

 Mortality and Longevity

 Aging and Retirement

2020 Living to 100 Discussant Comments 3A: New Developments in Mortality Modeling



Discussant Comments Session 3A: New Developments in Mortality Modeling

Cynthia Thomas Edwalds, FSA, ACAS, MAAA

First of all, let me congratulate these two authors, Zhiping Huang and Nhan Huynh, on two excellent and mathematically challenging papers. One of the aspects of being a discussant at the Living to 100 Symposium that I particularly enjoy is that I have to read and understand papers like these. I knew that I would find these papers interesting, but I might have been too lazy to actually read them otherwise.

While both papers deal with mortality modeling, their focus and their approaches are quite different, so I will discuss them separately.

Longevity Risk

Before I begin, I want to make a few comments about longevity risk. In her paper, with credible citation, Dr. Huang defines longevity risk as the risk that the overall survival probability of a reference population is higher than expected. In my opinion, it is important to distinguish between individual longevity risk and aggregate longevity risk. The definition given by Dr. Huang is correct for aggregate longevity risk, but without individual longevity risk, there would be no aggregate mortality risk.

Individual longevity risk is the risk that an individual will be faced with excess cash outflows due to living longer than expected. Mechanisms such as defined benefit pension plans, Social Security, and payout annuities offer relief to individuals who can use these mechanisms to transfer some or even most of their longevity risk. These mechanisms all operate by aggregating the longevity risk of individuals, thereby reducing the overall risk thanks to the Law of Large Numbers.

In a static mortality environment, this would be adequate, and the aggregate risk would continue to shrink as more individuals are added by these aggregators. However, mortality is not static, and these aggregators are faced with the aggregate longevity risk defined by Dr. Huang.

The Application of Affine Processes in Cohort Mortality Risk Models – Zhiping Huang; UNSW Australia, Michael Sherris, FSA, FIAA, FIA, UNSW Australia, Andres M. Vllagas Ramirez, Ph.D.; UNSW Australia; Jonathan Ziveyi, Ph.D. UNSW Australia.

To set the table for my discussion of Dr. Huang's paper, let me show you an equation that ought to be familiar to you from your days of studying Life Contingencies:

$$\bar{a}_x = \int_0^{\omega} e^{-\mu_{x+t} - \delta t} dt$$

This equation is not in her paper, but it is actually the starting point for her work. It also helps me to feel better by looking at an equation that is familiar.

This equation for continuous life contingent annuities is fairly simple, and much of what we studied as we prepared for the Life Contingencies exam involved simple deterministic assumptions about μ and δ . However, in reality both μ and δ represent complex stochastic processes. In order to grasp the risk represented by this equation, it is necessary to model both of these stochastic processes.

Much work has been done to model the stochastic process δ for interest rate paths, and separate work has been done to model the stochastic process μ for mortality. Dr. Huang observes that continuous time affine models have been used for both interest and mortality, and using these models together allows for an integrated framework for assessing the risk of life contingent annuities.

In particular, she proposes the Arbitrage-Free Nelson-Siegel (AFNS) Model for modeling the stochastic force of mortality. This model has closed form solutions for survival probabilities, which is very advantageous because it makes it possible to use these formulas in spreadsheets. There are three components of the survival probability formulas: a level term, a slope term, and a curvature term. The formulas are otherwise not very intuitive.

She fits the AFNS model to several other candidates to cohort mortality data for the US from the Human Mortality Database, using mortality rates at ages 50-100 for birth cohorts from 1883 to 1915. This focus on retirement and near-retirement ages is appropriate since her goal is to measure longevity risk. She compares the models and finds that the AFNS model outperforms the others in out-of-sample forecasting, using Mean Absolute Percentage Error to assess performance.

The use of cohort mortality data is appropriate, particularly since the prior work with affine mortality models has focused on their use with cohort mortality data. Cohort mortality models are advantageous for studying longevity risk because the patterns of mortality within a cohort are likely to be related and the long-term mortality trajectories are likely to be consistent. This makes them particularly useful for assessing long term obligations like annuities and pensions. However, the disadvantage of using cohort mortality models is that the cohorts that were used in this study were born a long time ago in order to include deaths at age 100. The out-of-sample data used for comparing the models was the mortality data for the 1916 birth cohort. How relevant is this data for current pensioners, many of whom were born in the Baby Boom following World War II?

Multi-Population Longevity Models: A Spatial Random Field Approach – Nhan Huynh; Mike Ludkovski, Ph.D., University of California at Santa Barbara; Howard Zail, Elucider, LLC

Dr. Huynh's paper is exciting, because there are many applications for joint modeling of the mortality of distinct populations in insurance. A few that come to mind right away are male versus female mortality, insured mortality versus population mortality, and company mortality versus industry mortality. Another application that would take advantage of the multi-population (instead of only two population) modeling capability of this technique would be modeling multi-level preferred mortality together with standard mortality.

