Understanding the Volatility of Experience and Pricing Assumptions in Long-Term Care Insurance

Sponsored by
Society of Actuaries' Long-Term Care Insurance Section

Prepared by
Roger Loomis
Christopher Churchill
J. Michael Crooks
Julius Heim
Lyle Semchyshyn
Actuarial Resources Corporation of Kansas
May 2014

ARC
SOFTWARE ■ CONSULTING

Contents

Acknowledgements 2
Disclaimer 3
Introduction 4
Discussion about risk and uncertainty 4
Process Risk 7
Parameter Risk 20
Interest Rate Risk 28
Key Risk Measures 30
Other Product Designs 33
Pricing Margins and Rate Increases 41
Summary 44
Conclusions 45
Appendix 1: LTC Volatility Assumptions 47
Appendix 2: Beta Distribution 62
Appendix 3: Incidence Rate Scenario Factors 65
References 70

Acknowledgements

We offer our sincere thanks to the Society of Actuaries (SOA), the Long-Term Care Insurance Section, and the project oversight group for their guidance and feedback on this project:

Steve Schoonveld, Chair

James Berger

Sivakumar Desai

Bob Hanes

David Hippen

Perry Kupferman

Alex Silva

Barbara Scott, SOA Research Administrator

Steven Siegel, SOA Research Actuary

We also offer special thanks to the following individuals for their contributions to the ideas and production of this research: Melissa Boudreau, Rachel Brewster, Nate Campbell, Sam Gutterman, Stuart Klugman, and Brett Sundet.

Disclaimer

The opinions expressed by the authors are their own and do not represent any official position or opinion of the Society of Actuaries or its members. The Society of Actuaries makes no representation or warranty to the accuracy of the information.

This report was prepared by Actuarial Resources Corporation of Kansas (ARC), and is published by the Society of Actuaries (SOA). The information published in this report was developed from hypothetical products and assumptions. The SOA and ARC do not recommend, encourage, or endorse any particular use of the information provided in this report.

Any distribution of this report must be in its entirety. Nothing contained in this report is to be used in any filings with any public body, including, but not limited to state regulators, the Internal Revenue Service, and the U.S. Securities and Exchange Commission.

Actuarial Resources Corporation of Kansas (ARC), its directors, officers and employees, disclaim liability for any loss or damage arising or resulting from any error or omission in this analysis. The report is to be reviewed and understood as a complete document.

Introduction

As carriers have learned over the last two decades, Long-Term Care insurance (LTCi) is a risky business. The intent of this paper is to illustrate how these risks can be better understood through modeling the liabilities using a Monte Carlo simulation approach. Specifically, the paper will address the following questions:

1- How risky is a typical LTCi block of business?
2- How can an actuary differentiate between poor experience due to the inherent volatility in the underlying LTCi business and poor experience due to incorrect assumptions?
3- How can these risks be exacerbated or mitigated by product design modifications?
4- What implications does this have for pricing margins and triggers for rate increases?

There are also a number of items that are outside the scope of this paper.

1- This paper does not establish actuarial or industry standards for the underlying assumptions, nor should it be interpreted as doing this. The morbidity and mortality assumptions in this paper are illustrative only.
2- This paper does not define what specific pricing margins are appropriate for any specific product. The point is to illustrate how the question can be approached, not to provide the final answer.
3- This paper does not set a specific, numeric benchmark for regulatory approval of rate increases.

Discussion about risk and uncertainty

Risk vs. Uncertainty

Following the convention first established by Frank Knight in 1921,[1] the term "risk" should be used when we have some understanding of a probability distribution that can be measured, but do not know what value the random variable will take in future trials. By contrast, "uncertainty" refers to unknowns involving possible future contingencies with probability distributions that cannot be objectively measured.

The models in this study focus on risk. The authors of this paper will do this by incorporating variables that are known to impact the results and for which we have a reasonable basis to estimate the underlying probability distribution. Since the future in the real world can, and probably will, contain contingencies that cannot be foreseen, actual results can be outside of the ranges these stochastic models suggest. The purpose of the stochastic models is not to determine the "real" probability distribution of an LTCi portfolio. Rather, it is to provide a framework to help us better understand the volatility of LTCi and to compare the volatility and riskiness of different products and product portfolios.

There are three main sets of risk our model will consider:

Process Risk: If we knew *a priori* the precise transition probabilities (e.g., incidence, recovery, lapse, and mortality probabilities), there would still be risk associated with how those probabilities

[1] *Risk, Uncertainty, and Profit*, Frank Knight (1921)

unfold in any future time period. This is process risk. For example, in casino games where the probabilities of every event are perfectly known, the only risk left is process risk; while the casino knows that the roulette ball will land on 00 one out of every thirty-eight spins *on average*, it does not know how many times it will land on 00 in the *next* thirty-eight spins.

When random variables in a model capture only process risk, the variance of those random variables is called process variance. This variance establishes the boundaries on our ability to precisely forecast the future.

Parameter Risk: The level of confidence we have in a model's underlying parameters should be included in the analysis. For example, if we perform an experience study and observe that 10 out of 1,000 policies lapsed in a given year, we could say that the historical lapse rate is 1%. However, that does not tell us what the underlying *probability* of lapsing actually is; the actual probability of lapsing could be higher or lower than 1%, but still happen to produce 10 lapses in the 1,000 trials we observed. Just as the actuary should carefully evaluate what he thinks the best estimate of a given model parameter is, he should also express how confident he is in that estimate given the richness of the historical experience, likelihood of changed conditions, and so forth. For the example above, he could say that based upon the experience his best estimate for the underlying probability of lapse is 1% and that, after examining the likelihood function associated with the data, he is 95% confident that the actual probability generating these results is between 0.7% and 1.4%.

Economic Scenario Risk: This is the risk associated with unknown future interest rates and other economic scenarios. It is a type of parameter risk, but given its unique characteristics and the key role interest rates play in the performance of LTCi, we consider it separately.

Prediction Intervals

A prediction interval[2] is the range in which projected values are expected to fall, with a stated level of confidence.[3] The stochastic models described in this paper can make prediction intervals on any number of operational and financial metrics. The first model described incorporates only process risk, and is useful for understanding how much volatility is inherent in a specific metric for a given block of business and reporting period. Models that incorporate parameter risk and economic scenario risk will then be layered on top of the process risk model. These models give a more realistic view of the risks insurers face.

Modeled Risks

Long-Term Care insurance is among the riskiest insurance products sold. It is driven by more assumptions (incidence rates, recovery rates, lapse rates, death rates, utilization rates, inflation scenarios, interest rates) than most insurance products and those assumptions could change

---

[2] Actuaries generally prefer to think of their models as projections, not predictions. Nevertheless, the term "prediction interval" is well-established in the greater forecasting community and will be used here.

[3] In *The Signal and the Noise,* Nate Silver "encourages readers to think carefully about the signal and the noise and to seek out forecasts that couch their predictions in percentage or probabilistic terms. They are a more honest representation of the limits of our predictive abilities. When a prediction about a complex phenomenon is expressed with a great deal of confidence, it may be a sign that the forecaster has not thought through the problem carefully, has overfit his statistical model, or is more interested in making a name for himself than in getting at the truth." (Silver, Nate, *The Signal and the Noise: Why So Many Predictions Fail-but Some Don't* (pp. 404-405).)

significantly between current observations and what might eventually materialize before the insurance contract ends. Our basic understanding of many of these assumptions is limited by the relatively short time LTC insurance has been in existence, the limited availability of public data, and by the significant changes that have occurred in the market during that time. It is a long-duration product, featuring steeply increasing benefits typically funded by level premiums. This implies that it will have a large policy reserve, and that a significant part of the benefits will be funded by investment returns. This adds the risk of changes in the future investment environment. With permanent life insurance, the cash value, reserve, and death benefit all converge at the ultimate age, causing the risk to vaporize as it becomes financially irrelevant whether the policy ends due to death, surrender, or maturity. LTCi typically does not have a cash value. This gives policyholders a perverse use-it-or-lose-it incentive to go on claim, and causes profits to emerge in discrete windfalls when policies end due to death or lapse. This in turn makes the risks caused by anti-selection, and accurately predicting mortality and lapse especially significant to LTCi.

In principle, any uncertain element about the future could be included in a stochastic model provided you have a way to describe its probability distribution. The products described in this paper have been simplified to emphasize the main drivers of risk and as such have the following characteristics:

- Gross premium is equal to the valuation net premium; i.e., expenses and profit margins are ignored.
- The model includes net-level policy reserves and claim reserves, discounted at 3.5%.
- Assets are equal to reserves; i.e., no additional capital.

These and other simplifications we make will have an impact on the results of the model. This enables us to focus the analysis on specific risks without having to adjust for the impact of other aspects of the model also influencing the results. Where appropriate, we will address how these simplifying assumptions may have influenced the results.

Thus, the only stochastic elements in this model are the transition probabilities (incidence, recovery, lapse, and death) and economic scenarios. The transition probabilities will first be modeled to evaluate the process risk, and then parameter risk and economic scenario risk will be layered in.

Basic Stochastic Functionality

The model has monthly time steps. At the beginning of every time step, every policy is in one of two statuses:

- Active
- Disabled

If the policy is active, one of four things can then happen:

- The policy will lapse
- The policyholder will die
- The policyholder will go on claim
- The policy will remain active

If the policy is disabled, one of four things can then happen:

- The policyholder will die
- The policyholder will recover
- The policyholder will remain on claim
- The policy will end because the ultimate benefit level is reached

To visualize the process that repeats every month of the policy's life, imagine a line segment from 0 to 1 is divided into four segments (or three, if the policy is disabled), with the lengths corresponding to the probability of each possible state for the next period. The probabilities depend on age, sex, policy duration, etc. A random number is then drawn from a uniform distribution between 0 and 1. The number is compared to the line segment to determine the transition. If the policy lapses or dies, the modeling of that policy ends. Otherwise, the process is repeated for the next time step.

Depending upon the details of the specific policy, premiums are collected, benefits are paid, and reserves are established as if an actual policy were moving across time according to these transitions. This model is well-suited to accurately model path-dependent benefits because the specific history of the policy being modeled can be tracked. It is assumed that assets are always equal to the reserves.

After this process is finished for every policy in the block, the aggregate results show a plausible scenario of what the operational metrics and resulting financials might look like over time. If this process is repeated several times, the probability distributions of every operational metric and financial result in the model can be estimated. To illustrate this, we will examine several variables.

Process Risk

We will begin by analyzing process risk. This will be done by looking at different operational and financial metrics in a stochastic model that assumes the transition probabilities are known. Afterwards, parameter risk and economic scenarios will be layered into the analysis.

When analyzing financial and operational metrics, management should have an accurate understanding of how much process risk they are exposed to. Process risk is driven by several things:

- The probability distributions that drive mortality, morbidity, and lapse
- The benefits of the insurance contract
- The number of policies
- The demographics of the insured population
- The length of the reporting period

Management should have an accurate understanding of how much process variance is entailed in every operational and financial metric they analyze. If an observed metric is different than its expected value but still within the expected range, the discrepancy from the expected value can be attributed to process variance rather than incorrect assumptions or business issues.

Basic Policy: Comprehensive LTCi, 5% Compound Inflation, Indemnity Benefits

We will begin by analyzing a block of long-term care policies with the following characteristics:

- All policies are indemnity (i.e., pay the full cash benefit, regardless of care setting)[4]
- Original Benefit Amount: $3050 per month
- Elimination Period: 90 days
- Benefit Period: 5 years
- Inflation Protection: 5% Compound

The block contains 40,000 policies that are all sold on January 1, 2014. The policies have the following demographics:

Gender	Percentage
Male	26%
Female	74%

Issue Age Band	Percentage
50-59	20%
60-69	54%
70-79	26%

This policy design will be referred to in this paper as the "basic policy." We will analyze its risk, suggest ways the risk can be analyzed, and suggest key metrics for measuring risk. We will then look at the key risk measures of four other policy designs:

- Shorter benefit period with less inflation protection
- The basic policy with a return of premium rider added
- The basic policy with life combo-product feature
- A policy where premium rates and benefit levels are both indexed to an inflation index

Distribution of Lapses

Unlike many other operational and financial metrics in the model, the probability distribution of the total number of lapses for the first projection month has a simple closed-form solution. In this section we will show how to calculate that probability distribution and use it to calculate a 90% prediction interval for the total number of lapses. We will show how that distribution can be estimated using simulation. This will illustrate how simulation works, and will help provide a sense for how precisely it works in other situations where closed-form solutions are not available.

We assume the annual probability of lapsing in the first year is known to be exactly 5.5% for all 40,000 policies, which corresponds to a monthly rate of 0.4703%.[5] This means that the total number of lapses has a binomial distribution with N=40,000 and p=0.004703. The mean is 188.12,

[4] While some LTC policies are in fact indemnity policies as described here, most actual policies are expense incurred, meaning they will only reimburse actual expenses up to the daily maximum. Like any other risk, the parameter risk associated with the utilization rate for reimbursement policies could be included in stochastic models as well. Depending upon the interplay between interest, inflation, and salvage factor assumptions, and with the specific policy benefits, expense-reimbursement policies may have higher or lower risk than indemnity products.

[5] 1-(1-0.055)^(1/12) = 0.004703

and using the normal approximation of the binomial distribution, we can be 90% confident that the total number of lapses should be between 165 and 210.[6]

Solving the same problem through simulation, we simulate how many of the 40,000 policies lapse in the first month, and repeat that several times. In the first simulation, 176 policies lapse, because of the 40,000 random numbers that were selected, exactly 176 of them fell on the part of the line segment corresponding to lapse. This was repeated for 200 simulations, with observed lapse counts of 185 lapses, 197 lapses, 215 lapses, etc.

Each of those simulations had an equal probability of occurring. The average number of lapses across the 200 simulations was 188.52. Of the 200 simulations, the central 180 resulted in lapse counts between 165 lapses and 213 lapses. This implies that we can be about 90% confident that the real number of lapses observed will be between 165 and 213, provided our parameter assumptions are correct.

To illustrate how well the simulation model works, we can compare a histogram of the simulated distribution to a histogram of the actual binomial distribution in Figure 1.

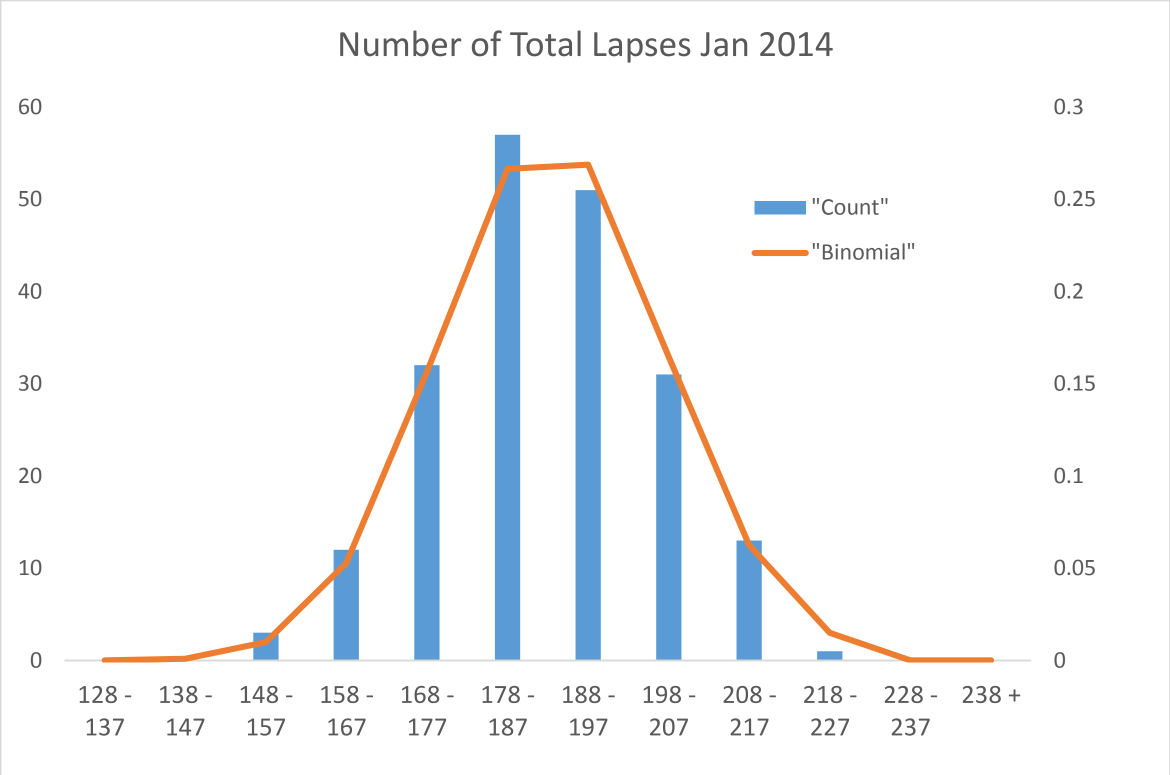

Figure 1: Simulation compared to binomial distribution

The implication of this exercise is that while it is true that the expected value of the number of lapses is 188.12, we can be only 90% confident that the actual number of lapses will lie somewhere

[6] $Standard\ Deviation\ = \sqrt{(0.004703)\ (1 - 0.004703)\ (40{,}000)} = 13.68 \quad F(0.95) = 1.645$
$90\%\ Confidence\ Interval = 188.12 \pm 13.68 \times 1.645 = [165.6, 210.6]$

in the range of 165 to 210; this is given we know the true probability of lapsing. The risk that lapses will not match the expected value, but that the deviation is caused by random fluctuations around perfectly known probabilities is what we are calling *process risk.*

One thing to note is that this range of 165 to 210 is a very narrow range relative to the 40,000 policies being modeled. Along with mortality over the same time period, this results in February premium having a 90% prediction interval ranging from $20.41 million to $20.43 million, or in other words, the expected value of premiums plus or minus 0.06%. Thus, premium forecasts are normally quite accurate.

Profitability over Time

We will now look at the probability distributions of several operational and financial metrics for one-month periods, usually near the beginning of the projection. We will then look at metrics for longer time periods and later in the projection. It is important to keep in mind that these distributions depend on the specific business mix being evaluated. The results will be different after the block has aged, and they would be different for other blocks as well.

While the variance of the number of lapses has a closed-form solution using the binomial distribution in this simple model, other metrics quickly become difficult to calculate without simulation. For example, the profit earned over a time period is a function of all of the transition probabilities. These transitions trigger changes in premiums, claims, and reserves, which in turn drive the change in profits.

Every component of every profitability metric that is a function of policy transitions is a random variable in its own right, the distribution of which can be estimated using simulation. In this section, we will look at how claims and reserves affect the aggregate process risk of earnings over a specific period.

Paid Claim Analysis

To illustrate some of the considerations needed, we will analyze paid claims for January 2015. The histogram in Figure 2 shows the distribution of paid claims for January, 2015.

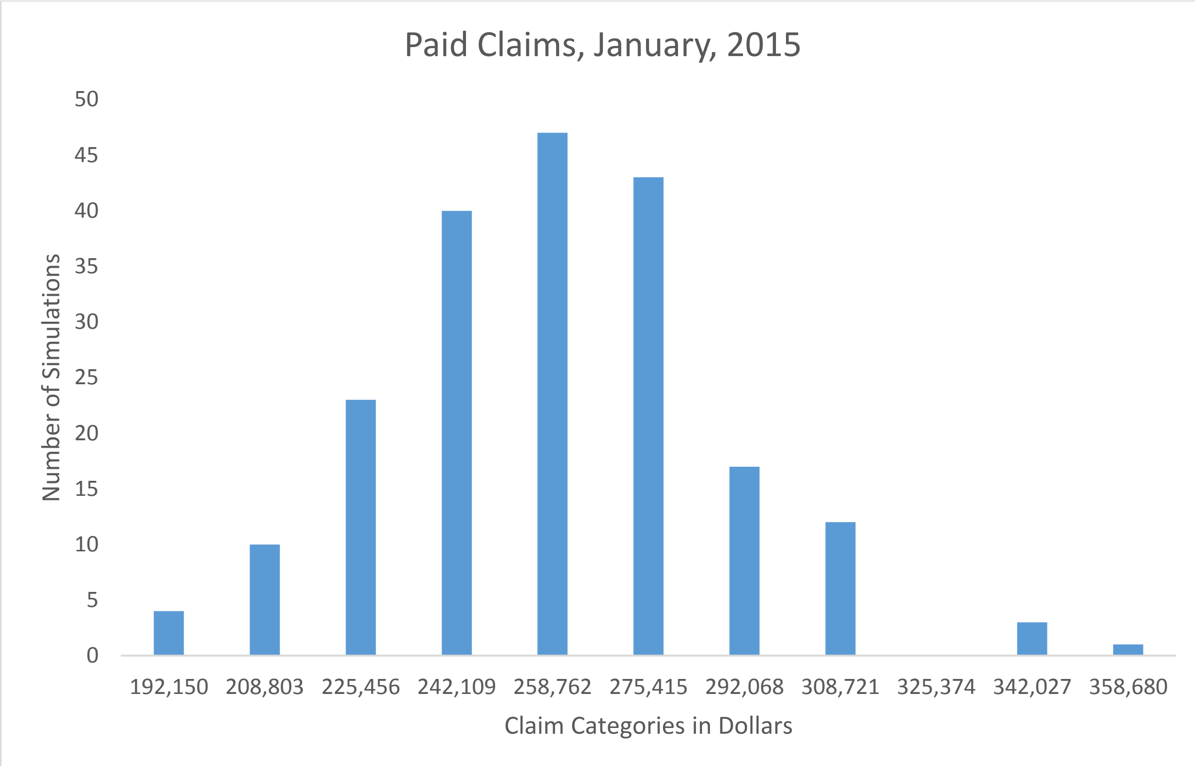

Figure 2: Histogram of Paid Claims, January 2015

The expected value of paid claims is $266,304, with a 90% prediction interval of $220,812 to $314,005. The sample standard deviation is $29,096.

