

Evaluating Life Expectancy Evaluations¹

Daniel Bauer, Michael V. Fasano, Jochen Russ, and Nan Zhu²

Abstract

The quality of life expectancy estimates is one key consideration for an investor in life settlements. The predominant metric for assessing this quality is the so-called *A-To-E Ratio*, which relies on a comparison of the actual to the predicted number of deaths. In this paper, we explain key issues with this metric: In the short run, it is subject to estimation uncertainty for small and moderately-sized portfolios; and, more critically, in the long run, it converges to 100% even if the underwriting is systematically biased. As an alternative, we propose and discuss a set of new metrics based on the *Difference in (Temporary) Life Expectancies*. We examine the underwriting quality of a leading US life expectancy provider based on this new methodology.

Keywords: Life Settlement; Life Insurance; Investment; Secondary Market; Life Expectancy Evaluation; A/E Ratio; Difference in Temporary Life Expectancies

¹ Parts of this article are taken from a newsletter article by the first and third author based on a presentation at the Fasano 9th Annual Life Settlement & Longevity Conference, and parts are taken from an old version of the paper “Adverse Selection in Secondary Insurance Markets: Evidence from the Life Settlement Market” by the first, third, and fourth author.

² Daniel Bauer is the Dai-ichi Life Insurance Company Endowed Professor in Actuarial Science and Risk Management in the Department of Economics, Finance, and Legal Studies at the University of Alabama (dbauer@cba.ua.edu). Michael V. Fasano is President of Fasano Associates, a medical underwriting firm and life expectancy provider (mfasano@fasanoassociates.com). Jochen Russ is managing partner of the Institute for Finance and Actuarial Science, an actuarial consulting firm based in Ulm, Germany and an adjunct professor for actuarial science at Ulm University (j.russ@ifa-ulm.de). Nan Zhu (corresponding author) is an Assistant Professor of Risk Management in the Department of Risk Management at Pennsylvania State University (nanzhu@psu.edu, phone: 814-863-8666, fax: 814-865-6284). We thank two anonymous reviewers whose comments helped to significantly improve the article. All remaining errors are ours.

1 Introduction

The life settlement market evolved from the viatical settlement market that was popular in the 1980s when the AIDS epidemic peaked. Unlike its predecessor that mainly targeted terminally ill life insurance policyholders with rather short estimated remaining lifespans, the life settlement market generally targets senior policyholders that are not terminally ill, yet still have below-average life expectancies. In a life settlement transaction, the original policyholder sells both the liability of future premium payments and the claim of the death benefit of a life insurance contract to a life settlement company for a lump sum payment. The life settlement company may further securitize a package of settled contracts in the capital market to general investors (Stone and Zissu, 2006; Braun, Gatzert, and Schmeiser, 2012).

A key ingredient when preparing an offer price for the transaction is an estimate for the (unknown) remaining lifetime of the insured. Such estimates are provided to the life settlement company by so-called *life expectancy (LE) providers*, who specialize in compiling individualized mortality forecasts in both primary and secondary life insurance markets. For investors, their return on each life settlement transaction is highly sensitive to the insured's remaining lifetime. As a consequence, the aggregate return for a portfolio of life settlement investments depends on the overall quality of the life expectancy estimates. More precisely, aggressive (too short) LE-estimates hurt investors as insureds end up living longer on average than projected by the life settlement company. On the other hand, LE-estimates are ordered by insureds' brokers and life settlement companies, who both potentially benefit from aggressive underwriting as shorter LE-estimates yield higher offer prices. This potential misalignment of interests is supported by empirical evidence of aggressive LE-estimates in early market years (Cook and Ezell, 2008, e.g.).

Objective measures for the quality of LE-estimates are therefore an important prerequisite for giving the life settlement market the credibility that general investors need. So far, the predominant metric for assessing this quality is an ex-post comparison of the *Actual* number of deaths to the predicted or *Expected* number of deaths, the so-called *A/E (A-To-E) ratio* (Actuarial Standards Board, 2013). If the actual number is close to the expected number, then the ratio of actual to expected deaths obviously will be close to 100%—which typically is interpreted as an indication of a high quality of medical underwriting. In this paper, we commence by explaining why such an interpretation can be deceiving. Using an illustrative example, we first show that relatively shortly after the point of underwriting, relatively few deaths occur, so that the metric is subject to estimation uncertainty for small to moderately sized portfolios. More importantly, we show that over time, the A/E ratio becomes increasingly less meaningful: Since eventually all lives that were expected to die (E) will indeed die (A), the A/E ratio will automatically be pulled towards 100%. Since existing portfolios are frequently made up of policies that were underwritten relatively recently and policies that were underwritten less recently, the extent to which the latter “bias” contributes to the overall underwriting evaluation is generally unclear. Thus, in many situations, the metric may not be suitable for revealing the true quality of the LE forecasts to analysts and investors.

As an alternative, we propose and analyze a set of new methodologies for measuring actual to expected performance via the *difference in life expectancies (DLE)*, the *difference in temporary life expectancies (DTLE)*, and the *implied difference in life expectancies (IDLE)*. Here the DLE is simply the difference in realized and estimated life expectancy, calculated as the average of realized lifetimes less the LE-estimates across a portfolio. We argue that while the first metric, DLE, has the most desirable properties in that it directly corresponds to the quantity of interest, its calculation is not possible until all individuals in the portfolio have died. The DTLE, on the other hand, is the difference of the temporary lifetime realized until the time of observation—that is, the minimum of the lifetime and the observation period—less its expected value under the provider’s estimates, which in the actuarial literature is usually referred to as *temporary life expectancy* (Bowers et al., 1997). Thus, an estimate for the DTLE metric can easily be obtained as the average of the differences in temporary lifetimes and temporary LE-estimates across a portfolio.

For a given portfolio, the DTLE converges to the DLE over time as the temporary life expectancy converges to life expectancy. Hence, unlike the A/E ratio, the informative value of the DTLE will improve as the portfolio ages and eventually converge to the metric of interest. Therefore, the DTLE overcomes the main issue of the conventional A/E metric. Furthermore, the standard errors and thus the uncertainty of the estimate are relatively small even early on when only few people have died. This latter aspect is associated with a drawback in the interpretation, however, since the information about deaths is sparse and is only incorporated over time. To elaborate, while a DTLE in the long run will be close to the DLE and so a DTLE of, say, 0.5 is easily interpreted, the significance of a DTLE of 0.1 early on is difficult to assess, since it is not clear if this value will result in a final value of, say, 0.5 years or 2 years.

