Article from

Reinsurance News

March 2019
Issue 116

DI for Dinner

Measuring Disability Income Insurance Volatility Using Survival Models

By Kai Kaufhold

A year ago, we started our SOA *Reinsurance News* series on predictive modeling using survival models with a casual introduction to survival models over lunch. In the second article, for afternoon tea, so to speak, we saw survival models being successfully applied to persistency within a book of life insurance business and discovered that this method also lends itself well to illustrating the drivers behind differences in persistency or mortality. In this third and final article of the series, we will wrap it up with a wholesome five-decrement dinner. My aim is to demonstrate that survival models are useful for predicting the outcomes of insurance business under multiple decrements (five, no less!) and then take it even one step further and show that we can use the method to quantify the volatility of a portfolio.

Most life insurance products combine multiple competing risks, such as death and lapse, or death, disability and lapse. For traditional actuarial models, this poses substantial challenges, because the actuary must make assumptions about the distribution of events during discrete time periods. By contrast, parametric survival models in continuous time entirely avoid that difficulty, because in each instance, each risk is acting simultaneously to all others.

Let's look at a case study of disability income risk. The challenge was not just to predict disability claims but to measure their volatility, to quantify by how much the predicted best estimate was likely to be wrong. To do this, I teamed up with a reinsurer, created a statistical model that described disability income insurance risk, and then used that model within a

Figure 1
Crude Hazard Rates for Active Death Against Attained Age and Policy Duration

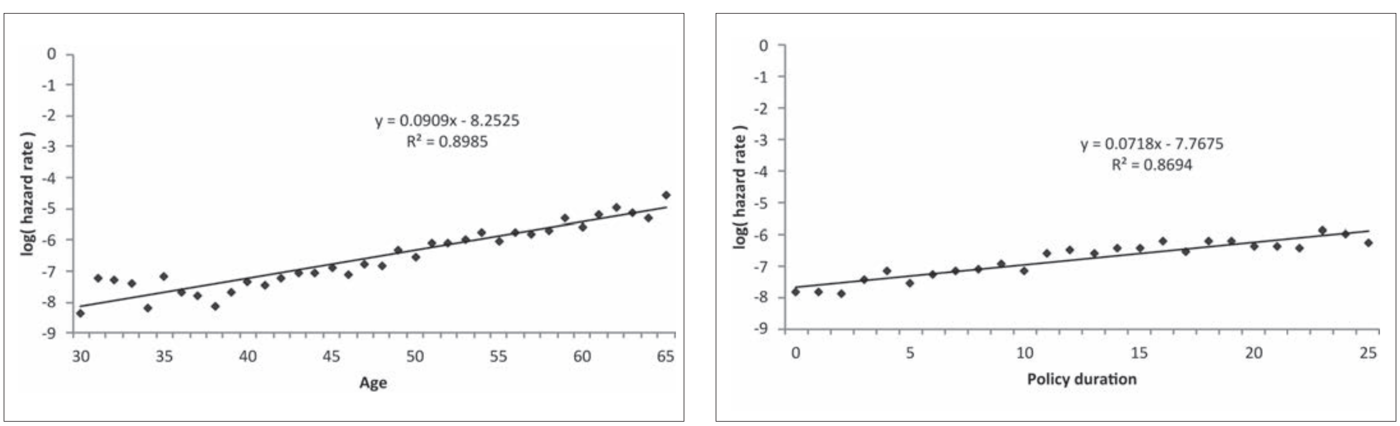

Source: Own calculations of time exposed to risk by age group and policy duration since inception.

