Insurance Mathematics and Economics xxx (xxxx) xxx

Contents lists available at ScienceDirect

Insurance Mathematics and Economics

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

Editorial

Longevity risk and capital markets: The 2024–25 update☆

This Special Issue of the *Insurance: Mathematics & Economics* contains 17 contributions to the academic literature all dealing with longevity risk and capital markets. Draft versions of the papers were presented at *Longevity 19: The Nineteenth International Longevity Risk and Capital Markets Solutions Conference* that was held in Amsterdam on 16–17 September 2024. It was hosted by the Research Centre for Longevity Risk at the University of Amsterdam and the Pensions Institute at Bayes Business School, City St George's, University of London.

Longevity risk and related capital market solutions have grown increasingly important in recent years, both in academic research and in the market we refer to as the Life Market, i.e., the capital market that trades longevity- and mortality-linked assets and liabilities. Mortality improvements around the world are putting increasing pressure on governments, pension funds, life insurance companies and individuals to deal with the longevity risk they face. At the same time, capital markets can, in principle, provide vehicles to hedge longevity risk effectively and transfer the risk from those unwilling or unable to manage it to those willing to invest in this risk in exchange for appropriate risk-adjusted returns or to those who have a counterbalancing risk that longevity risk can hedge, e.g., life offices and reinsurers with mortality risk on their books. Many new investment products have been created both by the re/insurance industry and by the capital markets. Mortality catastrophe bonds are an early example of a successful insurance-linked security. Some new innovative capital market solutions for transferring longevity risk include longevity (or survivor) bonds, longevity (or survivor) swaps, mortality (or *q*-) forward contracts and reinsurance sidecars (also called strategic reinsurance vehicles). The aim of the *International Longevity Risk and Capital Markets Solutions Conferences* is to bring together academics and practitioners from all over the world to discuss and analyze these exciting new developments.

As with the previous conferences, *Longevity 19 (L19)* consisted of both academic papers and more practical and policy-oriented presentations. There were four plenary sessions: (1) Hot topics in longevity, with Guy Coughlan (Chief Operating Officer at Clota Varde and member of Advisory Board at Longevity Solutions), Marine Habart (AXA) and Alexandre Boumezoued (Milliman), Sune Lehmann (Technical University of Denmark and Center for Social Data Science), and Tinhinane Talbi (Milliman); (2) Advances in the longevity risk transfer market, with Douglas Anderson (Club Vita), Paul Kitson (EY), and Leendert van Gastel (EIOPA); (3) Longevity issues in the Netherlands, the US, the UK and Australia, with Cara Dowell (RGA), Andrew Hunt (Pacific Life Re), and Andrea Staehr (Hannover Re); (4) The Netherlands, pension reforms, and opportunities for buy-out, with Rohit Mathur (Prudential Financial), Eilidh Brown (Willkie), Wichert Hoekert (WTW), and Marc Meertens (Head of Balance Sheet Management at NN Leven). A detailed summary of the sessions is provided as online supplementary material.

The academic papers that were selected by us as the editors of this Special Issue went through a refereeing process subject to the usual high standards of the *Insurance: Mathematics & Economics*. They cover the following themes: longevity indices, longevity risk transfers, reverse mortgages, pensions, annuities, mortality models and mortality forecasting, Covid-19 mortality, climate-related mortality, and the link between subjective well-being and mortality. We briefly discuss each of the 17 papers selected.

In 'Avoiding a longevity catastrophe: Harnessing longevity indices to mitigate individual, institutional and systemic longevity risks', Guy Coughlan considers the financial implications of an extreme increase in life expectancy for: (i) individuals with defined contribution pension plans and other forms of retirement savings; (ii) institutions such as defined benefit pension plans, insurance companies and reinsurers; and (iii) the financial system and economy as a whole. An extreme longevity scenario, as the IMF first acknowledged in 2006, is a long-term systemic risk that could impair the operation of the financial system with severe ramifications for the global economy. It also poses a significant risk to individuals who might live beyond the time that their retirement savings can support them. This paper explores one under-utilised way to mitigate these risks, viz., longevity index hedges, which can transfer longevity risk simply, rapidly and transparently away from where it is concentrated to a much broader set of organisations with appropriate levels of risk capital. For the market in these index hedges to grow requires a shared understanding of the hedges and their risk reduction potential by the insurance industry and regulators.