As a metric that combines both advantages, easy interpretability early on and meaningful results as time passes, we introduce the *implied difference in life expectancies (IDLE)*. More precisely, we define the (absolute) IDLE as the eventual deviation in DLE that results in a benchmark DTLE identical to the observed DTLE as of today. Therefore, an IDLE of around +/- one year indicates that the underwriter’s estimate will eventually turn out to be one year too short (aggressive) or one year too long (conservative), respectively, assuming the future evolution of survival rates matches the pattern observed in the past. Alternatively, instead of obtaining an absolute difference in eventual life expectancies, it is also possible to derive a relative percentage by which the final realized LE is expected to differ from its estimate given the observed value of DTLE today. We call this the *relative IDLE*. For instance, a relative IDLE of around +/- 10% indicates that the underwriter’s estimate will eventually turn out to be on average 10% too short or 10% too long, respectively, assuming the survival pattern remains steady. The latter notion has advantages in the context of heterogeneous portfolios, because clearly a one-year deviation has different implications depending on the length of the LE-estimate. Using this new metric, the results are easy to interpret, and properties—including asymptotic distributions and thus confidence bands—are easy to derive. In particular, we show that the IDLE will converge to the actual error of the estimated LE in the long run. This property should make it preferable to analysts and investors for assessing the quality of LE-estimates.

We apply the above new metrics using data from a large US LE-provider, Fasano Associates, in order to analyze its underwriting performance from beginning-of-year 2001 through end-of-year 2013. Our results show that underwriting has been spot on for the more recent time period: For lives that were underwritten from year 2006, both the DTLE and IDLE are very close to zero. For the entire underwriting portfolio, the metrics are slightly higher: The DTLE is less than two months, which corresponds to an estimated average deviation of the eventual DLE (absolute IDLE) of around eight months—or a relative IDLE of approximately 8%. We note that this is to a great extent caused by more aggressive underwriting in the early 2000s as evidenced by our analyses of the metrics over time. The metrics for the full portfolio are jointly influenced by both very good recent underwriting as well as lives too aggressively underwritten in the early years.

We note that our focus is on establishing a single metric that, much like the A/E ratio, seeks to assess the *average* underwriting quality of the estimates provided by an LE provider. Different approaches are possible, such as using several metrics (e.g. an A/E ratio for each year since underwriting) or testing for the accuracy of the shape of individual mortality curves. We comment on these alternative approaches in the conclusion.

Related Literature and Organization of the Paper

The comparison of actual and expected deaths has a long history in retrospective actuarial studies. Seligman and Kahn (1980) derive hypothesis tests for the correctness of the assumptions in state termination probabilities, using different (asymptotic) distributions of the numerator (actuals) in the A/E ratio corresponding to different cases. Rhodes and Freitas (2004) also derive approximate confidence bands for A/E ratios in the context of a Poisson model. We contribute to this line of literature by proposing a new set of metrics and analyzing asymptotic properties, including the calculation of asymptotic confidence intervals.

Various authors indicate the significance of medical underwriting results for life settlement investors (Bhuyan, 2009; Aspinwall, Chaplin, and Venn, 2009). In particular, Gatzert (2010) points to the underwriters' performance as a possible explanation of the gap between expected and realized investment returns in the life settlement market, and Januário and Naik (2016) provide empirical evidence on the relevance of “life expectancy estimation risk” for life settlement return (see also Braun, Gatzert, and Schmeiser (2012); Davó, Resco, and Barroso (2013); Giaccotto, Golec, and Schmutz (2017) for studies on life settlement investment performance). While there are general difficulties associated with composing and evaluating underwriting performance (Qureshi and Fasano, 2010), the specific problems associated with the bias of the A/E ratio may also lead to a misevaluation of underwriting performance by investors, and thus may exacerbate the issue of “life expectancy estimation risk.” Hence, our paper contributes by providing investors alternative—and as we argue superior—metrics, and by informing the marketplace about the results of one large underwriter based on this new approach.

The paper is organized as follows. In Section 2, we first provide analyses of the A/E ratio documenting the indicated deficiencies. In Section 3, we then introduce the novel DLE/DTLE/IDLE metrics as alternatives, discuss their properties, and present illustrative

calculations. In Section 4, we calculate empirical results using the Fasano data. Finally, we present our conclusions in Section 5.

2 A/E Ratios with Illustration

A/E ratios are commonly used in actuarial practice to assess the quality of LE forecasts by tracking the ratio of the actual number of deaths to the expected number of deaths over time (Seligman and Kahn, 1980). While in practice A/E ratios are typically reported based on the development over calendar years for a given portfolio, we focus here on the development over years since underwriting in order to better illustrate the drawbacks of the metric. More precisely, for each individual i in a portfolio of size N , we denote by $\tau_x^{(i)}$ the individual's realized remaining lifetime and by ${}_tq_x^{(i)}$ the individually estimated t -year mortality probability as seen from the time of underwriting, where in line with standard actuarial notation x is the individual's age (at the time of underwriting). The A/E ratio for the portfolio of all individuals that have been underwritten for at least t years is hence:

$$\frac{A}{E} = \frac{\sum_{i=1}^N 1_{\{\tau_x^{(i)} < t\}}}{\sum_{i=1}^N {}_tq_x^{(i)'}}$$

where $1_{\{\tau_x^{(i)} < t\}}$ is the indicator function that takes the value 1 if individual i died before time t and 0 otherwise. Hence, A is simply the number of people who have actually died before time t , whereas E is the expected number of deaths before time t under the assumption that the estimated mortality probabilities (corresponding to the estimated LEs) are correct.

To illustrate the deficiencies of the A/E ratio as a measure of the quality for life expectancy estimates, we consider a very simple hypothetical example. We assume an LE portfolio composed of 500 identical non-smoking male individuals aged 75 who have actual mortality probabilities twice as high as the probabilities from the standard 2008 Valuation Basic Mortality Table (2008 VBT).³ Moreover, we contemplate two biased LE providers with the first (A) being too aggressive in assessing LEs and the second provider (C) being too conservative. More precisely, we assume that each LE provided by (A) is two years too short whereas provider (C)'s LEs are two years too long. We further assume that the deviations originate only from an erroneous assessment of the "frailty factor" to be multiplied on the mortality probabilities from the VBT table. That is, we assume that there are no systematic deviations between the mortality table used by the underwriter and the "correct" table.

[Insert Figure 1 here]

The solid and dashed line in Figure 1 shows the corresponding development of A/E ratios for the aggressive LE provider (A) and the conservative LE provider (C), respectively, if the 500

³ We choose a multiplier ("frailty factor") of 200% since the life settlement market typically targets policyholders with below-average health status.

people die exactly according to their mortality probabilities (i.e., without any random fluctuation). Due to random fluctuation, the actual number of deaths will of course not evolve exactly according to the mortality probabilities. As a result, A/E ratios in reality will also deviate from these two lines. The dotted lines therefore show the 95% confidence intervals of the ratios. There are a number of contributions in the actuarial and statistical literature that provide analyses of the A/E ratio, including formulas for approximating asymptotic confidence intervals (see e.g. Rhodes and Freitas (2004) and references therein). However, here we rely on Monte Carlo estimates for the 2.5% and 97.5% percentiles of the A/E ratios for each year, based on 100,000 simulations of the random times of deaths in the portfolio as they should provide accurate estimates for the finite-sample intervals.