Figure 2
Crude Hazard Rates for Lapse Against Age and Policy Year

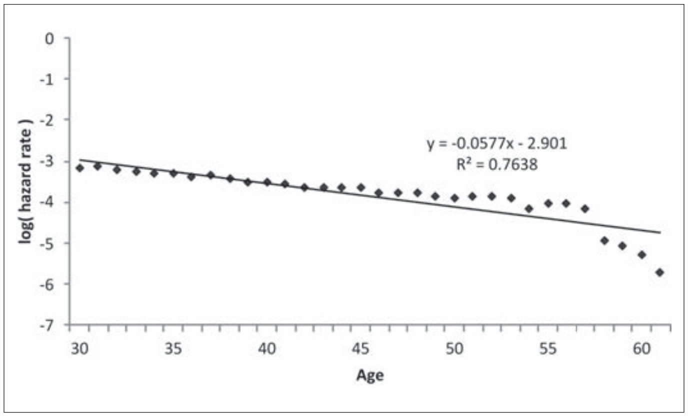

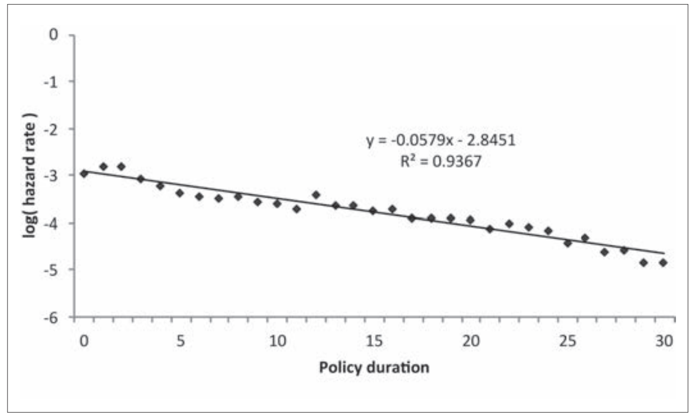

Source: Own calculations of time exposed to risk by age group and policy duration since inception.

Figure 3
Crude Hazard Rates for Incidence of Disability Against Age and Policy Year

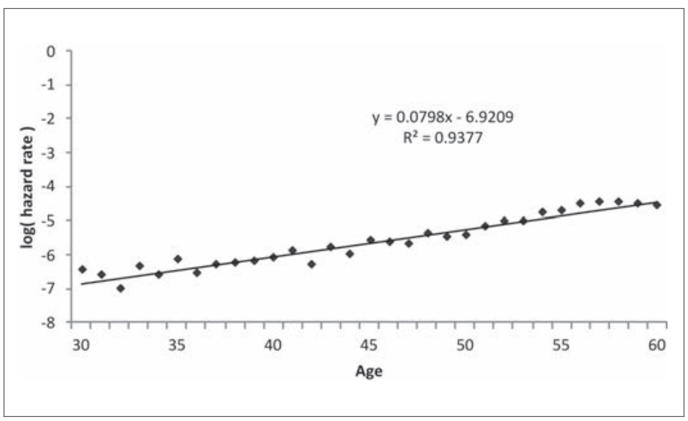

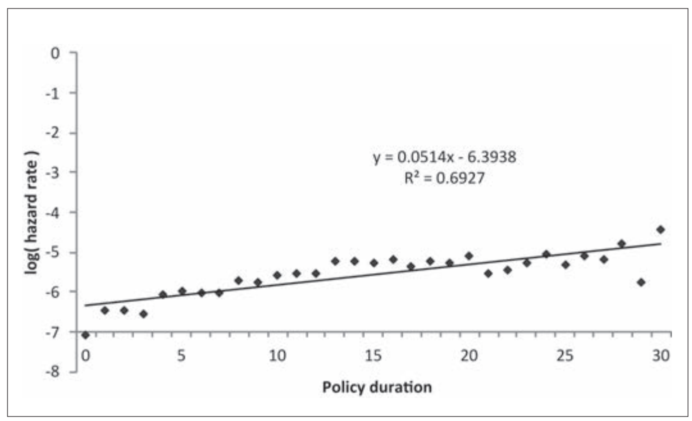

Source: Own calculations of time exposed to risk by age group and policy duration since inception.

Monte-Carlo simulation to measure volatility. We applied our method first to German disability income business (Berufsunfähigkeitsversicherung) and then to Australian individual disability income business.

Before designing the parametric survival models that we used for our predictive analysis, we reviewed the raw data to get an idea of the basic shape of the crude hazard rates. An active policyholder can die, lapse the policy or become disabled. Figures 1–5 include two charts—the left chart showing the age dependence of the respective hazard and the right-hand chart showing the hazard rates against policy duration. Note that in all cases, we have taken the logarithm of the crude hazard rates. I have included trend lines to indicate that on a logarithmic scale, a linear model should reasonably reflect the risk.

