictive samples for both the model parameters and the unobserved future infections. We repeat this procedure for every NUTS 2 region within France.

For scenario-based forecasting, each posterior predictive sample represents a complete trajectory of weekly new infections. We define a moderate, high, and severe influenza scenario by taking the point-wise 50%, 80%, and 95% quantile of the posterior predictive samples. Figure C.1 illustrates these three trajectories for weekly influenza in Pays de la Loire and Provence-Alpes-Côte d'Azur. The influenza scenarios for the remaining NUTS 2 regions can be found in Section E of the Online Appendix. Lastly, we set COVID-19 hospitalizations to zero in the forecasting period. However, users can create alternative scenarios to explore the potential impact of new respiratory infections such as COVID-19 on mortality, if desired.