There are two observations. First, the confidence intervals of the A/E ratios in the early years are very wide. In particular, the 95% confidence intervals for the two providers overlap for the first three years since underwriting. The intuition for this observation is straightforward: The probability of dying immediately after underwriting is relatively low, so the denominator in the A/E ratio is very small. Therefore, if some individuals die—or fail to die—early on, the impact on the ratio will be considerable. This indicates that the A/E ratio has to be interpreted with caution for *young* portfolios, especially those that are small or moderately-sized as in the example considered here.

Second and more strikingly, the A/E ratios for both the aggressive and the conservative LE provider are approaching 100%. For instance, after less than 20 years, the A/E ratios are well within the range between 90% and 110% for both providers with a high probability—even though both providers were systematically and significantly wrong in their original estimates. Again, the intuition is straightforward: Both, the actual number of deaths and the expected number of deaths will necessarily approach the number of people in the portfolio since, eventually, all individuals will die. Hence, both A/E ratios are “pulled” towards 100% in the long run, even if one provider was too aggressive and the other too conservative. This indicates that the A/E ratio has to be interpreted with caution for *old* portfolios.

[Insert Figure 2 here]

Of course, as the size of the portfolio N increases, the statistical uncertainty decreases and the results will be more accurate early on. For instance, in Figure 2, we plot the development of the A/E ratio and associated confidence intervals similarly as in Figure 1 but for a portfolio size of $N=50,000$, which roughly corresponds to our empirical application in Section 4 (the Fasano portfolio). We find that the widths of the confidence intervals for both A/E ratios become much tighter. Importantly, however, the long run bias is of systematic nature and is therefore not alleviated by the larger sample size.

Thus, we conclude that the A/E ratio is not meaningful to assess underwriting quality in the long run. Typical portfolios include both lives that were underwritten recently as well as lives underwritten further in the past. Therefore, it is difficult to assess the influence of this effect. The deficiencies for lives underwritten a relatively long time ago can be (partly) addressed within the A/E framework, particularly by using more than just one A/E-number (e.g. one A/E number for each year since underwriting). However, this may be inconvenient for

practitioners that would like an overall assessment of underwriting quality as well as for non-experts that will find many numbers difficult to process. Also, in practice, A/E analyses are frequently used to assess sub-portfolios, e.g. by gender, policy duration, impairment type, mortality ratings, etc. If deaths in such sub-portfolios are segmented by time to calculate several A/E numbers, many segments may be too small to generate statistically meaningful results.

3 DTLE/IDLE Methodology with Illustration

To determine a metric that overcomes some of the deficiencies of the A/E ratio, first assume that we knew the time of death for each individual in a given portfolio of LEs. In this case we could calculate the difference between the realized lifetime and the given LE-estimate for each life. For a given individual, of course this would ex ante be a random variable due to the intrinsic randomness of the remaining lifetime. In other words, even if all LE-estimates were exactly accurate, the realized lifetime will be larger than its expected value for some whereas it will be less than the life expectancy for others. Under the hypothesis that the LE-estimates are accurate, the average, however, should be close to zero. This is exactly the intuition of the *law of large numbers*, which states that the average of (mean-zero) random variables with bounded variance will converge to zero. On the other hand, if the average of these differences were positive, that would be a strong indication that the estimated LEs were too short on average—especially with a sufficiently large portfolio—whereas a negative average would indicate that estimated LEs were too long on average. We refer to the average of these differences as the *difference in life expectancies* (DLE).

However, in relevant situations, not all individuals will have died when one wants to assess the quality of an LE-underwriter. In particular, we can only observe times of death that occurred before the present time. For all other individuals we only know that they are still alive at this cut-off date but we do not know when they will die. To deal with this so-called *right-censoring* issue, we rely on an alternative concept from actuarial science and demographic research. More precisely, denote by $\bar{\tau}_x^{(i)} = \min(\tau_x^{(i)}, t_i)$ the *realized partial lifetime*, which is the right-censored version of the time of death, i.e. the number of years lived between time of underwriting and the cut-off date. Here, t_i is the (given) difference between the cut-off date and the time of underwriting for individual i . Its expected value, the so-called *temporary life expectancy*, is then given by the integral of the survival probabilities until time t_i (Bowers et al., 1997):

$$e_{t_i}^{(i)} = E[\bar{\tau}_x^{(i)}] = \int_0^{t_i} {}_tP_x^{(i)} dt,$$

where ${}_tP_x^{(i)}$ denotes the (unknown) “true” survival probabilities for individual i . Furthermore, we denote the estimated survival probabilities for individual i supplied by the LE provider by ${}_tp_x^{(i)} = 1 - {}_tq_x^{(i)}$, and denote the corresponding estimated temporary life expectancy by:

$$\hat{e}_{t_i}^{(i)} = \int_0^{t_i} {}_t p_x^{(i)} dt.$$

Then, under the assumption that the estimated temporary life expectancies are accurate on average, the difference in realized partial lifetime and estimated temporary life expectancy for individual i , $\bar{\tau}_x^{(i)} - \hat{e}_{t_i}^{(i)}$, is a zero-mean random variable with bounded variance. Thus, the average:

$$\text{DTLE} = \frac{1}{N} \sum_{i=1}^N [\bar{\tau}_x^{(i)} - \hat{e}_{t_i}^{(i)}],$$

which we will refer to as the *difference in temporary life expectancies* (DTLE), again converges to zero and is again asymptotically normally distributed according to Lyapunov's central limit theorem (Billingsley, 1995):

$$\frac{\text{DTLE}}{\frac{1}{N} S_N^{TR}} = \frac{1}{S_N^{TR}} \sum_{i=1}^N [\bar{\tau}_x^{(i)} - \hat{e}_{t_i}^{(i)}] \xrightarrow{d} N(0,1),$$

where $(S_N^{TR})^2 = \sum_{i=1}^N \text{Var}[\bar{\tau}_x^{(i)}] \approx \sum_{i=1}^N (\bar{\tau}_x^{(i)} - \hat{e}_{t_i}^{(i)})^2$. Clearly, as time passes by and all the t_i increase, the difference in temporary life expectancy will converge to DLE. Since, as described above, the latter is the “actual” metric we are interested in—ignoring estimation uncertainty—the long-run value of the DTLE will exactly reflect the desired information.

[Insert Figure 3 here]

To illustrate, we perform the DTLE analysis using the same hypothetical portfolio and assumptions of the two biased LE providers (A) and (C) as described in the previous section. Figure 3 shows how the DTLEs evolve over years since underwriting. Again, there are two observations. First, in the short run, the DTLE does not show a large deviation for either provider, and the confidence bands around the central estimates are also very narrow. While this may look like a deficiency, as indicated in the A/E analysis, the early deaths provide little usable information in view of the provider's quality. The point estimate for the A/E ratio ignores this and can take on values that heavily deviate from 100%; thus their statistical credibility is limited. The DTLE, on the other hand, reflects that there is little informational content in the early observations and is close to zero. If it is significantly different from zero, however, the result will be credible and likely reflect the provider's quality.

