SOCIETY OF ACTUARIES

Article from

Reinsurance News

November 2016

Issue 86

The Economic Value of Reinsurance

By Kai Kaufhold

Why buy reinsurance? To members of the Reinsurance Section that will sound like a rhetorical question. However, as sellers or buyers of reinsurance coverage, it makes sense to ask ourselves the "why?" question from time to time. What do life companies really need reinsurance for and how do they measure the value-for-money that they get from their reinsurance premium dollars?

REINSURANCE VALUE-ADDED

Reinsurers add value to the business of their clients in a number of ways, ranging from underwriting manuals, underwriting support and training, to actuarial support in product development and other areas. However, the original and most important purpose of reinsurance is risk management. Life companies buy reinsurance to limit their exposure to insurance risks. Therefore, the question of how much reinsurance to purchase is usually centered on a company's risk appetite: "What is the maximum loss per life that we are willing to sustain?", or more generally, "How does reinsurance help to reduce the volatility of our company's earnings?"

Ideally, a life company will be able to answer these questions within its enterprise risk management framework by defining its risk appetite in terms of economic capital and explicitly reflecting the impact of reinsurance in its economic capital model. In a research project[1] sponsored by the SOA's Reinsurance Section, the Financial Reporting Section and the Committee on Life Insurance Research, we set out to design a framework within which the risk-reduction impact of reinsurance can be quantified and expressed in terms of changes to a company's reserves and its economic capital. The key challenge was that we needed to develop a method by which we could ensure that the reserve margins and capital buffers reflected the riskiness of the business accurately, so that we could measure the impact of reinsurance. To accomplish that, we used a statistical tool called Parametric Survival Models to create portfolio-specific assumptions to calculate best-estimate reserves and carried out Monte-Carlo simulations to model the uncertainty associated with those reserves.

SURVIVAL MODELS

This statistical technique has been applied successfully by engineers and statisticians for decades. It is also widely used in the U.K. longevity market to assess the longevity risk associated with pension buy-ins, pension buy-outs and longevity swaps which are the different forms of longevity reinsurance by which pension plans and insurers in the U.K. transfer longevity risk to reinsurers. The beauty of using a statistical model is that it can be used to create portfolio-specific mortality (and lapse) assumptions and one can also measure the estimation error associated with the results. Importing this method to the life insurance practice area, our case study demonstrates how useful it is for life insurance valuation and capital management. Life insurers and reinsurers can use the framework outlined here to quantify the value of reinsurance in terms of the reduction in cost of capital which it achieves.

MULTI-DECREMENT CASE STUDY

We analyzed the mortality and lapse experience of a U.S. life insurer's term life business and created a survival model for the mortality behavior and lapses within this portfolio from experience data at the level of the individual. The study contained roughly 1.5 million records, of which around 15,000 related to deaths and 500,000 related to lapses and surrenders. The total face amount for all policies was around $600 billion and the study covered more than 10 million life-years of exposure. You will be thinking that this is actually a massive experience base to which not every company has access, and you would be correct. In order to test the sensitivity of results, we also applied the framework to random subsets of 25 percent, 50 percent and 75 percent of the total data and were able to fit models and run the simulations just as effectively on the smaller samples as on the entire block.

The process for building a set of survival models for the decrements death and lapse is simple:

1. First, pick parametric laws[2] in continuous time which match the general shape of the mortality and lapse experience in aggregate. These will be different for the two decrements.

2. Then, estimate the parameters of the hazard rate functions by maximum likelihood method.

3. Next, use the baseline survival model to identify different risk factors and quantify their impact.

As can be seen in Figure 1, the crude mortality hazard rates largely follow a log-linear pattern between the ages of 35 years and 75 years. This corresponds to the well-known Gompertz law for the force of mortality:

$$\mu_x = e^{\alpha+\beta x} \Leftrightarrow \ln(\mu_x) = \alpha + \beta x$$

with intercept α and slope β. In step 2 we estimate the parameters α and β and then define risk classes which have different adjustments to the intercept and slope of the baseline mortality law.