In 'Contract structure and risk aversion in longevity risk transfers', David Landriault, Bin Li, Hong Li and Yuanyuan Zhang develop an economic framework for optimal longevity risk transfer between a buyer and a seller with different risk aversions. The authors compare static (long-dated, pre-committed) and dynamic (short-dated, rolled) longevity swaps in a Stackelberg game. They find that static contracts are preferred when the buyer is more risk averse, while dynamic contracts are preferred when the seller is more risk averse. For the capital-market setting, they extend the benchmark by introducing seller-side ambiguity about the mortality distribution and robust max-min

☆ David Blake and Richard MacMinn are co-founders of the *Longevity Risk and Capital Markets Solutions Conferences*. We are extremely grateful to Marilyn Parris-Bell and Mike Mansfield for their superb organization of the conference. Cairns and Blake have a paper published in this Special Issue; the refereeing process for this paper was handled by the IME's General Editor.

https://doi.org/10.1016/j.insmatheco.2026.103309

Available online 9 October 2026
0167-6687/© 2026 Elsevier B.V. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

valuation. Even moderate ambiguity can eliminate the market for static swaps, while dynamic designs remain viable. The authors then extend the analysis to index-based swaps with basis risk: relative to indemnity swaps, optimal loadings are lower and gains are smaller for both parties, though the static–dynamic preference pattern is unchanged.

In 'Reverse mortgages with long-term care cost coverage: The role of health state granularity in valuing prepayment options', Tian-Shyr Dai, Liang-Chih Liu, Shang-Hung Tsao, and Sharon S. Yang point out that the growing elderly population has intensified interest in financial products such as reverse mortgages (RM) with long-term care (LTC) cost coverage, which provide funding solutions for senior homeowners facing income shortfalls and rising care costs. Their study advances existing research on reverse mortgage contracts with long-term care cost coverage (RM LTC) by systematically examining how the granularity of health state information influences contract valuation, particularly in the presence of prepayment options and market risks. They develop a numerical framework to determine the fair annuity rate of RM LTC contracts. The framework jointly captures stochastic house prices, interest rates, and multiple-health-state transition matrices calibrated with real-world disability data. Their results show that finer health state granularity improves the precision of prepayment decisions, increasing the value and therefore the cost of embedded options while lowering fair annuities. Conversely, relying on less detailed health information leads to an underestimation of option value and results in higher annuities offered to homeowners. By clarifying these mechanisms, the study underscores the pivotal role of health state granularity in valuing RM LTC contracts and provides insights into the design of home-equity-based financing for long-term care.

In 'Life care reverse mortgages: Monitoring the net cashflows of a new hybrid insurance product', Giovanna Apicella, Emilia Di Lorenzo, Giulia Magni and Marilena Sibillo argue that combining retirement income and long-term disability protection is a well established concept in the literature. They present a new connection between the lifetime annuity provided by a reverse mortgage and long-term care (LTC) insurance, jointly managed in a unique bundled product. Indeed, they define the RM (considered as a source of lifetime income) and LTC as two different regimes of a jointly managed insurance product, which they call the 'Life care reverse mortgage' (LCRM). From an actuarial perspective, the authors design the underlying structure of the LCRM and the inherent framework for the ex-ante estimation of a time-dependent profit/loss function, that provides a measure of the expected annual net cashflows for a life insurer holding a portfolio of LCRMs. They then perform a numerical application to illustrate the regime-switching mechanism on which the proposed LCRM insurance contracts are based and to quantify over time the lender's return for a pool of LCRM contracts through the designed time-dependent profit/loss function. Furthermore, they analyse the sensitivity of the portfolio profit/loss function to two sources of uncertainty: health patterns over time and house price dynamics.