Second, as soon as there is sufficient statistical content to reflect the underwriting quality, the DTLE will diverge from zero and converge to the “true” DLE value. In particular, after approximately 20 years, the central estimates for the metric correctly identify that both LE providers were off by exactly two years in opposite directions. There exists estimation uncertainty, but this is unavoidable since it is observationally equivalent for a provider to be systematically biased or to have “bad luck” in the sense that the finite sample did not evolve according to the “true” distribution.

Similar to the previous section, the confidence intervals will be considerably tighter for large portfolios. Figure 4 again shows an analogous illustration for a portfolio size of $N=50,000$. It is worth noting that here, however, the portfolio size does not only impact the estimation early on but also in the long run. More precisely, it is unlikely that for large portfolios, “good” or “bad” luck will have a considerable impact for the eventual DTLE result. Hence, in the medium to long run and for sizable portfolios, the DTLE provides meaningful and easy-to-interpret results for the underwriting quality of an LE provider.

[Insert Figure 4 here]

Although the small level of the DTLE in the early period after underwriting is associated with the absence of useful information due to a limited number of deaths, its interpretation is difficult. For instance, while in our example a value of DTLE close to 2 after about 20 years is easily interpreted, the significance of a DTLE of roughly 0.2 after 5 years or even of roughly 0.5 after 10 years is difficult to assess. As a solution that combines both attributes—stability in the medium to long run and an easy interpretation also in early years after underwriting—we introduce the *implied difference in life expectancies* (IDLE). Within the calculation of the IDLE, we solve for an eventual (absolute or relative) deviation in DLE that produces a hypothetical DTLE with the same value as the actual DTLE observed in the sample.

More precisely, based on the survival probabilities ${}_tP_x^{(i)}$ from the underwriter’s original LE-estimate, we can scale them to derive a hypothetical level of ${}_tP_x^{(i),*}$ that results in a new estimate with a consistent deviation across the entire portfolio.⁴ For instance, we could determine the ${}_tP_x^{(i),*}$ such that every LE-estimate is increased by d years or p percent of the original estimate. Then, under the assumption that the individuals die according to the adjusted distribution given by ${}_tP_x^{(i),*}$, we can calculate the benchmark DTLE (DTLE _{d} respectively DTLE _{p}) based on this hypothetical distribution:

$$\text{DTLE}_{d/p} = \frac{1}{N} \sum_{i=1}^N \left[\int_0^{t_i} {}_tP_x^{(i),*} dt - \hat{e}_{t_i}^{(i)} \right].$$

Now we determine the value of d (or p) such that $\text{DTLE}/\text{DTLE}_{d/p} = 100\%$, respectively, where the numerator (DTLE) is the one calculated from the realized deaths as in Figure 3. Therefore, the solved d (or p) provides an estimate of the final absolute (or relative) deviation of the provider’s estimates, and we refer to them as IDLE _{d/p} . Because of the monotonic transformation, the confidence bands of IDLE _{d/p} can be directly calculated as solutions of d (or p) that would result in the ratio $\text{DTLE}/\text{DTLE}_{d/p} = 100\%$, using the associated confidence bands of DTLE in the numerator.

[Insert Figure 5 here]

⁴ To be precise, the scaling is done by adjusting the mortality multiplier that is applied to the base table to match the modified LE-estimate. As pointed out in Section 2, we operate under the assumption that there are no systematic errors in the underlying mortality table.

Figure 5 plots the $IDLE_d$ for our hypothetical example.⁵ Of course, now the central estimates for the aggressive (A) and conservative (C) LE providers are *exactly* 2 years and -2 years, respectively, since these are the levels of the eventual DLEs we assume in setting up the calculation.⁶ Hence, ignoring statistical uncertainty, an observer can conclude very early on that the realized difference will probably result in a final DLE (or in a DTLE in late years) of approximately two years. Due to the normalization, the early results get inflated such that—similarly to the A/E ratio—the credibility is low and corresponding standard errors are large. Hence, there exists a certain “uncertainty principle” in that an easy-to-interpret early warning sign will necessarily have a large error associated with it. However, as with the A/E ratio, the error becomes less pronounced for large portfolio sizes N and, unlike the A/E ratio, results become *both* credible early on and meaningful as time passes.

4. Analyses on Fasano Associates Data

In this section, we assess the quality of the LE-estimates for the portfolio of a concrete LE-provider, Fasano Associates (Fasano), using the introduced DTLE/IDLE metrics. Our data consist of 53,947 distinct life expectancy evaluations underwritten by Fasano between beginning-of-year 2001 and end-of-year 2013.⁷ For each record, we are given an LE-estimate that is the result of the underwriting process. Beyond these LE-estimates, we are also given realized times of death for the individuals that had died before January 1st 2015.

Two practical issues arise before specific analyses can be performed. First, we need to establish reasonable assumptions on the IBNR (*Incurred But Not Reported*) rate, since LE providers are typically not informed about an underwritten individual’s death and have to rely on public data that are often inaccurate, incomplete, or delayed (Behan, 2012). In this paper we use a constant annual IBNR rate of 7%, which was recommended by a number of LE underwriters based on Social Security Master Death File data. Second, we need to fix a consistent evaluation methodology for all LEs in the portfolio, which of course were furnished at different points in time.⁸ Here, we always use the original LE-estimates (i.e. we do not try to adjust for the fact that under current methodology, different LE-estimates would be given for some individuals) and we use Fasano’s most current proprietary mortality table to generate mortality distributions around such point estimates. More precisely, we solve for a mortality multiplier (frailty factor) that has to be applied to this proprietary table in order to match the LE-estimate that was originally provided by Fasano. We use the resulting mortality probabilities in our analyses.

[Insert Table 1 here]

⁵ We only provide $IDLE_d$ for the illustrative example since we assume homogeneous policyholders. We analyze both (absolute and relative) metrics in the following section when looking at the Fasano portfolio.

⁶ To clarify, the central estimates will be *exactly* ± 2 years in the limit. Since here the estimate is derived based on 100,000 repetitions of a Monte Carlo analysis, some estimation uncertainty remains.

⁷ We use the earliest observation in the analyses when the same individual has been underwritten several times.

⁸ For instance, the Actuarial Standards Board (2013) recommends the so-called adjusted-to-current method, although interpretations of and opinions on this approach differ in the industry.

Table 1 shows the metrics of interest for the entire Fasano portfolio as of January 1st 2015, together with their 95% confidence intervals that are determined based on the normal approximation discussed in Section 3. From the table we first observe that the DTLE is positive and at a level of less than two months ($0.144 \times 12 = 1.73$). While a deviation in this range on average certainly may appear innocuous on its own, we recall from the previous illustrations that the DTLE is always small initially because of the limited number of deaths. Since the entire portfolio we study here contains LE-estimates for lives that were underwritten both recently and not so recently, the interpretation of this number is not immediately clear. We recall that the definition of the IDLE addresses this issue with interpretation and can be consistently interpreted throughout the period the individual is observed. We solve for the IDLE considering both the absolute and relative deviations. Table 1 shows that for the entire portfolio, the LE-estimates are on average too short by approximately 8 months or a relative deviation of 8%, +/- roughly one month or 1 percent, respectively.