IDENTIFYING RISK FACTORS

For both the force of mortality and the lapse hazard model, we include a number of different risk factors which influence the mortality or lapse outcomes. It is important that we include as many statistically significant risk factors as possible to ensure that we do not underestimate the estimation error. Take sex as an example for a risk factor: fitting a model for aggregate unisex mortality likely gives a smoother fit and smaller estimation error than if we fit curves for males and females separately. The apparent better fit in aggregate, however, is useless, because it introduces distribution risk. While the unisex table might work for the exact business mix of policies within the experience data, the sex distribution may shift due to changing new business sales or simply because the men lapse and die at higher rates than the insured women. The same problem arises with any set of risk factors which have a significant impact on mortality. For our case study, we found that we had to differentiate by sex, duration, smoking status, underwriting class, product type (10-year term, 20-year term, etc.), face amount band, and whether a policy was rated at issue or was accepted as standard.

FIGURE 1: FORCE OF MORTALITY—CRUDE MORTALITY HAZARD RATES

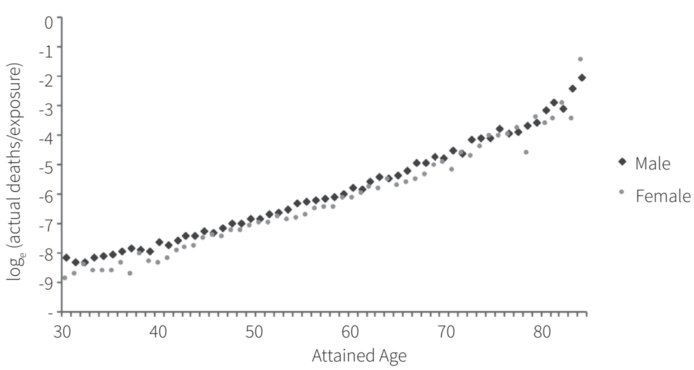

Source: Kaufhold and Lennartz (2016), aggregate count of deaths over entire study period within each one-year age bracket divided by the time for which the lives in each age bracket were exposed to risk.

MEASURING UNCERTAINTY

We capture the variability of results by applying a stochastic Monte-Carlo simulation in two steps. First, we take the parametric model that describes the best-estimate mortality and lapse behavior of the portfolio and give the parameters a little shake. In other words, we randomly perturb the parameter set in a way that is consistent with the experience data. The perturbed parameters then describe a mortality and lapse behavior that is

a little different from the best-estimate, but that also could have happened this way. So, we have created an alternative scenario consistent with the experience data.

Within this perturbed scenario, we can calculate the survival curves and stochastically determine whether an insured life survived until the end of the level term period, whether they died or whether their policy lapsed. To do so, we simply draw a random number between 0 and 1, and then use the survival curve to check which remaining life-time this randomly drawn probability corresponds to. Since it is a two-decrement model, we need two random probabilities to get the corresponding random times until death and until lapse. If both are longer than the remaining time until the end of the level term period, we have a survivor. Otherwise, we count the event as a lapse or a death, whichever happened first. By going through the entire list of lives, and letting them randomly survive, lapse or die, we can add up what the total present value of claims would have been in our first perturbed scenario.

These steps are repeated many times to get a probabilistic distribution of total claims that reflects the uncertainty associated with mortality and lapse. To this distribution, we can then apply different reinsurance structures and study their impact.

First dollar quota share reinsurance has no impact on the riskiness of the retained business, which has the same profile as the gross business before reinsurance. Excess of retention reinsurance, however, changes the risk profile of the claims occurring in any given period.

