PRICING LIFE INSURANCE USING BIVARIATE TEMPERATURE-MORTALITY SEASONAL HIDDEN MARKOV MODELS

Samuel Darwin Dona Kouton, Karim Barigou, Thai Nguyen

LIDAM Discussion Paper ISBA
2026 / 19

UCLouvain

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

ISBA

Voie du Roman Pays 20 - L1.04.01

B-1348 Louvain-la-Neuve

Email : lidam-library@uclouvain.be

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

Pricing life insurance using bivariate temperature-mortality seasonal hidden Markov models

Samuel Darwin Dona Kouton* Karim Barigou† Thai Nguyen‡

May 29, 2026

Abstract

We study the dependence between weekly mortality and temperature through a non-homogeneous Seasonal Hidden Markov Model (SHMM) with trend. The framework jointly models mortality and temperature while allowing for regime switches, explicit seasonality, smooth long-term trends, and time-varying transition probabilities in latent dynamics. We show that the general identifiability and strong consistency results of Touron (2018) apply to our Poisson- and Gaussian-emission specifications. We estimate the model on weekly Italian mortality and temperature data using two- and three-state specifications. Empirically, the SHMM improves fit relative to non-switching benchmarks and yields an interpretable state-dependent characterization of the temperature–mortality relationship. We then develop a change-of-measure framework for pricing mortality-linked insurance contracts that incorporates temperature-related mortality risk and latent regime uncertainty.

Keywords: life insurance; mortality modeling; regime-switching; seasonality; temperature; climate risk.

1 Introduction

The relationship between ambient temperature and mortality is nonlinear, seasonal, and state-dependent across time. Heat waves raise short-term mortality risks, while cold spells are associated with delayed increases in death counts; both interact with population vulnerability, baseline seasonality, and longer-term trends. For actuaries and demographers, capturing this dependence is crucial for short-term forecasting, scenario design, risk aggregation, and the pricing of temperature-sensitive mortality-linked products.

In this paper, we propose a *Seasonal Hidden Markov Model with trend* (SHMM) to jointly model weekly mortality and temperature. A discrete latent process governs regime changes, while the observable components incorporate seasonal patterns through Fourier terms, smooth time trends, and time-varying transition probabilities. Conditional on the latent state, mortality is modeled using either a Poisson or a Normal distribution, while temperature follows a state-specific Gaussian distribution. We consider two bivariate specifications, Normal-Normal and Poisson-Normal, and compare models with two- and three-state specifications.

*École d'Actuariat, Université Laval, Québec, Québec, G1V 0A6, Canada. Email: samuel-darwin-dona.kouton.1@ulaval.ca.

†Institute of Statistics, Biostatistics and Actuarial Science (ISBA), Louvain Institute of Data Analysis and Modeling (LIDAM), UCLouvain, Louvain-la-Neuve, Belgium. Corresponding author. Email: karim.barigou@uclouvain.be.

‡École d'Actuariat, Université Laval, Québec, Québec, G1V 0A6, Canada. Email: thai.nguyen@act.ulaval.ca.

This paper contributes to the literature in several ways. We first develop a non-homogeneous seasonal hidden Markov model with trend for the joint modeling of weekly mortality and temperature, allowing for seasonality, smooth long-run movements, and state-dependent regime changes. We then show that the general identifiability and consistency results available for seasonal and time-inhomogeneous hidden Markov models apply to our Poisson and Gaussian emission specifications. On the empirical side, using Italian weekly data, we find that the latent-state structure improves fit relative to non-switching benchmarks and yields a sharper interpretation of the temperature–mortality relationship. In particular, bivariate specifications that incorporate temperature produce more stable and interpretable latent regimes than mortality-only models. Finally, we derive a risk-adjusted change-of-measure framework within the SHMM and illustrate its use for pricing mortality-linked contracts under regime uncertainty. Our numerical results show that including temperature substantially reduces the sensitivity of premiums to transition-risk distortions, while a three-state specification allows the model to retain sensitivity to extreme mortality-climate episodes.

Our work sits at the intersection of three strands of research: environmental drivers of mortality, regime-switching/state-dependent mortality modeling, and actuarial pricing under model uncertainty. The use of aggregate climate indices (ACI) and related composite indicators has been explored in applied work (e.g., Pan et al. 2022) and in epidemiological studies, but their use in actuarial mortality forecasting remains limited. More broadly, the growing availability of high-frequency mortality data has motivated recent efforts to model intra-annual seasonal structure directly. For example, Bégin et al. (2025) introduce a seasonal overlay for age–period–cohort mortality models based on periodic splines and show, using daily mortality data from Quebec, that seasonal patterns are significant, persistent, and informative for the analysis of excess mortality. Recent contributions also incorporate temperature and other environmental covariates into mortality models in different ways: adaptations of Lee–Carter that include exogenous temperature factors improve forecasts over the baseline LC model (Seklecka et al. 2017, Barigou et al. 2025); extreme-value approaches capture tail dependence between deaths and temperature (Li & Tang 2022); and two-step frameworks combine epidemiological exposure–response estimates (e.g., DLNMs) with stochastic mortality models to separate temperature-attributable deaths from underlying demographic trends (Guibert et al. 2025, Robben & Barigou 2025). These approaches demonstrate the value of observable climate information for explanation and short-term attribution, but typically do not provide a unified, state-dependent probabilistic model that is directly usable for actuarial valuation and scenario generation.

Regime-switching models and hidden Markov frameworks have been used in actuarial and demographic contexts to capture abrupt shifts, excess-mortality episodes, and time-varying volatility (Hainaut (2012), Zhou (2019), Robben & Antonio (2024)). More recent work has begun to merge regime concepts with environmental risk: two-state and three-state regime specifications have been proposed to capture excess mortality due to extreme temperature and competing causes such as influenza (Robben & Antonio 2024, Robben et al. 2026). Our SHMM contributes to this literature by providing a fully joint model of temperature and mortality with seasonal, trend, and time-varying transition structure, enabling both (i) principled inference on state-dependent temperature–mortality links and (ii) straightforward generation of prospective scenarios that reflect both climatic variability and latent regime dynamics.

Risk-neutral pricing of life insurance contracts has been extensively investigated by many authors, see e.g. Bauer, Börger & Ruß (2010), Bauer, Bergmann & Kiesel (2010), Biffis et al. (2010), Cairns et al. (2006, 2008). In Hainaut (2012), the cost of regime switches is captured in a two-state Markov chain framework using change of pricing measures. However, a notable gap remains in the actuarial pricing literature: regime-switching or latent-state pricing approaches rarely condition regime dynamics on observed environmental drivers such as temperature, despite the relevance of these drivers for short- and medium-term mortality risk. Pricing studies typically tilt intensities or transition dynamics under an alternative measure, but they generally do so in models where regimes are purely latent and not jointly modeled with measurable covariates. By estimating a SHMM on Italian weekly data and deriving a change-of-measure framework that adjusts transition probabilities in a state-aware manner, our paper aims to bridge this gap and provide a tractable approach to pricing mortality-linked claims that explicitly accounts for temperature-driven risk and regime uncertainty.

The remainder of the paper is organized as follows. Section 2 discusses identifiability and consistency conditions for seasonal Hidden Markov models. Section 3 presents the joint temperature-mortality modeling framework. Section 4 applies the model to Italian data. Section 5 introduces the risk-neutral pricing framework based on SHMMs. Pricing results and sensitivity analyses are reported in Section 6. Section 7 concludes.

2 Identifiability and consistency of Seasonal Hidden Markov Models

In this section we recall the abstract identifiability and strong consistency results for seasonal and time-inhomogeneous hidden Markov models established by Touron (2018, 2019). Our contribution is to verify that the assumptions of these results are satisfied by the Poisson-mixture seasonal HMM introduced for mortality-temperature data. Thus, rather than reproving the general theory, we specialize it to our parametric framework.

2.1 Seasonal Hidden Markov Model (SHMM)

Let $S_t \in \{1, \dots, K\}$ denote the latent state at integer time $t \geq 1$ and let O_t take values in a Polish space $\mathcal{O}$ endowed with its Borel σ-algebra. The latent chain is first-order Markov with T-periodic transition matrices $Q(t) = (Q_{ij}(t))_{i,j=1}^K$ satisfying

$$Q(t+T) = Q(t) \qquad \forall t \geq 1, \qquad \Pr(S_t = j \mid S_{t-1} = i) = Q_{ij}(t).$$

Conditional on $S_t = k$, the observation O_t has emission distribution $\nu_{k,t}$, also T-periodic: $\nu_{k,t+T} = \nu_{k,t}$.

Definition 1 (Seasonal Hidden Markov Model)**.** *A pair of processes $(S_t, O_t)_{t\geq 1}$ is a seasonal hidden Markov model (SHMM) if: (i) (S_t) is Markov with T-periodic transition matrices; (ii) conditional on (S_t), the observations (O_t) are independent with T-periodic emission distributions $\{\nu_{k,t}\}$. Its law is determined by the initial distribution π of S_1, the cycle $\{Q(t)\}_{t=1}^T$, and the cycle of emissions $\{\nu_{k,t} : k = 1, \dots, K;\ t = 1, \dots, T\}$.*

Parametric structure. For the transition mechanism we use a multinomial-logit parametrization with trigonometric polynomials:

$$\log \frac{Q_{ij}(t)}{Q_{iK}(t)} = \sum_{l=0}^{d} \left(a_{ij,l} \cos \frac{2\pi l t}{T} + b_{ij,l} \sin \frac{2\pi l t}{T} \right), \qquad j = 1, \dots, K-1. \tag{1}$$

The corresponding collection of coefficients is denoted by β. Emission parameters are collected in δ, and the full parameter is $\theta = (\pi, \beta, \delta)$, ranging over a parameter set Θ, later assumed compact.

2.2 Identifiability Assumptions

Identifiability ensures uniqueness of model parameters given the distribution of a sufficiently informative block of observations. We state the conditions needed for seasonal (time–inhomogeneous) HMMs.

(A1) Positivity and full rank. Each transition matrix $Q(t)$ has strictly positive entries and full rank (i.e. non-singular).

(A2) Cycle irreducibility. The product $Q(1)Q(2)\cdots Q(T)$ is irreducible (every state can reach every other state within one cycle), implying a unique cycle-stationary distribution π^*.

(A3) Linear independence of emissions. For every $t \in \{1,\dots,T\}$, the family $\{\nu_{k,t}\}_{k=1}^K$ is linearly independent in the vector space of finite signed measures on $(\mathcal{O},\mathcal{B}(\mathcal{O}))$. Equivalently, no emission distribution can be written as a nontrivial linear combination of the others at that time.