A principle objective of this study is to suggest metrics that can be used to describe the relative riskiness of a block of business. A useful statistic might be the Relative Standard Error, which is the sample standard error (i.e., estimate of the standard deviation) divided by the sample mean.[7] In this case, the Relative Standard Error for the Random Variable "January 2015 Paid Claims" is:

$$\text{RSE}=\frac{s}{\overline{X}}=\frac{29{,}096}{266{,}304}=10.93\%$$

One way in which this analysis could be used is to define pricing margins and triggers for rate increases more objectively. If claim experience materializes within the 90% prediction interval, the experience could be considered moderate with respect to expectations, or simply "moderate." Within this range, there is no statistical evidence that any of our underlying assumptions are wrong, and any deviation from the mean can be attributed to process variance.

The Central Limit Theorem implies that this paid claims random variable has a Gaussian distribution. The 95th percentile of the standard normal distribution is 1.64. Thus, if we multiply

[7] If we were dealing with the entire population rather than a sample, the corresponding statistic is called either the Relative Standard Deviation or the Coefficient of Variation.

the Relative Standard Error by 1.64, we get the threshold that distinguishes moderate experience from experience that is beyond moderate.

We will call this value the *95th Percentile Claims Margin*:

$$PCM_{95} = RSE \times 1.64 = 17.92\%$$

This implies that for the block of policies and month in question, if claims are more than 17.92% higher or lower than expected, the experience is beyond moderate.

Another way to interpret this is through a hypothesis-testing paradigm. The null hypothesis is that the best-estimate assumptions are correct. Under that assumption, we can be 95% certain that claims for this month will be less than 117.92% of the expected value. Claims higher than that constitute statistical evidence that the underlying assumptions are wrong.

There are a couple of caveats here. First, the factor $PCM_{95} = 17.92\%$ is specific to this unique block of business for this month only. For blocks with other sizes, demographics, and benefits, the factor could be quite different.

Second, it is important to recognize how this differs from hypothesis testing. Under hypothesis testing, a single test is created and then performed. If the results of that single trial fall beyond the critical value, the null hypothesis is rejected. In contrast, with the process described here, the PCM_{95} is being continuously monitored. Thus, if the null hypothesis is correct, we would expect the results to be higher than the CM_{95} five percent of the time. So if this metric is being continuously monitored as a block of business matures, occasional periods with extreme experience are to be expected.

One could argue that if the probability of being beyond the CM_{95} is 5% for one month, the probability of being beyond it two months in a row is $0.05 \times 0.05 \times 100\% = 0.25\%$. This approach is only valid if the two months are statistically independent. That is unlikely, given the fact that most of the individual claims that caused the high claims the first month will still be on claim the second month.

To illustrate, the null hypothesis for January would be rejected in 12 of the 200 scenarios. If we repeat the procedure for the next month, there are nine scenarios in which the null hypothesis would be rejected. Of the nine scenarios that had disappointing results in February, eight also had disappointing results in January.

It is important to understand that this correlation of one month's results with the next is driven by the nature of the benefits; if there is an adverse month where process variance causes a high number of new claims, it will take several months for that bad luck to be absorbed into the system so that you no longer see elevated paid claims.

Incurred Loss Ratio Analysis

An alternate metric to consider is incurred loss ratio, defined as incurred claims (paid claims plus the change in the claim reserve) divided by earned premium. Using this metric, the same process as above could be implemented.

The loss ratio for this month is expected to be 16.79%, with a standard deviation of 5.07%. This implies that the prediction interval for the incurred loss ratio is $16.79\% \pm 1.645 \times 5.07\% = 16.79\% \pm 8.34\%$.

This result might be surprising; the Actual-to-Expected value for the loss ratio could be as high as 149% and still be considered moderate experience with the variance attributed only to process variance. In comparison, the Actual-to-Expected ratio for paid claims for the exact same block was only 114.74%.

The difference between these two numbers is the claim reserve. Every time a policy goes on claim or goes off of claim, a claim reserve is established and then released. These claim reserves can be quite large. The claim reserves ensure that the full liabilities are acknowledged on the balance sheet, as intended. A side effect is that they tend to have a high process variance, thus making earnings relatively unpredictable.

To illustrate the phenomenon, the following graphs show paid claims and the total of paid claims plus the increase in claim reserves, which is the definition of incurred claims.

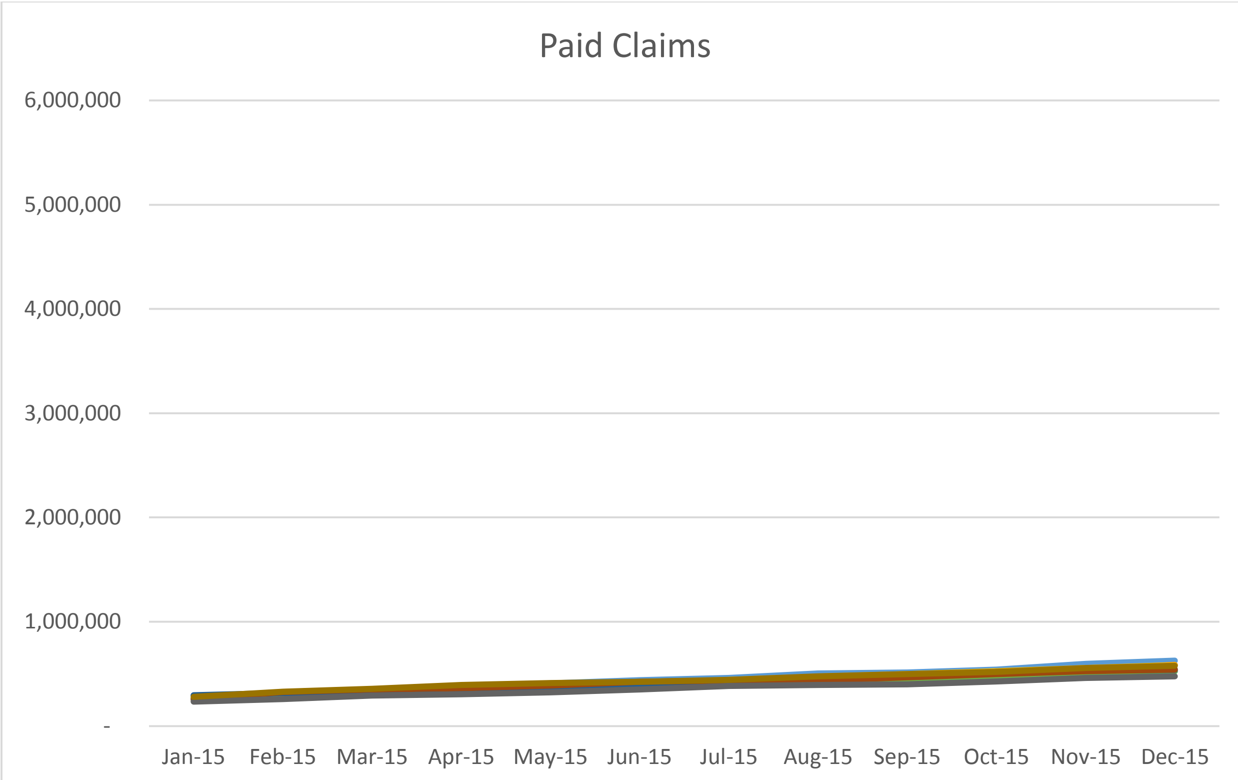

Figure 3: 10 Simulations of projection of future paid claims has low variance month-to-month

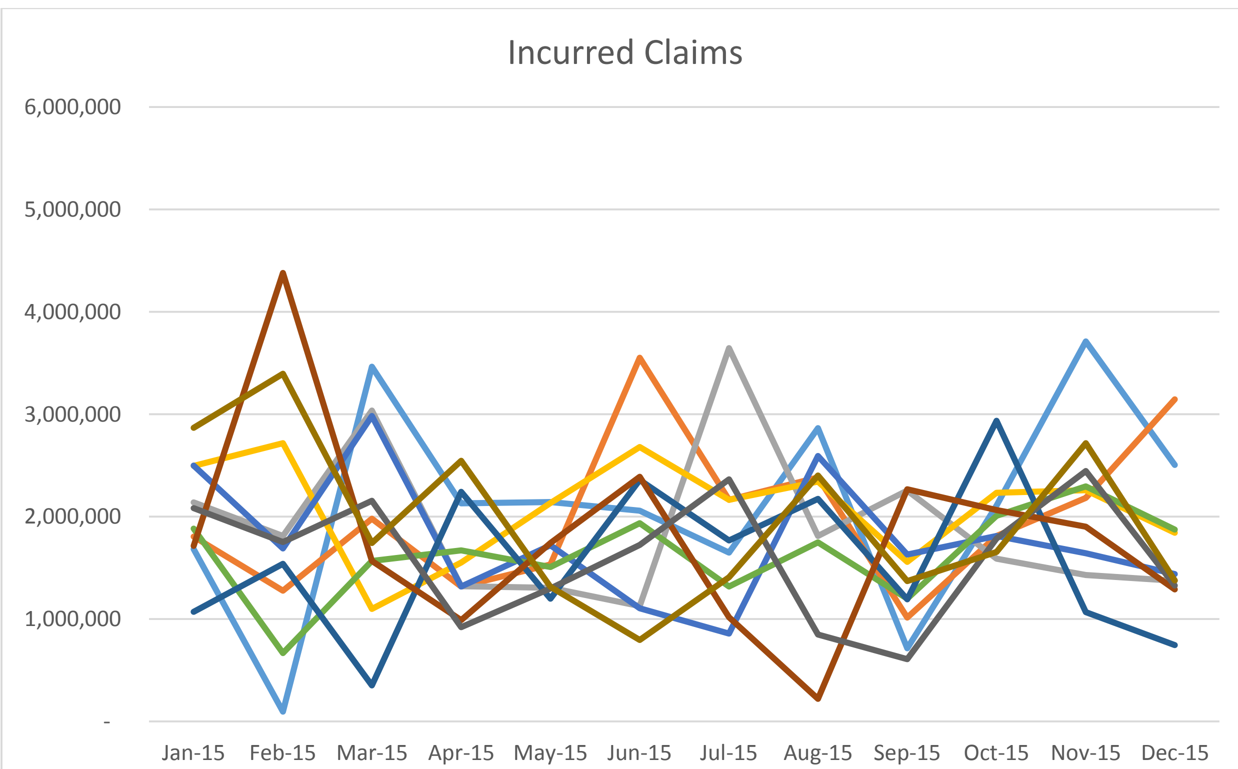

Figure 4: 10 Simulations of future incurred claims has high variance, driven by variance in claim reserves

The shifting of the claim reserves tends to dominate the claim payments. The extent of the domination depends upon the specific characteristics of the block for the time period being examined.

As with paid claims, we can use the 95th-percentile loss ratio margin ($PLRM_{95} = RSE \times 1.64 = 49.53\%$) to differentiate moderate experience from extreme experience. In this particular case, if experience is more than 149.53% of the 16.79% loss ratio we are anticipating (i.e., a loss ratio greater than 25.10%), the experience for the month is extreme.

It still holds that even if the assumptions are true, process error alone will cause about one of every twenty months to show extreme experience. Thus, although a single month of high loss ratios should raise a flag, it is not enough to statistically prove that the underlying assumptions are not valid.

In contrast with the paid claims statistic, there is not an inherent positive correlation across months with regards to the loss ratios. The question is, how many extreme months should be required before the null hypothesis that the best-estimate assumptions are correct is rejected?

If the null hypothesis is correct, about once every 20 months one should expect an extreme result. After that happens, there is about a 5% chance that the subsequent month will be extreme, too. Again, if the null hypothesis is true, we would expect the process variance to produce two extreme months in a row $0.05^2 = 0.25\%$ of the time, or about once every 400 months, or 33.3 years. Given the long-term nature of these contracts, process error is more likely to cause that to happen at some time over the life of the block than not. However, the probability of the process error causing this to happen three times in a row is $0.05^3 = 0.0125\%$ of the time, which is about once every 8,000 months (666.7 years).

Profit

How much process variance is on this product's bottom line?

In this model, Profit is defined as:

$$Profit = Premium + Inv\_Inc - Inc\_Pol\_Res - Inc\_Clm\_Res - Claims$$

Since we are using net premiums as the basis for our modeling, our expectation for profits should be that they are zero. However, we used deterministic methods to develop the premiums for our product that did not exactly match up with the probabilistic values in the Monte Carlo simulation. We adjusted for this at a macro level so that overall results were appropriate. But at a micro level, over short time spans, the emergence of claims and reserves do not exactly match up. This does not impact our analysis of the variability of profits other than to shift the mean away from zero.

For January 2015, the expected profit is $2.8M (million), with a standard deviation of $1.0M. This implies the 90% prediction interval for profits is:

$$PI(Profit)_{95} = \overline{X} \pm 1.645 \times s = [\$1.16M, \$4.46M]$$

In other words, we are only 90% certain that profits will be somewhere between $1.16M and $4.46M. Further, we should expect that due to process variance alone, the profit should be outside of that range 10% of the time.

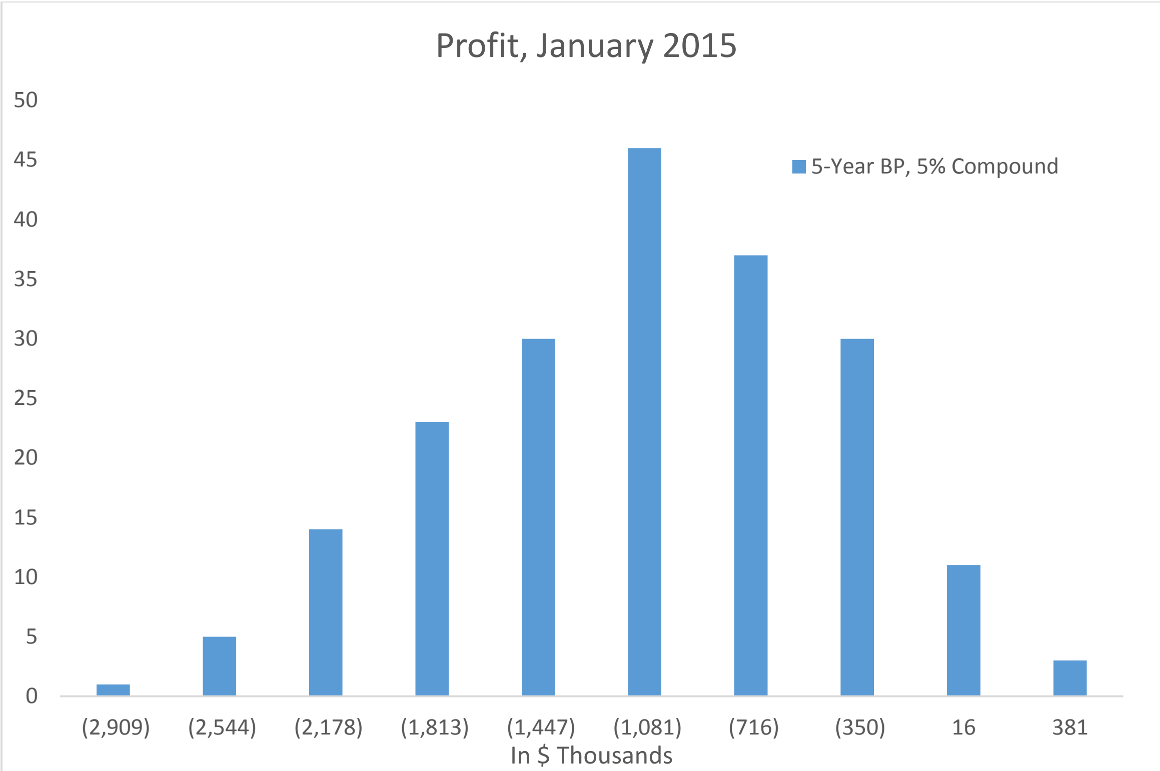

Figure 5: Profit in January, 2015

The total variance of profits is the sum of the variance of the underlying components plus the covariance. The model assumes that all of the underlying transition probabilities are independent. Even with that assumption, there is a correlation of these pieces on the income statement. For example, if claims are higher, then the claim reserves are also likely to be higher. The question is, which components of the income statements have the largest impact on the total variance?

The total variance of the total profit is given by:

$$VAR(Profit) = VAR(Premium) + VAR(Inv\_Inc\ ) + VAR(Inc\_(Pol\_Res\ )\ ) \\ + VAR(Inc\_(Clm\_Res\ )\ ) + VAR(Claims) + Covariance$$

In this formula, “Covariance” represents the total of all of the covariance terms in the full formula. The simulation model can be used to calculate estimates of terms in this formula, except covariance. We subtracted the variance of profit from the variance of the underlying components to back into the covariance. The graph below shows the standard deviation of each component of total variance for January 2015.

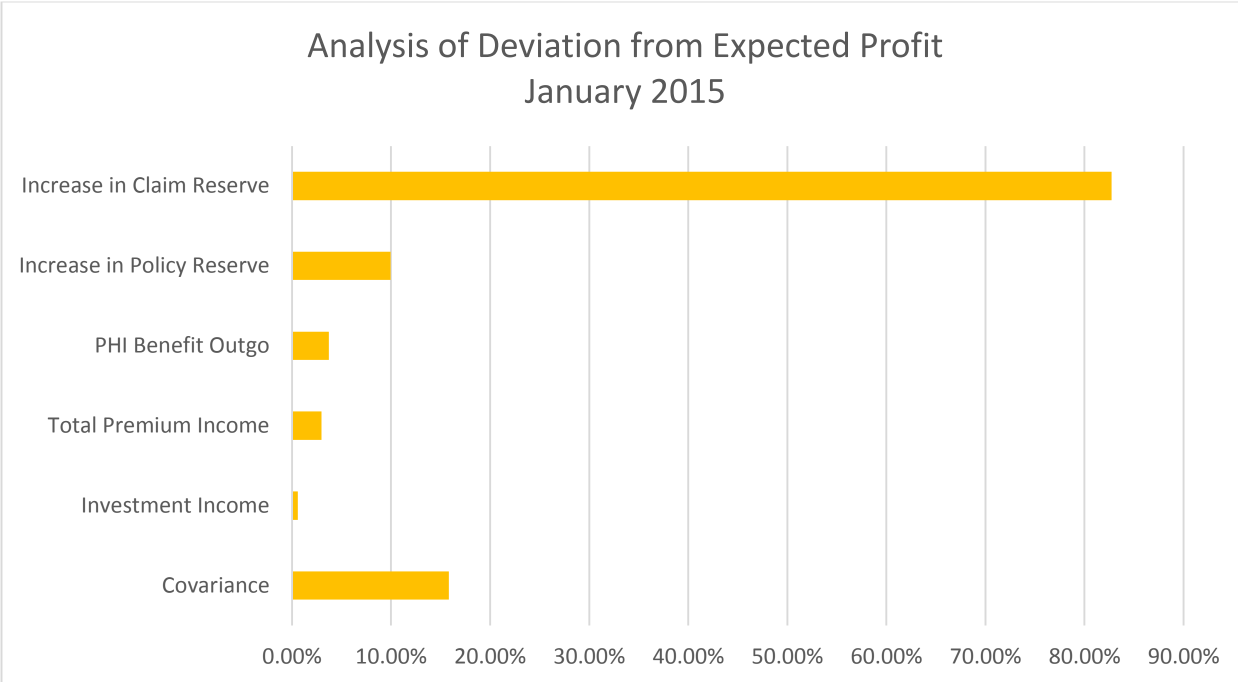

Figure 6: Analysis of Deviation from Expected Profit January, 2015

83% of the total standard deviation is due to the change in claim reserve, with the next highest element being the covariance, followed by increase in policy reserve, which comes in at 10%. Reserves tend to have a naturally high level of process variance.

Different Time Periods

The results we have looked at so far were all for a specific time period—January 2015, which is both the 13th month of the policies' lives, and the 13th month of the projection. The results can be substantially different if we look at longer reporting periods.

To illustrate this effect, we will look at the model's profit margin random variable across different time periods. This random variable is simply the time period's profit divided by its earned premium. A deterministic model could tell us the expected value of the profit margin for different time periods, but we need a stochastic model to see the risk. The process risk for the profit margin is proportional to its standard deviation.

The following table and graph show the standard deviation of this metric for four different time periods for 2014: the first month, quarter, half, and the full year:

Reporting Period	SD of Profit/Prem
Jan-14	2.69%
1Q 2014	1.49%
1H 2014	1.04%
2014	0.77%

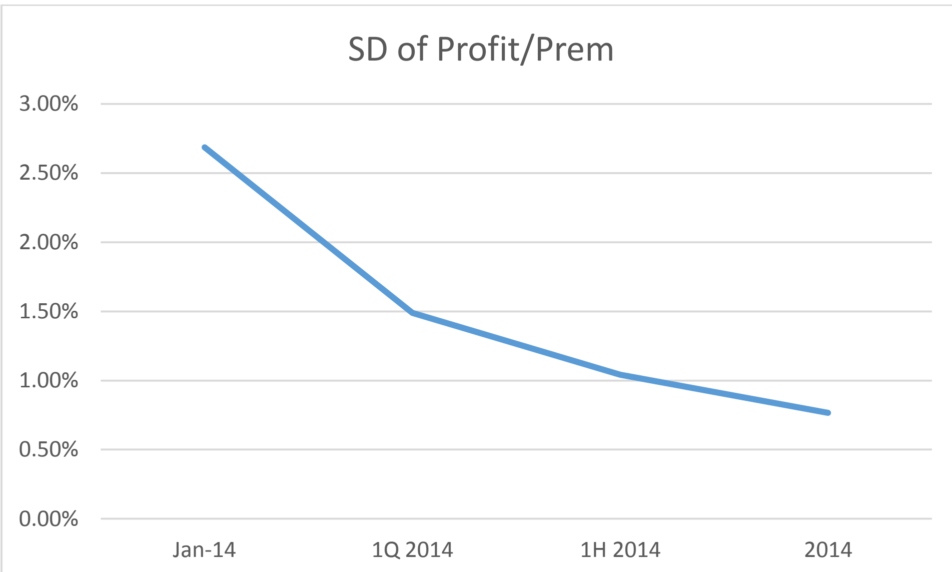

Figure 7: Standard Deviation of Profit Margin

When we increase the reporting period from monthly to quarterly, the standard deviation decreases by almost 50%. Going from quarterly to annually, it declines by about 50% again. These results are expected and should make intuitive sense. In general, if a sample size is increased by a factor of four, the relative standard error (RSE)[8] decreases by 50%.

