

ARTICLE TYPE: RESEARCH ARTICLE

Cross-Population Mortality Dependence and Longevity Risk: A Multivariate Copula Framework

George Papadopoulos and Anna Lindström

Independent Researchers, Thessaloniki, Greece and Stockholm, Sweden

Received:	19 June 2026	Accepted:	15 September 2026
Published:			

Abstract

Longevity-risk diversification across populations depends critically on whether mortality improvements remain correlated in adverse scenarios. This paper develops a multivariate copula framework for cross-population mortality dependence using a simulated panel of six populations observed over ages 55–90 and 45 calendar years. Population-specific mortality trends are estimated with a common-factor model, after which dependence among standardized mortality innovations is modelled using Gaussian, Student-t and factor-copula specifications. The data-generating process includes a shared longevity trend, region-specific factors and episodic joint mortality-improvement shocks, allowing the effect of tail dependence on annuity liabilities to be measured against a known benchmark. Results show that Gaussian dependence overstates diversification when extreme longevity improvements occur simultaneously, whereas the Student-t and factor-copula models reproduce joint longevity shocks more accurately. The impact is modest for one-year mortality forecasts but substantial for the upper tail of 30-year annuity present values. Bootstrap and rolling-window analyses indicate that dependence uncertainty becomes a material component of longevity capital when portfolios span highly related populations. The study offers a reproducible framework for integrating mortality forecasting and dependence modelling in longevity-risk management.

Keywords: Mortality; longevity risk; multi-population models; copulas; annuities; dependence

JEL Classification: C32; C53; G22; J11

1. Introduction

Longevity risk arises when mortality improves more rapidly than expected, increasing the duration and value of pension and annuity liabilities. Because insurers and pension funds often hold exposures to several demographic groups, cross-population diversification is potentially important. That diversification, however, depends on how mortality innovations co-move. If extreme improvements are shared across populations, a portfolio that appears diversified under Gaussian dependence can remain highly exposed in the longevity tail.

Multi-population mortality models address common trends explicitly. Extensions of the Lee–Carter framework, including coherent and common-factor approaches, allow populations to share long-run mortality dynamics while retaining population-specific deviations. These models are well suited to forecasting mean mortality paths, but the residual dependence that remains after fitting the common trend can still matter for joint liability distributions.

The present paper connects multi-population mortality modelling with copula-based dependence analysis. The objective is not to propose another central mortality trend model; rather, it is to examine how the dependence of

mortality innovations affects longevity-risk aggregation. A synthetic setting is used so that the true common trend, regional factors and episodic joint longevity shocks are known.

The study contributes by separating three sources of cross-population co-movement: deterministic common trend, persistent regional dependence and tail co-movement in innovations. This distinction helps explain why a mortality model can fit historical means well yet still understate portfolio tail risk if the dependence layer is too restrictive.

2. Mortality Forecasting and Cross-Population Dependence

The Lee–Carter model represents log mortality as an age pattern plus an age-specific sensitivity multiplied by a period index. Li and Lee (2005) extended this idea to coherent mortality forecasting, allowing related populations to share a common mortality development. Other approaches use product-ratio decompositions, common factors or multivariate time-series structures. The common objective is to prevent implausible long-run divergence among populations with related demographic experience.

Dependence matters beyond the conditional mean. Suppose two populations share a long-run period factor but also experience unusually strong mortality improvements in the same years. If those joint innovations are more frequent than Gaussian dependence implies, the upper tail of annuity values will be understated. This issue is conceptually similar to joint loss modelling in insurance, except that the risky event is unexpectedly low mortality rather than high claim severity.

Copulas provide a natural separation between population-specific innovation distributions and their dependence. A Student-t copula permits symmetric joint extremes, while factor copulas provide a scalable way to represent dependence among many populations. For longevity portfolios, dependence can be assessed directly on mortality innovations or on derived liability shocks. The former is statistically more transparent, whereas the latter is closer to the economic quantity of interest.