The extent to which the data in one population influences the model for the mortality of the other population is determined by the lengthscale parameter, θ . A large value of θ means that the models for the populations will be very similar, while a small value of θ means that the models will be very independent. θ is estimated by maximum likelihood, along with all the other parameters of the joint mortality model.

One of the illustrations in the paper that particularly caught my eye was Figure 7, which illustrated the model discussed in Section 5.2. The heat map in Figure 7(a) shows the projected relative mortality for males and females between the ages of 70 and 84 in Denmark for the years 2020 to 2060 when the genders are modeled separately, while Figure 7(b) shows this projected relative mortality when the genders are modeled jointly. The color scales for these two heat maps are different, so you need to look closely to see what is happening.

In Figure 7(a), the projected mortality for males at almost all of these ages is lower than the projected mortality for females for an entire decade, in spite of the fact that there is absolutely no historical data in which male mortality is lower than female mortality. In Figure 7(b) female mortality remains comfortably below male mortality for all years and ages in the projection, and the range of fluctuation of the mortality ratio is quite small.

It would have been advantageous if we had been able to use this joint modeling technique when we were developing the 2015 VBT mortality table. We modeled male and female experience mortality separately, and we modeled male and female mortality improvement separately. We had considerable difficulty with mortality crossovers in the resulting draft tables, which we corrected manually so that the eventual published tables would meet what we called “heuristic mortality constraints”.

As exciting and useful as this paper is, however, there are limitations to its usefulness. The model fitting process is computationally expensive, so the models in the paper were all fit using ages 70 – 84 only. On top of that, Dr. Huynh suggests limiting any model fitting process to 5 distinct populations, although the method can theoretically be applied to more. The amount of time for the model solution to converge is an important consideration.

There is also a theoretical limitation to this modeling process. The method relies on the assumption of covariance stationarity, which constrains the shape of the resulting model. In particular, it means that the lengthscale parameters must be the same for all ages throughout the model. Dr. Huynh suggests segmenting models by age, with about 15 consecutive ages in each model. This is a reasonable solution, but it creates issues with potential discontinuities at the boundaries between the age segments. I also wonder if confining the model to such a small number of ages in each segment results in an implicit Gompertz-Makeham assumption for the underlying mortality of each segment.

Congratulations to the Authors

Both of these papers are excellent and practical. I am honored to have the opportunity to discuss them at this conference. I congratulate both Drs. Huang and Huynh for their contributions to our literature on modeling mortality at advanced ages.

About The Society of Actuaries

With roots dating back to 1889, the [Society of Actuaries](#) (SOA) is the world's largest actuarial professional organizations with more than 31,000 members. Through research and education, the SOA's mission is to advance actuarial knowledge and to enhance the ability of actuaries to provide expert advice and relevant solutions for financial, business and societal challenges. The SOA's vision is for actuaries to be the leading professionals in the measurement and management of risk.

The SOA supports actuaries and advances knowledge through research and education. As part of its work, the SOA seeks to inform public policy development and public understanding through research. The SOA aspires to be a trusted source of objective, data-driven research and analysis with an actuarial perspective for its members, industry, policymakers and the public. This distinct perspective comes from the SOA as an association of actuaries, who have a rigorous formal education and direct experience as practitioners as they perform applied research. The SOA also welcomes the opportunity to partner with other organizations in our work where appropriate.

The SOA has a history of working with public policymakers and regulators in developing historical experience studies and projection techniques as well as individual reports on health care, retirement and other topics. The SOA's research is intended to aid the work of policymakers and regulators and follow certain core principles:

Objectivity: The SOA's research informs and provides analysis that can be relied upon by other individuals or organizations involved in public policy discussions. The SOA does not take advocacy positions or lobby specific policy proposals.

Quality: The SOA aspires to the highest ethical and quality standards in all of its research and analysis. Our research process is overseen by experienced actuaries and nonactuaries from a range of industry sectors and organizations. A rigorous peer-review process ensures the quality and integrity of our work.

Relevance: The SOA provides timely research on public policy issues. Our research advances actuarial knowledge while providing critical insights on key policy issues, and thereby provides value to stakeholders and decision makers.

Quantification: The SOA leverages the diverse skill sets of actuaries to provide research and findings that are driven by the best available data and methods. Actuaries use detailed modeling to analyze financial risk and provide distinct insight and quantification. Further, actuarial standards require transparency and the disclosure of the assumptions and analytic approach underlying the work.

Society of Actuaries
475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
www.SOA.org