In 'The Big Thaw: Unfreeze defined benefit pension with cash balance plans', Yijia Lin and Tianxiang Shi point out that the recent resurgence of defined benefit (DB) pensions through cash balance (CB) plans has sparked discussions and captured the attention of employers, unions, and benefits experts. However, there remains limited understanding of how unfreezing DB pensions impacts a firm. To address this gap, the authors explore the cost and benefit tradeoffs associated with pension unfreezing using a model calibrated to market data within a stochastic framework. Their analysis reveals that the extent to which restarting DB pensions affects a firm's total pension cost depends on factors such as pension surplus magnitude, employee salaries, firm financial stability, and Pension Benefit Guaranty Corporation (PBGC) premiums. In particular, with a high funding surplus and elevated employee salaries, unfreezing a CB plan can yield greater cost efficiency for companies than pension risk transfer, especially in a high interest rate environment. The authors' findings provide valuable insights for pension sponsors contemplating the reopening of DB pensions through CB plans, which could potentially counteract the current trends in pension de-risking.

In 'Rethinking the annuity puzzle: The role of loss aversion and money-back guarantees', Sebastian Hallstein, Daniel Liebler and Raimond Maurer argue that consumption and saving models typically suggest accumulation in life-only annuities that is at odds with actual investments. Milevsky and Salisbury (2022) add another dimension to this annuity puzzle by documenting a shift from life-only annuities to those offering money-back guarantees to the heirs of the annuitant. The authors provide further evidence for a balanced coexistence of both annuity types; however, standard utility frameworks that rely solely on consumption and bequest utility fail to reproduce this relationship. To overcome the dominance of life-only annuities, they incorporate loss aversion into individual preferences. Specifically, they assume aversion to dying before the total payouts exceed the premiums paid, creating a natural obstacle to life-only annuities. Using realistically calibrated life cycle models with access to liquid savings and both life-only and refundable annuities, they show that loss aversion can flexibly control the demand for annuities, thereby enabling a balanced relationship between the annuity types. While the authors' findings shed further light on behavioral drivers behind the annuity puzzle, their utility framework does not simultaneously establish the empirical relationship across annuity types as well as the overall demand for annuities.

In 'Robust parameter estimation for the Lee-Carter family: A probabilistic principal component approach', Yiping Guo and Johnny Siu-Hang Li point out that although the impact of outliers on stochastic mortality modelling has been examined, previous studies on this topic focus on how outliers in the estimated time varying indexes may be detected and/or modelled, with little attention being paid to the adverse effects of outliers on estimation robustness, particularly that pertaining to age-specific parameters. In this paper, the authors propose a robust estimation method for the Lee-Carter model, through a reformulation of the model into a probabilistic principal component analysis with multivariate t distributions and an efficient expectation maximization algorithm for implementation. The proposed method yields significantly more robust parameter estimates, while preserving the fundamental interpretation for the bilinear term in the model as the first principal component and the flexibility of pairing the estimated time-varying parameters with any appropriate time-series process. Thay also extend the proposed method for use with multi-population generalizations of the Lee-Carter model, allowing for a wider range of applications such as quantification of population basis risk in index-based longevity hedges. Using a combination of real and pseudo datasets, they demonstrate that the superiority of the proposed method relative to conventional estimation approaches such as singular value decomposition and maximum likelihood.

In 'An empirical study of small populations based on Taiwan's data', Hsin-Chung Wang and Jack C. Yue argue that the estimation of mortality and disease incidence rates for small populations poses significant challenges for insurers in meeting regulatory compliance requirements, such as those imposed by the European prudential framework Solvency II. These difficulties primarily arise from data scarcity and the substantial volatility inherent in estimates derived from limited observations. To alleviate these issues, several methodologies have been proposed that effectively enlarge the underlying data set – for instance, through multi-population mortality models or the incorporation of reference populations – to reduce the stochastic variation in small-sample estimates. Such approaches, however, rest critically on the assumption that the reference population exhibits comparable demographic and epidemiological characteristics to those of the target population. A frequently adopted strategy involves pooling data from geographically adjacent districts; nevertheless, neighboring areas often differ markedly in their socioeconomic and demographic compositions, which may induce pronounced heterogeneity in age-specific mortality and incidence patterns. In this study, the authors propose a data-driven framework for identifying suitable reference populations to enhance the

accuracy of mortality estimation in small populations. Specifically, cluster analysis is employed to group small populations together with their neighboring districts, thereby constructing reference populations using techniques grounded in the heterogeneity measures proposed by Lee (2003). The study areas are partitioned into clusters exhibiting intra-group homogeneity and substantial inter-group heterogeneity in mortality rates. The proposed methodology is empirically validated using mortality data at the county and township levels in Taiwan. Based on the constructed reference populations, life tables are established, and life expectancy estimates for small populations are subsequently derived.