The paper adopts both perspectives. Copulas are fitted to mortality innovations and then propagated through annuity valuation to evaluate portfolio capital. This allows the statistical dependence model to be judged by its financial consequences rather than by goodness-of-fit alone.

$$\log m_{x,t}^{(i)} = a_x^{(i)} + b_x^{(i)} k_t^{(i)} + \varepsilon_{x,t}^{(i)} \quad (1)$$

$$k_t^{(i)} = \beta_i K_t + \gamma_i R_{g(i),t} + u_t^{(i)} \quad (2)$$

$$PV^{(i)} = \sum_{s=1}^T v^s \square_s p_x^{(i)} \quad (3)$$

3. Synthetic Mortality Panel and Estimation Strategy

The synthetic panel contains six populations grouped into three regions. Ages 55–90 are observed for 45 years. Log mortality follows a declining common period factor, age-specific sensitivities and a region factor. Population-specific innovations are heavy-tailed and include five system-wide longevity-improvement shocks. These shocks lower mortality simultaneously in all populations but with heterogeneous intensity.

Population models are estimated in two stages. First, age schedules and period factors are extracted using singular-value decomposition under standard identification constraints. Second, period-factor dynamics are modelled with random walks with drift and regional adjustments. Standardized innovations are then transformed to pseudo-observations for copula estimation.

Three dependence models are compared: Gaussian copula, Student-t copula and a one-factor Student-t copula with population-specific loadings. Forecasts are generated for 30 years, and mortality paths are converted to cohort survival probabilities. A stylized immediate annuity at age 65 is then valued under each simulated path using a fixed discount curve.

4. Evaluation Design

Model evaluation is performed at both the mortality and liability levels. Mortality diagnostics include pairwise Kendall dependence, joint lower-innovation exceedances and calibration of the number of populations simultaneously experiencing unusually strong mortality improvement. Liability diagnostics include the mean annuity present value, 95% and 99.5% liability quantiles and portfolio diversification benefit across populations.

Forecast evaluation distinguishes one-year calibration from long-horizon liability risk. This is important because cross-population dependence has relatively little effect on individual one-year point forecasts but can accumulate substantially when many future survival probabilities are combined in a long-duration annuity.

Rolling-window fits are used to assess dependence stability. Bootstrap confidence intervals are obtained by resampling years in blocks to preserve low-frequency mortality dependence. Alternative age ranges and discount rates are evaluated as robustness checks.

Table 1. Dependence-model implications for joint longevity risk

Model	Joint longevity-shock error	99.5% liability index	Diversification benefit
Gaussian copula	0.31	1.0	0.142
Student-t copula	0.12	1.046	0.111
Factor-t copula	0.07	1.058	0.103

5. Results

The estimated common mortality trend is recovered accurately by all specifications because the central trend model is identical. Differences emerge in the residual dependence layer. Pairwise Kendall dependence is moderate for populations within the same region and lower across regions, but the simulated system-wide longevity shocks create more simultaneous extremes than a Gaussian copula predicts.

The Gaussian model therefore appears adequate in ordinary years yet materially understates the frequency of years in which four or more populations experience extreme mortality improvements simultaneously. The Student-t copula improves this diagnostic, while the factor-t model performs best because it combines heavy-tailed joint shocks with heterogeneous population loadings.

At the liability level, mean annuity values are nearly identical across dependence models. This is expected because dependence does not materially change marginal mortality forecasts. The difference appears in portfolio quantiles. The 99.5% liability measure under the factor-t model is approximately 5.8% above the Gaussian benchmark, and diversification benefit falls from 14.2% to 10.3%. Thus, the principal effect of dependence modelling is on capital and diversification rather than on expected liabilities.

The effect grows with forecast horizon. At five years the capital difference is small, while at 20–30 years it becomes material because common longevity shocks propagate through many survival probabilities. This finding is particularly relevant for long-dated annuity and pension exposures.




Figure 1. Estimated common mortality period factor




Figure 2. Pairwise innovation dependence

6. Robustness and Sensitivity

Changing the age range to 60–89 slightly increases dependence estimates because older-age mortality innovations share a stronger common component in the simulation. The ranking of copula models is unchanged. Using a lower discount rate increases all annuity values but also magnifies the absolute capital difference between Gaussian and heavy-tailed dependence.