It is worth pointing out that many companies do monthly reporting in order to pay closer attention to how the experience for the quarter is developing. When doing so, it has to be remembered that the signal-to-noise ratio[9] is much lower for monthly results compared to quarterly results. For example, if it is expected that quarterly results have a loss ratio of 70% plus or minus 10%, it should be expected that the monthly results will have a loss ratio of 70% plus or minus 17%.[10]

This diversification across time works in the other way as well. If the quarterly results are expected to have a loss ratio 70% plus or minus 10%, then the annual results should be expected to have a loss ratio of 70% plus or minus 5%.

---

[8] Relative Standard Error is defined as standard deviation divided by mean.
[9] Using the term colloquially, as used by Nate Silver.
[10] If the sample size increases by a factor of 3, then the expected value and variance increase by a factor of 3. The standard deviation increases by a factor of $\sqrt{3} \approx 1.7$.

Present Value of Future Profits

The ultimate extension of diversifying a policy's profits across time is to look at the present-value of expected profits at issue. The 95% confidence interval of discounted profits divided by discounted premiums for this block is 0.15% plus or minus 0.79%. This is a relatively narrow range, and indicates that despite the bumpiness in the profits along the way, the process variance is limited over the lifetime of these policies. Also note that over the life of the policy the expected value of profits is much closer to zero, as it should be. The residual amount is due to sampling error between this sample and the sample used to set the premium originally.

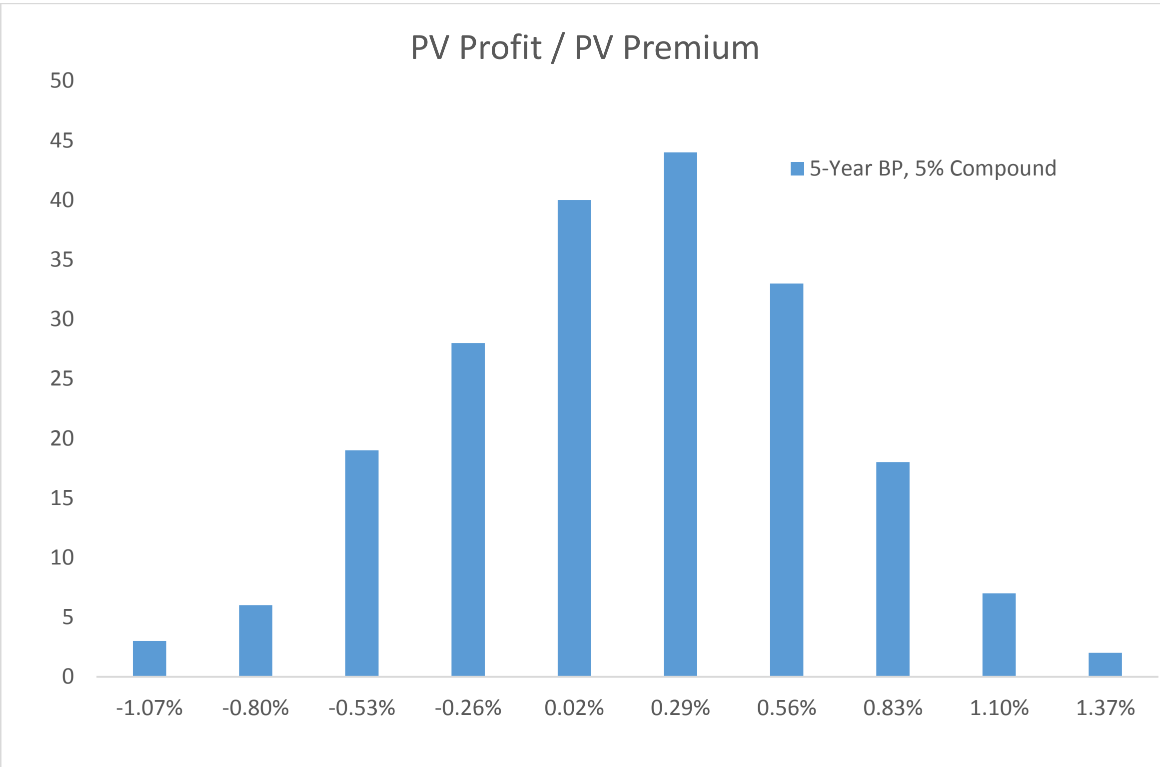

Figure 8: PV Profit/PV Premium

Parameter Risk

The above analysis of process risk shows that for a moderately sized block of business, process risk can cause significant variance in certain financial and operational metrics over short periods of time, but that over longer periods, that variance is largely diversified away. If the level of process risk is not correctly understood, there can be unreasonable expectations for smooth short-term results. Over the long term, process risk has a tendency to naturally even itself out.

In contrast, parameter risk is a much more serious issue. If parameter estimates do not correspond to the forces that actually drive morbidity in the real world, then actual results will differ from the model's predictions in a way that cannot be diversified.

Because the parameters plugged into forecasting models are merely estimates, the level of confidence we have in the parameters should be represented in forecasting models. In this section, we will provide an overview of how we incorporated parameter risk into this model and discuss some of our considerations.[11] We will then show the updated model's results and compare them to the results of the process-risk model.

For the illustrative purposes of this model, the parameter risk associated with claim incidence rates and lapse rates are explicitly modeled. The techniques described here may be used to model the parameter risk of any model assumption, including mortality, recovery, and claim utilization level. The focus should be on the assumptions with a high degree of uncertainty and to which the model's results are highly sensitive.

Parameter Risk and Sampling Risk

Parameter risk refers to the risk that the parameters in the forecasting model are not appropriate for what we are trying to model. This could be driven by either sampling risk or data bias. Sampling risk is the result of differences between the sample and the population. Data bias is the risk that the parameters are wrong because they are not trended properly.[12]

LTC is prone to both types of parameter risk. Because so much of the risk is concentrated at the extreme ages for which there is little historical (sample) data, sampling risk can be especially high at those ages. Changes that have occurred to things like underwriting standards and claims adjudication can cause experience data to vary from that underlying the current population. Censoring of the data which is naturally occurring due to policies only having been in force for a relatively short amount of time is another factor that adds to sample risk.

LTC is susceptible to data bias as well. We do not know what future mortality and morbidity will look like. If mortality continues to improve, will the additional years of life be generally healthy, or will they be additional poor health years, prone to longer periods of care? We simply do not know.

To incorporate parameter risk into forecasting models, there needs to be a way to incorporate a set of several plausible transition probabilities, weighted by the relative likelihood that any one of them represents the "real" unobserved forces of morbidity and mortality.

---

[11] See Appendix 1 for a detailed explanation of the model's assumptions.

[12] This risk classification scheme is from the Casualty Actuary Society's study note, "A Note on Parameter Risk" by Gary Venter and Rajesh Sahasrabuddhe,

The model in this paper extensively uses Beta distributions to model parameter risk. The Beta distribution is well suited to model sampling risk, because its shape is very similar to the likelihood function associated with the sample.[13]

Parameter Risk and Model Misspecification

When parameterizing a model, the true parameter risk can be hidden by model misspecification. This can create the unexpected impression that adding statistically significant covariates to a model will cause the total parameter risk to go up.

To illustrate this, assume that we had 10,000 units of exposure in our experience period, and observed exactly 100 claims. What is the underlying *probability* of going on claim that generated this experience? The true probability could be described as a random variable itself. If we assume the 10,000 exposure units are homogenous, the probability of going on a claim P has a Beta distribution with parameters 100 and 9,900 (i.e. $P \sim$Beta (100,9900)). This implies our best estimate of the actual probability of going on claim is 1.0% and the standard deviation of this estimate is 0.10%.[14]

Upon further investigation, we discover that the data set has exactly 5,000 males and 5,000 females. We then notice that 80 of the 100 claims were generated by females. This clearly implies that our original model was misspecified.

Specifying the model correctly, we see that the force of morbidity for the females is Beta (80,4920), implying that we can be 95% confident that the true incidence rate is really 1.6%, with a standard deviation of 0.18%. For males, the morbidity is given by Beta (20,4980), which has a mean of 0.40% and a standard deviation of 0.09%. For females, correcting the model misspecification causes the standard deviation of the estimated incidence rate increased from 0.10% to 0.18%, an 80% increase.

This may seem paradoxical; why would correctly specifying a model cause the parameter risk to increase? The answer is that the actual parameter risk did not change; we have simply replaced an erroneous model that significantly understated this risk by a more refined model that properly recognizes it.

Parameter Risk: Basic Incidence Rates

The model's incidence rates and related factors were calculated from our set of hypothetical experience using Generalized Linear Modeling (GLM). The model resulted in a set of base incidence rates and factors that are similar in structure to what many companies use. We could have arrived at essentially the same incidence rates by dividing claims by exposure for the various cells and then smoothing the results. We used GLM instead for two reasons:

---

[13] See Appendix 2 for a description of the Beta distribution and why it is useful for describing parameter risk.
[14] As explained more fully in Appendix 2, if we observe α successes and β failures from a series of independent and identically distributed Bernoulli trials with an unknown probability of success P, then P has a Beta distribution with a mean of $\frac{\alpha}{\alpha+\beta}$ and a variance of $\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$.

1. The GLM gives objective statistical indication about whether or not the various factors considered in the model are significant

2. The GLM provides standard errors and correlations between the various factors, which helps determine the level of sample risk inherent in the data

Our data consisted of about 2 million exposure years of LTC claims data. The GLM provided us with the following attained age rates and multiplicative adjustment factors for duration and gender:

Attained Age	Estimated Rate	Lower CL 5%	Upper CL 95%
< 40	0.00031	0.00008	0.00112
40-49	0.00018	0.00008	0.0004
50-59	0.00068	0.00056	0.00083
60-64	0.00124	0.00105	0.00147
65-69	0.00224	0.00199	0.00253
70-74	0.00553	0.00507	0.00603
75-79	0.0142	0.01324	0.01523
80-84	0.0339	0.03175	0.0362
85-89	0.06729	0.06207	0.07296
90+	0.1199	0.10116	0.1421

DURATION	Estimated Rate	Lower CL 5%	Upper CL 95%
1	0.56404	0.50212	0.63359
2	0.75192	0.67886	0.83284
3	0.91277	0.83118	1.00236
4	1.02683	0.93854	1.12342
5	1	1	1

Gender	Estimated Rate	Lower CL 5%	Upper CL 95%
Female	1	1	1
Male	0.70224	0.65875	0.7486

To interpret the results of the GLM, we begin with the assumption that the model is in fact specified correctly—that we have the correct parameters in the model, and the correct functional form. If those assumptions are correct, then we can be 90% confident that the actual probability of incurring a claim is within the range given.[15] For example, we can be 90% confident that the true probability of a claim incurring for a female age 70-74, in policy duration 5 or higher, is between

[15] Technically, these confidence intervals apply to the factors themselves. Only in the base cases (any age; duration 5+; female) can you directly make a probability statement about the incidence rate itself.

0.00507 and 0.00603. To correctly interpret these results, a couple of issues relating to covariance need to be understood.

Sample Size

These metrics can overstate the sampling risk because they are based on the assumption that the experience of insureds among various cells are independent. In reality, the probability of claims by people 65-69 and 75-79 has bearing on the probability of a claim incurring by people in the range between. So the additional confidence we receive from the information given by the surrounding cells should be incorporated into the model.[16]

To recognize this additional confidence, the standard error of each cell of the GLM was divided by a somewhat arbitrary factor of 1.41 ($\sqrt{2}$). This level was chosen based on the idea that the two adjacent cells provide as much weight to the estimate as the cell itself.[17]

Using these adjusted standard errors and the means from the incidence rate table above, you can create distributions of the *probability* of incurring a claim (e.g., the probability of going on a claim q is a random variable).

Covariance

The GLM model tells us the mean and standard error for each of the attained-age rates. The question remains, is there a correlation between the rates? For example, if the "true" probability of a 75-79-year-old incurring a claim is actually near the upper end of its 90% confidence interval, does that tell us that the true probability of an 80-84-year old incurring a claim is more likely to be near the top of his confidence interval as well? We assumed that there should in fact be a positive correlation between the factors; the incidence rate scenarios should reflect scenarios where the overall rates are high or low. We assumed, for the purposes of this study, that the correlation factor is 90% across all attained ages.

Based on the mean and standard errors from the GLM, along with a correlation factor of 90%, we created a table of incidence rate risk factors. The factors for the first 200 simulations are in Appendix 3.

Tail Risk

We had limited data for attained ages greater than 90, so when running the GLM, all ages for this band were combined. The average age for the population in the 90+ category was 94. In order to get rates for ages above this, we extrapolated the data by fitting the data to a logistic curve, as shown in figure 9.

---

[16] While recognition of the sequential nature of the categories could be achieved by using ordinal categories or even a continuous factor for age, the relationship between incidence rates across ages is not often well-described by a linear component within the GLM. Because we have elected to use a non-ordinal categorical model design, we need to explicitly adjust for this deviation from the model design's implicit assumption that there is no concept of "adjacent" categories.

[17] In practice, the actuary should verify to the extent possible that the resulting distribution of incidence rates by period for any given age category matches the distribution implied by the historical observed data, and set the adjustment factor accordingly.

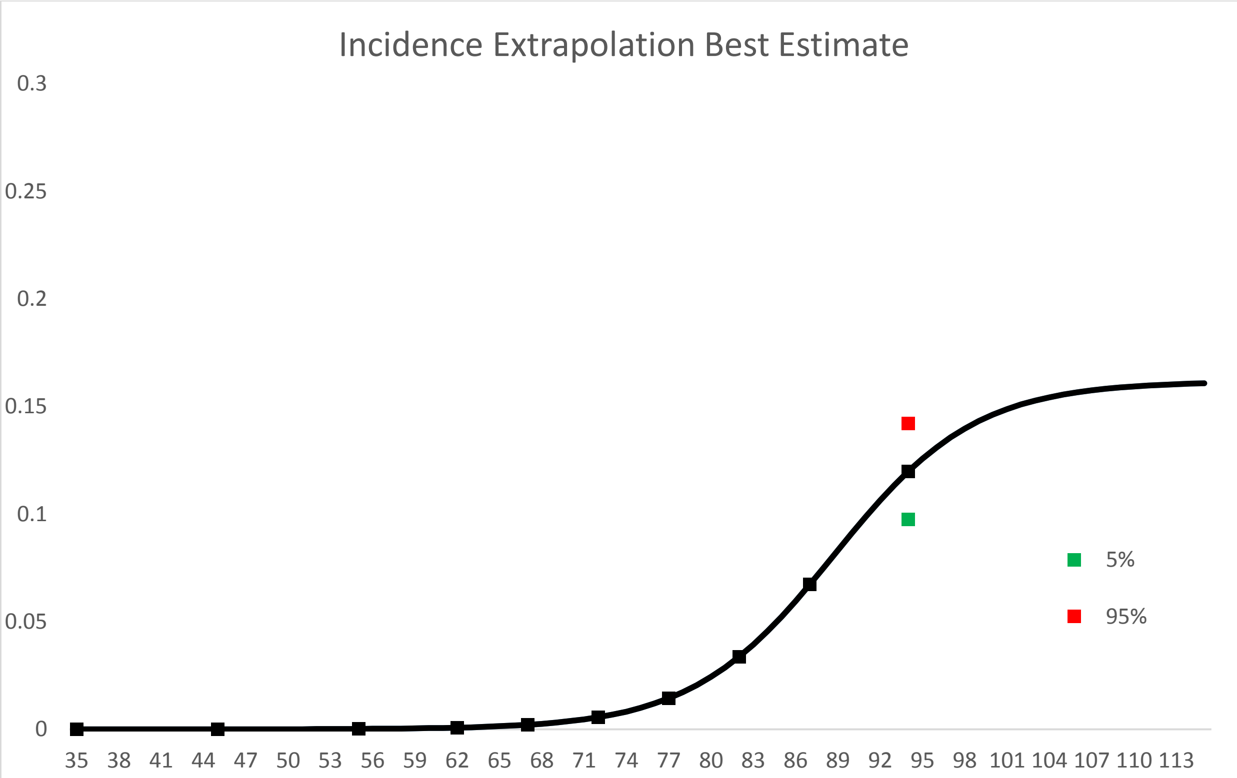

Figure 9: Logistic curve fit to best estimates of incidence rates; 5% and 95% of 90+ data point included.

The GLM results show that for the 90+ category, one can be 90% confident that the rate for that age is between 10.1% and 14.2%. This is a fairly large range. Furthermore, logistic extrapolation can be quite sensitive to the values at the end of the tail.

We decided to create a total of 19 equally likely incidence rate extrapolation scenarios. This was done by assuming the actual incidence rate for the 90+ cohort is normally distributed, with the standard error that was used to create the 90% confidence interval described above. For every 5th percentile of the normal distribution, we extrapolated the incidence rate. Thus, the incidence rate curves are more heavily weighted towards the middle.

The extrapolated curves are shown in figure 10:

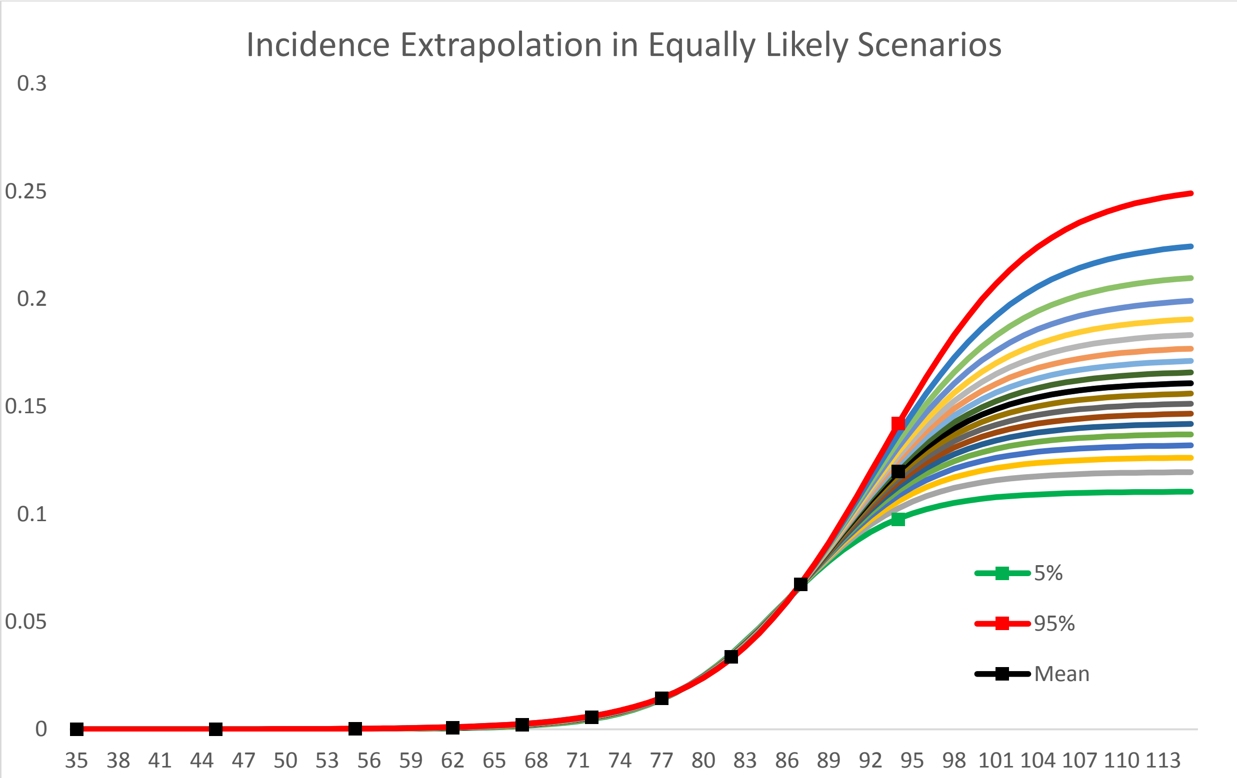

Figure 10: extrapolating 19 incidence rate curves

These curves illustrate that the uncertainty regarding rates in the 90+ cohort of our data is amplified when it is extrapolated to give plausible rates for the higher ages.

These curves were then converted to factors that we call "extrapolation risk factors."

Parameter Risk: Lapse Rates

Our lapse rates were derived from a study of historical lapses that did not incorporate a GLM. Originally, we used professional judgment to determine both the lapse rates themselves, and our level of confidence in the rates. We determined that the lapses grade down from 5.5% to 1% over 14 years, and that we were 90% confident that the actual lapses would be in the range of the best estimates plus or minus 20%.

To illustrate the implications of this, we fit the lapse rates to a set of Beta distributions, one for each of the 14 rates. To parameterize the Beta distributions, we solved for the distributions' alpha and beta that met the following criteria:

1. The mean of the Beta distribution is equal to the best estimate of the underlying probability.
2. The Beta distribution's cumulative distribution function at 5% matches 80% of the best estimate, which is the lower bound of the stated confidence interval (i.e., $F_X(5\%) = P(X < 0.8 \times E(X))$).