It is worth pointing out that the considered LE provider—and more generally the industry—has been adapting its underwriting technology over the past decade. Therefore, the overall deviation in LE-estimates might be heavily attributed to a more aggressive underwriting in early years, but does not represent the current underwriting performance. Indeed, differences for lives underwritten early on will disproportionately impact the DTLE, simply because they are larger in absolute value because of the longer history. Thus, we also separate the entire portfolio into two sub-portfolios based on the time of underwriting (before and after January 1st 2006), and calculate the same metrics for each sub-portfolio (also displayed in Table 1). We observe from the results that for lives that were underwritten more than a decade ago, the average bias in the underwriting is more pronounced ($IDLE_d$ of approximately one year, or $IDLE_p$ of 13.5%), whereas for the more recent sub-portfolio, the underwriting is fairly accurate: The point estimates for all metrics are very close to zero, and their respective confidence intervals all contain the zero point.

[Insert Figure 6 here]

Figure 6 shows the development of the DTLE for the Fasano data from two distinct perspectives. The left-hand panel shows the evolution of the DTLE and its confidence band at the end of each calendar year from year 2001 to year 2014, using all individuals that had been underwritten by the respective cutoff date. Therefore, the values at the end of year 2014 are the ones in Table 1. We observe that the figure exhibits a “smirk” shape. The DTLE for early years is (significantly) positive although the value is relatively low. Again, the latter is not surprising since the DTLE is naturally small over the early years. However, the observation that it is significant and positive corroborates that underwriting has been more aggressive for these early lives. While the DTLE first decreases to levels close to zero, it eventually increases to slightly below two months, a level that is disproportionately influenced by the early LE-estimates as explained above.

To better illustrate the underwriting performance over time, as within our illustrations in earlier sections, we also plot the DTLE results using years since underwriting on the x -axis in

the right-hand panel of Figure 6. As a consequence, the underlying dataset differs for each year (t). More precisely, the results for one year include all available data until end-of-year 2013—since we have reported deaths until beginning-of-year 2015 so that we observed all these lives at least for one full year. The results for, say, 10 years, however, include only the portion of the data that have been observed for more than 10 years, i.e. lives underwritten before beginning-of-year 2005. In particular, the results for low values of t include the most recent underwriting data and, thus, are the most relevant estimates for assessing the underwriter’s current quality, whereas the results for high values of t are solely based on lives underwritten in earlier years and hence measure the quality of the underwriting methodology used at that time. The DTLE is close to 0 for low values of t , but then increases to levels close to 1 year for $t=12$. The latter corroborates insights from Table 1 and the left-hand panel of Figure 6 that early underwriting (around 2001-2002) was more aggressive, although now the level (DTLE \approx 1) is more accurate. The finding that results for low values of t are close to zero confirms that underwriting has improved over the years, although as discussed the DTLE is difficult to interpret in this range.

[Insert Figure 7 here]

[Insert Figure 8 here]

Finally, Figures 7 and 8 depict the development of $IDLE_d$ (Figure 7) and $IDLE_p$ (Figure 8) over both, calendar year (left panel in each figure) and years since underwriting (right panel in each figure). Hence, one can interpret the figures as today’s best estimate for the (absolute or percentage) deviation in the eventual life expectancy that can only be observed in the future, when all insured will have deceased. In other words, under the assumption that the (today unknown) DLE will turn out to be this value in the future, the expected DTLE today coincides with the observed value for DTLE. Again, the results for 2014 in the left-hand panels coincide with the results for the entire portfolio from Table 1.

From both figures, the results for early calendar-years in the left-hand panels and the results for a large number of year-since-underwriting (t) in the right-hand panels confirm that LE-estimates were more aggressive in the early 2000s. For instance, the percentage of the increases amounts to roughly +10% for $t=9$ —meaning that on average estimates compiled pre-2006 were approximately 10% too short. However, results that include more recent underwriting experience show that this is no longer the case. Indeed, results for $t=1$ —which as discussed include the entire portfolio experience—demonstrate that both $IDLE_d$ and $IDLE_p$ are close to zero, whereas the results for t from 2 to 5 suggest that the underwriting has even been slightly on the conservative side in some periods. It is also worth noting that the confidence intervals are rather tight for small t due to the large sample size, suggesting that the estimations of IDLE, and furthermore, the assessment of the quality of the provider’s LE-estimates, are both credible and informative.

5. Conclusion

The quality of life expectancy estimates is of utmost importance for participants in the life settlements market. But the question of how to measure underwriting quality is also of certain relevance in primary insurance markets. In order to assess the quality of LE-estimates, typically so-called A/E ratios are being used. In the present paper, we explain that this metric has substantial shortcomings, both, for “young” and “old” portfolios. In particular, an A/E ratio for an old portfolio is artificially “pulled” towards 100% even if LE-estimates were systematically and significantly too short or too long. Therefore, under certain circumstances this metric can be misleading or even meaningless.

In this paper, we propose several alternatives. The number that is directly linked to the quality of the LE-estimates is the difference in life expectancies (DLE). Since this number cannot be observed until all insured in a portfolio have eventually died, a natural alternative is the difference in temporary life expectancies (DTLE). This number has the desired properties for old portfolios but its interpretation for young portfolios is not straightforward. We therefore propose the implied difference in life expectancies (IDLE) as an intuitive estimate based on today’s data for the DLE that will eventually materialize. We apply our methodology to the database of an LE-provider and find that this underwriter’s current methodology appears to be of high quality although it may have been slightly conservative in some recent years and more aggressive in very early years.

Of course, aggregating all the underwriting experience to a single number (be it A/E ratio, DTLE, or IDLE) is associated with a loss in information. Aggregating the information to multiple numbers, or even curves, will clearly reveal more information of the underwriting quality. However, processing multiple numbers/curves requires more expertise. Given the practice (and possibly the requirement) to have a single number that reflects underwriting performance, we argue that IDLE is a better aggregate measure than the A/E ratio (or DTLE) in terms of easy interpretation and explanatory power, for both older and younger portfolios. Furthermore, several IDLEs for sub-portfolios will be more informative to non-experts than several A/E ratios since all IDLEs are measured in the same way. When evaluating the performance of an aggregate, heterogeneous portfolio composed of policyholders with a wide range of LE-estimates, we recommend $IDLE_p$ over $IDLE_d$ since the former has a consistent interpretation irrespective of the size of the estimate. Moreover, it can directly be interpreted as the systematic deviation in the “frailty factors” that are the result of an LE provider’s underwriting process.