Naturally, the most important decrement is the incidence of disability, shown for the German case study in Figure 3, which increases with age and policy duration.

Switching to disabled lives, there are two ways benefit payments can terminate, except reaching the end of the benefit period, of course: by the disabled person either dying or going back to work. Mortality increases more moderately by age for disabled lives than for active lives. We also see in the right-hand chart of Figure 4 that disabled mortality is highest just after the disability occurs and decreases over time.

And finally, the chance that disabled people return to work declines both with age and with time since the disability occurred, as can be seen in Figure 5.

Figure 4
Crude Hazard Rates for Disabled Deaths Against Age and Duration of Disability

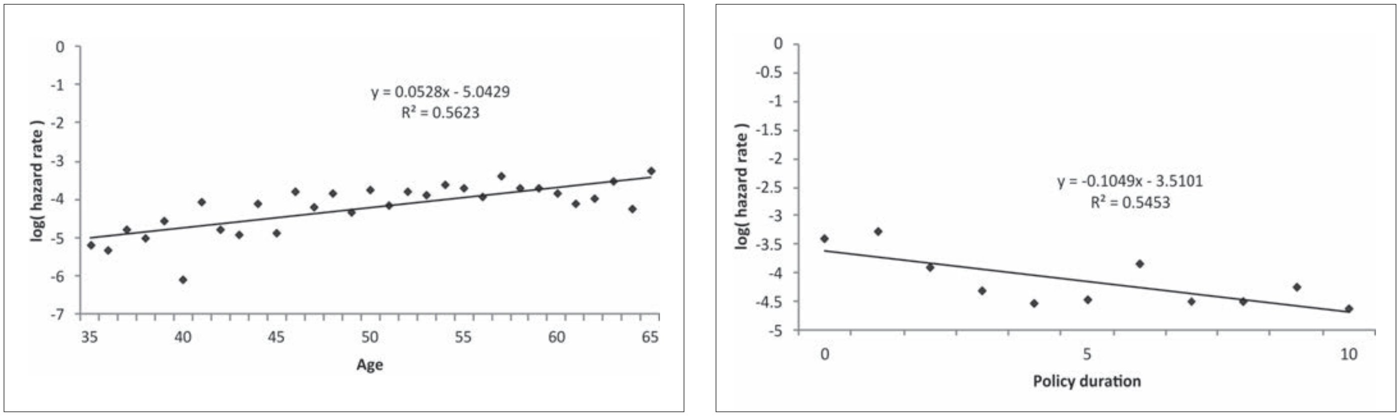

Source: Own calculations of time exposed to risk by age group and time since the date of disability.

Figure 5
Crude Hazard Rates for Reactivation Against Age and Duration of Disability