After the Beta distributions were solved, we could observe the Sample Size (SS) parameters of the Beta distributions.[18] The following graph shows the relationship:

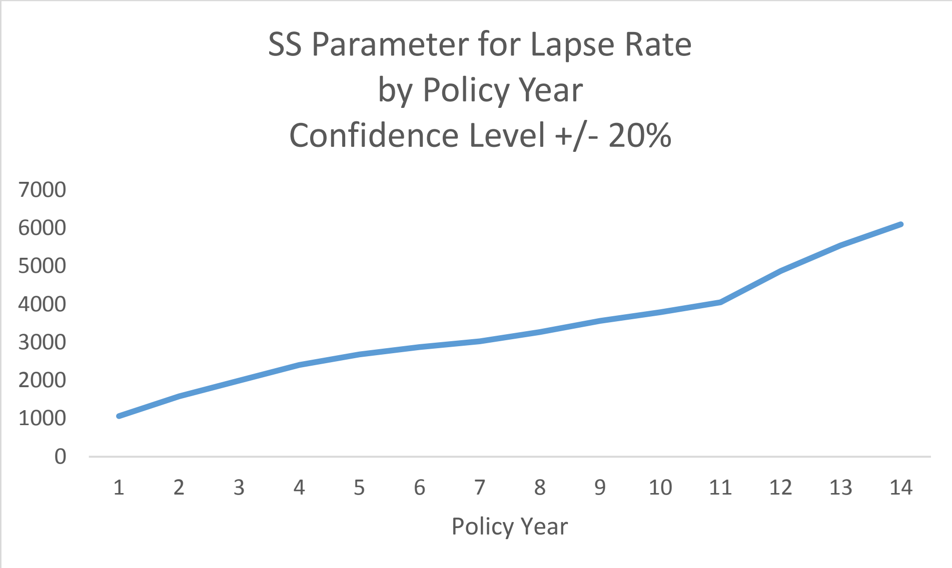

Figure 11: Sample Size (SS) Parameter for Lapse Rate by Policy Year

This graph illustrates that as the lapse rate decreases, a progressively larger sample size is needed to justify a constant confidence level. For example, if one observed that 5.5% of 1,000 policies lapse, the prediction interval for future lapses will be 5.5% plus or minus 20%. However, if you require a ±20% prediction interval on a lapse rate of 1%, you would need to observe that 1% of 6,000 policies lapse. This makes intuitive sense, as the confidence can be expected to depend more on the actual number of observed lapses than on the number of opportunities to lapse.

Note that the reason the SS is in the range of 1,000 to 6,000 is not because the lapse estimates are only based on between 1,000 and 6,000 observations for each age. Rather, in this case, the Beta distribution is being used to capture uncertainty about the future; even though our best estimate for the ultimate lapse rate is 1% for each policy, there are likely correlations between the lapse behaviors of the policies. For example, if there is a macro-economic event that causes one person to lapse, it is likely that many others will lapse, too.

After reviewing the implications of stating our 95% confidence level as a constant percentage of the best estimate, we decided that expressing our confidence as a constant SS for all ages better represented how we saw the uncertainty. As a result, we decided that a constant SS of 3,000 for all policy years is consistent with the original 20% level, but with a more appropriate weighting across policy years.

A constant SS of 3,000 for all ages is the same level of confidence that we would have if the rates were determined by 3,000 observations for each age. The 90% confidence interval for this assumption is illustrated in Figure 12.

[18] The Sample Size (SS) parameter refers to the sum of Alpha and Beta in the Beta distribution. SS can be thought of as representing the sampling error associated with having SS exposures in the historical sample. See Appendix 2 for more details on the Beta distribution.

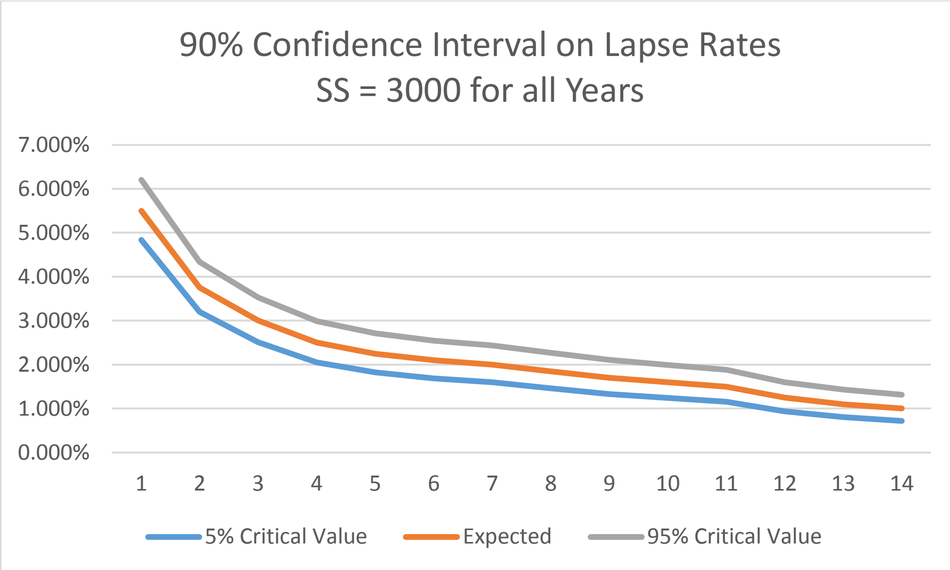

Figure 12: 90% Confidence Interval on Lapse Rates for SS = 3000

This graph shows that the overall width of the confidence interval is about the same for all ages, which better captures our intuitive level of confidence in the respective rates.

Parameter Risk Results

The same model used to quantify process risk was run, using the same product, demographics, and best-estimate assumptions as above. The only difference is that in this model, parameter risk was added. This was done through the following three actions at the beginning of each simulation:

1. Randomly choose a base incidence rate table
2. Randomly choose a set of extrapolation risk factors
3. Choose a lapse table by randomly drawing from the Beta distribution described above

The added variance in this model compared to the process-risk model is attributable to parameter risk.

By adding these parameter risk elements to the model, the standard deviation of the PV of Profits / PV of Premium random variable increases by a factor of 4.7, as shown in the graph below.

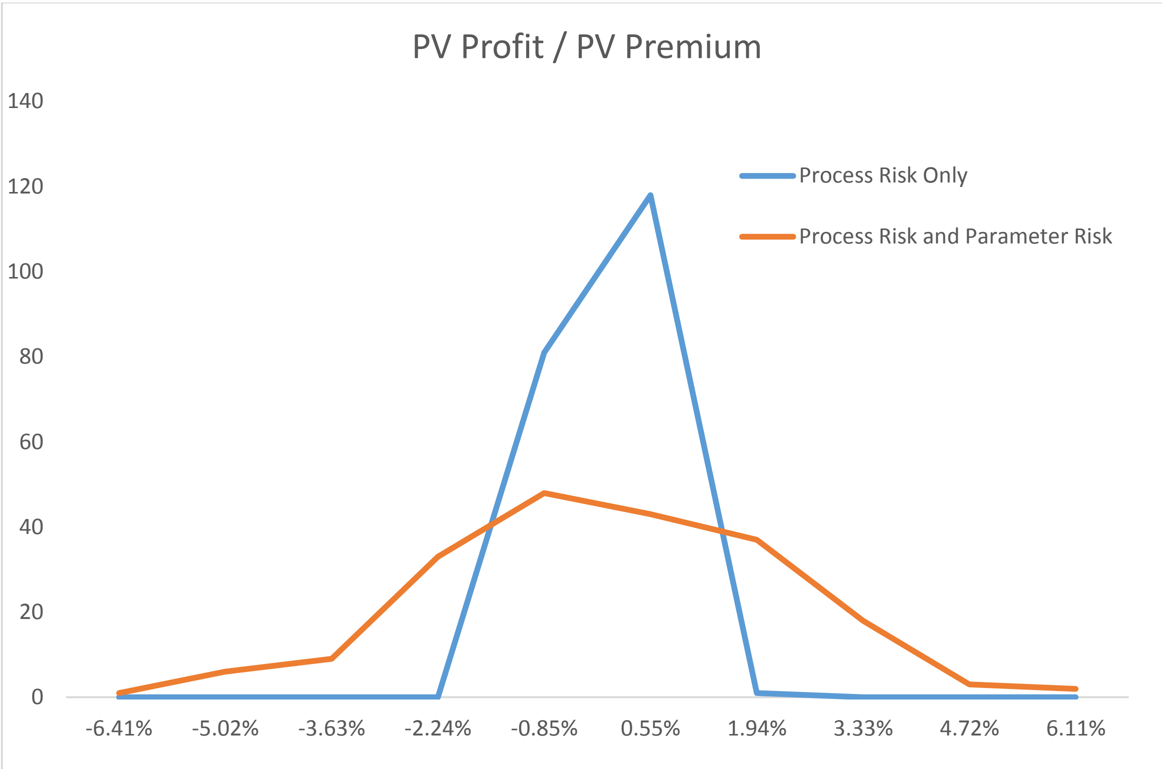

Figure 13: PV Profit/PV Premium

It is important to bear in mind that this added variance in the results is not driven by the possibility that the future will be inherently different than the past. Rather, it is driven by recognizing the level of uncertainty we have about the model's parameters, which were derived from a sample of historical data that might not be representative of the actual range of possible outcomes from the underlying insurance processes.

Parameter Risk and Credibility

Another way of thinking about these results uses the language of credibility theory. If our assumptions were 100% credible, the only risk we would face is process risk. If we find our assumptions are not 100% credible, that uncertainty can be expressed in the forecast by explicitly modeling the parameter risk as was done here.

Interest Rate Risk

In the base model, nearly half of benefits are funded by earned interest (rather than net premiums). Given the fact that uncertainty regarding future interest rates cannot be diversified away, the risk and uncertainty about this aspect of the product is crucial.

The scope of this paper is limited to just touching on the relative magnitude of this issue, and to give direction on how it can be analyzed. The results could be substantially different from company to company depending upon the company's investment strategy, the assets already in the portfolio, the specific characteristics of the liabilities, and what the current and prospective interest rate curves look like at the time of model projection.

For relatively short liability horizons, interest rate risk can be mitigated by asset portfolio management. However, asset portfolio management can become less effective when the duration of the liabilities is very long. Another option for effectively managing interest rate risk is the design of the LTCi product itself.

In the models discussed previously (process risk, process risk with parameter risk), we assumed the assets earned 3.5%, equal to the valuation interest rate. We will now assume that the company invests in 10-year bonds, and make the simplifying assumption that this can be replicated by taking a 10-year rolling average of the 10-year treasury par yield. In both cases, the asset portfolio is equal in size to the reserves.

The model uses 200 interest rate scenarios that were generated from the American Academy of Actuary's interest rate scenario generator.[19]

Adding the interest rate risk has a dramatic effect on the overall risk, and the 95% prediction interval on PV Profits / PV Premiums has a range of -20% to + 25% as seen in the following figure:

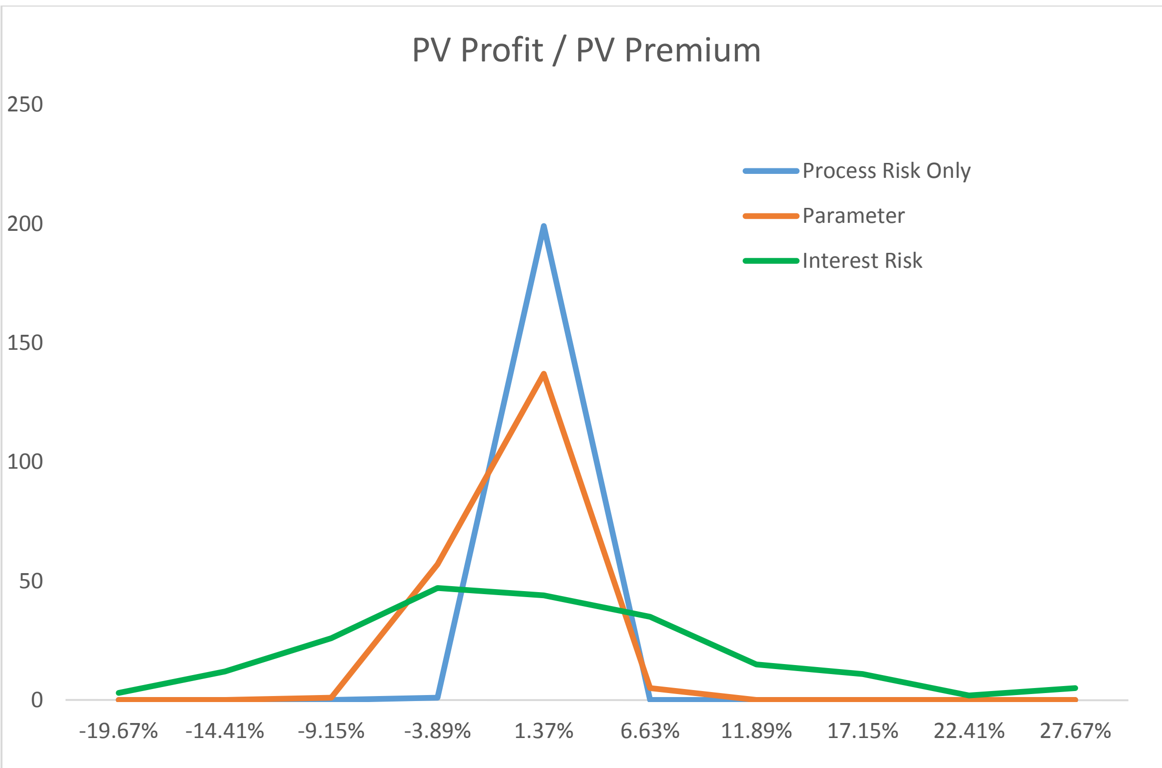

Figure 14: PV Profit/PV Premium

The uncertainty of interest rates over time is not a risk that can be diversified across policies, and asset portfolio management strategies can't deliver both low risk and high returns. Thus, LTC products will either need very large margins (i.e. prices based on very low interest rates), or product designs that transfer some of the investment risk (including upside potential) to policyholders. This could be done through adjusting the premiums or benefits to reflect investment returns.

[19] See the section on Economic Scenario Assumptions in Appendix 1 for more details.

Key Risk Measures

We suggest the following risk measures for evaluating the riskiness of a block of business, which will later be demonstrated in the comparison of the products modeled in this study. The values of these metrics for a specific block of business are driven by many things, including all model assumptions, demographics of the insured population, specific benefits of the policies, and the company's investment strategy. The values of the metrics for the study's sample products are for illustrative purposes only and should not serve as benchmarks for these product designs.

In principle, all of the metrics can be calculated from any stochastic model, regardless of which model assumptions are stochastic. For consistency, we will describe which stochastic elements are most suitable for each metric.

Loss Ratio Margin (LRM)

This measure is useful for setting expectations for incurred claims on a year-by-year basis. It is a possible candidate for triggering rate increases.

Stochastic elements:

- Process Risk
- Parameter Risk

This is a simple statistic based on the random variable Loss Ratio:

$$LR = \frac{Incurred\ Claims}{Earned\ Premium}$$

The 95th Percentile Loss-Ratio Margin (LRM_{95}) is the margin required to be added to the expected loss ratio ($E(LR)$) so that we can be 95% confident that the actual loss ratio will be less than the expected loss ratio plus the margin:

$$P(LR < E(LR) + LRM_{95}) = 0.95$$

Stated another way, this is similar to a one-tailed hypothesis test:

H_0: Assumptions that effect incurred claims are correct

H_1: Actual morbidity is higher than assumptions

Thus, LRM_{95} is simply 1.64 standard deviations of the loss ratio random variable, LR:

$$LRM_{95} = 1.64 \times s_{LR}$$

Thought of in another way, assume moderate experience for a specified period is defined as achieving a loss ratio in the prediction interval, $E(LR) \pm LRM_{95}$. The loss ratio materializing higher than that prediction interval then constitutes statistical evidence that the pricing assumptions are too low.

If the pricing assumptions are correct, process variance and parameter variance can be expected to cause the loss ratio to be higher than that range about once out of every 20 years. However, the loss ratio exceeding that interval two years in a row would constitute significant evidence that a rate increase should be considered.

Just as the incurred loss ratios typically increase substantially as policies age, so does the loss ratio margin. The loss ratio margin for the basic product by policy year is shown here:

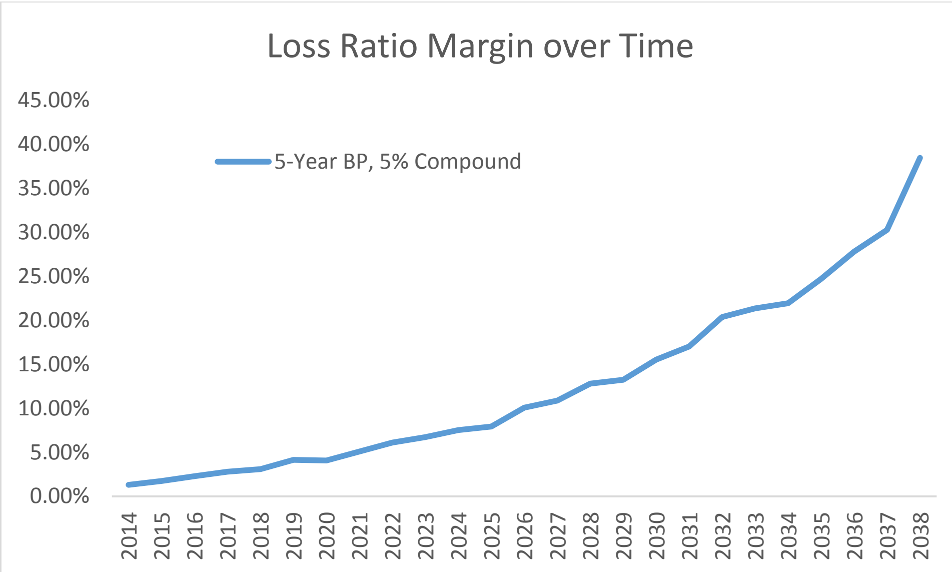

Figure 15: Loss Ratio Margin over Time

This can be thought of as the required addition to the loss ratio that is needed in order to differentiate moderate experience from more than moderate. Note that the margin increases in a shape similar to the loss ratio to which it is added.

Alternatively, this metric could be converted to a multiplicative factor by dividing it by the expected loss ratio. The multiplicative version of the margin looks like this:

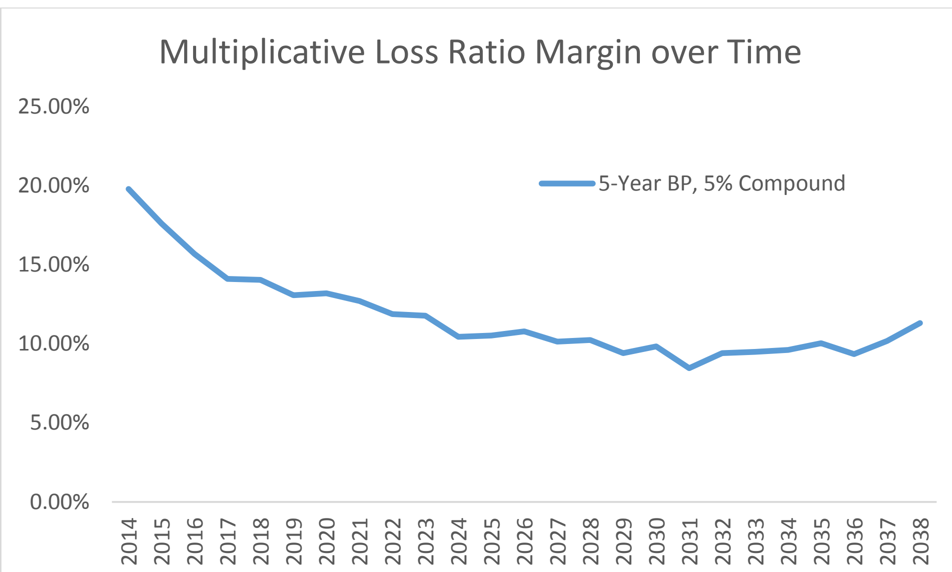

Figure 16: Normalized Loss Ratio Margin over Time

As can be seen here, the margin is a higher percentage in earlier years, but overall is relatively close to a constant percentage of the loss ratio itself.

Profit Ratio Margin (PRM)

As this paper has demonstrated, interest rate fluctuations can be a critical risk in Long-Term Care insurance. If a product depends upon investment returns in order to generate the revenue required to pay claims, a rate increase might be necessary if the investment yield is too low. Or insurers could be required to recognize this risk in their pricing and either mitigate it through product design and/or investment strategies, or charge for it appropriately.

In this context, the income earned on invested assets should be considered when defining moderate experience.

Stochastic elements:

- Process Risk
- Parameter Risk
- Interest Rate Risk

The calculation is similar to the Loss Ratio Margin calculation described above. The difference is that the random variable in question is the profit ratio:

$$PR = \frac{Profit}{Earned\ Premium}$$

The 95th Percentile Profit-Ratio Margin (PRM_{95}) is the margin required so that we can be 95% confident that the actual profit plus the margin will be greater than or equal to the expected profit:

$$P\big(PR + PRM_{95} \geq E(PR)\big) = \ 0.95$$

Thus, PRM_{95} is 1.64 standard deviations of the profit ratio random variable, PR:

$$PRM_{95} = 1.64 \times s_{PR}$$

As with LRM_{95}, moderate experience for the specified period is defined as achieving a profit ratio in the prediction interval, $E(PR) \pm PRM_{95}$. The profit ratio materializing lower than the prediction interval constitutes statistical evidence that the pricing assumptions are wrong.

Standard Deviation of Lifetime Loss Ratio: SDLR

This metric is useful for evaluating the overall riskiness of a portfolio.

Stochastic elements:

- Process Risk
- Parameter Risk

This is calculated by running several simulations and calculating the Lifetime Loss Ratio (LLR, defined as the present value of claims divided by the present value of premiums) for each scenario. All discounting is at the valuation interest rate. We then calculate the standard deviation.

For the base product we have been analyzing $\sigma_{LR} = 2.25\%$. So assuming that we can obtain the valuation interest rate as a portfolio return, 2.25% is one standard deviation of the lifetime loss ratio. This implies that a margin of $1.64 \times 2.25\% = 3.69\%$ is enough to be 95% certain that the actual lifetime loss ratio will be lower than the expected lifetime ratio plus the margin.

Standard Deviation of Lifetime Loss Ratio Discounted at Short Rate: $SDLR_{sr}$

This is calculated identically to the SDLR described above but with one difference. While the SLDR is discounted at the valuation interest rate, $SDLR_{sr}$ is discounted at the short rate specific to each scenario.

Thus, it has the following stochastic elements:

- Process Risk
- Parameter Risk
- Interest Risk

The specific product design, set of interest scenarios, and investment strategy can have a significant impact on the riskiness of the product. SDLR disregards the interest risk, while $SDLR_{sr}$ highlights it.

For the base product, $SDLR_{sr} = 9.78\%$. The difference between this and SDLR emphasizes how large interest rate risk is.