Block-bootstrap intervals for 99.5% liability capital are wider than intervals for mean liabilities, confirming that dependence uncertainty is concentrated in the tail. The factor-t model remains above the Gaussian benchmark in more than 90% of bootstrap replications. A placebo simulation with no system-wide longevity shocks removes most of the capital difference, demonstrating that the effect is genuinely driven by joint extremes rather than by ordinary correlation.

Alternative factor specifications with two global factors yield only minor improvements because the generating process contains one dominant global source plus regional components. In real data, however, model selection should allow for richer structures and for time variation in dependence.

7. Discussion

The results show why mortality diversification should not be inferred solely from historical correlation. Common trend models can eliminate long-run divergence while leaving residual tail dependence that remains important for solvency. A liability portfolio may therefore be well diversified in expectation but less diversified in extreme longevity scenarios.

The analysis also provides a practical modelling hierarchy. First, coherent mortality forecasting should capture common long-run trends. Second, residual innovations should be checked for cross-population dependence and tail concentration. Third, the dependence model should be validated using liability-level diagnostics because statistically small differences can become economically meaningful after long-horizon survival aggregation.

The synthetic design is intentionally stylized and does not include migration, cause-of-death interactions, pandemic shocks or cohort effects. Those features could alter both marginal forecasts and dependence. Nevertheless, the framework is sufficiently general to accommodate richer mortality models because the copula layer is applied to standardized innovations rather than tied to a specific mortality specification.

8. Conclusion

Cross-population dependence is a material component of longevity-risk aggregation. In the controlled panel studied here, Gaussian dependence reproduces ordinary co-movement but understates simultaneous extreme mortality improvements and therefore overstates diversification. Student-t and factor-t copulas provide more accurate joint-shock calibration and produce higher, more realistic tail liability measures.

The findings support an integrated approach in which mortality forecasting and dependence modelling are evaluated jointly. For insurers and pension funds with exposures to related populations, tail dependence should be treated as a distinct source of model risk rather than assumed to be captured automatically by a common mortality trend.

9. Dependence Modelling in Longevity Portfolios

Cross-population longevity dependence can arise from several distinct sources: shared medical technology, public-health policy, socioeconomic convergence, environmental exposure and common measurement practices. A statistical dependence parameter does not reveal which mechanism dominates. Nevertheless, separating common deterministic or stochastic trends from residual joint extremes is useful because the two components have different implications for hedging. A common trend may be partially hedgeable with broad longevity instruments, while residual population-specific tail dependence may remain as basis risk.

The choice of dependence scale is also important. Modelling innovations in period indexes is parsimonious, but it assumes that age-specific dependence is adequately summarized by the factor representation. Direct dependence

modelling across ages would be more flexible but vastly more dimensional. A practical compromise is to model a small number of mortality factors and then test residual age-specific dependence at economically important ages. This hierarchical approach mirrors factor-copula strategies in finance.

Liability aggregation can amplify small dependence differences because survival probabilities are products of many annual terms. Even if one-year mortality innovations differ only slightly across copula models, persistent joint improvements can shift long-horizon annuity quantiles. This nonlinearity explains why liability-level validation is essential. Statistical fit at the mortality-innovation level should be considered necessary but not sufficient for longevity-capital applications.

Hedging effectiveness is similarly dependence sensitive. A hedge based on another national population will perform well only if the relevant innovations remain dependent in stress states. A Gaussian correlation estimated over the full sample can overstate hedge reliability when dependence weakens or changes shape in extreme longevity scenarios. Tail-sensitive dependence measures can therefore contribute directly to basis-risk assessment.

10. Model Risk, Validation and Governance

Longevity models are long-horizon models, so parameter and structural uncertainty often dominate ordinary process risk. Dependence parameters add another source of uncertainty that should be propagated into capital rather than fixed at point estimates. Bootstrap distributions in this paper show that tail-capital intervals widen meaningfully when copula parameters are re-estimated, especially for smaller population groups.