Note that Assumptions (A1)-(A2) prevent degeneracies in latent dynamics and guarantee that each state influences long-run mixtures of observations. Linear independence (A3) is the analogue of distinct emission distributions in classical HMM identifiability; seasonal variation does not weaken the requirement but only demands it hold at each phase of the cycle.

Theorem 1 (Identifiability up to time-indexed label permutations)**.** *Let $\theta^* = (\pi^*, Q^*(t), \nu^*_{k,t})$ satisfy (A1)-(A3). If another parameter set $\tilde{\theta} = (\tilde{\pi}, \tilde{Q}(t), \tilde{\nu}_{k,t})$ yields the same joint distribution for $(O_1,\dots,O_{T+2})$, then there exist permutations $\sigma_1,\dots,\sigma_T$ of $\{1,\dots,K\}$ with $\sigma_{T+1} = \sigma_1$ such that*

$$\tilde{\pi}_k = \pi^*_{\sigma_1(k)}, \qquad \tilde{\nu}_{k,t} = \nu^*_{\sigma_t(k),t}, \qquad \tilde{Q}_{kl}(t) = Q^*_{\sigma_t(k),\sigma_{t+1}(l)}(t).$$

Sketch. Factor the joint distribution of the observation block into transition and emission tensors. Apply uniqueness of minimal tensor factorizations under linear independence (cf. Allman et al. (2009)) combined with positivity (A1)–(A2) to recover parameters up to permutation. See Touron (2018, 2019) for details. □

Thus, SHMM parameters are identifiable modulo a consistent relabeling of states across time indices within a cycle.

2.3 Strong Consistency of the Maximum Likelihood Estimator

Let $(O_1,\dots,O_n)$ be data from the SHMM with true parameter θ^*. For initial distribution π, denote the (incomplete-data) likelihood by $L_{n,\pi}(\theta)$ and define

$$\hat{\theta}_{n,\pi} := \arg\max_{\theta\in\Theta} L_{n,\pi}(\theta).$$

For each parameter $\theta \in \Theta$, each state $k \in \{1,\dots,K\}$, and each season index $t \in \{1,\dots,T\}$, assume that the emission distribution $\nu^\theta_{k,t}$ admits a density (or probability mass function) with respect to a fixed dominating measure η on $\mathcal{O}$. We denote this by

$$f^\theta_{k,t}(o) := \frac{d\nu^\theta_{k,t}}{d\eta}(o), \qquad o \in \mathcal{O}.$$

To state the strong consistency result, we impose the following additional regularity assumptions.

(A4) Parameter identifiability. Transition parameter β and the normalized emission parameter δ are identifiable (cf. Theorem 1) up to latent state permutation.

(A5) Uniform positivity. $\exists\,\alpha > 0$ such that $Q_{ij}(t) \geq \alpha$ for all i, j, t and all $\theta \in \Theta$.

(A6) Continuity. For every i, j, t and every $o \in \mathcal{O}$, the maps

$$\theta \mapsto Q_{ij}(t), \qquad \theta \mapsto f^{\theta}_{k,t}(o)$$

are continuous.

(A7) Uniform bounds on emissions. For each (k, t) and all $o \in \mathcal{O}$,

$$0 < \inf_{\theta \in \Theta} f^{\theta}_{k,t}(o) \leq \sup_{\theta \in \Theta} f^{\theta}_{k,t}(o) < \infty.$$

(A8) Log-integrability. Writing

$$c_t(o) := \inf_{\theta \in \Theta} \sum_{k=1}^{K} f^{\theta}_{k,t}(o), \qquad d_t(o) := \sup_{\theta \in \Theta} \sum_{k=1}^{K} f^{\theta}_{k,t}(o),$$

we have $\mathbb{E}_{\nu^*_{k,t}}[-\log c_t(O_t)] < \infty$ and $\mathbb{E}_{\nu^*_{k,t}}[\log d_t(O_t)] < \infty$ for all k, t.

We remark that Assumption (A5) prevents vanishing transitions that could fragment the latent state space. (A7)-(A8) rule out degeneracies in the emission distributions, permitting uniform control of log-likelihood increments and application of ergodic/LLN arguments over seasonal cycles.

Theorem 2 (Strong consistency of the MLE)**.** *Under (A1)-(A8), for any initial distribution* π *there exists a sequence of permutations* $(\sigma^{(n)})_{n \geq 1}$ *of the latent state labels (applied cycle-wise) such that*

$$\sigma^{(n)}\left(\hat{\theta}_{n,\pi}\right) \xrightarrow[n \to \infty]{a.s.} \theta^*.$$

Sketch. Decompose the log-likelihood into additive contributions over cycles; use uniform bounds (A7)–(A8) and positivity (A5) to apply a subadditive ergodic argument controlling deviations. Identifiability (A4) ensures uniqueness of limit points up to label permutations. See Touron (2018, 2019) (pp. 1058-1062) for full details. □

Remark 1 (Permutation resolution)**.** *In practice, one imposes deterministic ordering constraints (e.g., ordering states by average mortality intensity or by seasonal peak timing) to anchor labels and interpret regime-specific parameters consistently across samples and resampling procedures.*

We have stated interpretable assumptions guaranteeing (i) identifiability of seasonal HMM parameters up to label permutations, and (ii) strong consistency of the MLE. These results justify the use of likelihood-based estimation (EM / Baum-Welch) and the interpretation of latent regimes as structurally distinct mortality-temperature environments. In the next section we will verify these assumptions for Poisson emission distributions tailored to weekly mortality counts with exposure factors.

2.4 Poisson Mixture Emission Distributions

We now consider the SHMM with transitions parameterized by (1) and with emission laws given by finite mixtures of Poisson distributions. For each latent state $k \in \{1, \dots, K\}$ and each time t, the conditional law of the weekly count O_t is

$$\nu_{k,t}(y) = \sum_{m=1}^{M} p_{km} \, \mathcal{P}\big(E_t \mu_{km}(t); y\big), \qquad y \in \mathbb{N}, \tag{2}$$

where:

- $E_t > 0$ is the known exposure at time t;
- $p_{km} > 0$, with $\sum_{m=1}^{M} p_{km} = 1$, are the mixture weights in state k;
- $\mathcal{P}(\lambda; \cdot)$ denotes the Poisson law with parameter λ;
- $\mu_{km}(t) > 0$ is the component intensity.

We use the multiplicative decomposition

$$\log \mu_{km}(t) = \lambda_{km} + S_k(t) + T_k(t),$$

where λ_{km} is a baseline level, $S_k(t)$ is a trigonometric polynomial of degree d,

$$S_k(t) = \sum_{l=1}^{d} \left(\gamma_{k,2l-1} \cos \frac{2\pi l t}{T} + \gamma_{k,2l} \sin \frac{2\pi l t}{T} \right),$$

and $T_k(t)$ is a trend term depending on parameters α_k.

In our model, the trend therefore enters through the emission intensity, rather than additively at the observation level. This specification is natural for count data, since it preserves the integer-valued nature of the observations and ensures positivity of the intensities.

The relevant identifiability issue is the uniqueness of the decomposition of the log-intensity into baseline, seasonal, and trend components. Indeed, for any constant c_k,

$$\lambda_{km} \mapsto \lambda_{km} + c_k, \qquad T_k(t) \mapsto T_k(t) - c_k$$

leaves $\mu_{km}(t)$ unchanged. To remove this ambiguity, we impose a normalization on the trend component; for instance, we may require that

$$T_k(t_0) = 0$$

at a fixed reference time t_0. Any equivalent centering condition over one cycle would also be suitable.

Throughout this subsection we assume the following uniform bounds:

$$\lambda_{km} \in [\lambda_{\min}, \lambda_{\max}], \qquad S_k(t) \in [s_{\min}, s_{\max}], \qquad T_k(t) \in [\tau_{\min}, \tau_{\max}],$$

$$p_{km} \in [p_{\min}, 1], \qquad p_{\min} > 0, \qquad E_t \in [E_{\min}, E_{\max}],$$

and we assume that the parameter $\theta = (\beta, p, \Lambda, \Gamma, \alpha)$ ranges over a compact set. We also impose:

(A9) For every k, one has

$$\lambda_{k1} < \lambda_{k2} < \cdots < \lambda_{kM}.$$

(A10) For every k, m, one has $p_{km} > 0$.

Assumptions (A9)-(A10) ensure that each state-dependent emission law is represented as a mixture with exactly M distinct Poisson components.

Lemma 1. *Let $\lambda_1, \ldots, \lambda_n > 0$ be pairwise distinct. Then the Poisson measures*

$$\mathcal{P}(\lambda_1), \ldots, \mathcal{P}(\lambda_n)$$

are linearly independent.

Proof. Reordering if necessary, assume $0 < \lambda_1 < \cdots < \lambda_n$. Suppose

$$\sum_{i=1}^{n} a_i \mathcal{P}(\lambda_i) = 0.$$

Evaluating at $\{y\}$, $y \in \mathbb{N}$, gives $\sum_{i=1}^{n} a_i e^{-\lambda_i} \frac{\lambda_i^y}{y!} = 0$, hence $\sum_{i=1}^{n} a_i e^{-\lambda_i} \lambda_i^y = 0$. Dividing by $e^{-\lambda_n} \lambda_n^y$ and letting $y \to \infty$, we obtain $a_n = 0$. Repeating the argument successively yields $a_{n-1} = \cdots = a_1 = 0$. ☐

We now verify the regularity assumptions required for consistency and discuss the structural assumptions needed for identifiability.

Proposition 3. *For the Poisson-mixture SHMM defined above, Assumptions (A5)-(A8) hold on every compact parameter set satisfying the uniform bounds stated above and (A10).*

Proof. We verify the assumptions one by one.

Verification of (A5). Under the multinomial-logit parametrization (1), each linear predictor is a continuous function of θ on a compact set, hence is uniformly bounded. Therefore each transition probability $Q_{ij}(t)$ is bounded below by a constant $\alpha > 0$, uniformly in i, j, t and $\theta \in \Theta$.

Verification of (A6). For fixed i, j, t, the map $\theta \mapsto Q_{ij}(t)$ is continuous by construction. Likewise, for fixed k, t and $y \in \mathbb{N}$,

$$f_{k,t}^{\theta}(y) := \nu_{k,t}(y) = \sum_{m=1}^{M} p_{km} \exp\big(-E_t \mu_{km}(t)\big) \frac{(E_t \mu_{km}(t))^y}{y!},$$

which is continuous in θ, since it is formed from sums, products, exponentials, and trigonometric polynomials.

