Article from:

Product Matters

June 2015 – Issue 92

An Efficient Statistical Estimator for Validating Life Expectancy Reports in the Life Settlements Market

by Gordon Gillespie

Gordon Gillespie is an actuarial consultant based in Berlin, Germany. He has been offering quantitative risk management expertise to insurers, banks and other financial institutions since 2006. You can reach him at info@gillespie-aktuar.de.

Introduction

Investors in the life settlements market require a quality assessment of life expectancy reports. They generally rely on "actual to expected" analyses based on historical life expectancy reports, mostly issued or authorized by the medical underwriters who provided the reports. These analyses purport to and seemingly do show that the life expectancy estimates provided by the medical underwriters were reasonably accurate, or at least not statistically inconsistent. However, as we will show, a mere actual to expected analysis is inconclusive, if not misleading.

This article will present an alternative validation methodology, which makes much better use of the available mortality information. In its simplest form, it reduces the testing to the estimation of a single parameter which can be considered as a measure for a certain kind of systematic over or underestimation of the life expectancies.

Investors are thereby able to perform their own analyses, even for rather small portfolios with shorter histories, and to draw statistically valid conclusions concerning the life expectancies on which they have based their pricing and management of future cash flows.

The Role of Medical Underwriters in the Life Settlements Market

The basic principles underlying a hypothetical life settlement transaction include the following:

- An elderly insured person with a medical condition resulting in a substandard state of health no longer requires the protection provided by the life insurance policy.
- "Selling back" the policy to the insurer is not an attractive option for the insured, since the insurer's offer, the cash surrender value, lies far below the "fair" value.
- The investor makes an offer much more lucrative for the insured, however, at a price point that is still considerably below the fair value.
- If the insured accepts the investor's offer the policy stays in effect with the investor becoming the new beneficiary, in return for paying the offer price to the insured and the future premiums due to the insurer.

Obviously, investment in a single policy bears far too great a risk of financial loss. Therefore, investors usually purchase entire portfolios with at least 30 to 40 insured persons which creates a more predictable and less risky block of business.

This reassuring message to investors holds true only if the life expectancy estimates for the insured persons provided by medical underwriters and taken as key parameters for pricing the policies are not too "optimistic," i.e., more favorable than the "true" life expectancies. So, why rely on these estimates? Why not simply base one's pricing on accessible and generally accepted mortality tables, such as the Valuation Basic Table(s) developed and published by the Society of Actuaries?

The main reason is that, in the typical case of life settlements, we are dealing with an insured persons with severe medical conditions and a more substandard population as a result. The basic table underestimates the mortality for such a group, leading to an offer price below the fair value, equally unattractive to the insured as the surrender value. Therefore an appropriate estimate is needed for the degree by which the mortality rates in the basic table are to be increased for pricing purposes. Providing such estimates is the main service medical underwriters have to offer in the life settlements market.

The Medical Underwriters' Methodology

The methodologies employed by expert medical underwriters can differ in various respects—for instance, the amount and type of medical and socio-economic data taken into account. For the purpose of this article, however, we can neglect the differences and regard the methodologies alike, simply as functions taking biometrical and other relevant parameters as input and rendering a *mortality factor* as output, by which the mortality rates of the basic table are to be multiplied in order to derive the appropriate mortality rates.

In the past, mortality factors well above the 100 percent standard were common, and it was not unusual to see numbers as high as 500 percent for example. The impact of such figures can be seen in the following diagram, showing the mortality probability distributions, of a cohort of female insureds age 75, with the VBT 2008 as the basic table:

Mortality Probability Distributions for a cohort of Female, Age 75, for different Mortality Factors based on VBT 2008

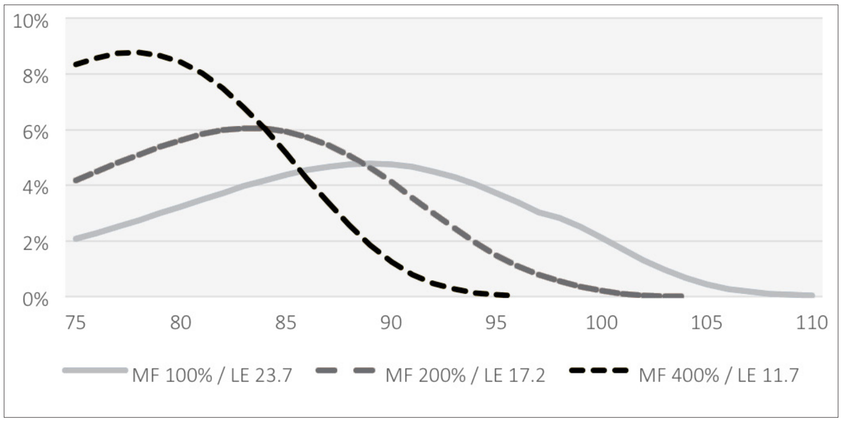

The diagram above shows that applying a constant mortality factor greater than 100 percent to the mortality rates of an insured person for all future years not only reduces the life expectancy estimate but also the "longevity risk." If one also takes into account that many policies traded in the life settlements market have flexible premiums, potentially growing year by year at ever higher rates, it is clear that overestimated mortality factors can lead to significant overpricing of policies for sale.