Other Product Designs

The risk metrics described above were all for the basic policy with a 5-year benefit period and 5-year compound inflation protection. Some product designs are intrinsically riskier than others. To illustrate, we will compare the risk metrics of four additional policy designs:

- 2-year benefit period with 3% simple inflation protection
- Life/LTC combo product
- Product with a return of premium rider
- Product with benefits and premium incrementing according to the CPI

Across all of these models, all of the other assumptions (demographics, incidence rates, lapse rates, etc.) will stay the same. This is done in order to isolate the effect that benefit design has on the risk metrics. In the real world, different product designs would likely have different assumptions to reflect different underwriting requirements, expected sales demographics, unique policy incentives, etc. While those assumptions have an effect on overall riskiness as well, analyzing that aspect of the riskiness is outside of the scope of this analysis.

For all of these products, the premium and valuation claim costs were adjusted to reflect the benefits of the specific plan.

2-Year Benefit Period, 3% Simple Inflation

The first alternate product design we will consider is a policy with smaller benefits. What happens if we reduce the benefit period from 5 years to 2, and reduce the inflation protection from 5% compound to 3% simple?

When you analyze the present value of claims, these two product designs have about the same level of risk. This is seen by comparing the SDLR and $SDLR_{sr}$ of these two products:

Product	SDLR	$SDLR_{sr}$
5-Year BP, 5% Compound	2.25%	9.78%
2-year BP, 3% Simple	2.32%	8.56%

An interesting result is that when you look at Loss Ratio Margin (LRM) on a year-by-year basis, the product with the 2-year BP shows less risk:

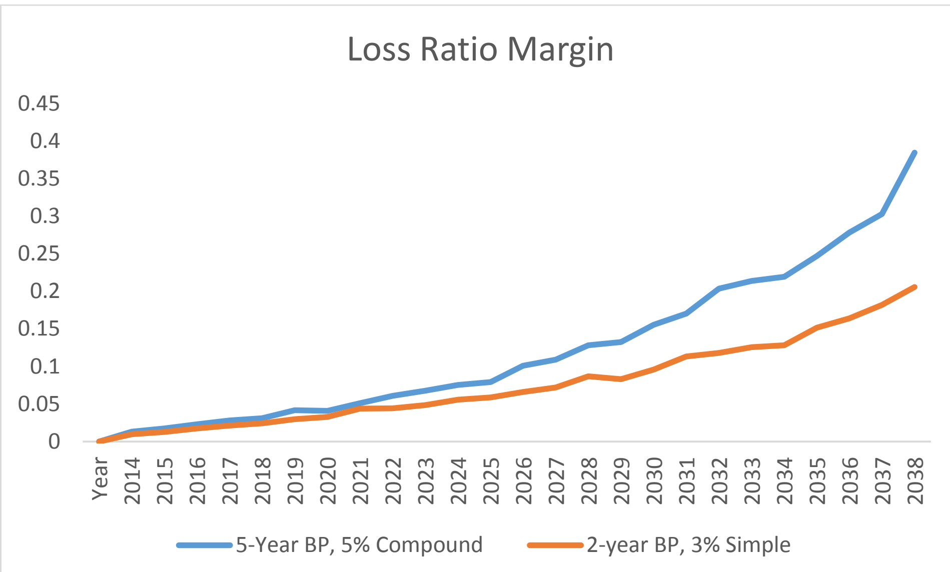

Figure 17: Loss Ratio Margin for 2-Year Benefit Period

Deeper analysis indicates that the divergence of the LRMs is driven by the claim reserves. The product with the longer benefit period also has a daily benefit that is growing at a higher rate. Over 10 years the 5% compounding generates benefits 62.9% higher compared to 30% higher for 3% simple interest. After 20 years the difference is 165% to 60%. These higher benefits are reflected in the claim reserves. So the same claim event causes a much bigger shock for the one compared to the other. These bigger shocks in turn cause the risk to be higher for the 5-year product.

Return of Premium (ROP) Rider

This product has the same benefits as the first product described, including a 5-year benefit period and 5% compound inflation protection. However, with the ROP rider, there is one additional benefit: upon death, all premium paid is returned, less benefits paid.[20]

[20] A return-of-premiums-less-benefits-paid is a path-dependent benefit, meaning that you do not know how much the death benefit is without knowing how much in LTC benefits had already been paid at the time of death (i.e., without knowing the path taken to death). Path-dependent benefits are difficult to model

Adding the return-of-premium rider to the product marginally lowers the risk according to the SDLR metric, but the gains are reversed when interest is considered:

Product	SDLR	$SDLR_{sr}$
5-Year BP, 5% Compound	**2.25%**	**9.78%**
With Return of Premium	**1.41%**	**9.88%**

This should be expected. Without the ROP rider, several individual policyholders receive little or no benefit because they die before receiving more in benefits than they paid in premiums. By adding this rider, those scenarios receive the ROP benefit. Returning the premium of the policies that did not receive claim benefits lowers the overall variance of benefits.

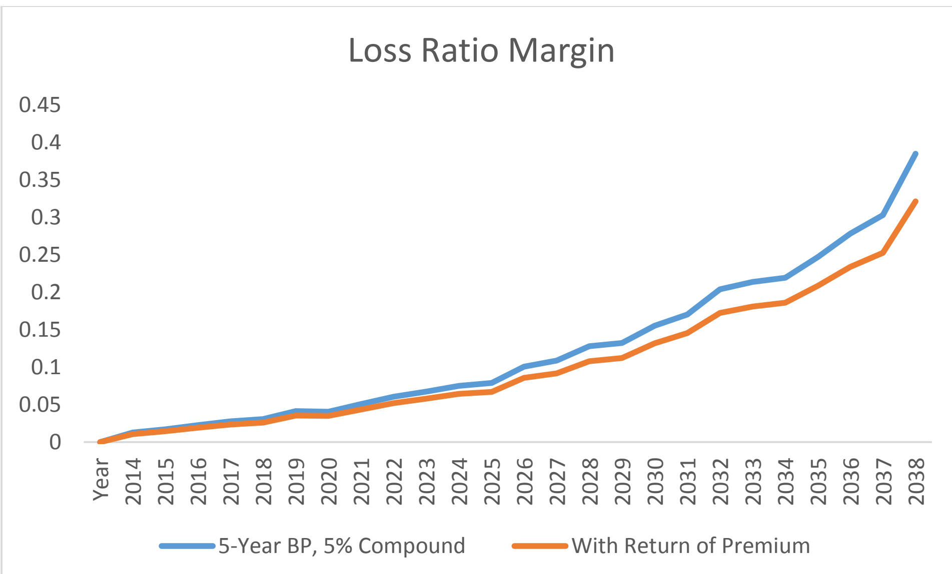

Figure 18: Loss Ratio Margin with Return of Premium

However, adding this additional benefit makes the product much more suseptible to fluctuations in interest rates. Thus, when looking at the $SDLR_{sr}$ metric, the lower risk caused by adding this benefit is offset by a higher investment risk.

Remember that in the baseline product approximately half of the benefits were paid from premiums and the other half from interest earned on reserves. For an ROP product the benefit paid becomes a function of the premium. Higher premiums correspond to higher benefits. This makes an ROP product extremely sensitive to interest rates. This can be seen in the graph below.

One other thing should be mentioned. Our simplifying assumptions do not include expenses or profit margins. In the real world ROP benefits are based on gross, not net, premiums. This means that in addition to funding excess claims, interest earnings on reserves need to be able to fund

deterministically without making some sweeping assumptions, but are straightforward to model in simulation models such as this one.

commissions, overhead, and profit margins as well. So the difference in the risk would be even greater than illustrated here.

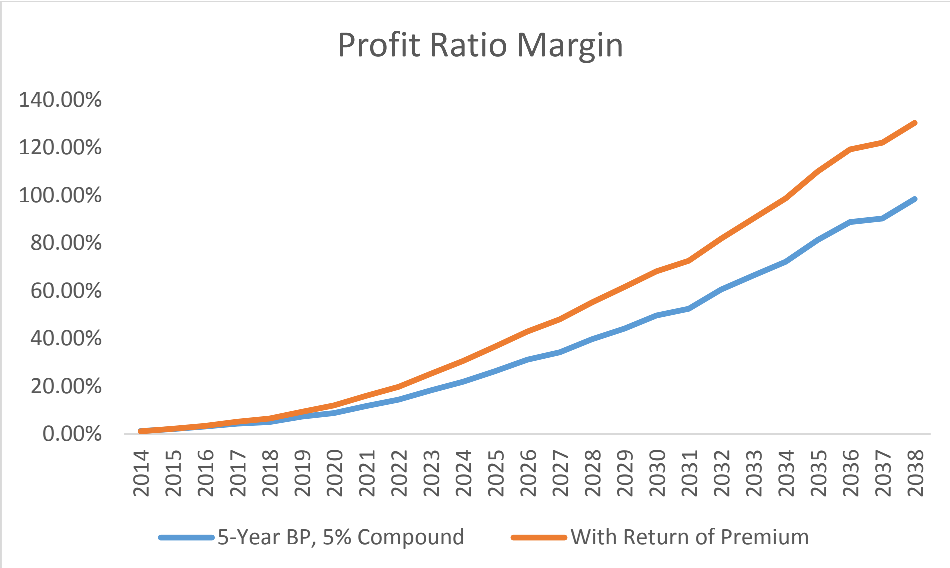

Figure 19: Profit Ratio Margin with Return of Premium

Combo Products

For the combo product, we begin with exactly the same LTC product we have been analyzing, but add one feature: at the time of death, the maximum lifetime benefit minus any claims already paid will be returned as a death benefit.

It should be intuitive that this type of combo product is less risky than a stand-alone LTCi product because there is a natural hedge between life insurance and stand-alone LTCi.[21] In stand-alone LTCi, when the policyholder dies no benefits are paid and the entire policy reserve is released to profit. While this effect emerges smoothly in deterministic models, simulation models reveal that this can have a high process variance when policyholders with policies having large reserves either die or continue on at each time step.

With LTC/life insurance combo products, death results in a substantial benefit being paid to the policyholder. Thus, as long as the insured does not lapse, he will receive either a death benefit or a LTC benefit. The uncertainty about whether a benefit will be paid is replaced with the substantially smaller uncertainty of when the benefit will be paid.

[21] See "Quantification of the Natural Hedge Characteristics of Combination Life or Annuity Products Linked to Long-Term Care Insurance" by Linda Chow et al. http://www.soa.org/research/research-projects/ltc/research-2012-03-quant-nat-hedge.aspx

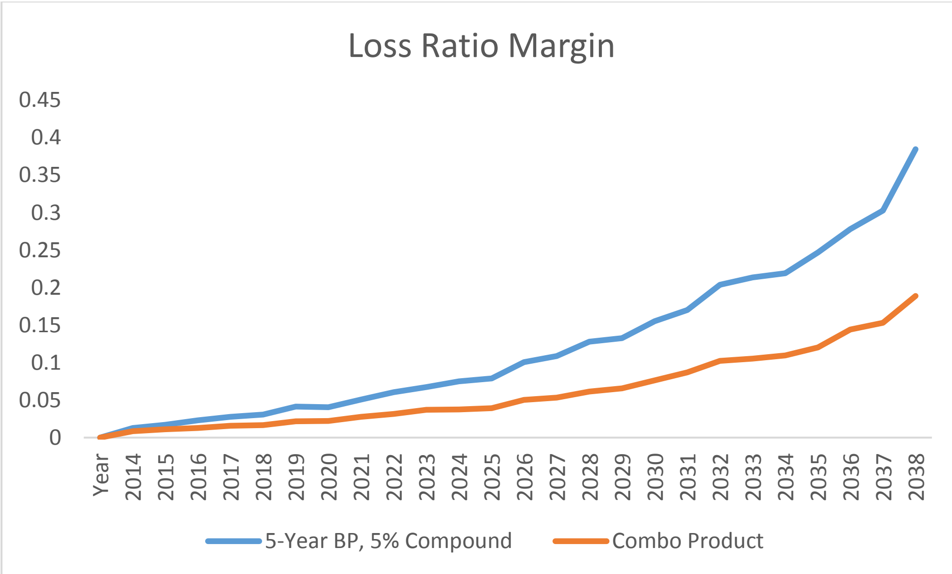

Figure 20: Loss Ratio Margin with Combo Product

This reduction in risk is substantial. The following table shows the loss ratio margin by year:

Loss Ratio Margin (LRM)			
	(A)	(B)	(A)/(B)
Year	5-Year BP, 5% Compound	Combo Product	
2014	1.29%	0.87%	67.75%
2015	1.73%	1.14%	66.21%
2016	2.30%	1.29%	56.28%
2017	2.79%	1.57%	56.36%
2018	3.09%	1.68%	54.27%
2019	4.14%	2.17%	52.41%
2020	4.08%	2.24%	54.94%
2021	5.09%	2.77%	54.39%
2022	6.08%	3.18%	52.39%
2023	6.74%	3.71%	55.11%
2024	7.53%	3.75%	49.72%
2025	7.91%	3.95%	49.96%
2026	10.08%	5.04%	50.02%
2027	10.89%	5.34%	49.05%
2028	12.81%	6.13%	47.85%

We see here that for most years, the loss ratio margin is about half of the margin for the identical product without the death benefit.

The reduction in risk for this product is perhaps understated in these metrics. In the real world, we might expect very few lapses because there are no benefits paid on lapse. Reducing the lapse rate assumption for the combo product would reduce the Loss Ratio Margin even further. It is also possible that regulators would require non-forfeiture benefits to be paid on combo products. In this case, people might have more incentive to lapse since they would be able to walk away with some benefit. None of this was reflected in our modeling because, for the purposes of this study, we wanted to focus on only the effects of changing the benefit design.

However, the reduction in risk associated with claims has a tradeoff. Because of the additional benefits paid at death, the duration of the liabilities increase, leaving the overall profitability more susceptible to interest rate risk. This can be seen by looking at the PRM:

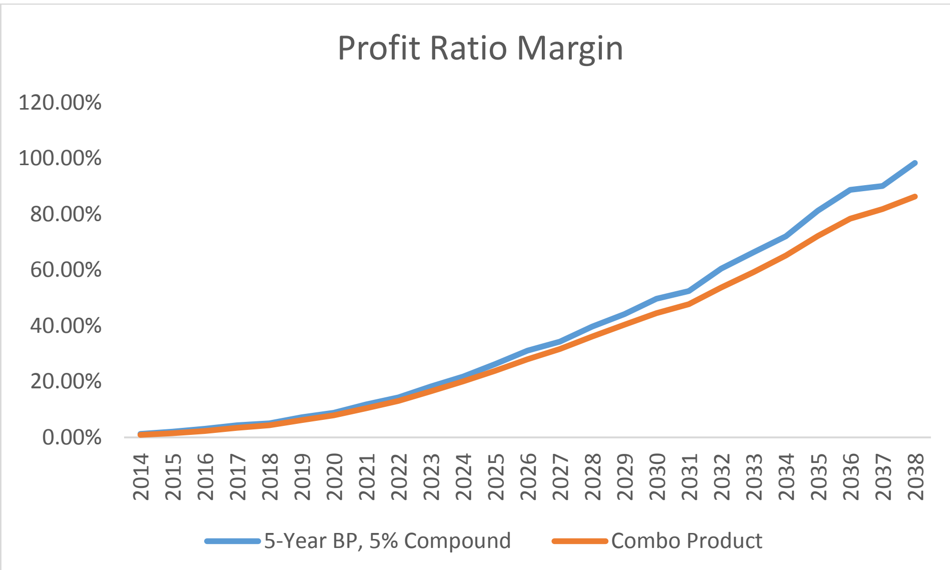

Figure 21: Profit Ratio Margin with Combo Products

The lifetime metrics tell the same story: increasing the duration of the liabilities makes the product more susceptible to interest rate risk.

Product	SDLR	$SDLR_{sr}$
5-Year BP, 5% Compound	**2.25%**	**9.78%**
Combo Product	**0.29%**	**10.74%**

Benefits and Premium Indexed to CPI

Like the base product, this product has a 5-year benefit period. Rather than having a constant inflation protection, in this product the maximum daily benefit each year increases with the Consumer Price Index (CPI). The monthly premium also increases with the CPI. Assuming that LTC costs and the policyholder's ability to pay the premium both increase in proportion to the CPI, this policy has the attractive feature of offering the policyholder the right amount of inflation protection, regardless of the future inflation scenario. But how risky is it from the insurance company's perspective?

Across the model's 200 economic scenarios, there are a wide variety of inflation scenarios. This causes there to be a higher variance in both incurred claims and earned premium. Because the effect of inflation on these two model variables is highly correlated, the change in overall risk is largely mitigated, resulting in an overall Loss Ratio Margin that is only marginally higher than the basic product.

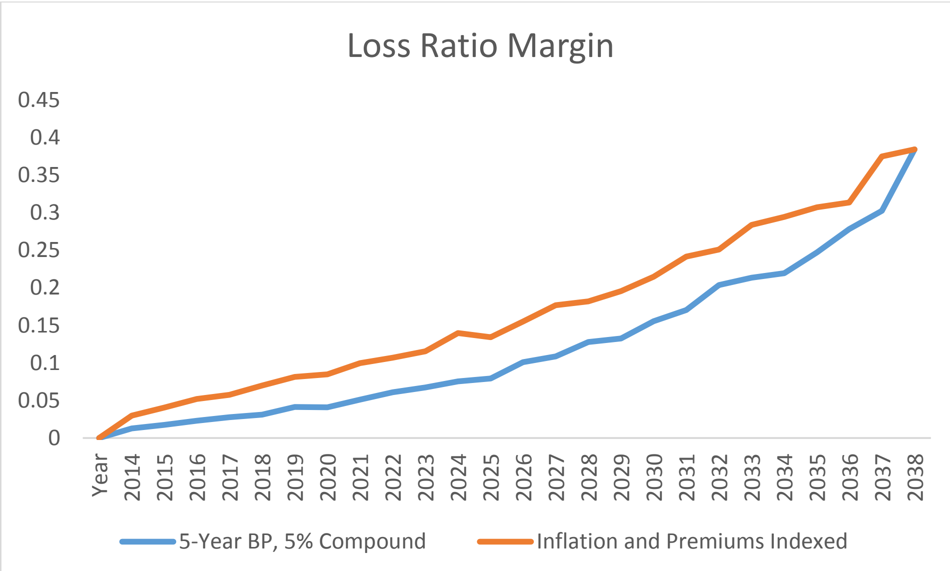

Figure 22: Loss Ratio Margin with Inflation and Premium Indexed

In contrast, when we look at the Profit Ratio Margin for this product, we see a dramatic reduction in risk:

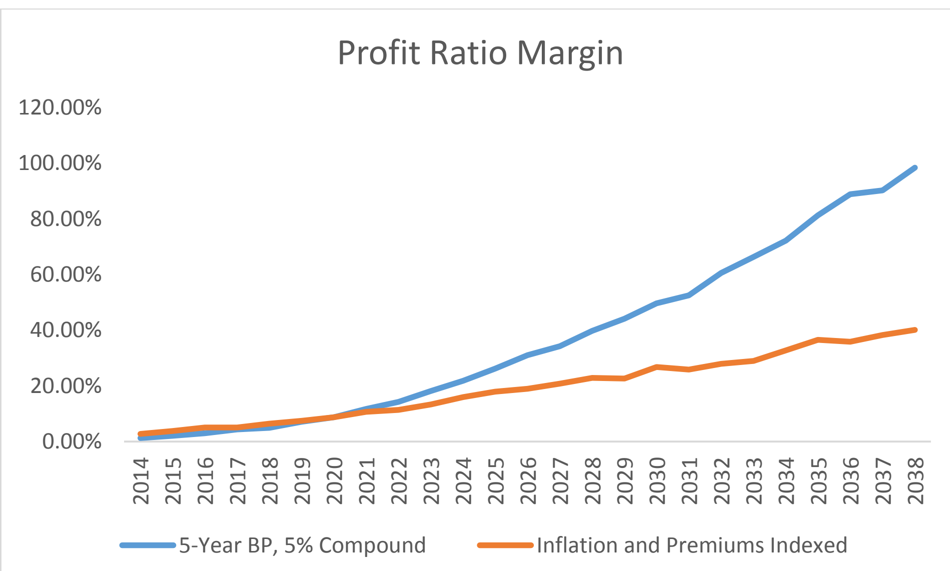

Figure 23: Profit Ratio Margin with Inflation and Premium Indexed

This reduction in risk is driven by the correlation between investment returns and the inflation rates. The scenarios with higher inflation generally have higher market returns, while scenarios with lower inflation have lower market returns. When compared to the base case, losses associated with low interest rates are offset by lower benefit increases. But in the scenarios with higher inflation, higher returns are obtained on the assets, and higher premiums are collected.

The lifetime metrics show this reduction in risk just as dramatically:

Product	SDLR	$SDLR_{sr}$
5-Year BP, 5% Compound	**2.25%**	**9.78%**
Inflation and Premiums Indexed	**8.56%**	**2.31%**

It should be pointed out that a smarter investment strategy could potentially lower the risk further. For example, investing only in bonds that were indexed to the CPI would theoretically eliminate all inflation risk, and only leave risk associated with the spread between the CPI and the yield of the index bonds changing.[22]

It should also be pointed out that in reality there would be tradeoffs between the two inflation approaches. With 5% compounding you normally have to project to what extent actual inflation in the benefits is going to keep up with the maximum daily benefit. This is normally done by projecting a salvage amount, which represents the amount that the actual reimbursement is less than the maximum daily benefit. That adds variability to the results, which is not showing up in our modeling because we chose the simplifying assumption that the benefits were indemnity based. So while our results show the indexed option to have a higher loss ratio margin, the two are likely to be much closer together in reality.

Product Design Summary

The following table summarizes the risk metrics of the five products:

Product	SDLR	$SDLR_{sr}$
5-Year BP, 5% Compound	**2.25%**	**9.78%**
2-year BP, 3% Simple	**2.32%**	**8.56%**
With Return of Premium	**1.41%**	**9.88%**
Combo Product	**0.29%**	**10.74%**
Inflation and Premiums Indexed	**8.56%**	**2.31%**

We see that by adding a death benefit, either in the form of a ROP rider or a Combo Product death benefit, the risk associated with the benefits decreases. This is because we are replacing a benefit that varies greatly in both timing and amount with one that varies primarily with timing only. Not only is the intrinsic riskiness of these products lower, they also could provide positive incentives for the policyholder; for example removing the "use it or lose it" feature that might incentivize

[22] To see this, consider the basic prospective valuation formula: $V = PV(Ben) - PV(NP)$. If the inflation for the year is i, then the benefits and net premium should both increase by i. Thus, the reserve after the bump in inflation would be: $PV\big(Ben(1+i)\big) - PV\big(NP(1+i)\big) = [PV(Ben) - PV(NP)](1+i) = V(1+i)$. If the reserves were invested in inflation-indexed bonds, then the assets supporting the reserves would increase by the needed $i\%$, regardless of how large or how little it is.

people to file a claim who otherwise would not. The extent to which these products would be required to provide non-forfeiture benefits might limit the impact of this benefit.