Our metrics should be of interest to all LE-providers when assessing their own quality of underwriting in order to identify weaknesses and improve their methodology. Also, investors in life settlements should consider using our metrics when assessing the quality of LE-providers, since the explained bias in A/E-ratios for old portfolios might conceal differences between LE-providers and might make LE-providers look better than they are. Moreover, when assessing the performance of their own portfolios, our metrics might be of use to investors. This is of particular interest when calculating the net asset value and the expected future return of a portfolio or fund: Given today’s information, IDLE is a natural and intuitive estimate for the DLE that will eventually materialize in the portfolio when all insureds have died. Stressing all original LE-estimates by the value of IDLE and calculating the net asset

value, the expected future return (or a return distribution) results in values that “automatically” correct for deficiencies in the original LE-estimates.

A possible caveat to this statement is the implicit assumption that imperfections in the estimates are “homogeneous” in the sense that they only result from wrong frailty factors but that there are no systematic deviations between the shapes of the mortality curves (this is also assumed in our illustrations in Sections 2 and 3). If we allow for arbitrary deviations between the “correct” mortality curve and the curve used in underwriting, it will be impossible to make an inference of the quality of eventual LEs and it will be impossible to assess underwriting quality with a single number. From a statistical perspective, it is possible to consider alternative hypotheses. For instance, one possible hypothesis is the adequacy of the entire provided mortality curve for each individual in the portfolio, for which it is possible to construct formal statistical tests, such as the Kolmogorov-Smirnov (K-S) test for the goodness-of-fit, or a dynamic version of Pearson’s test derived based on a counting process framework. However, adequacy of the full curve (rather than solely the LE) and not allowing for unsystematic deviations may be too restrictive as a quality criterion.

When considering the average quality of the LEs and when positing that past underwriting performance is a guide to future underwriting performance, our metrics and particularly the IDLE overcome deficiencies of the industry standard and provide a more accurate picture of the underwriter’s quality.

References

Actuarial Standards Board (2013). “Life Settlements Mortality.” Available at http://www.actuarialstandardsboard.org/wp-content/uploads/2014/02/asop048_175.pdf.

Aspinwall, Jim, Geoff Chaplin, and Mark Venn (2009). “Life Settlements and Longevity Structures: Pricing and Risk.” New York: John Wiley and Sons.

Behan, Donald F. (2012). “Measurement of Incurred but Unreported Deaths in Life Settlements.” ARC 2013.1 Proceedings.

Billingsley, Patrick (1995). “Probability and Measure.” 3rd Edition. New York: John Wiley and Sons.

Bhuyan, Vishaal B. (2009). “Life Markets: Trading Mortality and Longevity Risk with Life Settlements and Linked Securities.” New York: John Wiley and Sons.

Bowers, Newton L., Hans U. Gerber, James C. Hickman, Donald A. Jones, and Cecil J. Nesbitt (1997). “Actuarial Mathematics.” Second Edition, Schaumburg, IL: Society of Actuaries.

Braun, Alexander, Nadine Gatzert, and Hato Schmeiser (2012). “Performance and Risks of Open-End Life Settlement Funds.” *Journal of Risk and Insurance*, 79: 193-230.

Cook, David, and Glenn Ezell (2008). “Life Settlement Characteristics and Mortality Experience for Two Providers.” Available at

<http://us.milliman.com/insight/research/insurance/pdfs/Life-settlement-characteristics-and-mortality-experience-for-two-providers/>

Davó, Nuria Bajo, Carmen Mendoza Resco, and Manuel Monjas Barroso (2013). “Portfolio Diversification with Life Settlements: An Empirical Analysis Applied to Mutual Funds.” *The Geneva Papers on Risk and Insurance—Issues and Practice*, 38: 22-42.

Gatzert, Nadine (2010). “The Secondary Market for Life Insurance in the U.K., Germany, and the U.S.: Comparison and Overview.” *Risk Management and Insurance Review*, 13: 279-301.

Giaccotto, Carmelo, Joseph Golec, and Bryan P. Schmutz (2017). “Measuring the Performance of the Secondary Market for Life Insurance Policies.” *Journal of Risk and Insurance*, 84: 127-151.

Januário, Afonso V., and Narayan Y. Naik (2016). “Testing for adverse selection in life settlements: The secondary market for life insurance policies.” *Available at SSRN*.

Qureshi, A. Hasan and Michael V. Fasano (2010). “Measuring Actual to Expected Accuracy for Life Settlement Underwriting.” *Reinsurance News*, 68: 23-26.

Rhodes, Thomas E. and Stephen A. Freitas (2004). “Advanced Statistical Analysis of Mortality.” Proceedings of the 14th International AFIR Colloquium.

Seligman, Edward J., and Sheldon Kahn (1980). “Testing for Significant Differences Between Actual and Expected Results.” *Transactions of Society of Actuaries*, 32: 585-600.

Stone, Charles A., and Anne Zissu (2006). “Securitization of Senior Life Settlements: Managing Extension Risk.” *Journal of Derivatives*, 13: 66-72.

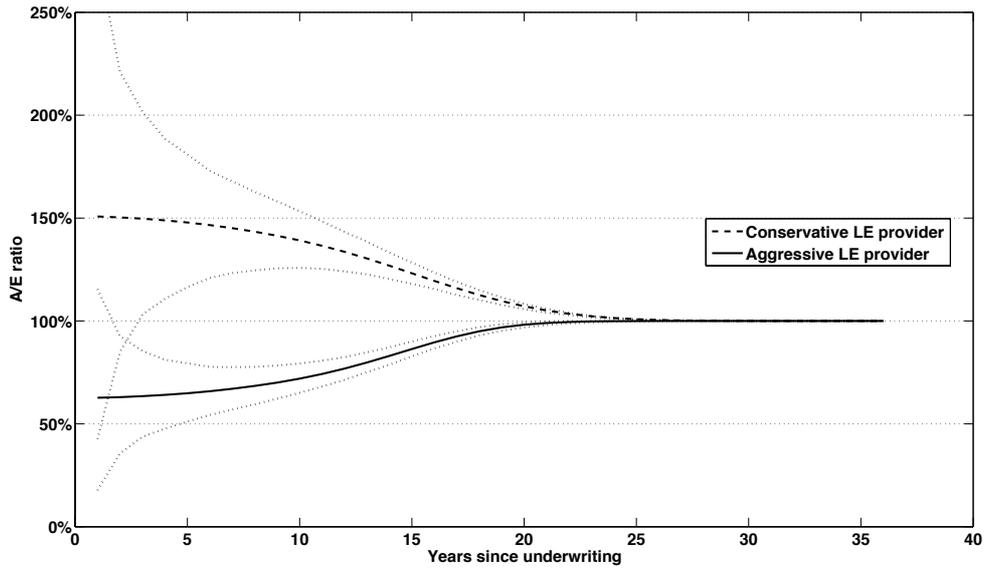


Figure 1: The development of the A/E ratios for the aggressive LE provider (A) [solid line] with associated 2.5% and 97.5% percentiles [dotted lines] and the conservative LE provider (C) [dashed line] with associated 2.5% and 97.5% percentiles [dotted lines] in the hypothetical example. $N = 500$.