Verification of (A7). By the uniform bounds on E_t, λ_{km}, $S_k(t)$, and $T_k(t)$, there exist constants $0 < \underline{\lambda} \leq \overline{\lambda} < \infty$ such that

$$\underline{\lambda} \leq E_t \mu_{km}(t) \leq \overline{\lambda} \qquad \text{for all } k, m, t, \theta.$$

Hence, for each fixed $y \in \mathbb{N}$, all Poisson parameters appearing in (2) lie in the compact interval $[\underline{\lambda}, \overline{\lambda}]$. Since

$$\lambda \longmapsto e^{-\lambda} \frac{\lambda^y}{y!}$$

is continuous on $[\underline{\lambda}, \overline{\lambda}]$, it attains a strictly positive minimum and a finite maximum there. Because the weights satisfy $p_{km} \geq p_{\min} > 0$, it follows that

$$0 < \inf_{\theta \in \Theta} f_{k,t}^{\theta}(y) \leq \sup_{\theta \in \Theta} f_{k,t}^{\theta}(y) < \infty.$$

Thus (A7) holds.

Verification of (A8). Define

$$c_t(y) := \inf_{\theta\in\Theta} \sum_{k=1}^{K} f^{\theta}_{k,t}(y), \qquad d_t(y) := \sup_{\theta\in\Theta} \sum_{k=1}^{K} f^{\theta}_{k,t}(y).$$

By the bounds obtained in the proof of (A7), $c_t(y)$ and $d_t(y)$ are positive and finite for every $y \in \mathbb{N}$. Moreover, they are bounded above and below by positive multiples of Poisson probabilities with parameters in $[\underline{\lambda}, \overline{\lambda}]$. Consequently, $\log d_t(y)$ and $-\log c_t(y)$ grow at most linearly in y and in $\log(y!)$. Since under the true model O_t has a finite Poisson-mixture distribution, these quantities are integrable under each $\nu^*_{k,t}$. Therefore (A8) holds. □

The next proposition isolates the nondegeneracy conditions that are not automatic consequences of the Poisson-mixture parametrization and must therefore be imposed separately.

Proposition 4. *Assume that, for every $t \in \{1, \ldots, T\}$,*

1. *the transition matrix $Q(t)$ has full rank;*
2. *the family $\{\nu_{k,t}\}_{k=1}^{K}$ is linearly independent.*

Then Assumptions $(A1)$-$(A4)$ hold.

Proof. Since the transition probabilities are obtained from a multinomial-logit parametrization, all entries of $Q(t)$ are strictly positive. By assumption, each $Q(t)$ also has full rank, so (A1) holds. The strict positivity of every $Q(t)$ implies that the product $Q(1)\cdots Q(T)$ is itself strictly positive; in particular, it is irreducible, and (A2) follows.

Assumption (A3) is exactly the second hypothesis of the proposition. Finally, once $(A1)$-$(A3)$ hold, identifiability of the SHMM up to permutation of latent states follows from Theorem 1, which gives (A4). □

The previous proposition isolates the assumptions that must be verified, or imposed, to apply the abstract identifiability theorem. In the present Poisson-mixture setting, linear independence of emissions is closely related to the linear independence of Poisson laws with distinct parameters.

Indeed, for each fixed t, every emission law $\nu_{k,t}$ is a finite linear combination of Poisson laws with parameters among the set

$$F_t := \left\{ E_t \exp\left(\lambda_{km} + S_k(t) + T_k(t)\right) : 1 \le k \le K,\ 1 \le m \le M \right\}.$$

Writing the distinct elements of F_t as $\tilde{\lambda}_1(t), \ldots, \tilde{\lambda}_{p(t)}(t)$, where $p(t) \le KM$, we may rewrite

$$\nu_{k,t}(y) = \sum_{j=1}^{p(t)} B_{kj}(t)\, \mathcal{P}(\tilde{\lambda}_j(t); y),$$

for a suitable coefficient matrix $B(t)$. By Lemma 1, the family $\{\mathcal{P}(\tilde{\lambda}_j(t))\}_{j=1}^{p(t)}$ is linearly independent as soon as the $\tilde{\lambda}_j(t)$ are pairwise distinct. Therefore, whenever $p(t) \ge K$ and $B(t)$ has rank K, the emission family $\{\nu_{k,t}\}_{k=1}^{K}$ is linearly independent. In applications, this condition can either be imposed directly, or regarded as a generic nondegeneracy condition on the emission parameters.

Combining the two previous propositions yields the following consequence.

Corollary 1. *Assume the compactness and boundedness conditions above, together with (A9)-(A10). If, in addition, for every $t \in \{1,\dots,T\}$, the matrix $Q(t)$ has full rank and the family $\{\nu_{k,t}\}_{k=1}^{K}$ is linearly independent, then the assumptions of Theorem 2 are satisfied. Consequently, the maximum likelihood estimator is strongly consistent up to permutation of the latent states.*

The results recalled from Touron (2018, 2019), together with the preceding verifications, provide the theoretical justification for likelihood-based inference in our seasonal Poisson-mixture hidden Markov model. In particular, once the nondegeneracy conditions ensuring identifiability are satisfied, the model is identifiable up to permutation of latent states and the maximum likelihood estimator is strongly consistent. We note that the case of the seasonal Gaussian-mixture hidden Markov model was already studied in Touron (2018).

3 Mortality and temperature with a seasonal hidden Markov model

We model joint weekly dynamics of mortality and temperature with a seasonal hidden Markov model (SHMM), in which both transition probabilities and state-dependent emission means may vary over time. Let $S_t \in \{1,\dots,K\}$ denote the latent regime at week $t = 1,\dots,T_{\text{obs}}$. The observable vector is

$$O_t := \big(O_t^{(1)}, O_t^{(2)}\big) \in \mathcal{O} := \mathbb{R}_+ \times \mathbb{R},$$

where $O_t^{(1)}$ is the mortality measure for age group x in region r at week t (either log-force of mortality or death counts), and $O_t^{(2)}$ is the corresponding weekly temperature for the same region and week. Denote by $E_t > 0$ the population exposure (adjusted for week length) when $O_t^{(1)}$ is a count.

3.1 Seasonal, time-varying transition probabilities

The latent process (S_t) is a first-order, time-inhomogeneous Markov chain with T-periodic transition probabilities, where $T = 52.18$ reflects the average number of weeks per year:

$$\Pr(S_t = j \mid S_{t-1} = i) = Q_{ij}(t), \qquad Q(t+T) = Q(t), \qquad 1 \le i,j \le K, \tag{3}$$

$$Q_{ij}(t) = \frac{\exp\big(P_{ij}(t)\big)}{1 + \sum_{\ell=2}^{K} \exp\big(P_{i\ell}(t)\big)}, \qquad j \ge 2, \tag{4}$$

$$Q_{i1}(t) = \frac{1}{1 + \sum_{\ell=2}^{K} \exp\big(P_{i\ell}(t)\big)}. \tag{5}$$

Here $P_{ij}(t)$ is a trigonometric polynomial of known degree $d \ge 1$,

$$P_{ij}(t) \;=\; \beta_{ij,0} \;+\; \sum_{l=1}^{d} \Big\{ \beta_{ij,(2l-1)} \cos\Big(\tfrac{2\pi l}{T}\, t\Big) \;+\; \beta_{ij,2l} \sin\Big(\tfrac{2\pi l}{T}\, t\Big) \Big\}, \tag{6}$$

so that regime persistence and switching probabilities vary smoothly and periodically over the calendar year. The parameter collection $\beta := \{\beta_{ij,\cdot}\}_{i,j}$ governs seasonal transitions.

3.2 State-dependent emission distributions

Conditional on the state $S_t = k$, the distribution of $O_t = (O_t^{(1)}, O_t^{(2)})$ is

$$\nu_{k,t} = \nu_{k,t}^{(1)} \otimes \nu_{k,t}^{(2)}. \tag{7}$$

Conditional on $S_t = k$, we assume the mortality and temperature components to be independent. Dependence between the two margins is therefore captured through the latent regime process.

Mortality component. Let $\mu_k(t)$ denote the state-specific mortality intensity, with log-scale decomposition

$$\log \mu_k(t) = \lambda_k + T_k^{(1)}(t) + S_k^{(1)}(t), \tag{8}$$

where λ_k is a baseline level, $T_k^{(1)}(t) = \alpha_k^{(1)} a(t)$ is a linear trend term, and $S_k^{(1)}(t)$ is a trigonometric polynomial of degree d:

$$S_k^{(1)}(t) = \sum_{l=1}^{d} \left\{ \gamma_{k,2l-1} \cos\left(\tfrac{2\pi l}{T} t\right) + \gamma_{k,2l} \sin\left(\tfrac{2\pi l}{T} t\right) \right\}.$$

We consider two observation models:

$$\text{(i) Log-force model:} \quad \nu_{k,t}^{(1)} = \mathcal{N}\left(\log \mu_k(t),\ \zeta_k^2\right), \tag{9}$$

$$\text{(ii) Count model:} \quad \nu_{k,t}^{(1)} = \text{Poisson}\big(E_t \mu_k(t)\big). \tag{10}$$

In the count specification, E_t scales the intensity according to exposure; in the log-force specification, ζ_k^2 is a state-specific variance.

Temperature component. Following Touron et al. (2018), we model the temperature margin by a Gaussian emission with state-specific trend and seasonality:

$$\nu_{k,t}^{(2)} \ = \ \mathcal{N}\Big(m_k + T_k^{(2)}(t) + S_k^{(2)}(t),\ \sigma_k^2 \Big), \tag{11}$$

where

- m_k is a state-specific location,
- $T_k^{(2)}(t) = \alpha_k^{(2)}\, a(t)$ is a linear trend in temperature,
- $S_k^{(2)}(t)$ is a trigonometric polynomial of degree d,

$$S_k^{(2)}(t) \ = \ \sum_{l=1}^{d} \Big\{ \delta_{k,(2l-1)} \cos\Big(\tfrac{2\pi l}{T}\, t\Big) \ + \ \delta_{k,2l} \sin\Big(\tfrac{2\pi l}{T}\, t\Big) \Big\}.$$

We remark that the product structure maintains interpretability in the sense that conditional on $S_t = k$, mortality and temperature are independent. State dependence induces dependence in the joint observations. Also, weekly periodicity is set to $T = 52.18$; in practice T can be fixed at 52 or 53 depending on the calendar convention, with negligible impact on inference. Finally, additional covariates (e.g., humidity, influenza proxies, holidays) may enter additively in (8) and/or (11) without altering the SHMM structure.