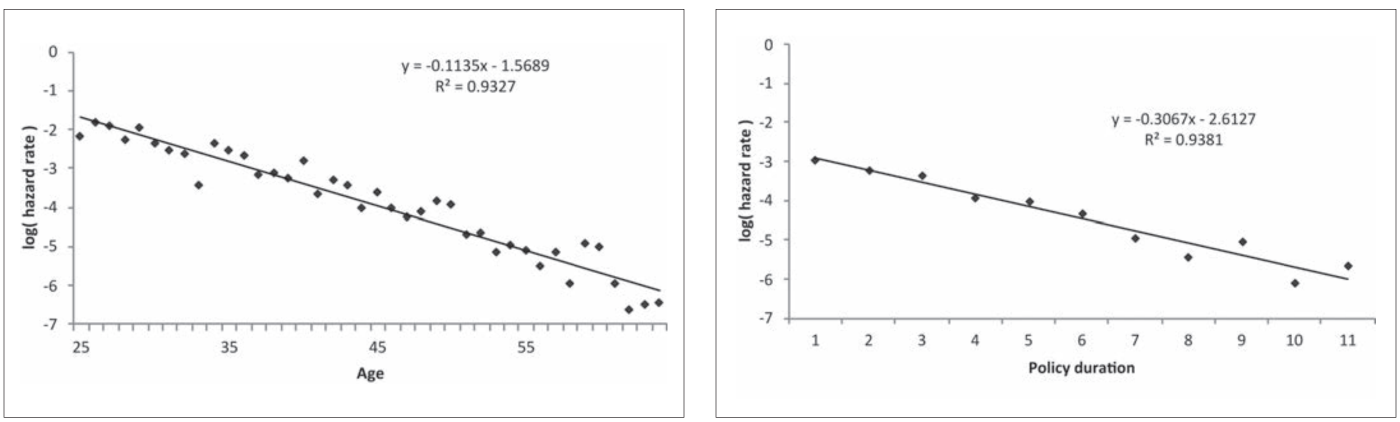

Source: Own calculations of time exposed to risk by age group and time since the date of disability.

An important feature of the disabled life models is that we have much fewer data, because we are limited to disabled lives. This explains the greater variability of results and relatively low scores for the R^2 statistic on disabled deaths. By contrast, the reactivation rates are much more tightly bunched around the log-linear trend line.

Having identified the basic shape for our parametric hazard rate function as a simple linear exponential function, equivalent to the Gompertz law of mortality, we can now use the maximum likelihood method to fit parameters and identify additional risk factors that might have an impact on the respective hazards, just as we saw in the previous article on survival models in the July issue of *Reinsurance News*. In our case study for the German disability portfolio, we limited the models to include only age, duration and gender as risk factors and thereby graduated a set of assumptions that was directly comparable to the German industry tables for disability risk.

We now use the five hazard rate functions for the five different decrements to predict the financial outcome of a disability income insurance portfolio and then run a Monte Carlo simulation, in which we go through the entire portfolio and simulate the outcome for each person. First, we "roll three dice"[1] to find out when each person lapses, dies or becomes disabled. All we need to do is check which happened first. If the first event predicted to happen is disability, If we acknowledge that incidence rates, lapse, death and termination rates are not deterministic, then we have to accept that the overall risk is 50 times higher.

we roll two more dice to decide whether the disabled person dies, goes back to work or remains disabled to the end of the benefit period.

The reason we can do this so easily is that we have analytical (continuous) expressions for the different hazard rates and thus the survival curves, which give us the cumulative probability of an event. By inverting the survival curves, we can use a randomly picked probability of, say, becoming disabled to calculate exactly when that event will take place, as illustrated in Figure 6.

Figure 6
Illustration of Disability Risk Simulation—Idiosyncratic Risk

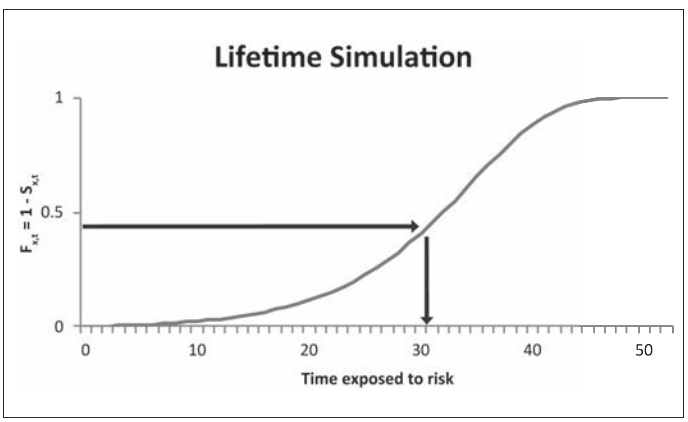

Source: Own calculations of cumulative risk of becoming disabled.

If we go through this process of rolling the dice for each life in the portfolio many times, we will get a distribution of disability claims that reflects idiosyncratic risk (i.e., the fact that disability, death, lapse and reactivation are all random events that will affect different individuals differently). This risk is often also referred to as process risk. An example of such a distribution generated for our German book of disability income risks is given in Figure 7.