Therefore, investors were shocked by news in 2008 that some of the leading medical underwriters had started to drastically lower their mortality factors, so that life expectancy estimates increased by 20 to 25 percent on average.[1]

Poor Validation Methods

Were the reductions of the mortality factors the consequence of necessary corrections to previously flawed methodologies or incorrect weighting schemes for certain medical conditions? Or were they merely due to the replacement of the VBT 2001 by the VBT 2008 as the "official" basic table as some medical underwriters have suggested?[2]

Quantitative analysis by the medical underwriters based on their comprehensive historical records could shed some light on these issues. However, to the author's knowledge, no such analyses, with statistically valid conclusions as to whether historical mortality factors were systematically overestimated, have been published. Instead, investors were presented with actual to expected analyses, which simply compared the actual number of deaths for a specific portfolio with the expected number of deaths to date according to life expectancy reports previously provided to the investors. Actual to expected ratios around 100 percent were then considered proof or, at least, an indication of a valid methodology used in the past.

With respect to such analysis, one market expert has commented: [3]

"You would think that the expected deaths used to determine a life expectancy provider's Actual to Expected ratio would be based on the actual LE estimates it gave to its clients. However, in the face of A/E ratios based on actual/historical data that are too low, some life settlement providers have adopted the practice of just lowering their "expected" deaths, ostensibly to reflect current methodologies and mortality tables, with the convenient benefit of making their adjusted A/E ratios higher and closer to 100%."

The Actuarial Standards Board, in an Exposure Draft of May 2013, has also drawn attention to a lack of rigorous estimation standards:

"The life settlements market has demanded actual-to-expected (A/E) results from the LE providers, but in the absence of specific guidelines and disclosures, practices for calculating A/E results have varied widely. A limited

CONTINUED ON **PAGE 22**

number of states require LE providers to file A/E ratios, but again, lack of specific guidelines has led to concerns with mortality tables and methodologies used."

Simple actual to expected ratios have little relevance and can actually be quite misleading, as the following diagram shows:

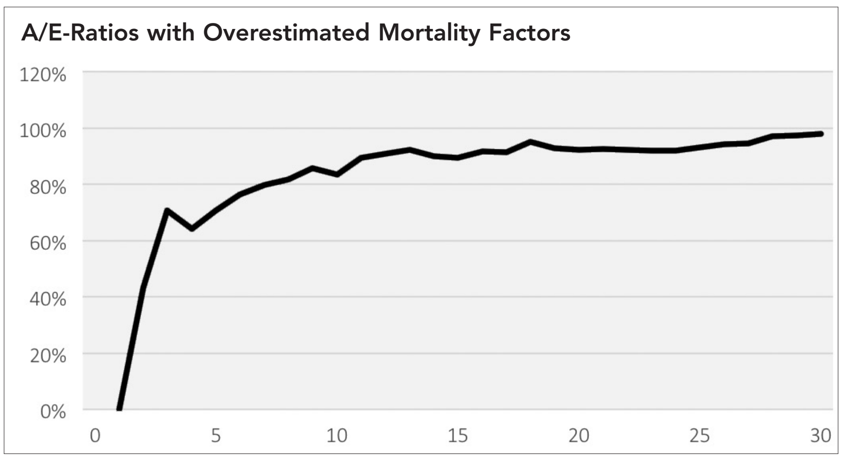

The graph depicts a series of actual to expected ratios for a fictitious portfolio of 48 men and 52 women whose policies were purchased eight to 10 years ago and whose ages at the time ranged from 60 to 90 years. Further, the mortality factors ranged from 150 percent to 300 percent of the 2001 VBT table. Finally, the ratios were simulated under the assumption that for each mortality factor the portion exceeding 100 percent is twice what it should have been in order to render the actual mortality rates (e.g., estimated factor 180 percent, correct factor 140 percent).