3.3 Parameterization and specifications used empirically

Let θ collect all parameters:

$$\theta = \Big(\beta,\ \{\lambda_k, \zeta_k\}_k,\ \{m_k, \sigma_k\}_k,\ \{\alpha_k^{(1)}, \alpha_k^{(2)}\}_k,\ \{\gamma_{k,\cdot}, \delta_{k,\cdot}\}_k\Big).$$

Two empirically tractable specifications are emphasized:

- Gaussian-Gaussian (log-force, temperature): $\nu_{k,t}^{(1)} = \mathcal{N}(\log \mu_k(t), \zeta_k^2)$ and $\nu_{k,t}^{(2)}$ as in (11).
- Poisson-Gaussian (counts, temperature): $\nu_{k,t}^{(1)} = \text{Poisson}(E_t\, \mu_k(t))$ and $\nu_{k,t}^{(2)}$ as in (11).

Both variants inherit the same seasonal and trend structure, enabling a unified comparison across $K \in \{2, 3\}$ latent regimes.

We remark that regimes capture qualitatively distinct mortality-temperature environments (e.g., temperate vs. stress seasons), with state-specific seasonal profiles and trends in each margin. This specification provides a clear actuarial interpretation and tractable likelihood-based estimation.

4 Empirical application: weekly mortality and temperature

We apply seasonal hidden Markov models to a bivariate weekly series of mortality and temperature. Our goals are threefold: identify latent climate-mortality regimes, improve characterization of seasonal mortality patterns, and compare two emission specifications-Gaussian for log mortality rates versus Poisson for death counts. The analysis focuses on Rome (NUTS 3 region ITI43) for individuals aged 65 and older, using data spanning 2013 to 2019.

4.1 Data

Weekly death counts were extracted from Eurostat and aggregated by sex and 5-year age groups at the NUTS 3 level across Europe. In the present application, we focus on Rome, corresponding to code ITI43, and restrict attention to individuals aged 65 and older. Weekly mortality rates are computed from weekly death counts and estimated weekly exposures.

Let $d_{x,y,w}^{(r)}$ denote the number of deaths in region r, age group x, calendar year y, and week w. Let $P_{x,y,0}^{(r)}$ denote the population exposure at the beginning of year y. We approximate weekly exposure by

$$E_{x,y,w}^{(r)} = \frac{P_{x,y,0}^{(r)} + P_{x,y+1,0}^{(r)}}{2 \times 52.18}, \qquad m_{x,y,w}^{(r)} = \frac{d_{x,y,w}^{(r)}}{E_{x,y,w}^{(r)}}, \tag{12}$$

where 52.18 is the average number of weeks per year. The restriction to ages 65+ is motivated by the greater stability of mortality signals in the multivariate setting.

Temperature data were obtained from the Copernicus Climate Data Store, using the E-OBS daily gridded dataset at 0.10° resolution (approximately 11 km). For each week w, let T_d denote the daily temperature on day d, and define the weekly minimum and maximum temperatures by

$$T_w^{\min} := \min_{d \in w} T_d, \qquad T_w^{\max} := \max_{d \in w} T_d.$$

To capture the thermal extremes most relevant for mortality, we construct a weekly extreme-temperature index T_w^{ext} defined by

$$T_w^{\text{ext}} = \begin{cases} T_w^{\min}, & \text{if } T_w^{\min} < Q_{0.5}\big(\{T_{w'}^{\min}\}_{w'}\big), \\ T_w^{\max}, & \text{otherwise,} \end{cases}$$

where $Q_{0.5}\big(\{T^{\min}_{w'}\}_{w'}\big)$ denotes the empirical median of the weekly minimum temperatures over the observation period. The threshold at the 50th percentile provides a simple rule for distinguishing cold-dominated from heat-dominated weeks.

This construction is motivated by the well-documented nonlinear relationship between temperature and mortality, with both cold and heat extremes exerting disproportionate effects on health outcomes (Analitis et al. 2008, Robben et al. 2026). It is also consistent with the logic of public-health warning systems, which are typically based on extreme-temperature thresholds rather than weekly average conditions (Pascal et al. 2006).

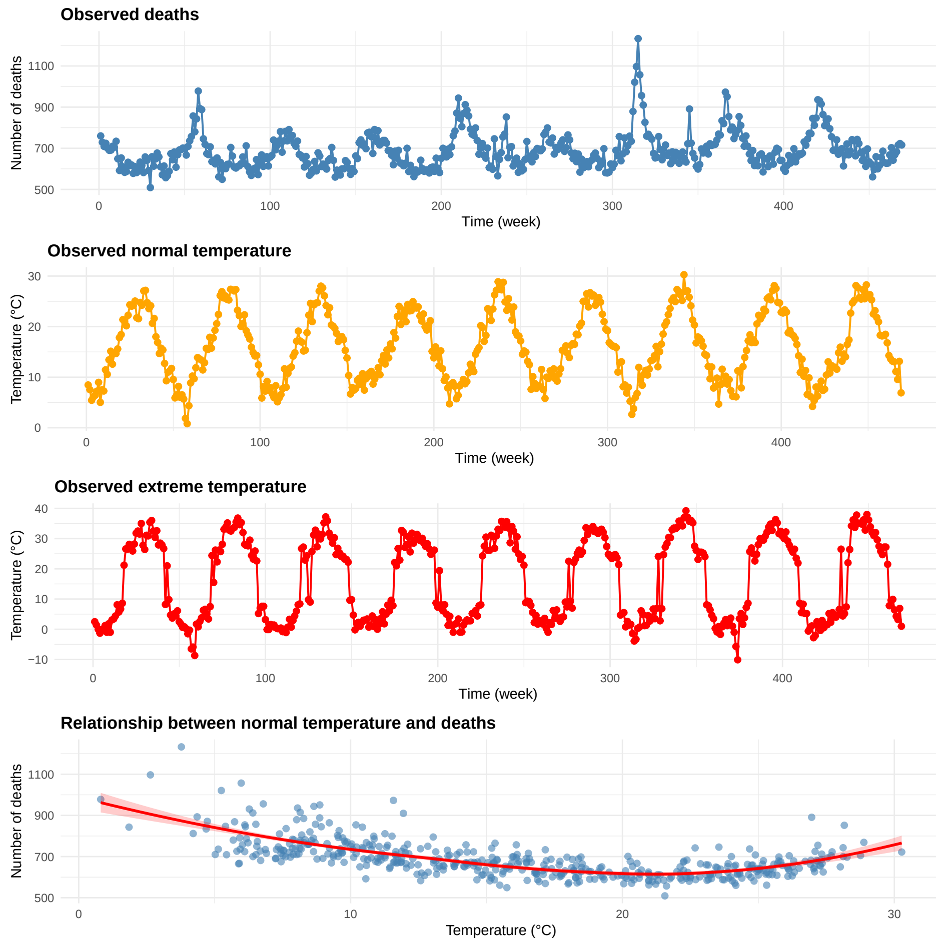

Figure 1: Joint time series of mortality and extreme temperature for ages 65+.

Visual inspection of the data (Figure 1) suggests that periods of elevated mortality tend to coincide with very low temperatures, occasionally exhibiting a lag of one to two weeks.

4.2 Model specifications

We estimate two bivariate SHMM specifications, each with $K \in \{2, 3\}$ latent regimes:

- **Normal-Normal**: the mortality margin is modeled by a Gaussian distribution for the log mortality rate, while the temperature margin is Gaussian; both margins include state-specific seasonal components (one harmonic) and linear trend terms.

- **Poisson-Normal**: the mortality margin is modeled by a Poisson($E_t \mu_k(t)$) distribution for weekly death counts, while the temperature margin is Gaussian with analogous state-specific seasonal and trend components.

In both specifications, transition probabilities follow the multinomial-logit formulation introduced in Section 3, with one seasonal harmonic (cosine and sine terms) for each origin state. All models are estimated by maximum likelihood using 150 random initializations per specification, in order to reduce sensitivity to local optima. Estimation is carried out with the R package `hmmTMB` (Michelot 2025).

4.3 Key fit metrics

Table 1 summarizes the main fit and diagnostic measures for the four fitted specifications. Since the Normal-Normal model is fitted to log mortality rates whereas the Poisson-Normal model is fitted to death counts, likelihood-based criteria such as the log-likelihood, AIC, and BIC should be interpreted primarily *within* each specification family rather than used for direct comparison across the two mortality scales. Cross-specification comparisons are therefore based mainly on predictive accuracy and residual diagnostics.

Table 1: Fit and diagnostic measures across model specifications.

Model	States	LogLik	AIC	BIC	MAPE	Pseudo-R^2	KS p-value
Normal-Normal	2	457.54	−861.08	−749.01	6.97	0.32	0.225
Normal-Normal	3	599.39	−1098.77	−891.24	7.12	0.33	0.880
Poisson-Normal	2	−3912.36	7874.71	7978.48	9.39	0.09	0.373
Poisson-Normal	3	−3653.44	7400.87	7595.95	7.39	0.34	0.147

Several patterns emerge. Within each specification family, moving from two to three latent states improves model fit substantially. In the Normal-Normal family, the three-state model yields markedly lower AIC and BIC values, a higher log-likelihood, and a KS p-value of 0.880 indicating adequate residual fit, against 0.225 for the two-state specification. The pseudo-R^2 improves modestly (0.32 to 0.33), while the MAPE increases only marginally (6.97 to 7.12). The Poisson-Normal family shows the same qualitative pattern, with the three-state model improving across all criteria, though its overall predictive performance remains weaker than that of the preferred Gaussian specification. Taken together, these results favour a three-state structure in both model families. The three-state Normal-Normal specification provides the most satisfactory overall fit and is therefore retained for the main interpretation.

4.4 Normal-Normal: Two-State Model

We begin by examining the two-state Normal-Normal specification to establish baseline regime structure.

Parameter	State 1	State 2
Mortality (log scale)		
Intercept (λ)	0.0654	0.0665
Sine (γ_1)	0.119	0.008
Cosine (γ_2)	0.191	-0.009
Trend ($\alpha^{(1)}$)	0.0002	-0.0029
Std dev (ζ)	0.081	0.070
Temperature		
Intercept (m)	7.85°C	21.14°C
Sine (δ_1)	-4.184	-5.466
Cosine (δ_2)	-6.366	-9.292
Trend ($\alpha^{(2)}$)	0.059	0.283
Std dev (σ)	2.374	2.492

The emission parameters reveal a clear thermal distinction between states. State 1 corresponds to lower baseline temperatures ($m = 7.85$°C) with higher baseline mortality ($\lambda = 0.065$) and pronounced seasonal amplitude (sine = 0.119, cosine = 0.191), while State 2 captures warmer conditions ($m = 21.14$°C) with slightly elevated baseline mortality ($\lambda = 0.067$) but markedly reduced seasonal variation. Both states exhibit similar thermal variability ($\sigma = 2.37$ vs. 2.49), suggesting comparable dispersion across the temperature range. Mortality trends remain near zero in State 1 ($\alpha = 0.0002$) and slightly negative in State 2 ($\alpha = -0.003$), indicating that once seasonal and thermal effects are accounted for, no meaningful linear drift is present over the observation window.