In 'Quantile-based interpretable neural network models: Mortality forecasting and actuarial simulations', Yang Qiao, Jinggong Zhang, Wenjun Zhu and Chou-Wen Wang introduce a class of quantile-based, interpretable neural network (NN) models for mortality prediction: the Lee-Carter neural network (LCNN) model, the Renshaw-Haberman neural network (RHNN) model, and the simple neural network (simpleNN) model that balances simplicity with predictive performance. These models preserve the linear interpretability of classic stochastic mortality models while harnessing the flexibility of NNs to capture complex nonlinear patterns in mortality data, thereby achieving enhanced predictive performance. Leveraging a composite loss function that integrates pinball, median anchoring, and quantile-crossing penalty terms, the authors estimate mortality quantiles and develop an efficient simulation scheme based on interpolation. This framework provides distributional insights essential for pricing, reserving and risk management. Extensive empirical analyses across multiple populations demonstrate that the proposed models consistently outperform traditional approaches in predictive accuracy. Their practical utility is illustrated through an application to longevity swap pricing.

In 'The future of mortality: Mortality forecasting by extrapolation of deaths curve evolution patterns', Matthias Börger, Martin Genz, and Jochen Ruß point out that a variety of mortality models can be used to project future mortality. However, the parameters of most of these models lack a clear demographic interpretation. Hence, the resulting projections may be demographically implausible in the sense that trends in key demographic statistics are not extrapolated in a reasonable way. When demographers make predictions on future mortality, they typically focus on one or few relevant demographic statistics related to certain aspects of the mortality evolution. However, they do not derive comprehensive mortality forecasts as required for actuarial purposes. This article aims to close the gap between these forecasting approaches. To this end, the authors establish a new deterministic mortality model which can be used for best estimate and scenario forecasts. They model the deaths curve, i.e. the age-at-death distribution, and derive forecasts based on the extrapolation of statistics that have a clear demographic interpretation. The four key statistics of the model are those from the classification framework of Börger et al. (2018). The design of the model makes sure that forecasts for the immediate future of the deaths curve are consistent with the most recent trends of all demographically relevant statistics. Moreover, expert opinions with respect to the future trends of certain demographically interpretable statistics can easily be incorporated – in particularly for the farther future where a pure extrapolation of historic trends might lead to implausible results. Finally, the authors present a possible implementation of the model and provide case studies that illustrate how the model can be applied.

In 'Scanning the horizon: Integrating expert knowledge into the calibration of stochastic mortality models', Richard G.A. Faragher, Arne Freimann and Jochen Ruß argue that expert knowledge from many different disciplines has the potential to inform on developments that could significantly increase or decrease human life expectancy. However, such knowledge is typically not considered in longevity risk management, since stochastic mortality models are generally only calibrated to historical mortality patterns, i.e., are fully data-driven. Following an interdisciplinary approach, the authors develop a methodology for how expert knowledge on the (uncertainty of the) future of human life expectancy can be integrated into the calibration of stochastic mortality models. They argue that this approach is particularly relevant if there are 'low probability / high impact' scenarios on the horizon, that are considered plausible by experts in their respective field but are 'virtually impossible' in models calibrated to historical data. Based on current research on treatments that might be effective in slowing down ageing, the authors motivate and propose an exemplary plausible scenario for the future development of human life expectancy. They assign a potential impact on life expectancy as well as a plausible probability of occurrence to the scenario and present a method for calibrating stochastic mortality models so that the resulting projections are in line with these parameters. In a case study, they analyse and compare the longevity risk in an exemplary annuity portfolio and show that this so-called 'driver-driven' calibration can lead to a structurally different assessment of longevity risk than the traditional 'data-driven' approach, especially with regard to tail risks.

