Insurance: Mathematics and Economics 68 (2016) 45–60

Contents lists available at ScienceDirect

Insurance: Mathematics and Economics

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

Statistical emulators for pricing and hedging longevity risk products

J. Risk*, M. Ludkovski

Department of Statistics & Applied Probability, University of California, Santa Barbara CA 93106-3110, United States

ARTICLE INFO

Article history:
Received September 2015
Received in revised form
February 2016
Accepted 24 February 2016
Available online 9 March 2016

Keywords:
Statistical emulation
Longevity risk
Life annuities
Valuation of mortality-contingent claims
Kriging
Gaussian processes

ABSTRACT

We propose the use of statistical emulators for the purpose of analyzing mortality-linked contracts in stochastic mortality models. Such models typically require (nested) evaluation of expected values of nonlinear functionals of multi-dimensional stochastic processes. Except in the simplest cases, no closed-form expressions are available, necessitating numerical approximation. To complement various analytic approximations, we advocate the use of modern statistical tools from machine learning to generate a flexible, non-parametric surrogate for the true mappings. This method allows performance guarantees regarding approximation accuracy and removes the need for nested simulation. We illustrate our approach with case studies involving (i) a Lee–Carter model with mortality shocks; (ii) index-based static hedging with longevity basis risk; (iii) a Cairns–Blake–Dowd stochastic survival probability model; (iv) variable annuities under stochastic interest rate and mortality.

1. Introduction

Longevity risk has emerged as a key research topic in the past two decades. Since the seminal work of Lee and Carter (1992) there has been a particular interest in building stochastic models of mortality. Stochastic mortality allows for generation of a range of future longevity forecasts, and permits the modeler to pinpoint sources of randomness, so as to better quantify respective risk. Longevity modeling calls for a marriage between the statistical problem of calibration, i.e. fitting to past mortality data, and the financial problem of pricing and hedging future longevity risk. At its core, the latter problem reduces to computing expected values of certain functionals of the underlying stochastic processes. For example, the survival probability for t years for an individual currently aged x can be expressed as a functional

$$P(t,x) = \mathbb{E}\left[\exp\left(-\int_0^t \mu(s, x+s)\,ds\right)\right], \tag{1}$$

where $\mu(s, x+s)$ is the force of mortality at date s for an individual aged $x+s$. In the stochastic mortality paradigm $\mu(s, x+s)$ is random for $s > 0$, and so one is necessarily confronted with the need to evaluate the corresponding expectations on the right-hand-side of (1).

The past decade has witnessed a strong trend towards complexity in both components of (1). On the one hand, driven by the desire to provide faithful fits (and forecasts) to existing mortality data, increasingly complex mortality models for $\mu(t,x)$ have been proposed. The latest generation of models feature multi-dimensional, nonlinear stochastic state processes driving $\mu(\cdot, x)$, see e.g. Cairns et al. (2009), Li et al. (2009), Lin et al. (2013), Barrieu et al. (2012) and Fushimi and Kogure (2014). These models are effective at calibration and emitting informative forecasts, but lack tractability in terms of closed-form formulas. On the other hand, sophisticated insurance products, such as variable annuities or longevity swap derivatives make valuation and hedging highly nontrivial, and typically call for numerical approaches, as closed-form formulas are not available, see e.g. Bacinello et al. (2011) and Qian et al. (2010). Taken together, pricing of mortality-linked contracts becomes a complex system, feeding multi-dimensional stochastic inputs through a "black box" that eventually outputs net present value of the claim.

These developments have created a tension between the complexity of mortality models that do not admit explicit computations and the need to price, hedge and risk manage complicated contracts based on such models. Due to this challenge, there remains a gap between the academic mortality modeling and the implemented models by the longevity risk practitioners. Because the aforementioned valuation black box is analytically intractable, there is a growing reliance on Monte Carlo simulation tools, which in turn is accompanied by exploding computational needs. For example, many emerging problems require *nested simulations* which can easily take days to complete. Similarly, many portfolios contain millions of heterogeneous products (see, e.g. Gan and Lin, 2015) that must be accurately priced and managed.

* Corresponding author.
E-mail address: risk@pstat.ucsb.edu (J. Risk).

http://dx.doi.org/10.1016/j.insmatheco.2016.02.006
0167-6687/

In this article we propose to apply modern statistical methods to address this issue. Our approach is to bridge between the mortality modeling and the desired pricing/hedging needs through an intermediate *statistical emulator*. The emulator provides a computationally efficient, high-fidelity surrogate to the actual mortality model. Moreover, the emulator converts a calibrated opaque mortality model into a user-friendly valuation "app". The resulting toolbox allows a plug-and-play strategy, so that the end user who is in charge of pricing/risk-management can straightforwardly swap one mortality model for another, or one set of mortality parameters for an alternative. This modular approach allows a flexible solution to robustify the model-based longevity risk by facilitating comparisons of different longevity dynamics and different assumptions.

Use of *emulators* is a natural solution to handle complex underlying stochastic simulators and has become commonplace in the simulation and machine learning communities (Santner et al., 2003; Kleijnen, 2007). Below we propose to apply such statistical learning within the novel context of insurance applications. In contrast to traditional (generalized) linear models, emulation calls for fully nonparametric models, which are less familiar to actuaries. To fix ideas, in this article we pursue the problem of pricing/hedging vanilla life annuities, a foundational task in life insurance and pension plan management. Except in the simplest settings, there are no explicit formulas for annuity values and consequently approximation techniques are already commonplace. Looking more broadly, our method would also be pertinent for computing risk measures, such as Expected Shortfall for longevity products, and in other actuarial contexts, see Section 8.