Figure 2 displays the data colored by the Viterbi-identified state sequence. The two regimes alternate in a clear seasonal pattern, with State 1 dominating winter months and State 2 prevailing through summer. The symmetric transition structure confirms that neither state is markedly more persistent than the other at the annual scale, with regime switches driven almost entirely by the seasonal harmonics in the transition probabilities.

4.5 Normal-Normal: Three-State Model

Building on the two-state results, we now introduce a third latent regime to assess whether additional structure improves model fit and interpretability. The transition matrix now exhibits richer seasonal dynamics.

The three-state model cleanly stratifies the climate-mortality relationship into three thermally distinct regimes. State 1 corresponds to the cool winter regime (baseline temperature ≈ 7.90°C), with a moderate mortality baseline ($\lambda = 0.0681$) and pronounced seasonal harmonics (sine = 0.081, cosine = 0.129), reflecting the typical winter excess mortality pattern. State 3 represents the warm summer regime (≈ 21.09°C), with the lowest baseline mortality ($\lambda = 0.0663$) and near-flat seasonal variation. State 2 stands out as the most mortality-intensive regime, exhibiting both the highest baseline mortality ($\lambda = 0.0804$) and the largest mortality variance ($\zeta = 0.0895$). Unlike the other two states, it is not strongly tied to a specific season: its weak seasonal harmonics (sine ≈ 0.003) indicate that the elevated mortality episodes it captures occur across a broad range of climatic conditions, in both winter and summer, albeit more frequently in winter. The high thermal variability ($\sigma = 3.241$) is consistent with this interpretation, reflecting a regime defined less by a characteristic temperature than by the presence of mortality excesses regardless of season.

Temperature trends are positive across all states, with the most pronounced warming signal

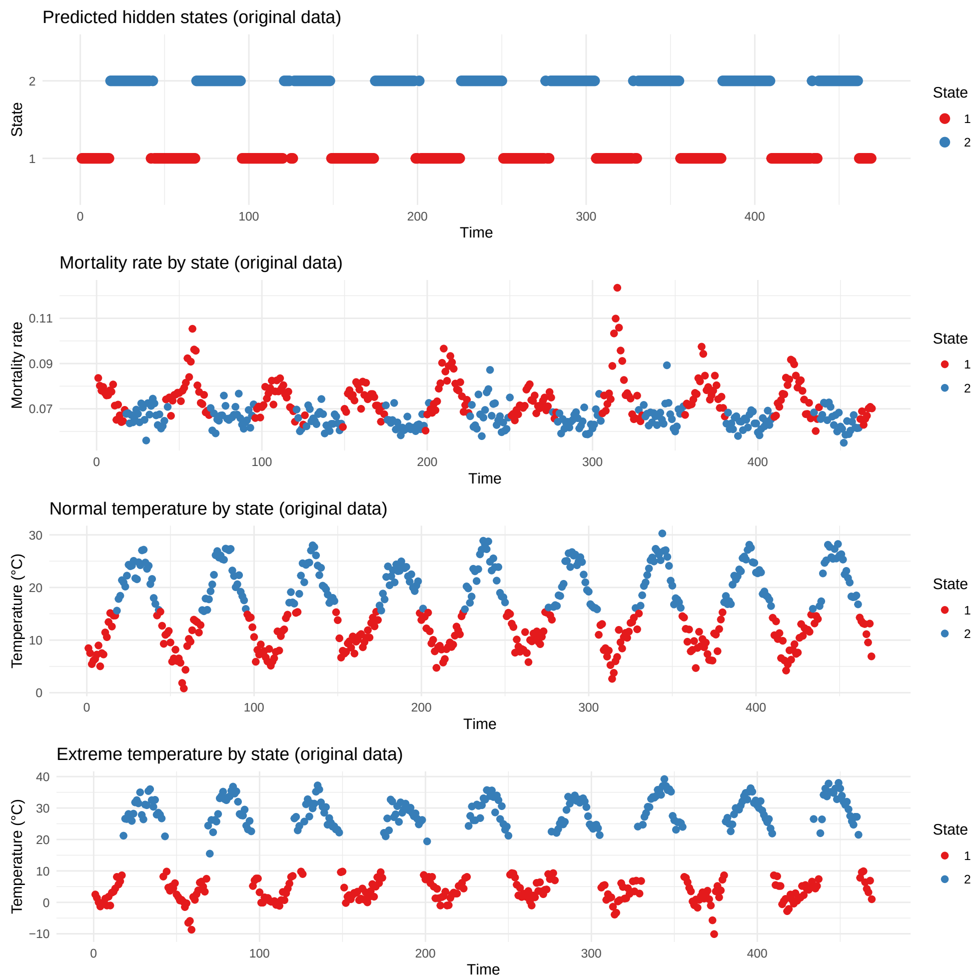

Figure 2: Weekly observations with points colored by inferred states for the 2-state Normal-Normal model.

Table 2: Emission parameters for the 3-state Normal-Normal model.

Parameter	State 1	State 2	State 3
Mortality (log scale)			
Intercept (λ)	0.0681	0.0804	0.0663
Sine (γ_1)	0.0810	0.0029	0.0075
Cosine (γ_2)	0.1288	0.2024	-0.0091
Trend ($\alpha^{(1)}$)	-0.00485	-0.01010	-0.00286
Std dev (ζ)	0.0496	0.0895	0.0695
Temperature			
Intercept (m)	7.90°C	4.18°C	21.09°C
Sine (δ_1)	-4.007	-3.675	-5.498
Cosine (δ_2)	-6.508	-4.817	-9.315
Trend ($\alpha^{(2)}$)	0.136	0.238	0.287
Std dev (σ)	1.805	3.241	2.500

in State 2 ($\alpha^{(2)} = 0.376$), indicating a systematic shift in transitional-season temperatures over the observation period. In contrast, mortality trends remain negligible in all three states ($|\alpha^{(1)}| < 0.0003$), confirming that once seasonal and climatic effects are accounted for, no residual linear drift is detectable over this seven-year window.

Figure 3 illustrates the identified state sequence. State 2 is the most persistent regime and accounts for the bulk of the observation period, consistent with its role as the transitional baseline. States 1 and 3 emerge in clear seasonal alternation, capturing cold winter episodes and warm summer periods respectively. The temporal segmentation aligns closely with the observed temperature series, confirming that the latent structure is strongly anchored to climatic conditions and mortality level rather than driven by latent dynamics alone.

4.6 Poisson-Normal: Three-State Model

For comparison, we present the three-state Poisson-Normal specification, which directly models death counts rather than log rates.

Emission parameters (selected)

Parameter	State 1	State 2	State 3
Mortality (Poisson, log-rate)			
Intercept (λ)	6.409	6.612	6.417
Trend ($\alpha^{(1)}$)	0.0122	0.0136	0.0112
Temperature			
Intercept (m)	8.00°C	15.43°C	21.53°C
Std dev (σ)	2.217	2.185	2.240

The Poisson specification yields a thermally ordered three-regime structure. State 1 corresponds to the coolest conditions (≈ 8.0°C), State 2 to an intermediate thermal regime (≈ 15.4°C), and State 3 to the warmest conditions (≈ 21.5°C). State 2 exhibits the highest mortality rate on the log scale ($\lambda = 6.612$), suggesting that intermediate thermal conditions are associated with elevated

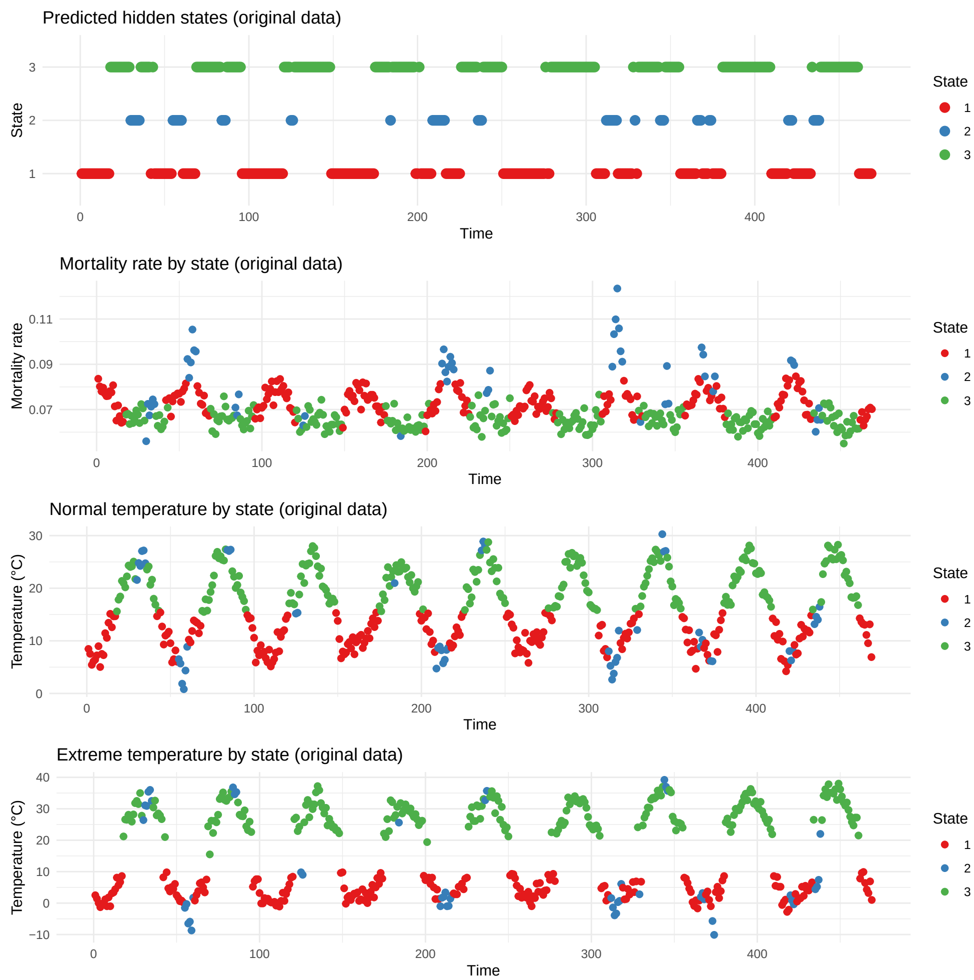

Figure 3: Real data colored by identified states for the 3-state Normal-Normal model.

mortality counts, possibly capturing transitional-season excess mortality similarly to the Normal-Normal specification. Thermal variability is comparable across all three states ($\sigma \approx 2.2$), offering less regime differentiation on this dimension than the Normal-Normal model. Mortality trends are positive and small in all states, with the largest increase in State 2 ($\alpha^{(1)} = 0.014$).