Figure 7
Simulated Distribution of Disability Benefits—Idiosyncratic Risk

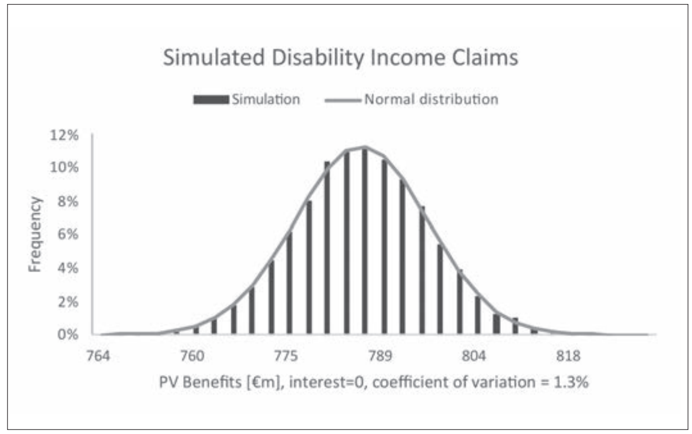

Source: Own calculations of random time of disability and contingent duration of disability claim for a portfolio of 140,000 lives. Monte-Carlo simulation with 50,000 runs. Coefficient of variation: 1.3 percent.

Within this same simulation framework, we can also incorporate estimation error by replacing the fitted parameters with a set of random parameters. Let's say a parameter has a maximum likelihood estimate that comes with a high standard error. Then the randomly "perturbed" new parameter should be farther away from the best-estimate parameter than for a parameter with a small standard error.[2] Figure 8 illustrates misestimation risk influencing the simulated survival curves.

Figure 8
Illustration of Disability Risk Simulation—Misestimation Risk

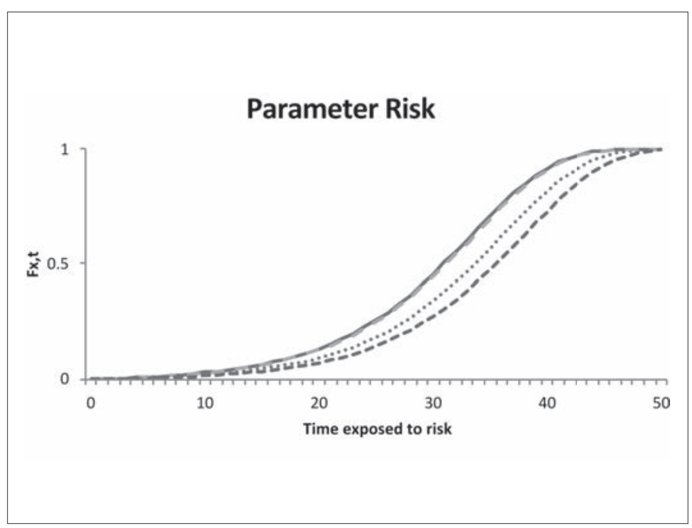

Source: Own calculations of cumulative risk of becoming disabled. Different cumulative distribution functions correspond to different sets of parameters, which have been randomly displaced from the best estimate in a way consistent with the experience data.

Rerunning the simulation 50,000 times, including misestimation risk, gives us a distribution as shown in Figure 9. To put it mildly, this distribution no longer has anything to do with a nicely behaved normal distribution. The simulated disability claims are heavily left-skewed, and volatility is 50 times higher than for the simulation without misestimation risk.

Figure 9
Simulated Distribution of German Disability Benefits—Idiosyncratic and Misestimation Risk

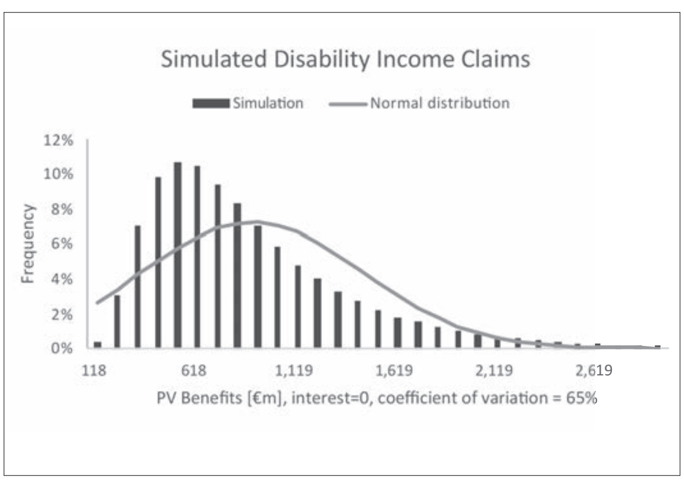

Source: Own calculations of random time of disability and contingent duration of disability claim for a portfolio of 140,000 lives. Monte-Carlo simulation with 50,000 runs. Coefficient of variation: 65 percent.