It is salient that for the product with inflation protection and premium indexed to inflation, the risk is high when the returns are discounted at a constant rate, but are low when discounted at the short rate. What is attractive about this is that this product lowers the risk of to both the insurer *and the insured*; assuming the cost of care tracks the CPI, this product design helps ensure that the policyholder will have adequate, but not too much, coverage in any given future scenario. Likewise, it provides her a constant premium, adjusted for inflation.

For the non-indexed products, the risk metrics that include interest returns are much higher. It should be understood that this high-risk metric is a function of the interest rate environment (as reflected in the model's interest scenarios), as well as the company's investment strategy and actual asset portfolio. The company's own investment strategy could be included in the model, which would give risk metrics specific to that company that are different than the ones shown here.

Pricing Margins and Rate Increases

This section contains the authors' opinions about how the modeling techniques and insights discussed could be applied to issues surrounding pricing margins and rate increases.

Step 1: Define Risk Tolerance Up Front

When filing products, insurance companies could describe to the regulators the risks they are willing to accept, as well as the risks that will be subject to rate increase. For example, they could say that risks involving claim incidence, claim severity, and portfolio returns are subject to rate increases, while all other risks, including risks associated with lapse rates, mortality rates, and business-mix,[23] are not. This could create a marketplace in which prospective policyholders would have trouble understanding exactly what they were buying. If each company chooses the risks that they are willing to underwrite and those that they are not how would this be communicated? And would regulators be equipped to deal with that much potential variation?

A better solution may be for the industry to come together to determine a uniform set of risks that they are willing to accept entirely as well as those risks that should potentially trigger rate increases given the proper circumstances. A model regulation could be drafted as the mechanism to put this solution in place. This would have the advantage to both consumers and regulators of creating a uniform playing field in which all carriers agree to be responsible for the same risks.

For the risks subject to rate increases, standards would need to be developed to determine what deviations from current assumptions could trigger a rate increase. For example, for incidence rate risk the requirement could be to accept all process risk, plus a moderate amount of parameter risk. The definition of moderate parameter risk would be explicitly stated (e.g., moderate parameter risk for incidence rates is defined as experience falling with the 90% prediction interval for the current pricing incidence rates).

[23] For example, females have higher LTC morbidity than males. If a company charges unisex premiums, it must sell enough policies to males in order to adequately subsidize the females. If it sells to an unfavorable business mix, it could be in trouble before the first claim is filed.

Step 2: Set Annual Performance Expectations

It is useful to monitor the emergence of experience against the expected prediction intervals. At the beginning of every calendar year, a company could run a simulation model to create a prediction interval of key risk measures. For example, the company would say, "Our expected loss ratio for 2020 is 25%, with a loss ratio margin (LRM) of 4%. Thus, we are 95% confident that given the model assumptions, our actual loss ratio will be less than 31.56%."[24] The expected loss ratio and the LRM would be objectively calculated using the block of business in-force at the beginning of the year, and stochastic pricing assumptions established in Step 1.

Similar prediction intervals could be created for other key operational and financial metrics. Making these prediction intervals at the beginning of every year should be useful for management to establish expectations. For rate increase purposes, the prediction intervals can then be used as the basis of a hypothesis test, the result of which indicates whether or not a rate increase should be considered.

Step 3: Monitor Experience

At the end of the year, it is beneficial to compare the actual loss ratio to the prediction interval that was forecasted a year earlier. If the experience is within the expected range, the experience is moderate. If the experience is outside of the range, a company could inform regulators that its loss ratio was higher than moderate.

If this only happens occasionally, it is not a problem; it is expected that if the model assumptions are all correct, actual experience on any given metric will exceed the critical value one out of every 20 years.

However, it is unlikely that mere chance will drive most metrics to exceed the critical value two years in a row.[25] That happening indicates that there is a problem that merits further investigation. If, for example, the loss ratio is higher than moderate two years in a row, a company could inform regulators that its loss ratio exceeded the critical value two years in a row and that a rate increase may be necessary.

Step 4: Investigate Assumptions

After two years of poor experience, a company could perform a new analysis on the underlying assumptions, incorporating the new data. The actual loss ratio exceeding the critical value two years in a row does not necessarily imply that premium is insufficient. Rather, it implies that there is something wrong with the model assumptions. The analysis should entail figuring out how the underlying model assumptions should be updated to reflect what has been learned about morbidity and the actual block of business since the original pricing.

When doing this, all model assumptions should be reviewed. For example, if lapse assumptions have been unfavorable but morbidity has been favorable, a rate increase based on new lapse assumptions should be tempered to reflect better morbidity.

[24] Z-score for 95th percentile = 1.64, so the critical value for the loss ratio is $25\% + 1.64 \times 4\% = 31.56\%$

[25] There are exceptions to this, such as paid claims metrics, which are correlated from year to year.

Step 5: Re-price Business

After the assumption analysis is complete, a company could go back and re-price business using the updated models and parameters. The premium could then be set to what it should have been initially, according to this analysis. Going forward these updated assumptions would become the current pricing assumptions driving the analysis in future years.

Summary

There are several benefits to using simulation models to analyze blocks of LTCi. It allows a company to set clear performance expectations with prediction intervals for business metrics and financial results. In addition, it allows modeling path dependent variables that cannot easily be modeled deterministically. This provides a way for companies to determine if the variation in their financial results is statistically credible or not. As a result, simulation models allow companies to better understand the risk of their business, design products that minimize risk, and set appropriate margins.

Process Risk

Some elements of the financial results have relatively low process variance (e.g., premiums), while others tend to have very high process variance (e.g., claim reserves). When monitoring emerging experience, it is important to compare the emerging results to more than just their expected values. The process risk prediction interval represents the boundaries of normal variation due to process risk. Actual results must fall outside of this range to be considered statistically significant deviations from expectations, resulting in a need for updated assumptions. Insurance companies should be expected to absorb all process risk.

Parameter Risk

Parameter risk measures the level of confidence in the model assumptions. This risk can be objectively evaluated by using statistical techniques to establish confidence intervals of model parameters, incorporating those distributions into forecasts, and evaluating those forecasts against emerging experience. It is reasonable to expect insurance companies to absorb a moderate amount of parameter risk. The amount of risk may vary by parameter and should be determined by the industry in concert with regulators.

Interest Rate Risk

Long-term care policies typically have long durations and generate large reserves. Because of this, investment returns are an important source of the funds needed to pay policy benefits. Future interest rate environments are uncertain and dramatically different outcomes are plausible. Insurance companies cannot diversify this risk across policies; however, they do have ways to mitigate it through product design and investment strategies. The level of interest rate risk associated with different product designs, and how well specific investment strategies deal with those risks, can be measured by including interest rate scenarios in the stochastic models.

Product Design

Product design can have a dramatic impact on the riskiness of a product. The stochastic models illustrated in this paper create prediction intervals on the operational and financial metrics that are specific to the product in question. Product designs that are less risky will naturally produce less variance in the prediction intervals.

This implies that objectively less-risky products can have smaller pricing margins, smaller triggers for rate increases, and smaller capital requirements. If this is recognized by regulators, companies will be incentivized to design less-risky products.

Defining Moderate Experience

Moderate experience needs to be objectively defined. This can be accomplished using the predication interval paradigm. To do so, stochastic models can be run that incorporate a moderate amount of parameter risk. These models will produce prediction intervals on all key operational and financial metrics. Actual results that fall within the prediction interval are, by definition, moderate. Adverse experience, defined as experience falling outside of the prediction interval, would constitute statistical evidence that the pricing assumptions should be revisited and that a rate increase might be appropriate.

Margins and Rate Increases

By using stochastic models to help design their products, companies can gain an understanding of the relative riskiness of various product designs. These same models can then be run to provide prediction intervals of emerging experience. The results of this analysis could be reported to regulators annually, and a criteria for allowing rate increases could be based upon them. For example, if results fell outside of the prediction interval two years in a row, that would indicate that the model assumptions should be reevaluated and a rate increase considered. There are many ways in which the criteria could be set; however, it is critical to establish what defines a statistically significant variation from expected results. This will create consistency within the regulatory environment and provide objective evidence to justify an insurer asking for an increase in rates.

Key Risk Measures

A number of key risk measures were explored and evaluated. The loss ratio margin, profit ratio margin, and standard deviation of the lifetime loss ratio were all described in detail. Each serves a somewhat different purpose and has its strengths and weaknesses. Several overarching principals emerged from the evaluations. Looking at monthly risk measures creates so much volatility that it is hard to learn anything useful from it. The prediction intervals created for monthly values are so broad that large variation from the expected value can still carry no statistical significance. In addition, the "increase in reserve" line on the income statement for traditional stand-alone LTCi is inherently volatile; it has a high statistical variance. Because of this volatility, care must be taken when analyzing experience metrics that incorporate changes in reserves.

Conclusions

To conclude, we will provide concise answers to the questions stated in the introduction.

How risky is a typical LTCi block of business?

Long-Term Care insurance is among the riskiest insurance products sold. It is driven by more assumptions (incidence rates, recovery rates, lapse rates, death rates, utilization rates, inflation scenarios, interest rates) than most insurance products and those assumptions could change significantly between current observations and what materializes before the insurance contract ends. Our understanding of many of these assumptions is not always robust. These products have not been available for very long and studies of emerging experience are confounded by changes that have been occurring in the industry over time. It is a long-duration product, featuring steeply increasing benefits typically funded by level premiums. This implies that it will have a large policy reserve, and that a significant part of the benefits will be funded by investment returns. All of these factors combine together to make LTCi cashflows extremely hard to predict; which is the definition of a risky financial services product.

How can an actuary differentiate between poor experience due to the inherent volatility in the underlying LTCi business and poor experience due to incorrect assumptions?

Key product assumptions such as incidence, recovery, death, and lapse can be expressed as probability distributions rather than deterministic rates. Using these probability distributions, a company's actual block of business can be forecast using Monte Carlo simulation. Variations from simulation to simulation in this type of model are known as process risk. Statistical modelling allows you to estimate probability distributions and prediction intervals for all key operational and financial metrics. If the actual experience that emerges is consistent with these models, there is no statistical evidence of incorrect assumptions. Results falling outside of the prediction intervals constitute statistical evidence that some of the model assumptions are incorrect.

How can these risks be exacerbated or mitigated by product design modifications?

Simulation models effectively illustrate the inherent volatility in LTCi and allow a company to compare the relative volatility of different product designs. Effective strategies for lowering risk through product design include sharing the risk with the policy holder and offering packages of benefits that collectively hedge the risk.

For example, a product design that has the premium and inflation protection tied to an inflation index can effectively hedge against interest rate risk. The inflation rate and interest rate tend to be correlated. Thus, low interest rate returns can be more easily absorbed, because they tend to occur in scenarios with lower growth in the benefit payments. Such products are mutually beneficial because they also lower the risk to the insured; the specific amount of inflation protection the insured receives is based on the actual increase in costs that he or she faces.

Products with "use it or lose it" benefits have higher risk because, on a policy-by-policy basis, the benefit payments are highly volatile; the benefit could be used in full, partially, or not at all. Further, it is difficult to set correct assumptions for these policies because of policyholder incentives. In contrast, if the product has a non-forfeiture benefit, then the value of the non-forfeiture benefit hedges the risk of the base LTCi benefit. Likewise, the acceleration of benefits rider in a combo product is a natural hedge against the base LTCi benefit.

What implications does this have for pricing margins and triggers for rate increases?

Insurance companies must be willing to accept all of the process risk associated with the products they offer. It is also reasonable to expect them to absorb a moderate amount of parameter risk. On an annual basis, stochastic models can be run that create prediction intervals of key operational and financial metrics. Actual results that fall within these prediction intervals are, by definition, moderate. If adverse experience falls outside of the prediction interval, product assumptions can be updated based on the new experience and a rate increase considered.

Using the stochastic models described in this paper, two key margins were suggested: the loss-ratio margin and the profit margin. These margins can be calculated based upon the specific risk characteristics of the block of business. If a product is intrinsically less risky, smaller margins can be justified.

Appendix 1: LTC Volatility Assumptions

The assumptions used in our modeling were developed primarily from the 1984-2007 SOA LTC intercompany experience study. We realize the weaknesses inherent in this data, yet chose to use publically available data so that the process involved could be completely open. The focus of this study was not to determine what the correct incidence or termination rates are for any given company or the industry as a whole. Rather, it was to illustrate how to incorporate uncertainty when making assumptions and then incorporate both risk and uncertainty in models.

Methodology

We used a variety of methods to set the various assumptions needed for our modeling. This was done to illustrate some of the many ways that reasonable assumptions could be determined. We make no claim that any of the methods we used were the best or the most appropriate; they are just some of many possibilities. We would expect actuaries involved in similar work to use their own judgment regarding the methods and techniques they believe are most appropriate for the data available and the particular circumstances they are addressing.

Rates, Probabilities, Risk, and Uncertainty

The term “rate” is associated with deterministic models and experience metrics. For example, a block of business might assume that the ultimate lapse rate is 1%, and might have an observed lapse rate of 1.2%. In contrast, the term “probability” is associated with stochastic models. For example, the underlying probability of a specified policy lapsing might be 1%, but a single policy will either lapse or it will not.

The term “risk” refers to the natural variance that occurs around correctly understood probabilities. For example, if the real probability of lapsing is 1% and there are 10,000 policies, elementary statistics indicates that you can be 95% confident in having between about 80 and 120 lapses with not knowing precisely how many lapses will be observed within that range is risk.

In contrast, “uncertainty” refers to the contingency of the underlying probabilities being unknown or different from what was presumed. Uncertainty can range from assumed probabilities being marginally wrong to fundamental unforeseen shifts in morbidity, the economic environment, and costs of care.

Uncertainty is generally a much bigger concern than risk. As illustrated in this project, risk has a tendency to offset itself over time, given that the assumption itself is correctly specified. Uncertainty, on the other hand, has a chronic effect, impacting the model in the same direction over the entirety of the projection.

To deal with uncertainty in our models, we begin with explicitly stating how confident we are that the model assumptions match the underlying probability that generated the experience. This can be done in multiple ways. For example, if the historical data has 100 lapses out of 10,000, we might conclude that the true underlying probability of lapsing that generated this experience is 1% plus or minus 0.2%. Of course, the true prospective uncertainty is greater than this because the future probabilities are not necessarily going to be equal to the probabilities that generated our historical experience.

Other methods can be used to express our level of confidence about the assumptions, such as determining boundaries based on a review of prior data and industry experience. For example, if an assumption had varied no more than ±5% over the last 120 months and had a standard deviation of 1% then it might be reasonable to assume that uncertainty was around ±3%. The rationale for this is that most of the time a value will fall within two standard deviations of its mean. So the mean should be somewhere in that smaller interval.

Incidence

Incidence rates were calculated using Generalized Linear Modeling (GLM) on a subset of the intercompany experience study data.

Data Filtering

The original data set underlying the pivot table in the file *1984-07 LTC Appendix D-B.xls* has 1,642,999 records representing a total of 44,054,975 policy-years of exposure.

We filtered this data to get a set of records commensurate with a hypothetical mid-sized company by only using records with the following characteristics:

Policy Type = "Individual",

Gender <> "unknown"

Benefit Period <> "unknown".

Underwriting = Full

Coverage Type = Comprehensive

EP = 90 days

This results in a data set with 83,717 records representing 2,050,374 exposure years.

GLM

We ran several models looking for factors that had a statistically-significant fit and that seem to make sense. The result is a table with incidence rates that vary by the attained-age category with adjustment factors for BP, Policy Duration, Gender, Marital Discount, and Region.

The basic model output is shown in Appendix 3.

Continuance

Continuance represents the expected number of claims that remain as time progresses. Its corollary in life insurance is survivorship with the claim termination rate taking the place of mortality. We chose to determine reasonable continuance rates from the data, leaving it for the model to solve for the termination rates. We could have chosen to solve for terminations instead.

We also had to decide which of the many factors impacting continuance rates we were interested in modeling. This was important because we did not want to develop assumptions that varied by factors we did not intend to use. For example, policy type is a factor that certainly has some impact on continuance rates, but since we planned to only model individual policies, we chose to ignore it. In the end we settled on gender, age group, and marital status as the factors for which we would vary our continuance assumption. We would have definitely used elimination period had it not been for the decision to only model 90-day eliminations.

Our basic strategy was to develop a base continuance rate for 90-day eliminations, and then create three sets of adjustment factors to take into account variations by gender, age group, and marital status. We chose to do the majority of the analysis within Excel. Our goal was to take the data from the 1984-2007 SOA study and produce a reasonable curve going through the data. Rather than simply creating tables we chose to create a mix of formulaic and tabular solutions. The sole purpose of this was to illustrate that there are many ways to approach the same problem.

Baseline Assumption – 90-Day Elimination Period

To fit the continuance data for a 90-day elimination period, we pulled the data from Appendix E of the 1984-2007 SOA Long-Term Care Intercompany study into Excel. Even within Excel, there are a number of approaches that can be taken to fitting data to a curve. For the baseline assumption, we thought it was important that the data fit fairly closely; we were also more concerned with the fit at earlier time periods than later in the tail. So we decided to fit the data using a weighted linear regression to a log transform of the data. We did this by setting up the equation for the linear regression in Excel and then we used Solver to determine the values of a and b that minimized the weighted squared error terms. There are a number of references online that can be found for this sort of analysis.

We graphed the results of our analysis so we could visually inspect the results. We found that a single exponential term was unable to closely fit the entirety of the curve. It could not drop fast enough while maintaining a tail thick enough to represent a reasonable continuance curve. Therefore we had to modify our design in some way. We decided to use two exponential regressions with one focused on the top of the curve and the other focused on the tail. The results fit the data quite well.

Thus, for the baseline continuance curve we ended up with two exponential equations. To choose where to change from one to the other, we looked to the point where the error in the first equation exceeded the error in the second. This provided us with a smooth transition between the two curves.

One note: such models are quite common in pharmaceutical research in which there are two mechanisms clearing the drug from the body. Here we have two mechanisms causing people to drop off a claim: recovery and death. Recovery tends to dominate in the first part of the curve while death becomes primary for long-term, chronic claimants.

The following graph shows the raw data from the SOA study along with the fitted data we derived.

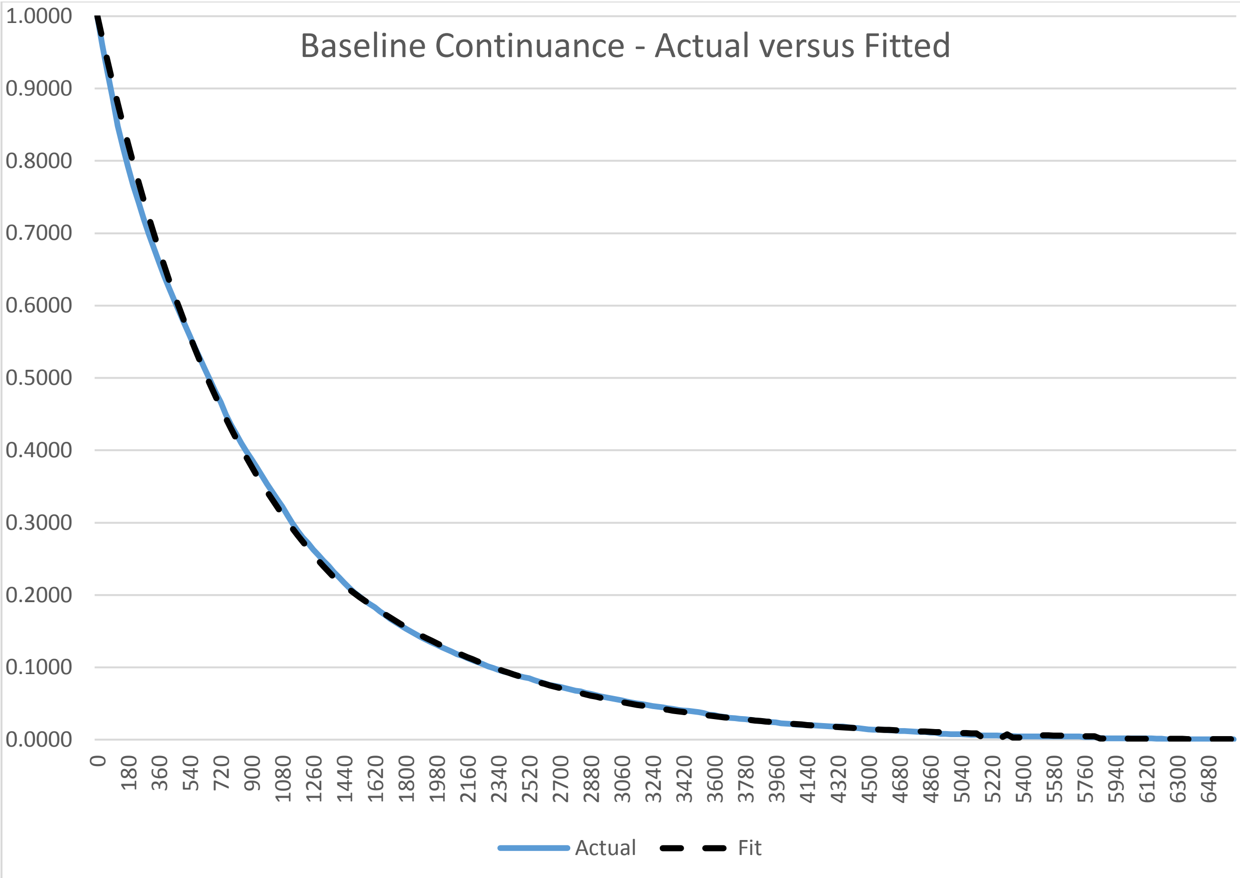

Figure 26: Baseline Continuance – Actual versus Fitted

Adjustment for Gender

Gender is an important factor in estimating claim continuance. There are several reasons for this. Females tend to stay on claim longer than men. This is primarily due to woman having lower mortality than men, but there are other considerations as well. During the early durations of a claim, women tend to have shorter continuance. This is due to women having a higher incidence of claims than men. So on average, women coming on claims are healthier and thus have shorter durations. Then as you get past about 180 days, men start to have shorter claims because the health status evens out and mortality starts taking over.