Figure 4 shows that while the Poisson-Normal model successfully identifies distinct thermal regimes, the state labels differ from the Normal-Normal ordering, and the boundaries between regimes appear less crisp-a consequence of the higher variance inherent in count-based modeling.

Summary. Across both emission specifications, the empirical results support a three-state representation of joint mortality-temperature dynamics, with regimes interpretable as cool, intermediate, and warm conditions. The three-state Normal-Normal model provides the clearest state separation and the best overall diagnostic performance among the fitted models. It is therefore retained as the reference specification for the actuarial pricing analysis developed in the next section.

5 Risk-Adjusted Pricing under the SHMM Framework

The seasonal hidden Markov model naturally leads to a *partial-information* pricing setting. Let

$$(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$$

be a filtered probability space, where the full filtration

$$\mathcal{F}_t = \sigma(S_s, O_s^{(1)}, O_s^{(2)},\, s \leq t)$$

contains both the latent regime process (S_t) and the observed mortality and temperature variables. In practice, however, only the observable processes are available, so pricing must be based on the smaller filtration

$$\mathcal{G}_t = \sigma(O_s^{(1)}, O_s^{(2)},\, s \leq t) \subset \mathcal{F}_t.$$

Because the latent regime is unobserved and mortality risk is not fully hedgeable, the market is incomplete. Consequently, no unique valuation measure is singled out by no-arbitrage arguments alone. We therefore work with a risk-adjusted probability measure $\widetilde{\mathbb{P}}$, assumed equivalent to $\mathbb{P}$, under which actuarial cash flows are valued. This perspective is consistent with the literature on risk-adjusted valuation of life-insurance liabilities and mortality-linked contracts; see, among others, Bauer, Börger & Ruß (2010), Bauer, Bergmann & Kiesel (2010), Biffis et al. (2010), Cairns et al. (2006, 2008), Hainaut (2012). In particular, our approach is related in spirit to the state-dependent change-of-measure constructions used in regime-switching mortality and catastrophe pricing (Hainaut 2012). In the present paper, however, we adopt a parsimonious specification in which the valuation adjustment acts only on the latent transition kernel.

For a payoff X_τ payable at time τ, the time-t value is given by

$$V_t = \mathbb{E}^{\widetilde{\mathbb{P}}}[B(t,\tau)\, X_\tau \mid \mathcal{G}_t]\,. \tag{13}$$

Here $B(t, \tau)$ denotes the discount factor from τ to t. In the numerical application below, we assume a constant annual risk-free rate r, so that $B(t, \tau) = (1 + r)^{-(\tau-t)}$.

In this paper, we adopt a parsimonious risk-adjustment scheme that modifies only the latent transition dynamics, while leaving the state-conditional emission distributions unchanged. This choice isolates the effect of regime-persistence risk on valuation and is particularly natural in an incomplete-market setting where the latent climate–mortality regimes are the main source of non-hedgeable structural uncertainty.

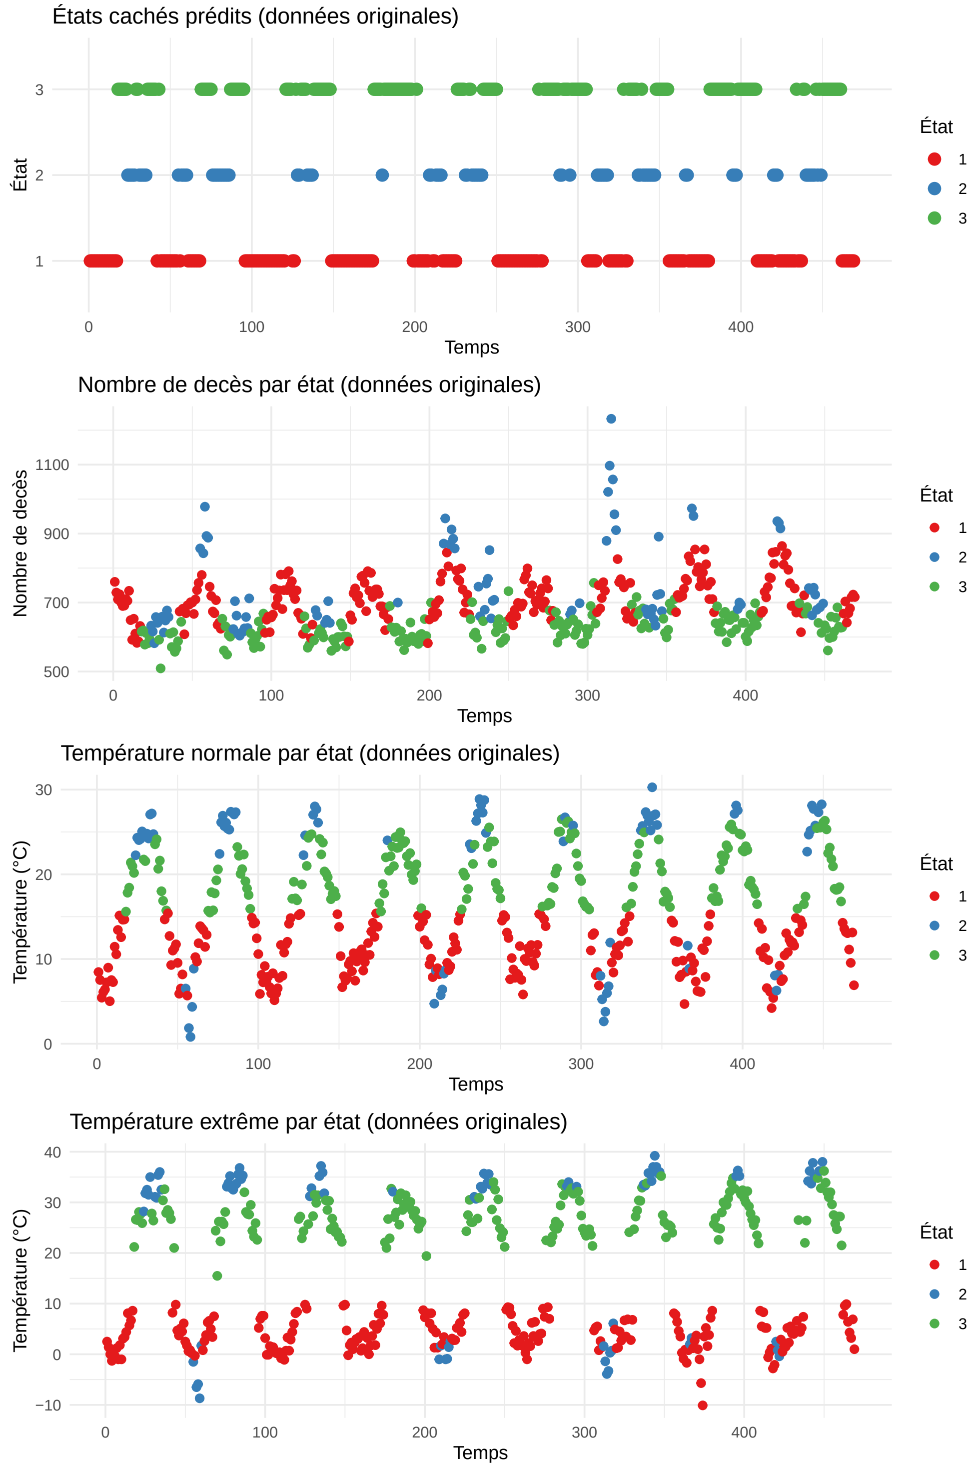

Figure 4: Real data colored by identified states for the 3-state Poisson-Normal model.

Under the historical measure $\mathbb{P}$, the latent chain evolves according to the estimated seasonal transition kernel

$$Q(t) = \big(Q_{ij}(t)\big)_{i,j=1}^{K}, \qquad Q_{ij}(t) = \mathbb{P}(S_{t+1} = j \mid S_t = i).$$

Under the pricing measure $\widetilde{\mathbb{P}}$, we replace $Q(t)$ by a tilted transition kernel

$$\widetilde{Q}_{ij}(t) = \frac{Q_{ij}(t)\,\eta_{ij}(t)}{\sum_{m=1}^{K} Q_{im}(t)\,\eta_{im}(t)}, \qquad \eta_{ij}(t) > 0, \tag{14}$$

where $\eta_{ij}(t)$ is a transition-risk adjustment factor. This transformation preserves the Markov structure while reweighting the probabilities of persistence and switching between regimes.

For ex ante valuation at policy inception, let $\widetilde{\boldsymbol{\alpha}}_t = (\widetilde{\alpha}_t^{(1)}, \ldots, \widetilde{\alpha}_t^{(K)})$ denote the forecast state-probability vector under $\widetilde{\mathbb{P}}$, with initial distribution $\widetilde{\boldsymbol{\alpha}}_0$. These probabilities evolve according to

$$\widetilde{\boldsymbol{\alpha}}_{t+1} = \widetilde{\boldsymbol{\alpha}}_t\, \widetilde{Q}(t). \tag{15}$$

The corresponding risk-adjusted mortality intensity is then defined by

$$\widetilde{\mu}_t = \sum_{k=1}^{K} \widetilde{\alpha}_t^{(k)}\, \mu_k(t), \tag{16}$$

where $\mu_k(t)$ denotes the state-specific mortality intensity estimated under the SHMM.

For annual actuarial pricing, the weekly model is aggregated into one-year death probabilities. Let $\widetilde{\mu}_{t,w}$ denote the weekly risk-adjusted mortality intensity for policy year t and week w. We define the corresponding one-year death probability at attained age $x+t$ by

$$\widetilde{q}_{x+t} = 1 - \exp\left(-\sum_{w=1}^{52} \widetilde{\mu}_{t,w}\,\Delta_w\right),$$

where Δ_w denotes the week-length fraction of a year. The risk-adjusted survival probabilities are then

$${}_t\widetilde{p}_x = \prod_{s=0}^{t-1} (1 - \widetilde{q}_{x+s}).$$

As an illustration, consider a level term insurance contract issued at age x, with maturity n years, benefit C, and constant annual interest rate r. Its single premium under $\widetilde{\mathbb{P}}$ is

$$\Pi^{\widetilde{\mathbb{P}}} = C \sum_{t=0}^{n-1} (1+r)^{-(t+1)}\, {}_t\widetilde{p}_x\, \widetilde{q}_{x+t}. \tag{17}$$

This formulation highlights the role of transition-risk adjustment in pricing. Since the state-conditional mortality dynamics are left unchanged, differences between historical and risk-adjusted premiums arise entirely from the modified persistence and switching behavior of the latent climate–mortality regimes. In the next section, we investigate the numerical sensitivity of premiums to such transition distortions under several SHMM specifications.

