ARTICLE IN PRESS

Insurance: Mathematics and Economics ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Insurance: Mathematics and Economics

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

Multi-population mortality models: A factor copula approach

Hua Chen [a,*], Richard MacMinn [b], Tao Sun [a]

[a] Temple University, United States

[b] Illinois State University, United States

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Multi-population mortality model
Factor copulas
Maximum entropy principle
Mortality/longevity risk pricing
Mortality/longevity risk hedging

ABSTRACT

Modeling mortality co-movements for multiple populations have significant implications for mortality/longevity risk management. A few two-population mortality models have been proposed to date. They are typically based on the assumption that the forecasted mortality experiences of two or more related populations converge in the long run. This assumption might be justified by the long-term mortality co-integration and thus be applicable to longevity risk modeling. However, it seems too strong to model the short-term mortality dependence. In this paper, we propose a two-stage procedure based on the time series analysis and a factor copula approach to model mortality dependence for multiple populations. In the first stage, we filter the mortality dynamics of each population using an ARMA–GARCH process with heavy-tailed innovations. In the second stage, we model the residual risk using a one-factor copula model that is widely applicable to high dimension data and very flexible in terms of model specification. We then illustrate how to use our mortality model and the maximum entropy approach for mortality risk pricing and hedging. Our model generates par spreads that are very close to the actual spreads of the Vita III mortality bond. We also propose a longevity trend bond and demonstrate how to use this bond to hedge residual longevity risk of an insurer with both annuity and life books of business.

1. Introduction

In recent years, a few mortality models with multiple populations have been proposed and developed. Modeling mortality dependence for two or more populations is still, however, in its infancy. Multi-population mortality models are typically structured assuming that the forecasted mortality experiences of two or more related populations are linked together and do not diverge over the long run. This assumption might be justified by the long-term mortality co-movements and thus be applicable to longevity risk modeling. It seems, however, too strong an assumption to use in modeling short-term mortality dependence. Here, we propose a copula-based multivariate model to capture mortality dependence for multiple populations. In the first stage, we filter the mortality dynamics of each population using the ARMA–GARCH process with heavy-tailed innovations. In the second stage, we model the residual risk using a one-factor copula model that is widely applicable to high dimension data. We then use our mortality model to price the Swiss Re Vita III bond and a longevity bond in the spirit of the Kortis bond to analyze the hedging problem of an insurer with both life and annuity books of business.

In order to have a thorough assessment of mortality/longevity risk, the correlation or co-integration between mortality improvements of different populations have to be evaluated. Cairns et al. (2011) provide a detailed discussion on why correlations among multiple populations must be taken into account. First, from the natural hedging perspective, a typical life insurer may want to hedge mortality risk from insured lives with longevity risk from annuitants. These two insured groups, however, may have different but correlated patterns of mortality experience. Second, sometimes the mortality data may not be available or sufficiently reliable for a small population due to a small number of deaths, a limited number of calendar years of data, age range or simply poor quality of data, thus causing highly inaccurate parameter estimates. Jointly modeling the small population and a larger linked population allows the small-population mortality forecasts to be consistent with those of the larger population. Third, with the rapid development of mortality securitizations, the payoffs of almost all mortality bonds are contingent on a weighted mortality index

* Correspondence to: Department of Risk, Insurance, and Health Management, Temple University, 1801 Liacouras Walk, 625 Alter Hall, Philadelphia, PA 19122, United States.

E-mail addresses: hchen@temple.edu (H. Chen), richard.macminn@ilstu.edu (R. MacMinn), tao.sun@temple.edu (T. Sun).

http://dx.doi.org/10.1016/j.insmatheco.2015.03.022
0167-6687/

based on multiple populations.[1] We need to understand mortality correlations in order to better evaluate the overall risk and the final payoff to investors. Last but not least, entities seeking to hedge their exposures to mortality/longevity risk using capital market instruments need to determine hedge ratios that minimize basis risk between their own population and the population associated with the hedging instruments.

A few two-population stochastic mortality models have been proposed. Li and Lee (2005) present an augmented common factor model for a group of populations, imposing a common mortality change by age but allowing each its own age pattern and level of mortality. Li and Hardy (2011) consider four extensions to the Lee–Carter model to incorporate mortality dependence. Dowd et al. (2011) develop a gravity mortality model for two-related but different sized populations, where gravity effects bring the state variables of the small population toward those of the large population in a manner consistent with biological reasonableness. A similar model was developed by Jarner and Kryger (2011). Cairns et al. (2011) introduce a general framework for two-population mortality modeling; they employ a mean-reverting process that permits different short-run trends in mortality rates but parallel long-run improvements. Zhou et al. (2013a) propose a two-population Lee–Carter model with transitory jumps under this framework. Yang and Wang (2013) and Zhou et al. (2013b) apply a co-integration analysis to investigate the long-run equilibrium in multi-country mortality data and use a vector error correction model (VECM) for mortality forecasts. A common feature of these models is that they are constructed in a way that mortality forecasts in different populations do not diverge in the long run. This assumption might be justified by the long-term mortality co-movements and so be applicable to longevity risk modeling, but it seems too strong to model the short-term mortality dependence.