What does this mean? In simple terms, disability income risk is perfectly well-behaved as long as we can assume that we know the disability incidence rates as well as all other decrements exactly. Then the law of large numbers applies, and it is possible to predict disability claims quite accurately (standard deviation < 2 percent of mean). If we acknowledge that incidence rates, lapse, death and termination rates are not deterministic, then we have to accept that the overall risk is 50 times higher.

My reinsurance colleagues and I wondered whether this phenomenon applied only to German disability business (say it with me: "*Berufsunfähigkeitsversicherung*," aka "BU") or whether disability income risk showed the same profile in other countries. We carried out the same analysis for a portfolio of Australian individual disability income insurance business and were able to confirm that the distribution of claims shows the same pattern, if not worse.

Figure 10
Simulated Distribution of Australian Disability Benefits—Idiosyncratic and Misestimation Risk

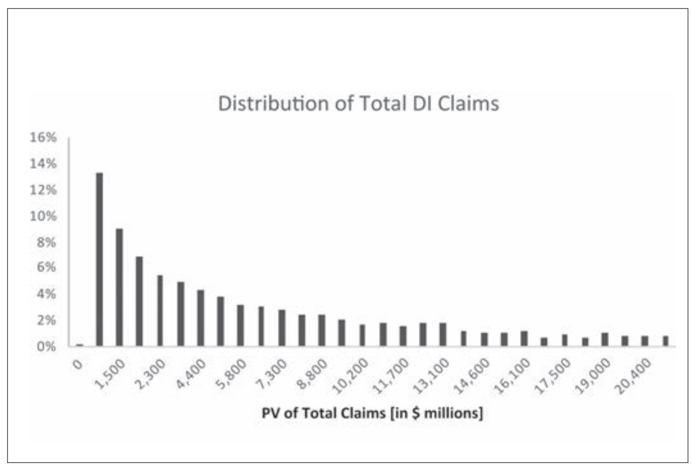

Source: Own calculations of random time of disability and contingent duration of disability claim for a portfolio of 111,000 lives. Monte-Carlo simulation with 5,000 runs. Coefficient of variation: 143 percent.

The total claims distribution for Australian disability income risks shown in Figure 10 is even more left-skewed than the German BU results and has a coefficient of variation that is twice as high. There are several reasons things would be worse for the Australian portfolio that we analyzed. For example, the Australian portfolio showed greater heterogeneity between short-term and long-term disability benefits and different occupational classes. We were also able to measure annual lapse spikes in the Australian DI portfolio that may have led to anti-selective lapses, which would not be present in the German-level premium disability business to this extent.

From a reinsurer's perspective, all this shocking news about the riskiness of disability income insurance business is, of course, scary but at the same time is the best-possible sales argument. A single life insurance company has no way of handling a large portfolio of disability income risk on its own. It needs the support of a well-diversified, financially strong reinsurance partner who can withstand the potentially catastrophic results of disability income business. Our results prove that disability risk in and of itself is frightfully difficult to get right, even if you make no mistakes.

See you in Las Vegas[3]—if you are interested in rolling some dice or discussing DI over dinner. ■

Kai Kaufhold, Aktuar (DAV), is a partner, lead for Longevity and Prediction Consulting with NMG Consulting. He can be contacted at Kai.*Kaufhold@nmg-group.com*.

ENDNOTES

1 In this example, rolling the dice symbolizes drawing uniformly distributed random numbers between zero and one.

2 For the interested practitioner, S.J. Richards gives an in-depth introduction to mis-estimation risk in his paper: Mis-estimation risk: measurement and impact, *British Actuarial Journal*, 21(3), pages 429–475 (including discussion).

3 ReFocus 2019 will take place March 10–13, 2019, in Las Vegas and is jointly sponsored by ACLI and SOA.