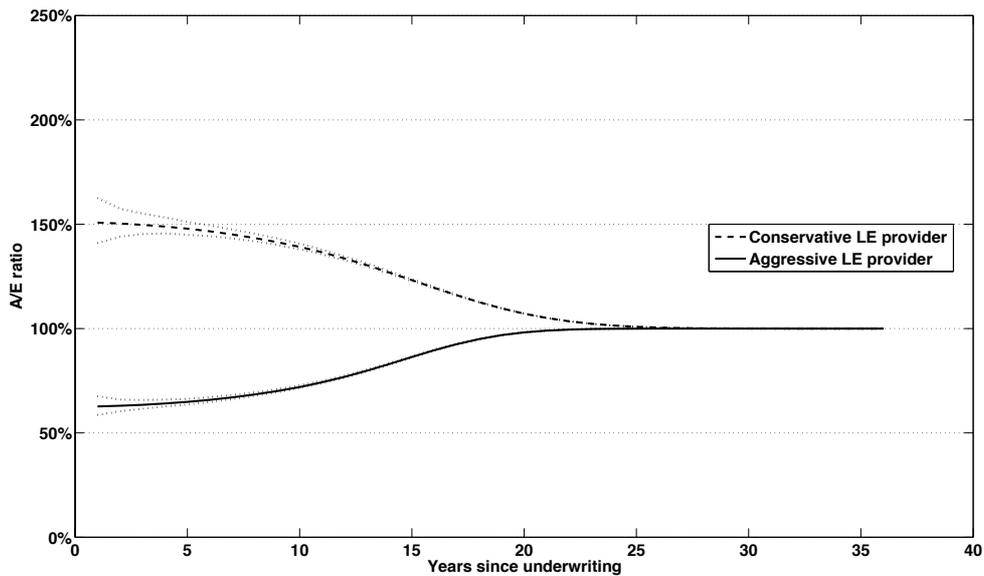


Figure 2: The development of the A/E ratios for the aggressive LE provider (A) [solid line] with associated 2.5% and 97.5% percentiles [dotted lines] and the conservative LE provider (C) [dashed line] with associated 2.5% and 97.5% percentiles [dotted lines] in the hypothetical example. $N = 50,000$.

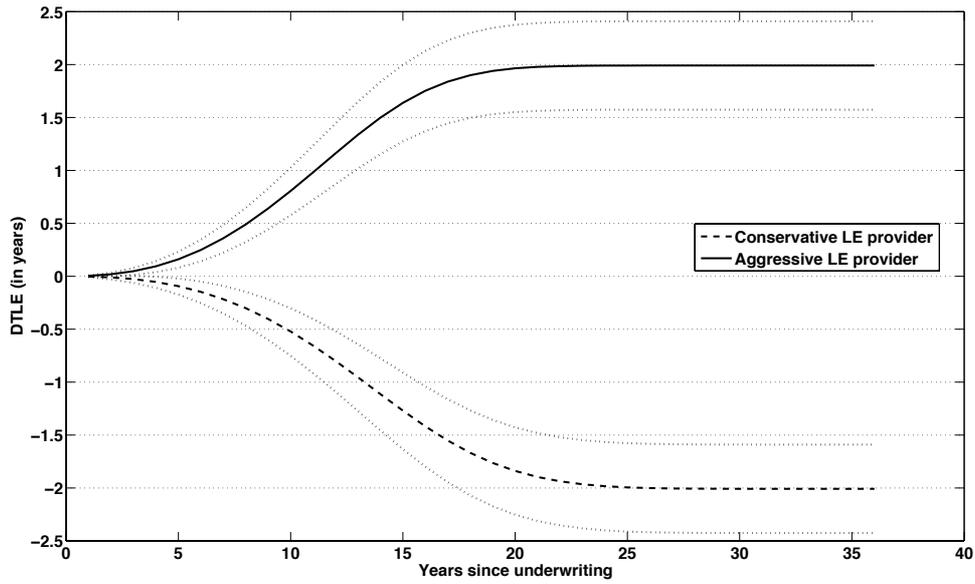


Figure 3: Development of DTLE for the aggressive LE provider (A) [solid line] with associated 2.5% and 97.5% percentiles [dotted lines] and the conservative LE provider (C) [dashed line] with associated 2.5% and 97.5% percentiles [dotted lines] in the hypothetical example. $N=500$.

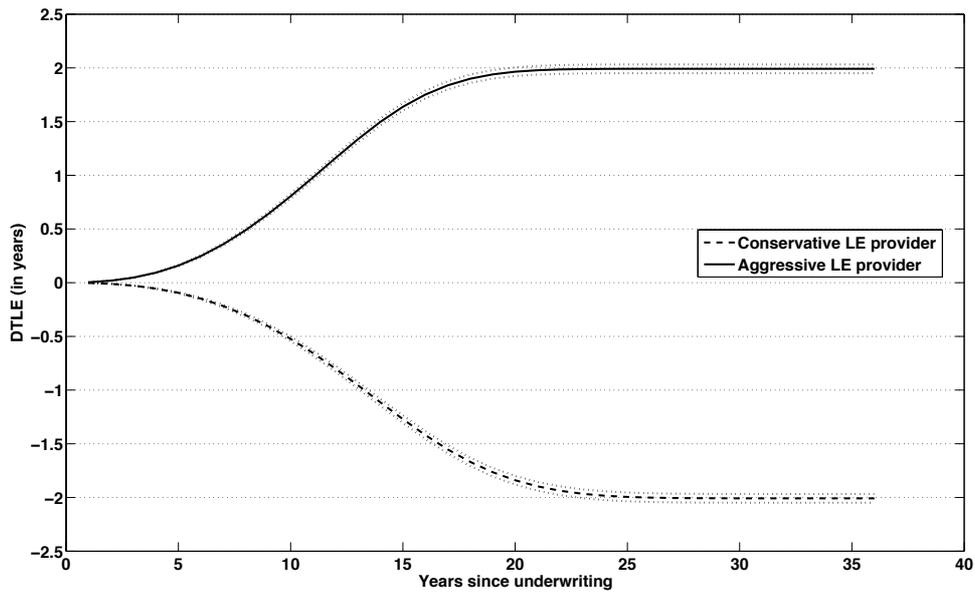


Figure 4: Development of DTLE for the aggressive LE provider (A) [solid line] with associated 2.5% and 97.5% percentiles [dotted lines] and the conservative LE provider (C) [dashed line] with associated 2.5% and 97.5% percentiles [dotted lines] in the hypothetical example. $N=50,000$.