6 Numerical Results and Sensitivity Analysis

In this section, we present numerical results for the valuation of mortality-linked contracts under the Seasonal Hidden Markov Model. In particular, we adopt a partial change of measure that tilts only the latent transition kernel, while leaving state-conditional mortality and temperature emissions unchanged.

6.1 Baseline Premiums

Let $S_u \in \{1, \ldots, K\}$ denote the latent regime at week u. Under the historical measure $\mathbb{P}$, transitions follow the estimated seasonal kernel $Q(u)$, and each regime is associated with a state-specific weekly mortality intensity $\mu_k(u)$. Let $\boldsymbol{\alpha}_u = (\alpha_u^{(1)}, \ldots, \alpha_u^{(K)})$ denote the forecast state-probability vector under $\mathbb{P}$, evolving according to

$$\boldsymbol{\alpha}_{u+1} = \boldsymbol{\alpha}_u Q(u), \qquad \boldsymbol{\alpha}_0 \text{ given.}$$

The corresponding marginal weekly mortality intensity is

$$\mu_u = \sum_{k=1}^{K} \alpha_u^{(k)} \, \mu_k(u).$$

Annual death probabilities q_{x+t} are then obtained by aggregating the weekly intensities over policy year t.

Recall that for a level term death benefit C issued at age x, with term n years and flat annual risk-free rate r, the single premium under $\mathbb{P}$ is given by

$$\Pi = C \sum_{t=0}^{n-1} v^{t+1} \, {}_t p_x \, q_{x+t}, \qquad v = (1+r)^{-1},$$

where q_{x+t} is the one-year death probability obtained by aggregating weekly intensities over policy year t, and the survival probabilities satisfy

$${}_{t+1}p_x = {}_t p_x (1 - q_{x+t}), \qquad {}_0 p_x = 1.$$

Below, we consider three calibrated specifications under the historical measure $\mathbb{P}$ (i.e., without transition tilting), estimated for Rome (ITI43), ages 65+, over the period 2013-2019:

- Univariate 2-state SHMM (mortality only)
- Bivariate 2-state SHMM (mortality + temperature)
- Bivariate 3-state SHMM (mortality + temperature with an extreme regime)

The reference contract is a term life insurance with issue age $x = 65$, maturity $T = 40$ years, benefit $C = 100{,}000$, and constant interest rate $r = 2.25\%$.

Table 3: Baseline single premiums under $\mathbb{P}$ (no transition tilt)

Model	States	Premium (\$100,000)	Premium per \$1
Univariate (mortality only)	2	74,259.45	0.7426
Bivariate (mortality + temp)	2	63,992.22	0.6399
Bivariate (mortality + temp)	3	65,119.71	0.6512

Table 3 shows that introducing temperature is associated with a substantial reduction in the baseline premium (-13.8% relative to the univariate model). This pattern is consistent with improved state identification when temperature is included as an informative observable, thereby reducing regime ambiguity that would otherwise be absorbed by the latent dynamics alone. Moreover, moving from two to three states increases the premium moderately (+1.8% relative to the 2-state bivariate model). This increase is consistent with the emergence of an additional regime associated with elevated mortality and greater variability. Relative to the 2-state specification, the 3-state model appears better able to represent tail mortality episodes associated with unusually adverse climatic conditions, resulting in a higher aggregate risk assessment.

6.2 Premium Sensitivity under Transition Risk

We now analyze the sensitivity of premiums to transition dynamics by moving from $\mathbb{P}$ to $\widetilde{\mathbb{P}}$ through distortions of the latent transition matrix. Recall that under the valuation measure $\widetilde{\mathbb{P}}$, only the transition kernel is modified:

$$\widetilde{Q}_{ij}(t) = \frac{Q_{ij}(t)\,\eta_{ij}(t)}{\sum_{m=1}^{K} Q_{im}(t)\,\eta_{im}(t)}, \qquad \eta_{ij}(t) > 0.$$

In the sensitivity analysis, we do not impose an explicit parametric form on the distortion factors $\eta_{ij}(t)$. Instead, we specify target persistence probabilities $\widetilde{Q}_{ii}$ and recover the off-diagonal entries by normalization.

For the univariate two-state model, Table 4 reports premiums across a grid of target persistence values $(\widetilde{Q}_{11}, \widetilde{Q}_{22})$. The observed range is [0.7007, 0.7518], corresponding to a relative amplitude of approximately 7.29%. This pronounced sensitivity is consistent with the absence of an observable anchoring variable: regime identification relies entirely on latent dynamics, so valuation is highly dependent on persistence assumptions. In particular, greater persistence of the high-mortality regime leads to higher premiums, whereas more rapid switching reduces exposure to adverse mortality conditions.

Table 4: Premium rate per \$1 of benefit (univariate 2-state model)

$\widetilde{Q}_{11} \backslash \widetilde{Q}_{22}$	0.95	0.80	0.65	0.50	0.35	0.20	0.05
0.95	0.7388	0.7259	0.7048	0.7224	0.7100	0.7007	0.7117
0.80	0.7494	0.7392	0.7391	0.7366	0.7315	0.7289	0.7281
0.65	0.7515	0.7442	0.7421	0.7386	0.7372	0.7345	0.7337
0.50	0.7511	0.7463	0.7458	0.7436	0.7407	0.7402	0.7368
0.35	0.7512	0.7497	0.7462	0.7439	0.7411	0.7408	0.7397
0.20	0.7516	0.7495	0.7469	0.7448	0.7438	0.7424	0.7417
0.05	0.7518	0.7490	0.7474	0.7464	0.7450	0.7436	0.7426

For the bivariate two-state model, Table 5 shows that premiums remain confined to the interval [0.6359, 0.6398], corresponding to a very small relative amplitude of about 0.61%. The inclusion of temperature substantially stabilizes valuation by linking latent regimes more closely to observable climatic conditions, so that transition uncertainty has only a limited effect on premiums. This pattern is consistent with a marked reduction in regime ambiguity once temperature is incorporated into the observation process.

For the bivariate three-state model, results are reported for two representative values of the intermediate-regime persistence $\widetilde{Q}_{22}$. Tables 6 and 7 correspond respectively to $\widetilde{Q}_{22} = 0.95$ and $\widetilde{Q}_{22} = 0.50$. Across both tables, premiums range from 0.6364 to 0.6501, corresponding to a relative amplitude of about 2.13%. While higher than in the 2-state bivariate case, this variability remains well below the univariate benchmark. The persistence of the extreme regime $\widetilde{Q}_{33}$ appears to be the main driver of valuation differences: when this regime becomes more persistent, premiums increase, consistent with more sustained exposure to high-mortality climatic episodes.

Table 5: Premium rate per \$1 of benefit (bivariate 2-state model)

$\widetilde{Q}_{11} \backslash \widetilde{Q}_{22}$	0.95	0.85	0.75	0.65	0.55	0.45	0.35	0.25	0.15	0.05
0.95	0.6393	0.6397	0.6395	0.6393	0.6396	0.6398	0.6397	0.6397	0.6397	0.6398
0.85	0.6387	0.6388	0.6389	0.6388	0.6391	0.6393	0.6391	0.6393	0.6392	0.6395
0.75	0.6374	0.6376	0.6384	0.6386	0.6388	0.6385	0.6388	0.6385	0.6392	0.6390
0.65	0.6367	0.6373	0.6376	0.6377	0.6383	0.6382	0.6386	0.6385	0.6387	0.6387
0.55	0.6364	0.6369	0.6374	0.6378	0.6382	0.6381	0.6382	0.6386	0.6385	0.6386
0.45	0.6365	0.6369	0.6372	0.6375	0.6377	0.6378	0.6380	0.6382	0.6383	0.6384
0.35	0.6364	0.6367	0.6371	0.6375	0.6375	0.6376	0.6378	0.6380	0.6382	0.6382
0.25	0.6361	0.6366	0.6367	0.6374	0.6376	0.6376	0.6374	0.6379	0.6380	0.6380
0.15	0.6359	0.6363	0.6367	0.6369	0.6373	0.6374	0.6377	0.6377	0.6377	0.6380
0.05	0.6359	0.6362	0.6366	0.6368	0.6369	0.6372	0.6375	0.6376	0.6376	0.6378

Table 6: Premium rate per \$1 of benefit (bivariate 3-state model, $\widetilde{Q}_{22} = 0.95$)

$\widetilde{Q}_{11} \backslash \widetilde{Q}_{33}$	0.95	0.85	0.75	0.65	0.55	0.45	0.35	0.25	0.15	0.05
0.95	0.6447	0.6453	0.6465	0.6477	0.6469	0.6460	0.6473	0.6461	0.6474	0.6476
0.85	0.6455	0.6469	0.6479	0.6478	0.6474	0.6476	0.6484	0.6479	0.6482	0.6488
0.75	0.6451	0.6484	0.6483	0.6486	0.6486	0.6489	0.6491	0.6487	0.6490	0.6488
0.65	0.6463	0.6479	0.6481	0.6489	0.6496	0.6494	0.6491	0.6494	0.6489	0.6490
0.55	0.6457	0.6484	0.6483	0.6487	0.6488	0.6494	0.6491	0.6492	0.6495	0.6494
0.45	0.6463	0.6475	0.6482	0.6486	0.6494	0.6496	0.6498	0.6497	0.6495	0.6497
0.35	0.6451	0.6477	0.6488	0.6490	0.6493	0.6493	0.6497	0.6497	0.6495	0.6497
0.25	0.6452	0.6474	0.6478	0.6494	0.6490	0.6493	0.6495	0.6496	0.6497	0.6498
0.15	0.6456	0.6475	0.6489	0.6487	0.6496	0.6499	0.6496	0.6500	0.6499	0.6501
0.05	0.6458	0.6472	0.6484	0.6492	0.6498	0.6497	0.6496	0.6501	0.6501	0.6500

The aggregate results are summarized in Table 8. In the two-state specification, the inclusion of temperature reduces premium dispersion by more than 90% relative to the univariate benchmark, suggesting that climate information provides a strong anchoring mechanism for latent regimes. The three-state model reintroduces moderate variability, which is consistent with a richer representation of extreme mortality-climate episodes. Overall, three main insights emerge. First, incorporating temperature substantially stabilizes premiums by reducing regime ambiguity. Second, introducing a third regime allows the model to reflect tail mortality risk associated with unusually adverse climatic conditions. Third, transition distortions provide a transparent way to incorporate forward-looking views on regime persistence without altering the historically estimated state-conditional mortality dynamics.