The paper is organized as follows: In Section 2 we introduce the emulation problem and review the mathematical framework of stochastic mortality. Section 3 discusses the construction of emulators, including spline and kriging surrogates, as well as generation of training designs and simulation budgeting. The second half of the paper then presents four extended case studies on several stochastic mortality models that have been put forth in the literature. In Section 4 we examine a Lee–Carter model with mortality shocks that was proposed by Chen and Cox (2009); Section 5 studies approximation of hedge portfolio values in a two-population model based on the recent work by Cairns et al. (2014). Section 6 considers valuation of deferred annuities under a Cairns–Blake–Dowd (CBD) (Cairns et al., 2006) mortality framework. Lastly, in Section 7 we consider variable annuities and their future distributions for risk measure analysis, using stochastic interest rate and the Lee–Carter framework.

2. Emulation objective

We consider a stochastic system with Markov state process $Z = (Z(t))$. Throughout the paper we will identify Z with the underlying stochastic *mortality factors*. In Section 2.2 we review some of the existing such models and explicit the respective structure of Z. Typically, Z is a multivariate stochastic process based on either a stochastic differential equation or time-series ARIMA frameworks. For example, Z may be of diffusion-type or an auto-regressive process.

In the inference step, the dynamics of Z are calibrated to past mortality data that reflect as closely as possible the population of interest. In the ensuing valuation step, the modeler seeks to evaluate certain quantities related to a functional $F(T, Z(\cdot))$ looking into the future. Here F maps the stochastic factors to the present value of a life insurance product at a future date T, or alternatively the actuarially fair value of a deferred contract, common in longevity risk, that starts at T. Our notation furthermore indicates that F potentially depends on the whole path $\{Z(t), t \geq T\}$, such as

$$F(T, Z(\cdot)) = \exp\Big(-\sum_{t=T}^{\infty} h(Z(t))\Big), \tag{2}$$

for some $h(z)$. Given F, a common aim is to compute its expected value based on the initial data at $t = 0$,

$$\mathbb{E}\left[F(T, Z(\cdot)) \mid Z(0)\right]. \tag{3}$$

Another key problem is to evaluate the quantile $q(\alpha; F(T, Z(\cdot)))$, eg. the Value-at-Risk at level α of F. Other quantities of interest in actuarial applications include the Expected Shortfall of F, $\mathbb{E}[F(T, Z(\cdot)) \mid F(T, Z(\cdot)) \leq q(\alpha; F(T, Z(\cdot))), Z(0)]$ and the correlation between two functionals, $Corr(F_1(T, Z(\cdot)), F_2(T, Z(\cdot))|Z(0))$.

Our initial focus is on (3) which is a fundamental quantity in pricing/hedging problems. When $T > 0$, the evaluation of (3) can be broken into two steps, namely first we evaluate

$$f(z) \doteq \mathbb{E}[F(T, Z(\cdot))|Z(T) = z], \tag{4}$$

and then use the Markov property of Z to carry out an outer average,

$$\mathbb{E}[F(T, Z(\cdot))|Z(0)] = \int_{\mathbb{R}^d} f(z) p_T(z|Z(0)) dz,$$

where $p_T(z'|z) = \mathbb{P}(Z(T) = z'|Z(0) = z)$ is the transition density of Z over $[0, T]$. In addition to computing expected values from point of view of $t = 0$, computation of $f(z)$ is also necessary for analyzing the distribution of future loss in terms of underlying risk factors, e.g. for risk measurement purposes.

Crucially, because the form of $F(T, Z(\cdot))$ is nontrivial, we shall assume that $f(z)$ is not available explicitly, and there is no simple way to describe its functional form. However, since $f(z)$ is a conditional expectation, it can be sampled using a simulator, i.e. the modeler has access to an engine that can generate independent, identically distributed samples $F(T, Z^{(n)}(\cdot))$, $n = 1, \ldots,$ given $Z(0)$. However this simulator is assumed to be expensive, implying that computational efficiency is desired in using it.

Given an initial state $Z(0)$, a naive Monte Carlo approach to evaluate (3) is based on nested simulation. First, the outer integral over $p_T(z|Z(0))$ is replaced by an empirical average of (4) across $m = 1, \ldots, N_{out}$ draws $z^{(m)} \sim Z(T)|Z(0)$,

$$\mathbb{E}[F(T, Z(\cdot))|Z(0)] \simeq \frac{1}{N_{out}} \sum_{m=1}^{N_{out}} f(z^{(m)}). \tag{5}$$

Second, for each $z^{(m)}$ the corresponding inner expected value $f(z^{(m)})$ is further approximated via

$$f(z^{(m)}) \simeq \frac{1}{N_{in}} \sum_{n=1}^{N_{in}} F(T, z^{(m),n}(\cdot)), \quad m = 1, \ldots, N_{out}, \tag{6}$$

where $z^{(m),n}(t), t \geq T$ are N_{in} independent trajectories of Z with a fixed starting point $z^{(m),n}(T) = z^{(m)}$. This nested approach offers an unbiased but expensive estimate. Indeed, the total simulation budget is $\mathcal{O}(N_{out} \cdot N_{in})$ (where the usual big-O notation $h(x) = \mathcal{O}(x)$ means that $h(\cdot)$ is asymptotically linear in x as $x \to \infty$) which can be computationally intensive—for example a budget of 1000 at each sub-step requires 10^6 total simulations. As stochastic mortality models become more complex, models with $d = 3, 4, 5+$ factors are frequently proposed, and efficiency issues become central to the ability of evaluating (3) tractably.

For this reason, it is desirable to construct more frugal schemes for approximating (3). The main idea is to replace the inner step of repeatedly evaluating $f(z)$ (possibly for some very similar values of z) with a simpler alternative. One strategy is to construct