To adjust for gender, we once again chose to use a formulaic approach. But in this case, rather than coming up with the actual continuance, we came up with a multiplicative adjustment factor that can be applied to the baseline numbers. To do this, we normalized the female and male data through dividing by the overall continuance. There were a number of factors that needed to be considered here. The data by gender included all elimination periods. Consequently, we needed to normalize with all of the data. But we also had to consider the impact that the shorter elimination periods would have on the continuance by gender. We know that in early time periods, female claims have longer continuance. Thus, a 0-day elimination period has much different continuance by gender than a 90-day elimination period does. There is no perfect solution to this. We made the assumption that all of the claims started 60 days into the claim. This is reasonable given the mix of elimination periods included in the study.

Once the data was normalized, we had to decide how to fit the data to a curve. We chose once again to use Excel, but this time we decided that weighting the data was not as important. For this reason, we decided to use the built-in Excel array functions **linest()** and **logest()**. They do essentially the same thing that we did on the baseline using the Excel Solver. We first graphed the results to see what the raw data looked like. We found that the female data looked essentially linear whereas the male data was more of a curve. So we decided to try fitting the female with **linest()** while fitting the male with **logest().**

There was quite a bit of noise as the data thinned out at higher durations, but for the first 10 years, the equations appear to fit the data quite nicely. This is especially true for females who are experiencing higher than average continuance. For males, the curve fits closely for the first three years and then tends to understate the reduction. This was done on purpose to be reasonable yet conservative about the adjustment factors. The parameters feeding the formula were hand-selected by varying slightly the estimates **logest()** generated until the desired pattern emerged.

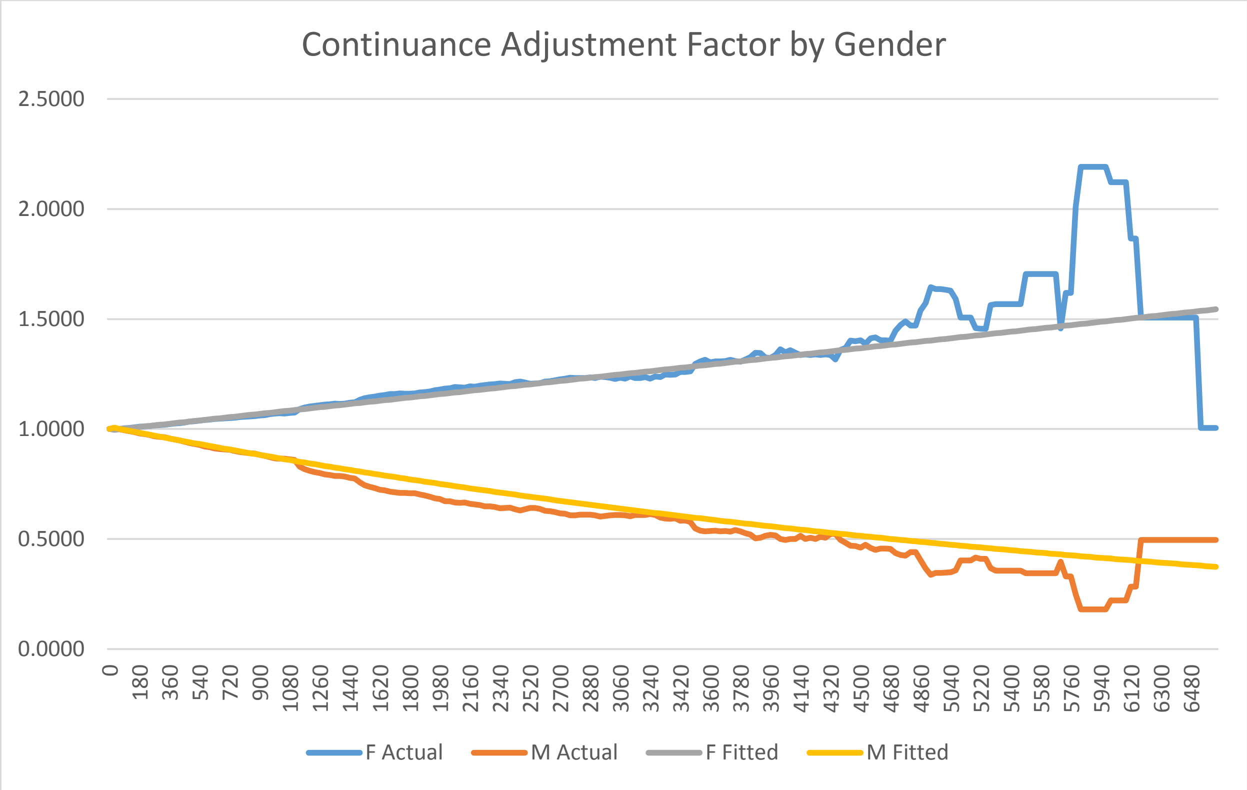

Figure 27: Continuance Adjustment Factor by Gender

Adjustment for Age Bands

Attained age date from Appendix E-3 is divided into five age bands: 55-64, 65-74, 75-84, 85-89, and 90+. In this case, it was determined that coming up with an adjustment factor formulae was not practical. The factors do not exhibit nice patterns by age, but move rather erratically. The best choice appeared to be to create a table with adjustment factors by age band. Because the data moved fairly erratically, we decided to use moving averages of the data to smooth out the table entries. We also decided to set the factor to a constant after 9 years. Past this, the factors started to move around in unreasonable ways due to the lack of data.

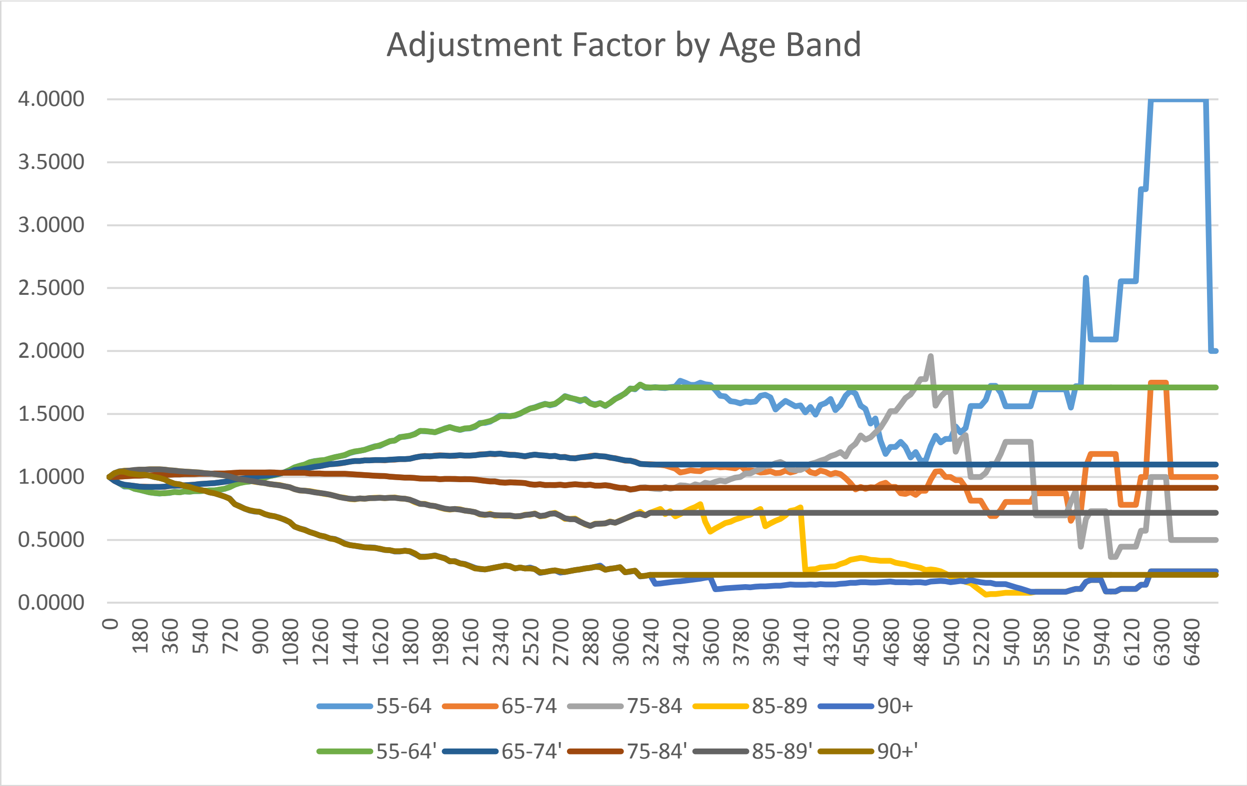

Figure 28: Adjustment Factor by Age Band

Adjustment Factor for Marital Discount

The marital discount for LTC comes from married couples having lower incidence rates and also shorter claim durations than unmarried policyholders. We normalized the data in Appendix E-7 using our gender specific factors. This showed a clear division between those with a marital discount and those without.

Interestingly, the male and female curves tend to track each other quite well up until about three years; then, male continuance jumps up significantly. Since we are only looking for a reasonable result, we chose to determine a single adjustment factor to be applied to the baseline continuance: marital or non-marital. We believe that this will achieve the goal of a reasonable continuance pattern when all of the factors are combined.

To illustrate one additional resource, we used the website "www.zunzun.com" to perform the curve-fitting for these final factors. Normally, fitting data to equations with no rational basis for the selection of which equation you use other than best fit will lead to over-fitting the data and possibly erroneous conclusions. However, in this case, we are again only looking to reasonably reproduce the curve in our model. Hence, any equation that gives a good fit to the data will be usable. The website also gives the ability to pick and choose from a large number of equations that the user would be able to select a reasonable choice, given the user has some idea of what the equation underlying the data should look like.

The following graph shows the result of normalizing the data for gender. The first letter M or U represents married or unmarried while the second letter F or M represents the gender.

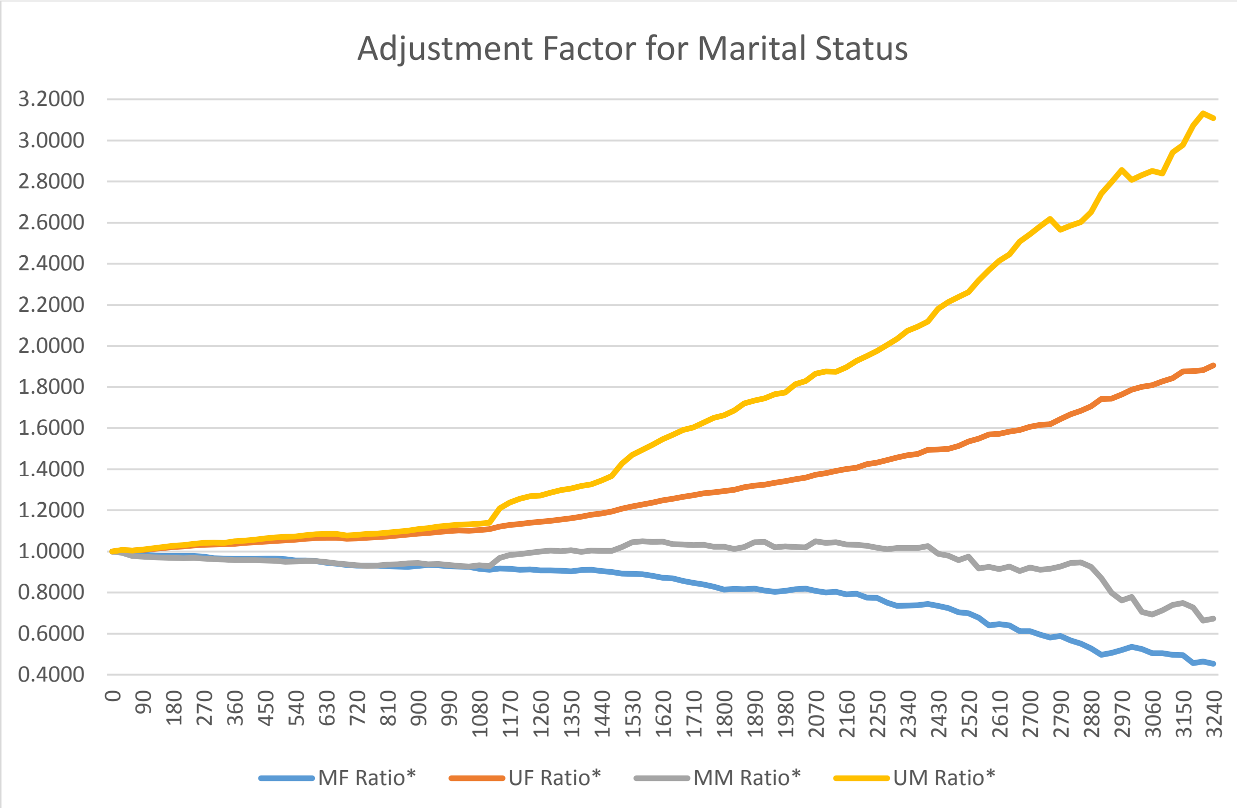

Figure 29: Adjustment Factor for Marital Status

For the non-marital factor, we chose to use a Weibull distribution where $y = a - b * exp(-c * x^d)$. We limited this factor to 2.0 at later durations. For the marital factor, we chose to use a negative exponential. It was one of the best fits and was consistent with prior choices. We limited the marital factor to 0.40 at later durations.

Final Combined Assumptions

Since three of our four factors are formulaic, it would be possible to try to combine each to create one complicated formula and one lookup table. However, for the sake of simplicity in modeling, editing, and debugging we chose to keep each of the factors separate. In our models, t is measured in months.

Baseline Formula

IF t < 48 THEN
  EXP(-0.0010807 * t * 30)
ELSE
  EXP(-0.0008678 * t * 30 – 0.2961)

Gender Adjustment Factor

IF SEX = 0 THEN ‘Sex = 0 is female; sex = 1 is male
  0.00083 * t * 30 + 0.994
ELSE
  1.01 * 0.99985 ^ (t * 30)

Marital Adjustment Factor

IF MARRIED = 1 THEN
  MAX(0.10576 * (1 – EXP(0.0005498 * t * 30)) + 1, 0.4)
ELSE
  MIN(0.056257 + 0.95904 * EXP(0.00000079013 * (t * 30) ^ 1.68771), 2.0)

Attained Age Factor

The attained age factors are simply in a lookup table. After nine years, the factors are held constant. For ease of use, the table was created using 30-day months. All of the other factors are consistent with this.

	Age Band				
Days	55-64	65-74	75-84	85-89	90+
0	1.0000	1.0000	1.0000	1.0000	1.0000
30	0.9742	0.9698	0.9996	1.0245	1.0298
60	0.9474	0.9480	1.0010	1.0417	1.0433
90	0.9282	0.9373	1.0041	1.0478	1.0399
120	0.9219	0.9325	1.0064	1.0505	1.0279
150	0.9063	0.9276	1.0084	1.0541	1.0223
180	0.8968	0.9235	1.0099	1.0572	1.0184
210	0.8864	0.9217	1.0110	1.0585	1.0161
240	0.8783	0.9202	1.0125	1.0596	1.0105
270	0.8721	0.9205	1.0138	1.0595	1.0002
300	0.8688	0.9209	1.0146	1.0600	0.9911
330	0.8701	0.9223	1.0162	1.0574	0.9800
360	0.8736	0.9272	1.0178	1.0530	0.9617
390	0.8807	0.9305	1.0192	1.0492	0.9465
420	0.8783	0.9314	1.0204	1.0461	0.9399
450	0.8841	0.9341	1.0215	1.0423	0.9288
480	0.8837	0.9369	1.0220	1.0405	0.9208
510	0.8864	0.9389	1.0234	1.0374	0.9082
540	0.8878	0.9432	1.0237	1.0343	0.8984
570	0.8894	0.9455	1.0250	1.0301	0.8895
600	0.8918	0.9489	1.0259	1.0266	0.8784
630	0.8935	0.9518	1.0266	1.0232	0.8691
660	0.8996	0.9571	1.0273	1.0178	0.8573
690	0.9066	0.9616	1.0282	1.0117	0.8456
720	0.9191	0.9681	1.0291	1.0040	0.8278
750	0.9380	0.9790	1.0323	0.9905	0.7836
780	0.9502	0.9851	1.0337	0.9810	0.7655
810	0.9585	0.9894	1.0346	0.9757	0.7475
840	0.9641	0.9956	1.0347	0.9695	0.7340
870	0.9728	1.0004	1.0347	0.9628	0.7271

900	0.9806	1.0058	1.0346	0.9559	0.7214
930	0.9952	1.0090	1.0354	0.9506	0.7057
960	1.0072	1.0129	1.0361	0.9435	0.6935
990	1.0146	1.0204	1.0345	0.9389	0.6861
1020	1.0252	1.0274	1.0331	0.9343	0.6748
1050	1.0378	1.0340	1.0330	0.9265	0.6608
1080	1.0543	1.0411	1.0329	0.9181	0.6401
1110	1.0755	1.0542	1.0327	0.9033	0.6064
1140	1.0889	1.0610	1.0328	0.8938	0.5905
1170	1.0990	1.0651	1.0323	0.8899	0.5775
1200	1.1174	1.0733	1.0283	0.8860	0.5627
1230	1.1245	1.0809	1.0273	0.8780	0.5488
1260	1.1307	1.0854	1.0276	0.8724	0.5350
1290	1.1335	1.0945	1.0250	0.8666	0.5262
1320	1.1460	1.1009	1.0244	0.8585	0.5120
1350	1.1560	1.1033	1.0251	0.8500	0.5038
1380	1.1635	1.1080	1.0255	0.8403	0.4890
1410	1.1735	1.1144	1.0258	0.8326	0.4687
1440	1.1913	1.1217	1.0231	0.8263	0.4590
1470	1.2021	1.1267	1.0200	0.8239	0.4531
1500	1.2077	1.1276	1.0174	0.8283	0.4479
1530	1.2164	1.1317	1.0154	0.8256	0.4416
1560	1.2282	1.1322	1.0126	0.8325	0.4384
1590	1.2411	1.1339	1.0093	0.8329	0.4364
1620	1.2500	1.1334	1.0077	0.8345	0.4300
1650	1.2660	1.1329	1.0064	0.8315	0.4209
1680	1.2840	1.1364	1.0018	0.8340	0.4183
1710	1.2898	1.1386	0.9996	0.8348	0.4108
1740	1.3141	1.1408	0.9970	0.8281	0.4091
1770	1.3211	1.1415	0.9950	0.8263	0.4134
1800	1.3260	1.1454	0.9934	0.8195	0.4093
1830	1.3411	1.1533	0.9908	0.8035	0.3917
1860	1.3661	1.1619	0.9872	0.7868	0.3647
1890	1.3644	1.1649	0.9859	0.7840	0.3651
1920	1.3601	1.1638	0.9870	0.7733	0.3699
1950	1.3541	1.1676	0.9854	0.7693	0.3753
1980	1.3712	1.1707	0.9826	0.7578	0.3650
2010	1.3828	1.1695	0.9838	0.7485	0.3502
2040	1.3947	1.1682	0.9849	0.7432	0.3303
2070	1.3831	1.1692	0.9831	0.7452	0.3282
2100	1.3760	1.1716	0.9826	0.7408	0.3144
2130	1.3844	1.1699	0.9822	0.7320	0.3060
2160	1.3883	1.1699	0.9826	0.7263	0.2892

2190	1.4026	1.1731	0.9790	0.7202	0.2736
2220	1.4236	1.1784	0.9748	0.7042	0.2694
2250	1.4337	1.1804	0.9701	0.6993	0.2644
2280	1.4439	1.1855	0.9649	0.7030	0.2721
2310	1.4604	1.1826	0.9646	0.6974	0.2795
2340	1.4825	1.1846	0.9567	0.6943	0.2876
2370	1.4828	1.1807	0.9546	0.6943	0.2949
2400	1.4811	1.1756	0.9568	0.6950	0.2905
2430	1.4897	1.1747	0.9559	0.6860	0.2733
2460	1.5054	1.1697	0.9538	0.6889	0.2802
2490	1.5223	1.1653	0.9505	0.6980	0.2728
2520	1.5434	1.1717	0.9419	0.7030	0.2730
2550	1.5511	1.1773	0.9376	0.7077	0.2686
2580	1.5660	1.1729	0.9414	0.6898	0.2421
2610	1.5795	1.1703	0.9347	0.6951	0.2437
2640	1.5738	1.1639	0.9368	0.7067	0.2517
2670	1.5802	1.1661	0.9355	0.7095	0.2577
2700	1.6090	1.1581	0.9388	0.6907	0.2489
2730	1.6388	1.1581	0.9328	0.6721	0.2454
2760	1.6278	1.1503	0.9387	0.6636	0.2525
2790	1.6181	1.1479	0.9414	0.6631	0.2588
2820	1.6082	1.1534	0.9388	0.6448	0.2642
2850	1.6143	1.1588	0.9359	0.6269	0.2722
2880	1.5850	1.1625	0.9372	0.6107	0.2786
2910	1.5710	1.1687	0.9308	0.6268	0.2860
2940	1.5851	1.1637	0.9310	0.6275	0.2835
2970	1.5663	1.1618	0.9333	0.6315	0.2626
3000	1.5878	1.1538	0.9287	0.6439	0.2689
3030	1.6183	1.1462	0.9197	0.6363	0.2759
3060	1.6379	1.1382	0.9128	0.6524	0.2829
3090	1.6645	1.1306	0.9126	0.6699	0.2427
3120	1.7014	1.1293	0.8995	0.6868	0.2488
3150	1.7003	1.1197	0.9044	0.7033	0.2548
3180	1.7330	1.1029	0.9131	0.7080	0.2127
3210	1.7145	1.0999	0.9128	0.6942	0.2151
3240	1.7109	1.0985	0.9120	0.7142	0.2212

Mortality

Appendix J-2 from the SOA LTC Intercompany Study demonstrates Mortality Rates by Attained Age Cohorts, split between male and female, compared to the 1994 Group Annuity Mortality (GAM) table, the 2000 Annuity Mortality Table, the 2001 Valuation Basic Table, and the 2008 Valuation Basic Table. Any of these tables could be used as a starting point for the mortality assumption development. We selected the 1994 GAM table because it appears to have the most consistent slope with regard to the LTC data from the Intercompany study. The comparison between the LTC data and 1994 GAM are smoothed to create the following adjustments.