7 Conclusion

We investigated how to jointly model weekly mortality and temperature in order to capture their seasonal, state-dependent interaction and to embed climate-related regime uncertainty into actuarial valuation. To this end we developed and analyzed a Seasonal Hidden Markov Model with Trend (SHMM), extending classical hidden Markov models to non-homogeneous, periodically varying transitions and regime–specific trend–season structures for mortality and temperature (or an extreme temperature index). On the theoretical side, we established identifiability and strong consistency conditions for these seasonal, time–varying models under Gaussian and Poisson (and Poisson mixture) emissions, adapting arguments for latent structure uniqueness and ensuring that maximum likelihood estimates are reliable in the presence of seasonal transition kernels and regime-

Table 7: Premium rate per \$1 of benefit (bivariate 3-state model, $\widetilde{Q}_{22} = 0.50$)

$\widetilde{Q}_{11} \backslash \widetilde{Q}_{33}$	0.95	0.85	0.75	0.65	0.55	0.45	0.35	0.25	0.15	0.05
0.95	0.6379	0.6395	0.6404	0.6405	0.6410	0.6405	0.6410	0.6410	0.6410	0.6407
0.85	0.6379	0.6389	0.6405	0.6409	0.6413	0.6413	0.6417	0.6417	0.6418	0.6420
0.75	0.6374	0.6397	0.6402	0.6414	0.6416	0.6416	0.6419	0.6425	0.6425	0.6426
0.65	0.6371	0.6394	0.6406	0.6413	0.6418	0.6420	0.6427	0.6426	0.6430	0.6429
0.55	0.6370	0.6401	0.6402	0.6412	0.6421	0.6423	0.6427	0.6433	0.6431	0.6435
0.45	0.6365	0.6389	0.6407	0.6413	0.6423	0.6425	0.6429	0.6433	0.6433	0.6437
0.35	0.6369	0.6395	0.6396	0.6414	0.6420	0.6424	0.6429	0.6432	0.6438	0.6440
0.25	0.6374	0.6390	0.6401	0.6414	0.6423	0.6427	0.6437	0.6433	0.6436	0.6441
0.15	0.6370	0.6385	0.6410	0.6419	0.6420	0.6434	0.6436	0.6436	0.6441	0.6442
0.05	0.6364	0.6388	0.6406	0.6415	0.6423	0.6430	0.6434	0.6435	0.6441	0.6447

Table 8: Premium dispersion under transition stress

Model	States	Range	Relative amplitude	Reduction vs. univariate
Univariate	2	[0.7007, 0.7518]	7.29%	–
Bivariate	2	[0.6359, 0.6398]	0.61%	-91.6%
Bivariate	3	[0.6364, 0.6501]	2.13%	-70.8%

dependent trends. Verification for Poisson emissions shows that bounded intensities and distinct seasonal profiles are sufficient for generic linear independence, giving a solid asymptotic foundation for applications.

Empirically, simulated experiments and real European weekly data (ages 65+) demonstrated that the SHMM recovers interpretable latent regimes: cold-excess periods, warm baseline conditions, and—when a third regime is introduced—an extreme or transitional regime with higher volatility. Mortality trend parameters were consistently negligible once temperature and seasonal harmonics were included, suggesting that short-horizon fluctuations are predominantly climate–season driven rather than linearly trending. Incorporating temperature improved numerical stability and eliminated anomalous age effects seen in mortality-only specifications, while the Normal–Normal (log-rate) emission structure produced clearer regime separation than the Poisson–Normal alternative.

In pricing applications we adopted a parsimonious change of measure that tilts only the latent transition matrix, isolating the impact of assumed regime persistence (duration of cold or extreme conditions) without altering state–conditional mortality or temperature distributions. Baseline term-death premiums fell slightly when temperature was included, reflecting improved predictability of winter excess mortality, and rose marginally with a third (extreme) regime due to added tail exposure. Sensitivity analyses showed a marked compression of premium dispersion (about an 85% reduction) when moving from a mortality-only to a bivariate (mortality + temperature) two-state model; the three-state variant reintroduced moderate variability useful for stress testing but remained substantially more stable than the univariate benchmark.

The current specification only captures contemporaneous temperature effects rather than distributed lag or non-linear exposure–response structures widely used in epidemiology. Future methodological extensions could integrate lagged non-linear components, Bayesian hierarchical pooling across regions, or mixture priors to manage structural uncertainty in the number of regimes and harmonics. From an actuarial perspective, a natural next step is to couple the SHMM with projected climate pathways (RCP / SSP scenarios) so that long-horizon mortality pricing and longevity capital can incorporate anticipated shifts in temperature distributions and regime persistence. Further enrichment with air quality, humidity, precipitation extremes, and multi-hazard indicators may capture synergistic influences and improve resilience of tariff structures to compound environmental shocks. Finally, extending this framework with forward-looking environmental scenarios and richer covariate sets promises to enhance its role in longevity risk management and climate-aware insurance product design.

Reproducibility. Code in `R` (`depmixS4`) with the custom Poisson exposure class and replication scripts are available at https://github.com/Dakopge2018/model_temps_mortality_using_hmm

Acknowledgement. This paper is based on the Master's thesis work of the first author, Samuel Darwin Dona Kouton. The authors gratefully acknowledge financial support from the Fonds de recherche AMF en gestion intégrée des risques des institutions financières (AMF-GIRIF), Université Laval.

References

Allman, E. S., Matias, C. & Rhodes, J. A. (2009), 'Identifiability of parameters in latent structure models with many observed variables', *The Annals of Statistics* pp. 3099–3132.

Analitis, A., Katsouyanni, K., Biggeri, A., Baccini, M., Forsberg, B., Bisanti, L., Kirchmayer, U., Ballester, F., Cadum, E., Goodman, P. G. et al. (2008), 'Effects of cold weather on mortality: results from 15 european cities within the phewe project', *American Journal of Epidemiology* **168**(12), 1397–1408.

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

Bauer, D., Bergmann, D. & Kiesel, R. (2010), 'On the risk-neutral valuation of life insurance contracts with numerical methods in view', *ASTIN Bulletin: The Journal of the IAA* **40**(1), 65–95.

Bauer, D., Börger, M. & Ruß, J. (2010), 'On the pricing of longevity-linked securities', *Insurance: Mathematics and Economics* **46**(1), 139–149.

Bégin, J.-F., Boudreault, M. & Landry, T. (2025), 'Modelling seasonal mortality: An age–period–cohort approach', *Insurance: Mathematics and Economics* p. 103162.

Biffis, E., Denuit, M. & Devolder, P. (2010), 'Stochastic mortality under measure changes', *Scandinavian Actuarial Journal* **2010**(4), 284–311.

Cairns, A. J., Blake, D. & Dowd, K. (2006), 'Pricing death: Frameworks for the valuation and securitization of mortality risk', *ASTIN Bulletin: The Journal of the IAA* **36**(1), 79–120.

Cairns, A. J., Blake, D. & Dowd, K. (2008), 'Modelling and management of mortality risk: a review', *Scandinavian Actuarial Journal* **2008**(2-3), 79–113.

Guibert, Q., Pincemin, G. & Planchet, F. (2025), 'Impact of climate change on mortality: An extrapolation of temperature effects based on time series data in france', *International Journal of Forecasting* .
URL: *https://www.sciencedirect.com/science/article/pii/S0169207025000639*

Hainaut, D. (2012), ‘Multidimensional lee–carter model with switching mortality processes’, *Insurance: Mathematics and Economics* **50**(1), 113–125.
URL: *https://doi.org/10.1016/j.insmatheco.2011.11.006*

Li, H. & Tang, Q. (2022), ‘Joint extremes in temperature and mortality: A bivariate pot approach’, *North American Actuarial Journal* **26**(1), 43–63.

Michelot, T. (2025), ‘hmmtmb: Hidden markov models with flexible covariate effects in r’, *Journal of Statistical Software* **114**(5), 1–45.
URL: *https://www.jstatsoft.org/article/view/v114i05*

Pan, Q., Porth, L. & Li, H. (2022), ‘Assessing the effectiveness of the actuaries climate index for estimating the impact of extreme weather on crop yield and insurance applications’, *Sustainability* **14**(11), 6916.

Pascal, M., Laaidi, K., Ledrans, M., Baffert, E., Caserio-Schönemann, C., Le Tertre, A., Manach, J., Medina, S., Rudant, J. & Empereur-Bissonnet, P. (2006), ‘Definition of temperature thresholds: the example of the french heat wave warning system’, *International Journal of Biometeorology* **50**(6), 369–376.

Robben, J. & Antonio, K. (2024), ‘Catastrophe risk in a stochastic multi-population mortality model’, *Journal of Risk and Insurance* **91**(3), 599–651.

Robben, J. & Barigou, K. (2025), ‘A penalized distributed lag non-linear lee-carter framework for regional weekly mortality forecasting’, *Arxiv Preprint Arxiv:2509.24087* .

Robben, J., Barigou, K. & Kleinow, T. (2026), ‘Granular mortality modeling with temperature and epidemic shocks: a three-state regime-switching approach’, *Insurance: Mathematics and Economics* p. 103250.
URL: *https://www.sciencedirect.com/science/article/pii/S0167668726000405*

Seklecka, M., Pantelous, A. A. & O’Hare, C. (2017), ‘Mortality effects of temperature changes in the united kingdom’, *Journal of Forecasting* **36**(7), 824–841.

Touron, A. (2018), ‘Consistency of the maximum likelihood estimator in seasonal hidden markov models’, *Statistics and Computing* .

Touron, A. (2019), Modélisation multivariée de variables météorologiques, PhD thesis, Université Paris-Saclay.
URL: *https://theses.hal.science/tel-02319170*

Touron, A., Hoang, T. T. H. & Parey, S. (2018), ‘Bivariate modelling of precipitation and temperature using a non-homogeneous hidden markov model’, *arXiv preprint arXiv:1810.09682* .

Zhou, R. (2019), ‘Modelling mortality dependence with regime-switching copulas’, *ASTIN Bulletin: The Journal of the IAA* **49**(2), 373–407.