According to this assumption the average life expectancy for each insured person at the time of purchase of his or her policy was 10.1 years. This figure is significantly greater than the average, 8.4 years, of the life expectancies featured in the hypothetical medical underwriting reports. Yet, seven years after the first purchase the actual to expected ratio has already reached a level of 80 percent and, in the following years, it continuously approaches the "perfect" 100 percent level.

This example demonstrates that a simple actual to expected analysis can make historical mortality factors appear much more accurate than they actually are, all the more so if the mortality factors are subsequently reduced for the purpose of the analysis (as hinted at in the above quote). Whether the latter is the case or not, in light of the methodological deficiencies of the actual to expected approach it should come as no surprise that most medical underwriters have presented ratios above the 90 percent level. Such figures cannot be taken as "statistical proof" of a valid underwriting methodology in the past.

Improved variations of the simple actual to expected approach have been developed.[4] These alternatives cannot fix the problem of the simple approach, though—namely not to take the entire *mortality distributions* associated with the mortality factors into account but only certain key figures thereof. By condensing the available information to, for instance, the number of deaths up to now and only comparing the expected and actual value with each other, too much information may be neglected.

It is often argued that statistical tests based on the *entire* information encoded in the mortality distributions (for the testing period) are so restrictive that they *have* to lead to a rejection of the medical underwriters' models.[5] Thus, it could be argued further that such tests are just as useless for validation purposes as simple actual to expected analyses, only for opposite reasons.

This argument would indeed have some merit if such tests were employed in order to verify, or rather falsify, the assumption of a perfect alignment between the actual mortality distributions and the ones implied by the mortality factors. Such a match is highly unlikely anyway. It is quite obvious that applying one factor to all future mortality rates for the insured in question is bound to lead to a more or less skewed or otherwise distorted mortality distribution.[6] One might concede that the mortality distributions implied by the medical underwriter's models only have to be in the general proximity of the actual mortality distributions.

Next, a statistical method is presented that provides investors with a reliable measure to satisfy that requirement.

An Alternative Method

How could an investor assess the quality of the historical life expectancy reports for the hypothetical portfolio of 48 men and 52 women? The investor will certainly not be able to assess each report, or rather each implicitly reported mortality factor *by itself.* The only information available for such an assessment would be that the respective insured person is still alive or, if not, when he or she died. These bits of information are clearly insufficient for any valid statistical conclusion.

Thus, some kind of connection has to be established between the mortality factors for *all* insured persons. One way of doing so is by introducing a parameter, α, which corresponds to a certain kind of systematic over- or underestimation of the mortality factors. Let $\mu_1, \mu_2, \ldots$ be the reported mortality factors, and assume that the true mortality factors are

$$\mu_{\alpha,k} = \alpha(\mu_k - 1) + 1 \ (k = 1, 2, \ldots)$$

Then, $\alpha = 100\%$ means that the reported mortality factors were correct, and $\alpha = 0\%$ means that, in contrast, the mortality rates given by the basic table directly applied to the portfolio. The assumption under which the series of actual to expected ratios depicted above was simulated corresponds to $\alpha = 50\%$.

This way of connecting the mortality factors can be criticized as being arbitrary. Indeed, it is not very plausible that all mortality factors were estimated with a systematic error expressible in such a simple manner by a single parameter. But the aim is not to develop a realistic model of how the mortality factors were systematically over or underestimated, if that was indeed the case. The aim is rather to develop a model that allows a statistically valid conclusion as to whether there was *some* kind of systematic estimation error based on the sparse mortality data available for the portfolio. And that is exactly what the proposed model does. Say, for instance, a low value for α is implied by the data, with a narrow margin of error and at a high level of confidence. Then the investor can justifiably claim that the life expectancy reports in question significantly overestimated the mortality factors, at least "on average." Perhaps more importantly, the investor will have a better basis for the modeling of future cash flows.

The classical maximum likelihood methodology offers a suitable estimator $\hat{a}$ for α. This estimator is determined as follows: Let T_k be the random time of death of the insured I_k, measured in full units, for instance months, viewed from the time of purchase $t_{0,k}$ of his or her policy. And let t_k and τ_k be the actual time of I_k's death and today, respectively, measured in the same full units and with respect to $t_{0,k}$, but viewed from some time in the future when all insureds will have deceased. Then $\hat{a}$ is the α-value for which the following product reaches its maximum:[7]

$$P = \prod_{t_k \leq \tau_k} P_\alpha(T_k = t_k) \prod_{t_k > \tau_k} P_\alpha(T_k > \tau_k)$$

$P_\alpha(T_k = t)$ being the probability of the event $\{T_k = t\}$ according to the mortality distribution of I_k implied by $\mu_{\alpha,k}$, as viewed from $t_{0,k}$.