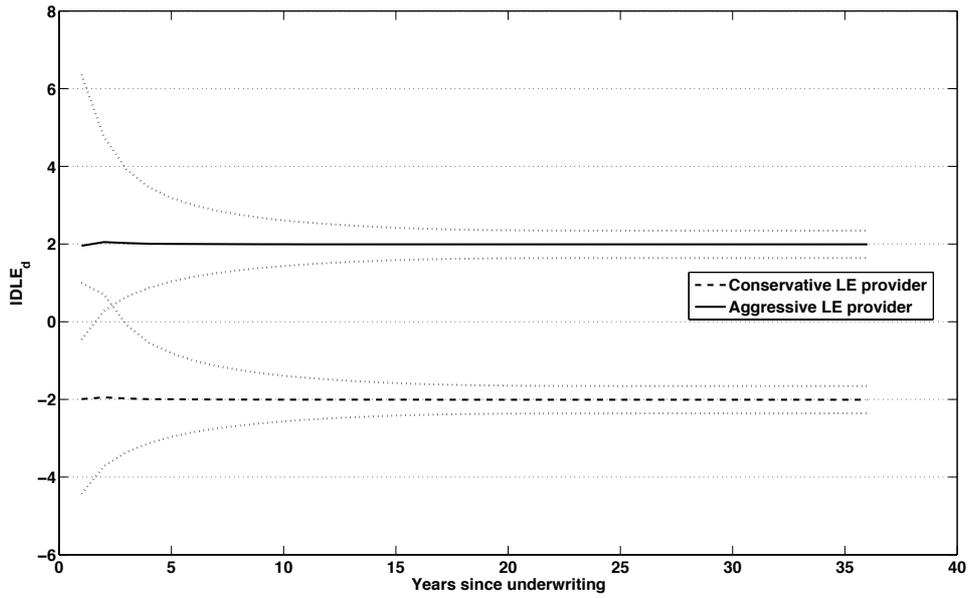


Figure 5: Development of IDLE for the aggressive LE provider (A) [solid line] with associated 2.5% and 97.5% percentiles [dotted lines] and the conservative LE provider (C) [dashed line] with associated 2.5% and 97.5% percentiles [dotted lines] in the hypothetical example. $N=500$.

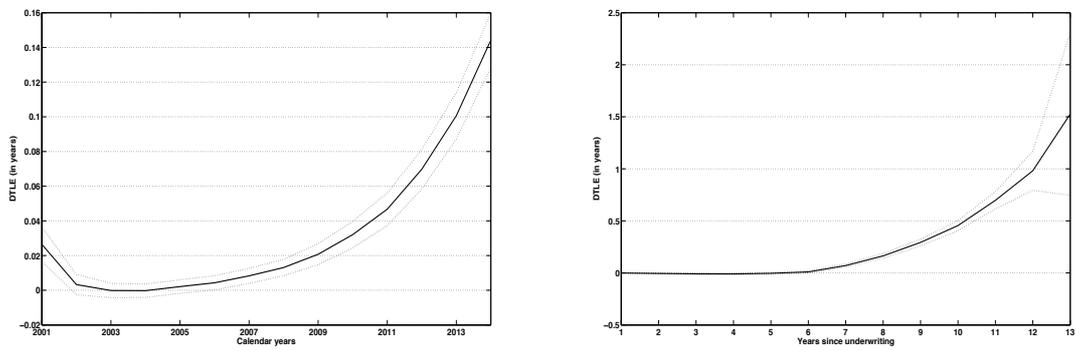


Figure 6: The development of DTLE for the Fasano dataset as a function of calendar year (left panel) or year since underwriting (right panel) (solid line) with associated 2.5% and 97.5% percentiles (dotted lines).

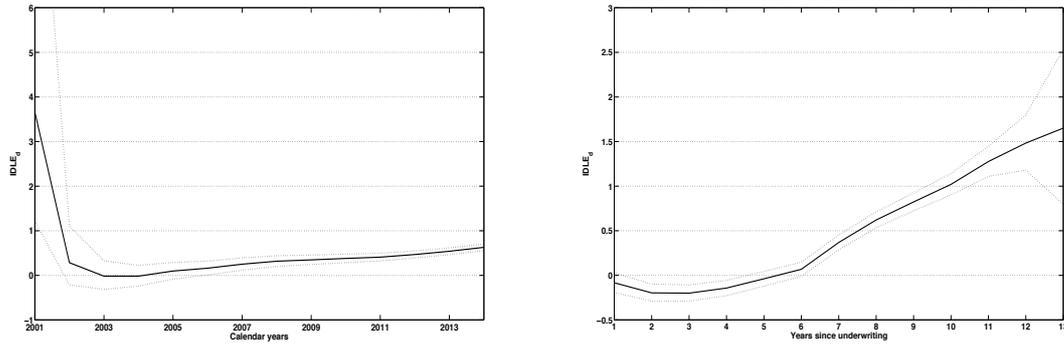


Figure 7: The development of $IDLE_d$ for the Fasano dataset as a function of calendar year (left panel) or year since underwriting (right panel) (solid line) with associated 2.5% and 97.5% percentiles (dotted lines).

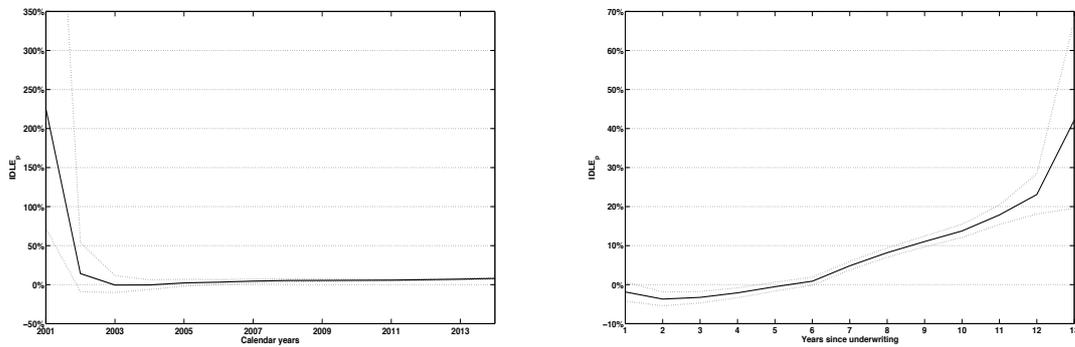


Figure 8: The development of $IDLE_p$ for the Fasano dataset as a function of calendar year (left panel) or year since underwriting (right panel) (solid line) with associated 2.5% and 97.5% percentiles (dotted lines).

	Entire Portfolio	Before January 1 st 2006	After January 1 st 2006
DTLE	0.144	0.423	-0.005
95% C.I.	(0.129, 0.160)	(0.386, 0.460)	(-0.019, 0.009)
$IDLE_d$	0.628	1.031	-0.034
95% C.I.	(0.556, 0.702)	(0.933, 1.131)	(-0.131, 0.064)
$IDLE_p$	0.080	0.135	-0.004
95% C.I.	(0.071, 0.089)	(0.122, 0.149)	(-0.016, 0.008)

Table 1: Summary statistics for DTLE, $IDLE_d$, and $IDLE_p$ and their associated 95% confidence intervals. The metrics are calculated as of January 1st 2015 for the entire Fasano portfolio, as well as two sub-portfolios that were underwritten before and after January 1st 2006, respectively.