In 'Prolonging life by vitagions: Modelling of mortality improvement shocks', Maria Carannante, Valeria D'Amato and Cinzia Di Palo argue that changes in life expectancy are commonly used to track shifts in mortality dynamics, yet they are hard to interpret, especially during so-called mortality shocks that arrive suddenly, are temporary, and run counter to the long-run trend. Purely extrapolative models cannot capture such changes. Therefore, the authors propose a forward-looking extension of the Lee-Carter framework, which integrates mortality improvements due to 'vitagions'. Vitagions are individual sources of mortality improvement, and projects how these might unfold in the future. To model mortality increases related to medical innovation, the authors consider extending the Lee-Carter model to capture heterogeneity in morality (i.e., a frailty-based model) and include a jump diffusion process estimated as a function of the relationship between medical progress and mortality trends to improve mortality projections.

In 'Beyond annual data: Mortality forecasting with mixed-frequency data', Runze Li, Rui Zhou and David Pitt argue that existing mortality models typically rely on annual data, with forecasts for a given year based on information available up to the end of the previous year. However, technological advances have enabled the collection and production of weekly and monthly death data, offering new opportunities to improve forecasting. Using only annual data overlooks the full range of available information. In this paper, the authors propose mixed frequency sampling (MIDAS) models to integrate monthly death counts (high frequency) with annual mortality rates (low frequency), enabling improved short-term prediction of annual mortality. Extending economic applications of MIDAS, which typically predict single variables such as GDP growth, their MIDAS framework accounts for age dependence unique to age-specific mortality modeling. The authors also evaluate different weighting functions, a core element in MIDAS that determines the relative importance of high-frequency data at different lags, and identify suitable weighting functions for mortality forecasting. Using U.S. mortality data, they demonstrate that their approach significantly improves prediction accuracy compared to models relying solely on annual data for short-term forecasting. These findings highlight the potential of MIDAS models as a useful tool for accurate and timely mortality forecasts.

In 'ADM's APPLE: The Accelerated Deaths Model with an application to the Covid-19 pandemic', Andrew J.G. Cairns and David Blake introduce the Accelerated Deaths Model (ADM) which builds on the hypothesis that, within a given age cohort, those who are less healthy are more likely to die if infected with Covid-19 than healthier people, leaving a pool of on-average healthier survivors. They use the term 'detrimental selection' which has two complementary aspects: the years of life lost by those who experienced an accelerated death; and the higher average life expectancy of survivors which we call their 'adjusted post-pandemic life expectancy' (ADM's APPLE). Their model represents a novel synthesis of recent advances in our understanding of mortality heterogeneity and the development of the Proportionality Hypothesis – both of which have improved our understanding of the Covid-19

pandemic. In particular, the authors identify an important positive relationship between mortality heterogeneity and accelerated deaths. They find, in the case of the Covid-19 pandemic in England, that the years of life lost by those who experienced an accelerated death, while significantly lower than pre-Covid life expectancy, was greater than reported in the media at the time. They also find that the increase in the mean life expectancy of survivors was very small. As a result, the impact on annuity providers (e.g., in terms of potentially higher annuity prices), pension schemes and life insurers was also very small. In contrast, the authors find that the impact on life expectancy of a general change in future mortality assumptions post-pandemic (i.e., the base mortality table and improvement rate) would be much greater. The ADM has potentially wide application, e.g., to other types of contagion and to climate-related deaths, where we would expect there to be a positive correlation between deaths and all-cause mortality (consistent with the Proportionality Hypothesis), but where the degree of detrimental selection might be different.

In 'On measuring Covid-19 excess mortality: Insights and challenges', Ayse Arik, Allen Klein, and Han Li examine all-cause excess mortality due to the Covid-19 pandemic across 16 countries by age and sex. Utilising the Short-Term Fluctuations data series from the Human Mortality Database, they analyse various measures of excess mortality on a weekly and yearly basis from 2020 through the end of 2023. They explore the strengths and limitations of different approaches to offer a comprehensive understanding of the pandemic's impact on diverse nations. Specifically, they evaluate two methodologies employed by the UK's Office for National Statistics, considering observed death counts both with and without adjusting for population sizes. They also apply a method based on annual mortality improvements and, as a final approach, implement a statistical model to weekly death counts. Each method defines excess mortality differently, allowing for comparison across alternative perspectives. Their analysis demonstrates largely consistent outcomes across different measures during the first two years of the pandemic, with significant variations in the last two years. These findings contribute to a nuanced understanding of various measures by highlighting their strengths and limitations.