A great advantage of the maximum likelihood estimator is that it not only renders a measure for a potential overestimation of mortality factors but also a measure for the reliability of that measure. For $\hat{a}$ is asymptotically normally distributed with mean α and variance $I(\alpha)^{-1}$, and $I(\alpha) = -E(\frac{d^2 \ln(P)}{d\alpha^2})$.[8] Together with an estimate of α statistical software packages will usually also provide an estimate of $I(\alpha)$.

The following diagram shows how well the asymptotic approximation works, even for a relatively small portfolio with a moderately long history as the one considered above:

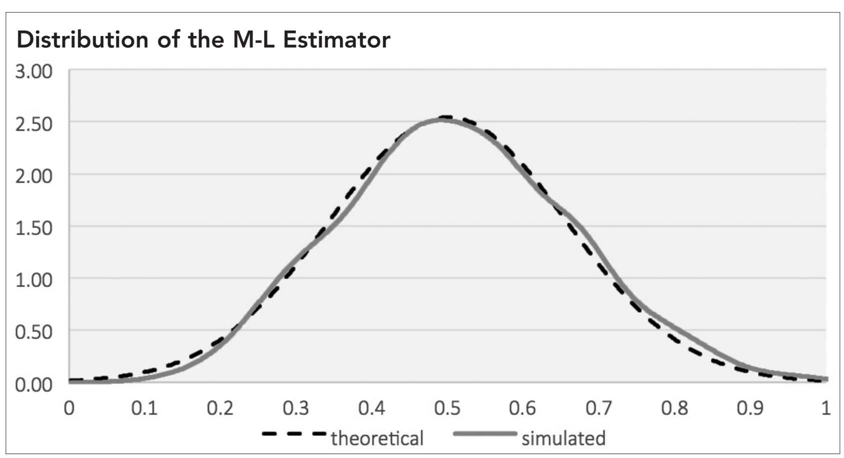

CONTINUED ON PAGE 24

The solid line shows the smoothed empirical probability distribution (or rather density) of $\hat{a}$, simulated for the hypothetical portfolio, with $\alpha = 50\%$. The first simulation run rendered $\hat{a} = 54.95\%$ and $\widehat{I(\alpha)^{-1}} = 15.69\%$.[2] The dotted line depicts the densitiy of the normal distribution with mean 0.5 and standard deviation 0.1569.

This example shows that the back testing method described above provides investors with estimates for the degree by which the historical mortality factors for their portfolios were possibly overestimated in a certain kind of way and "on average." Moreover, it also provides them with a reliable measure for the accuracy of those estimates.[9] In the hypothetical case the investor could conclude, at the 95 percent confidence level, that α is no greater than $54.95\% + 1.6449 \times 15.69\% = 80.76\%$.

Conclusion

For validating life expectancy reports, a mere actual to expected analysis is less than a valid substitute for quantitative studies which enable one to draw genuine statistical conclusions. This is due mainly to the loss of a notable portion of the information associated with any particular set of mortality factors, resulting from focusing solely on certain key figures of the mortality distributions involved.

The maximum likelihood estimator, $\hat{a}$, in contrast, uses the entire information and thereby allows investors to perform their own analysis. This analysis is intended to detect a systematic overestimation of mortality factors, if indeed such an overestimation did occur in the past.

The estimator, $\hat{a}$, and more sophisticated alternatives can prove to be powerful new risk management tools for investors in the life settlements market. ◻

ENDNOTES

1 *See http://www.lifepolicygroup.com/press/market-rocked-as-21st-services-changes-mortality-tables*

2 *See Siegert, Paul Evolution of Life Expectancies in the Life Insurance Secondary Market, Insurance Studies Institute, 2010.*

3 *Rebello, R. How Poor Actuarial Practices result in Multi-Million dollar losses for Life Settlement Investors, Colva Insurance Services*

4 *See Bauer, D., Russ, J. A New Methodology for Measuring Actual to Expected Peformance, 2012, http://www.ifa-ulm.de/downloads/DCLE.pdf*

5 *Ibid.: "Due to size of portfolio, deviations that would be considered small by practitioners would be statistically significant."*

6 *Consider the case of a patient upon whom a life-saving operation needs to be performed. Assume that the outcome of that operation will either be the patient's death or the patient's complete recovery. It is clear, in this case, that not all mortality rates are equally affected.*

7 *It is assumed that T_1,T_2,… are independent.*

8 *See e.g. Green, W.H. Econometric Analysis, 7th Ed., 2012.*

9 *With a further simulation study it can be shown that the estimates for the variance of do not vary too much, themselves.*