Validation should include historical backtesting where feasible, but the scarcity of extreme longevity events limits purely empirical assessment. Scenario analysis is therefore indispensable. Examples include accelerated cardiovascular improvement, a medical breakthrough affecting several countries, or divergent improvements caused by unequal adoption of treatment. Dependence models should be tested against such narratives to determine whether they generate plausible cross-population responses.

Model governance should also distinguish demographic judgment from statistical estimation. Long-run convergence assumptions, population grouping and cohort-treatment choices often reflect substantive demographic reasoning. These choices affect residual dependence and should be documented before copula calibration. Otherwise, the dependence model may absorb structure that belongs in the mortality trend model.

For portfolio applications, reporting only a single best-estimate dependence model is insufficient. Challenger specifications, parameter stress and alternative population groupings should be maintained. The spread in liability capital across plausible models provides a transparent measure of dependence-model risk.

11. Limitations and Future Research

The synthetic panel does not include cause-of-death information, migration or explicit cohort effects. Each of these can generate additional cross-population dependence. Cause-specific improvements may spread internationally at different speeds, while cohort effects can be shared among culturally or geographically related populations. Extending the copula layer to latent cause-specific or cohort factors would be a natural next step.

Pandemic and disaster mortality also require special treatment because they represent adverse mortality shocks rather than longevity improvements. A comprehensive life-insurance portfolio may face both short-term mortality catastrophe risk and long-term longevity risk. Dependence between these components across populations is not necessarily symmetric and may require regime-dependent or mixture copulas.

Future work could also connect dependence modelling to tradable longevity instruments. Evaluating hedge ratios and basis risk under Gaussian versus tail-dependent population relationships would provide a direct economic test of model relevance. Such an extension would move the framework from liability measurement toward active longevity-risk management.

References

Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741.

- Aas, K., Czado, C., Frigessi, A., & Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44(2), 182–198.
- Bedford, T., & Cooke, R. M. (2002). Vines—A new graphical model for dependent random variables. *Annals of Statistics*, 30(4), 1031–1068.
- Czado, C. (2019). *Analyzing Dependent Data with Vine Copulas*. Springer.
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66.
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions. *Journal of Econometrics*, 182(1), 119–134.
- Embrechts, P., McNeil, A., & Straumann, D. (2002). Correlation and dependence in risk management. In M. Dempster (Ed.), *Risk Management: Value at Risk and Beyond*. Cambridge University Press.
- Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3), 339–350.
- Joe, H. (2014). *Dependence Modeling with Copulas*. CRC Press.
- McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools* (2nd ed.). Princeton University Press.
- Nelsen, R. B. (2006). *An Introduction to Copulas* (2nd ed.). Springer.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2), 527–556.
- Cairns, A. J. G., Blake, D., & Dowd, K. (2006). A two-factor model for stochastic mortality with parameter uncertainty. *Journal of Risk and Insurance*, 73(4), 687–718.
- Lee, R. D., & Carter, L. R. (1992). Modeling and forecasting U.S. mortality. *Journal of the American Statistical Association*, 87(419), 659–671.
- Li, N., & Lee, R. (2005). Coherent mortality forecasts for a group of populations: An extension of the Lee–Carter method. *Demography*, 42(3), 575–594.
- Hyndman, R. J., & Ullah, M. S. (2007). Robust forecasting of mortality and fertility rates: A functional data approach. *Computational Statistics & Data Analysis*, 51(10), 4942–4956.
- Haberman, S., & Renshaw, A. (2011). A comparative study of parametric mortality projection models. *Insurance: Mathematics and Economics*, 48(1), 35–55.
- Villegas, A. M., Kaishev, V. K., & Millossovich, P. (2018). StMoMo: An R package for stochastic mortality modeling. *Journal of Statistical Software*, 84(3), 1–38.

Declarations

Funding: No external funding was used for the preparation of this manuscript.

Conflict of interest: The authors declare no competing interests.

Data and code availability: The numerical illustrations use simulated data and deterministic random seeds so that the analyses can be reproduced. Code and generated data may be deposited with the journal as supplementary material.