Lin et al. (2013) develop a jump diffusion mortality model for multiple countries where mortality dependence comes from common jumps and the correlation between idiosyncratic risks. Their approach focuses on analyzing correlations among the raw mortality data but the raw data exhibit noise such as autocorrelation or volatility clustering that could possibly confound the dependence structure. Their approach also entails the use of normalized multivariate exponential tilting to price the risks and that implicitly assumes that the risks follow a Gaussian copula. We believe this is problematic since Gaussian copulas lack tail dependence and are therefore inadequate to model the joint mortality events. Our approach is related to Lin et al. (2013), but our model uses a two-stage multivariate analysis based on a factor copula approach which overcomes the problems noted here and has other advantages that we subsequently discuss.

Our mortality model has two stages. The first stage estimates the conditional distribution of mortality rates for each population. In the seminal work of Lee and Carter (1992), an ARIMA model is used to forecast the time varying mortality factor. They assume that the error terms are white noise with zero mean and a small constant volatility; that assumption of homoscedasticity, however, is not realistic. Lee and Miller (2001) argue that the observed logarithm of central death rates is quite variable and the volatility is time varying. Recently, GARCH-related models have been used to model mortality rates (see, e.g. Gao and Hu, 2009; Giacometti et al., 2012; Chai et al., 2013). Along with this line of research, we use an ARMA–GARCH model to fit the mortality data of each population in order to remove autocorrelation and conditional heteroskedasticity from mortality time series. Instead of modeling morality jumps explicitly, we assume that innovations follow a heavy-tailed distribution. This is supported by empirical evidence provided in Giacometti et al. (2009, 2012) and Wang et al. (2013). Giacometti et al. (2009) observe that for some age groups the Lee–Carter model with normal inverse Gaussian (NIG) innovations produces a dominant fit compared to the Gaussian one. Giacometti et al. (2012) find that an AR–ARCH model with Student-t innovations is more suited to Italian data. Wang et al. (2013) conclude that the Renshaw and Haberman (2006) model with non-Gaussian innovations generates better forecasts for England and Wales, France, and Italy.

The second stage of our model is designed to capture mortality dependence among residuals revealed by the first stage. A very common and intuitive way to capture the dependence is to use copulas. Copulas have been studied in both actuarial science and finance to examine dependencies among risks (Frees and Valdez, 1998; Embrechts et al., 2003). Particularly, in mortality studies copulas have been applied to model the bivariate survival function of the two lives of couples (see, e.g. Frees et al., 1996; Carriere, 2000; Shemyakin and Youn, 2006; Youn and Shemyakin, 1999, 2001; Denuit et al., 2001). Surprisingly, we find no prior research using copula models in the multi-population mortality analysis. In addition, the majority of aforementioned papers use copulas in the Elliptical or Archimedean family, which usually have only one or two parameters to characterize the dependence between all variables, and are thus quite restrictive when the number of variables increases. To overcome this drawback and increase the flexibility of our model, we adopt factor copula models proposed by Oh and Patton (2014).[2] Their model has some advantages. First, it is based on a simple linear, additive factor structure for the copula and is particularly attractive for high dimension applications. Second, the factor copula permits separate development of the marginal distributions and the copula model. Hence, this method allows the use of any existing model for the univariate analysis with the subsequent focus on the copula model. Third, the factor copula provides more flexibility of the model according to the number of variables and available data. To the best of our knowledge, our model is the first to use the factor copula to fit mortality data for multiple populations.[3]

Pricing mortality-linked derivatives is challenging in an incomplete market. With sparse market price data, some prevalent pricing methodologies, such as the arbitrage free pricing method (Cairns et al., 2006; Bauer et al., 2010), the Wang transform (Dowd et al., 2006; Denuit et al., 2007; Lin and Cox, 2008; Chen and Cox, 2009), or the Esscher transform (Chen et al., 2010; Li et al., 2010), require the user to make one or more subjective assumptions to derive the market prices of risk. The pricing process becomes even more difficult when multiple risks are involved. For this reason, Li (2010), Kogure and Kurachi (2010) and Li and Ng (2011) advocate the use of the maximum entropy approach that "does not strictly require the use of security prices to predict other security prices". This is particularly important in today's market where there are a limited number of market participants and only a few mortality

[1] There are only two exceptions, i.e., the Tartan mortality bond sponsored by Scottish Re in 2006 and the Atlas IX mortality bond sponsored by SCOR Re in 2013. Their payoffs depend on US mortality data only.

[2] A one-factor copula model was first introduced in Vasicek (1987) to evaluate loan loss distributions and Li (2000) applies the Gaussian copula to multi-name credit derivatives. The model is generalized in Andersen and Sidenius (2004), Hull and White (2004), and Laurent and Gregory (2005). Extending Hull and White (2004)'s work, Oh and Patton (2014) use factor copula models to capture the dependence among firms in the S&P 100. Their paper presents a "simulated method of moments" approach to accommodate this high dimensional data.

[3] At the same time this article was written, Wang et al. (2013) proposed a two-stage procedure which is very similar to our model. They use an ARMA process in the first stage and time-varying copulas in the Elliptical or Archimedean family in the second stage.