OVERALL RESULTS

Applying the method described above to our term life portfolio in the case study, we found a number of interesting things:

1. The margin required for reserves at a certain confidence level depended on the business mix. It was different for the different products, with 10-year term requiring the greatest mark-up and longer-term products requiring a lower mark-up.
2. Different portfolio sizes required different levels of prudential provisions with smaller blocks needing a greater reserve buffer. This is totally unsurprising, but a good check to show that the method makes sense.
3. What did surprise us at first was that reserve margins were hardly affected by reinsurance. We expected to see that reinsuring large policies and thereby reducing the risk would change the risk profile of the business so much that the reserves on the retained portion would be a lot lower, relatively speaking, than the reserves without reinsurance. However, it turns out that reserves that reflect the present value of claims and premiums over an extended period of time are a lot less sensitive to the life insurer's level of retention, because volatile annual results are smoothed over time.
4. Excess reinsurance has a strong impact on the volatility of annual earnings, and therefore affects solvency capital requirements that serve the purpose of ensuring that the life company has enough capital to withstand short-term volatility. If the company has a large portion of short-term business, the reserve margins will be greater and will also be more sensitive to reinsurance.

FIGURE 2: OPTIMIZING RISK RETENTION—OPTIMAL RETENTION

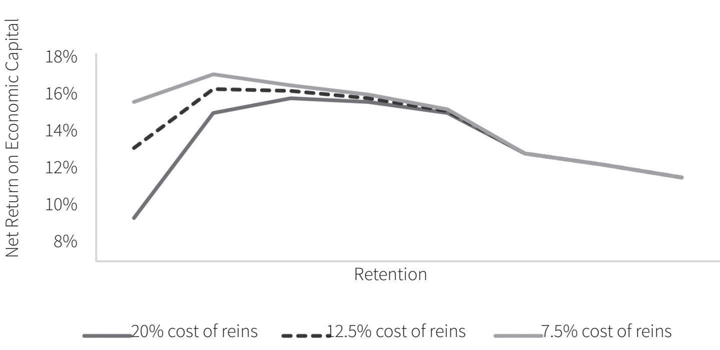

Source: Kaufhold and Lennartz (2016): Gross return on economic capital 12%, reinsurance increases return on EC, offset by increasing cost of reinsurance. Three scenarios show that optimal retention depends on the cost of reinsurance.

The purpose of the research project was to investigate the impact of reinsurance under modern reserving and solvency capital regimes.

CONCLUSION

The purpose of the research project was to investigate the impact of reinsurance under modern reserving and solvency capital regimes. In this respect, the key result was that reinsurance has a greater influence on capital levels than on reserve levels, and that reinsurance can be used to optimize the return on economic capital that properly reflects the riskiness of the business. An important byproduct of this project is that we had to develop a method for setting reserves that truly reflect the uncertainty associated with setting the mortality and lapse assumptions (estimation error), and the volatility of the business itself (adverse deviation). Our results showed that reserve levels will vary depending on the business mix of the company, and that it is therefore important for life insurers to carry out their own analysis to derive company-specific mortality and lapse assumptions and quantify explicit margins for uncertainty. The method is applicable for small- to medium-sized life companies, just as it is to large life insurers, and can be applied to any kind of insurance risk.

To find out more, please check out the SOA research report "Optimizing Risk Retention, Quantitative Retention Management for Life Insurers" available at *www.soa.org/Files/Research/Projects/research-2016-quantitative-retention.pdf.* The authors will be more than happy to answer any questions you may have regarding the case study and its results. Just drop us an email at *kai.kaufhold@adreservices.com*. ■

Kai Kaufhold, managing director, Ad Res, actuarial consulting firm in Cologne, Germany, is a member of the organizing committee of the Living to 100 Symposium and has studied longevity and mortality risk for more than two decades.

ENDNOTES

1 The SOA Research Report "Optimizing Risk Retention" can be accessed at https://www.soa.org/Files/Research/Projects/research-2016-quantitative-retention.pdf

2 The method is called Survival Model, because we estimate the parameters of the mortality law by maximizing the likelihood of future lifetimes (_t_i^)p_(x_i) (μ_(x_i+t_i))^(d_i) for each individual i , where (_t_i^)p_(x_i) is the probability of an individual aged x_i surviving t_i years, μ_(x_i) is the individual's force of mortality (a.k.a. mortality hazard rate) and d_i is a status variable which equals 1 if the individual has experienced death (or whichever decrement is being analysed) and 0 otherwise.