In 'The climate lifespan link: Investigating the influence of temperature change on human mortality', Yijia Lin, Yanxin Liu and Ruilin Tian extend the widely adopted Lee and Carter (1992) mortality model by incorporating temperature-related covariates to quantify the impact of climate risk on mortality. Their analysis reveals variation in mortality sensitivity to temperature across different ages. Based on U.S. data, they find that the elderly population is particularly vulnerable to the impacts of climate change. Specifically, seniors aged 65–70 are most susceptible to extreme cold, while those aged 65–75 are most vulnerable to hot weather. Additionally, the effects of cold and hot weather on mortality differ significantly across various contiguous U.S. divisions. Their proposed Temperature Lee-Carter (TLC) model projects that, under high-emissions scenarios such as SSP3–7.0 and SSP5–8.5, global warming could increase U.S. mortality rates among individuals aged 65 and above by 13.05% to 27.22% by 2099, with younger seniors (65–75) experiencing up to a 42.57% rise—posing substantial challenges for public health systems and the financial sustainability of life insurance markets. Their results suggest that the effect of global warming on mortality counteracts the current population aging trends, which have important implications for mortality risk management.

Finally, in 'Mortality risks, survival pessimism, and subjective well-being: Evidence from the Health and Retirement Study', Lisa Posey, Sharon Tennyson and Nan Zhu point out that the positive relationship between an individual's subjective well-being (SWB) and their future survival prospects has been well documented. Using data from the Health and Retirement Study (HRS), they present empirical evidence of how this relationship operates through various channels. They use a two-stage approach where objective survival probability estimates derived from a first-stage Cox model are used to construct the survival bias employed as the dependent variable in the second stage, with SWB and other demographic and health-related variables being covariates in both stages. Their findings reveal that even after controlling for objective health-related factors and potential private information on health and mortality, individuals' SWB remains as a significant factor to their objective mortality. Moreover, respondents' SWB also affects the bias in the survival estimation—measured as the disparity between their subjective and objective survival probabilities. In particular, a higher SWB is correlated with a reduction in survival pessimism. The authors provide new evidence that bias in survival estimates can distort consumption and saving decisions, and provide estimates of its impact on welfare across different levels of SWB.

Longevity 20 took place in Singapore on 8–9 September 2025; the *Journal of Risk & Insurance* will publish a special issue. *Longevity 21* took place in Rome on 8–9 September 2026. *Longevity 22* is planned for Toronto on 7–8 September 2027.

Supplementary material

Supplementary material containing both an update on developments in the Life Market and more details about the conference presentations is available online at https://ars.els-cdn.com/content/image/1-s2.0-S0167668726000995-mmc1.pdf or https://www.sciencedirect.com/special-issue/10322519XKV or doi:10.1016/j.insmatheco.2026.103309.

References

Börger, M., Genz, M., Ruß, J., 2018. Extension, compression, and beyond – A unique classification system for mortality evolution patterns. Demography 55 (4), 1343–1361.

Lee, R., Carter, L., 1992. Modeling and forecasting U. S. mortality. J. Am. Stat. Assoc. 87 (419), 659–671.

Lee, W.C., 2003. A partial SMR approach to smoothing age-specific rates. Ann. Epidemiol. 13 (2), 89–99.

Milevsky, M.A., Salisbury, T.S., 2022. Refundable income annuities: feasibility of money-back guarantees. Insur.: Math. &Econ. 105, 175–193.

David Blake[a,*], Andrew J.G. Cairns[b], Torsten Kleinow[c]

[a] *Faculty of Finance and Pensions Institute, City St George's, University of London, UK*

[b] *Faculty of Actuarial Mathematics and Statistics, Heriot-Watt University, UK*

[c] *Research Centre for Longevity Risk, University of Amsterdam, The Netherlands*

[*] Corresponding author.

E-mail addresses: D.Blake@city.ac.uk (D. Blake), A.J.G.Cairns@hw.ac.uk (A.J.G. Cairns), t.kleinow@uva.nl (T. Kleinow).