Attained Age Cohort	Smoothed Adjustment to 1994 GAM	
	Male	Female
20-29	30%	35%
30-39	35%	45%
40-49	35%	45%
50-59	35%	45%
60-69	35%	45%
70-79	50%	50%
80-89	65%	60%
90-99	90%	80%

In total, the SOA LTC data results in mortality significantly less than the 1994 GAM table, 54% for female and 50% for male. Based on a sample of actual rate review filings, it appears that companies are experiencing results near 80% of the 1994 GAM. For our model we scaled the results up to 80% of the 1994 GAM and utilized a wide variance to account for differing results. We believe this accounts for companies incorrectly recording deaths as lapses, as mentioned previously.

Appendix J-10 demonstrates Mortality by Underwriting Type, Issue Age Cohort, and Duration. This table may be utilized to create Select adjustments by duration. The table shows durations individually for years one (1) through ten (10) and it groups durations eleven (11) and above. As a result, it is assumed that ultimate mortality is reached by duration eleven (11). To determine the durational wear-off of full underwriting, we compare the mortality slopes for full underwriting against simplified and guaranteed issue policies. Assuming that ultimate mortality is reached at duration eleven (11), we normalize the mortality rates to 1.0 at this duration so that the slopes for varying underwriting types can be compared directly.

Duration	Issue Age Cohort			
	50-59	60-69	70-79	80-89
1	40%	70%	75%	100%
2	50%	80%	80%	100%
3	60%	85%	85%	100%
4	70%	90%	90%	100%
5	80%	95%	95%	100%

6	90%	100%	100%	100%
7	100%	100%	100%	100%
8	100%	100%	100%	100%
9	100%	100%	100%	100%
10	100%	100%	100%	100%
11+	100%	100%	100%	100%

Combining these two tables results in a Mortality table with Select and Ultimate values based on the 1994 GAM table. This table is utilized as the mean scenario in our model. The final results are provided in an Excel table along with this document.

One possible method to calculate the variance is to assume a distribution (for example the Beta distribution) and estimate the variance given the specific data points. Unfortunately, the Mortality tables in the SOA study appendix do not allow us to see the underlying data. Given that we only have summary snapshots of the results, and that Mortality data is likely underestimated in the study, we believe it is appropriate to create scenarios assuming the force of Mortality could vary by 20% in either direction. This allows us to fully explore the effect of Mortality results on the low and high end of the spectrum.

Lapse

The basis for lapse assumptions is Appendix F from the SOA LTC Intercompany Study. This provides multiple snapshots of results by issue age, duration, attained age, policy type, etc. As expected, the total lapse rate assumptions look quite high with duration one (1) above 9% and an ultimate rate above 2%. This is inconsistent with the experience we have reviewed and the assumptions we have seen used by most other companies. It is in line with the expectation that lapses are overestimated and mortality is underestimated within the SOA LTC Intercompany Study. It is worth noting that Appendix F states that each exhibit "excludes those companies who do not distinguish between deaths and terminations." It is still likely that companies that choose to track death and lapse independently do not do so with 100% accuracy for policies without death benefits.

Our model focuses on Individual policies with a 90-day elimination period. Consequently, we used Appendix F-5 as a starting point for our analysis. This particular slice of the data shows much more reasonable results for Individual policies with 90-day elimination periods. The values are smoothed below to attain the mean value for our model.

Duration	Lapse
1	5.50%
2	3.75%
3	3.00%
4	2.50%
5	2.25%
6	2.10%
7	2.00%
8	1.85%
9	1.70%
10	1.60%
11	1.50%
12	1.25%
13	1.10%
14	1.00%
15	1.00%
16	1.00%
17	1.00%
18	1.00%
19	1.00%
20	1.00%

Similar to Mortality, we did not use a specific distribution to calculate the variance of the Lapse assumption.

Economic Scenario Assumptions

We took a simplified approach to modeling assets for this project. While we recognize that the investments supporting LTC policies have risks associated with them depending on the nature and quality of the assets, these risks are not unique to LTC. Any asset portfolio has risks inherent to it; default, liquidity, reinvestment, and exposure to movements in interest rates are a few of these. Our focus is on the risks associated with uncertainty regarding future interest rates (i.e., the risk that the investment returns on reserves will not be sufficient to fund the policies future obligations), and how this can potentially be mitigated by policy design.

To explore this, we needed a set of economic scenarios that includes both future interest rates and future inflation rates.

We chose to use the AAA Scenario Generator (v 7.0.4) to create 200 pseudo-random economic scenarios of monthly asset returns over a 60-year time period, starting at third-quarter-end 2013. We modeled a simple asset portfolio and believe that 200 scenarios should provide a range of economic outcomes vast enough to enable us to adequately explore the risks of an LTC policy, as these risks relate to shifts in interest rates.

The methodology for setting the CPI sets the base year to be July 1980 (July time periods were chosen since the latest month of available historical data for CPI was July 2013). A curve was fit using each 1-year constant maturity treasury (monthly values) for the same time period (i.e., from July 1980 to July 2013). To achieve a better fit, the 1-year treasury was capped (7% through 12%, constraint added to Excel solver) and it was determined that a cap of 12% had the best fit. The fitted CPI rates using linear regression provided a slope = 0.345 and an intercept term = 1.45%.

Visually, the fitted CPI versus the actual CPI over that time period can be seen below.

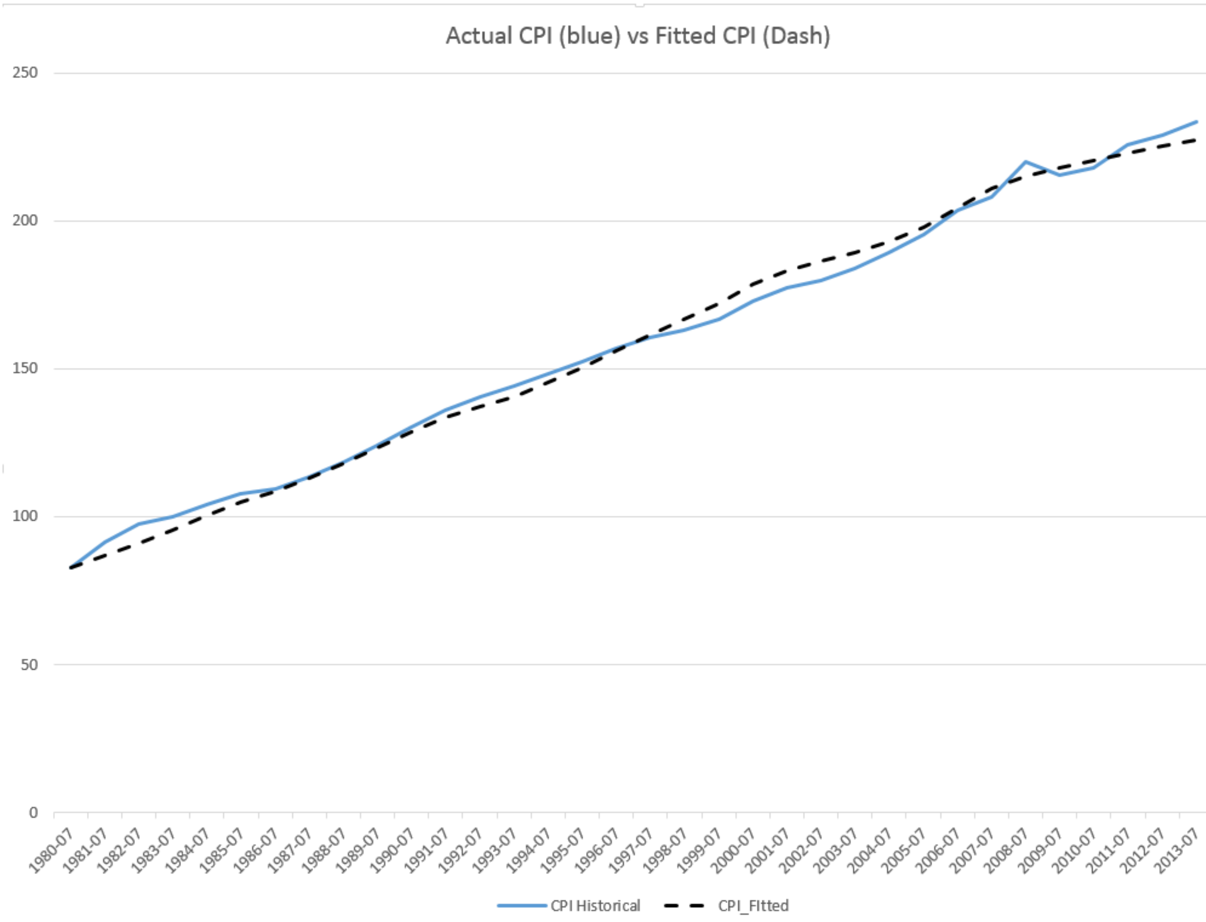

Figure 30: Actual CPI versus Fitted CPI

The source for historical CPI data is: ftp://ftp.bls.gov/pub/special.requests/cpi/cpiai.txt

The source for U.S. Treasury data is:
http://www.federalreserve.gov/datadownload/Output.aspx?rel=H15&series=153e3c14864f1d7d0117fb3b83f55d2a&lastObs=&from=&to=&filetype=csv&label=include&layout=seriescolumn

Rather than supplying the tables for these indices, we decided to use a formulaic approach within our model, which allows us to easily change the model to reflect any minor revisions to the slope or intercept parameters of each index, if necessary. This avoids the need to recreate the tables upon each revision.

Turning our attention to the asset part of our model, we linked our portfolio returns to be some combination of one or more of the Treasuries plus an implied spread.

Appendix 2: Beta Distribution

Beta Distribution Overview

The beta distribution provides an elegant way to express the confidence we have in our model's transition probabilities. A high-level overview of this distribution will be given here.

A beta-distributed random variable $\mathbf{X}$ with parameters α and β (i.e., $\mathbf{X} \sim \text{Beta}(\alpha, \beta)$) has the following key properties:

- Its domain is the interval [0,1]
- $E[\mathbf{X}] = \frac{\alpha}{\alpha+\beta}$

These properties suggest that it might be a good distribution to express our confidence level in the underlying probabilities that generated an experience set. For a given set of experience, we can compare the beta distribution to the likelihood function. In general, this exercise shows that the beta distribution and likelihood function have nearly identical shapes. This implies that the probabilities from the Beta distribution are consistent with the relative likelihood of each potential probability rate.

For example, say that we observe a set of 1,000 policies for a year, and that 50 of the policies lapse, while 950 persist. The lapse rate we observe here is simply $\frac{50}{1{,}000} = 5.0\%$. However, this does not tell us what the underlying force of lapsing was that produced this observation.

To evaluate this, we could look at the likelihood function. The likelihood function tells us how likely it is that any possible lapse probability produced this experience. For example, if x represents the specific 1,000 observations made, then the probability that the lapse rate 4.9% produced x is:

$$\mathcal{L}(0.049|x) = (0.049)^{50}(0.951)^{950} \approx 6.0436 \times 10^{-87}$$

If θ represents every possible rate of lapsing, then the likelihood function $\mathcal{L}(\theta|x)$ would look like this:

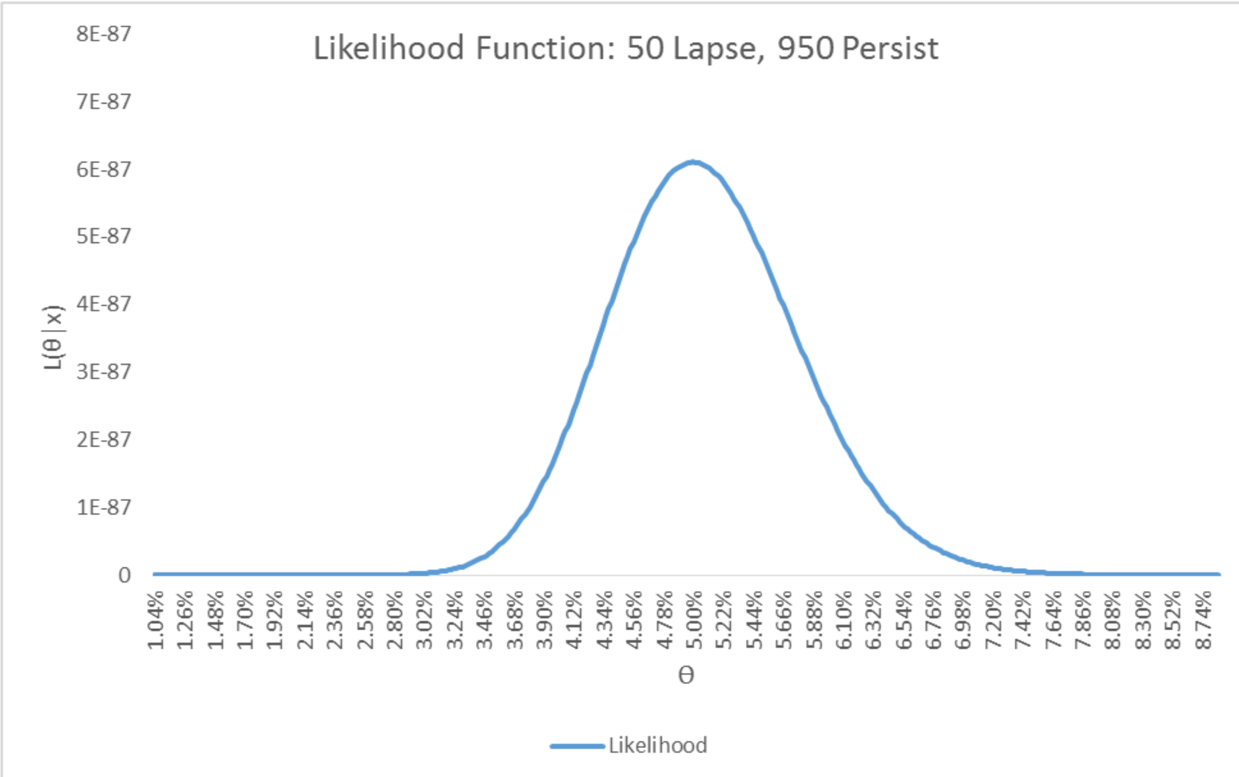

Figure 31: Likelihood Function

Taking the derivative of this likelihood function proves that the probability of lapsing most likely to result in 50 lapsing and 950 persisting is, in fact, 5.0% (i.e., the Maximum Likelihood Estimate (MLE) is 5.0%). The added value of looking at the entire likelihood function (rather than just the MLE) is that it allows us to see a range of probabilities that could have plausibly produced these results. In this case, we can see that anything from about 3.9% to about 6.7% could plausibly be the true underlying lapse rate.

It turns out that the Beta distribution **X**~Beta(50,950) has almost the same shape as the likelihood function. This is illustrated on the following graph:

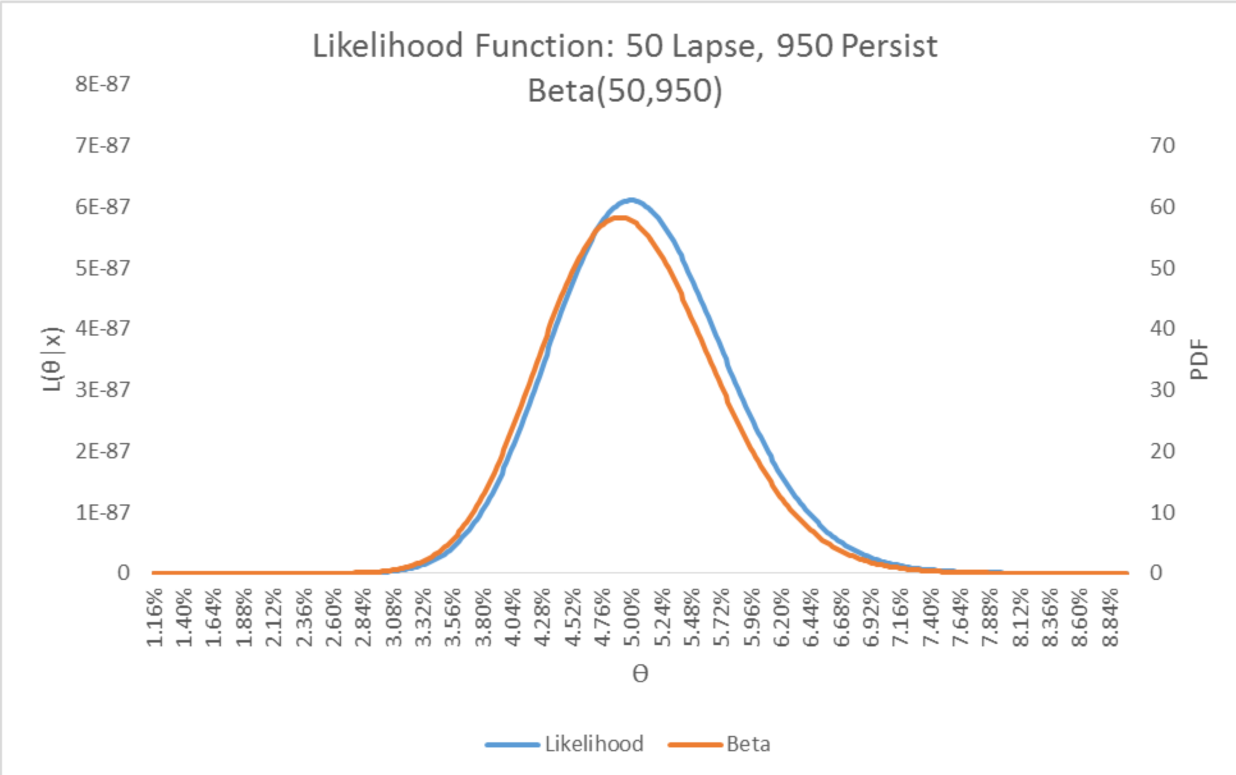

Figure 32: Likelihood Function and Beta Distribution

The likelihood function is not a pdf—the probability that *any* given probability would result in our 1,000 precise observations is incredibly small and the total area under a likelihood function is typically going to be much smaller than 1.0. If we could scale the whole likelihood function up so that the total area under the curve did equal one, it would be a pdf that would tell us the relative likelihood of each potential value. The Beta distribution is an elegant approximation of a scaled-up likelihood function.

Using the Beta Distribution to Model Parameter Risk

Actuaries are accustomed to developing single sets of tables to describe mortality rates, claim incidence rates, and so forth. For each of those rates, a second number can be added that represents the degree of certainty the actuary has in that particular rate. For example, an actuary might say the mortality rate for a 45-year-old is 0.00366 (i.e., $q_{45} = 0.00366$). However, that does not express his level of confidence in that rate.

The beta distribution can be used to describe a mortality rate that includes both the best estimate for the underlying probability and also a weighted range of other plausible values. To do this, two values are needed: α and β. However, this can be parameterized using numbers that are more meaningful. The two parameters we used in the report are *mean* and *sample strength* (SS) defined as:

$$Mean = \frac{\alpha}{\alpha + \beta}$$

$$SS = \alpha + \beta$$

The mean is the best estimate of the parameter. The sample strength is the number of observations supporting that number. So using the lapse example above, one could say there is a 5% lapse rate with $SS = 1{,}000$. This means that the confidence in the 5% lapse rate is what it would be if there are 1,000 trials supporting it.

For the mortality example, we would say that the mean mortality rate is $q_{45} = 0.00366$, with $SS_{45} = 55{,}000$, which can be thought of as 55,000 observations supporting our estimate that the mortality rate is 0.00366. If one did in fact observe a sample of 55,000 45-year-olds and saw that 0.366% of them died,[26] then the estimate of the underlying mortality rate that generated these results is beta distributed with $\alpha = 201.3$ and $\beta = 54{,}798.7$ (i.e., $X \sim \text{Beta}(201.3, 54798.7)$). By taking the inverse of this distribution, we see that we can be 90% confident that the real probability of death is between 0.00325 and 0.00409.

Just as actuaries need to use some smoothing and professional judgment to come up with the final best estimates of other model parameters, *aggregation* and professional judgment needs to be used to estimate the sample strength parameter; the sample strength parameter should reflect the aggregate confidence in the rate, taking into account all available sources of information, including the rates of related cells. As an example, suppose incidence rates are being estimated and somehow 10,000 exposures at age 75 are observed, only 100 at age 76, and then another 10,000 at age 77. Evaluating the rate at age 76 in isolation, results in an SS of 100, which would have a very high variance of parameter risk. What would make more sense though is to decide that the age 76 rate is the average of the neighboring rates, with an SS parameter equal to something like 10,000, reflecting the 10,000 exposures above and below it.

Appendix 3: Incidence Rate Scenario Factors

The following table shows the first two hundred incidence rate scenario factors. These were derived through simulation using the incidence rates and standard errors that came out of the GLM used to analyze historical incidence rates. A correlation of 75% was assumed between ages.

SIM	< 40	40-49	50-59	60-64	65-69	70-74	75-79	80-84	85-89	90+
1	0.959	1.184	1.037	1.030	1.037	1.031	1.013	1.017	1.006	1.015
2	0.858	0.971	0.985	0.970	0.969	0.976	0.998	0.986	0.980	1.019
3	1.135	1.146	1.015	1.009	1.015	0.997	0.987	1.009	1.005	1.001
4	1.807	2.033	1.083	1.082	1.069	1.052	1.041	1.032	1.039	1.111
5	0.587	0.862	0.951	0.930	0.938	0.952	0.998	0.966	0.966	0.921
6	0.738	0.793	0.986	0.986	0.961	0.991	0.983	0.992	0.979	0.964
7	3.789	2.372	1.181	1.159	1.104	1.052	1.087	1.049	1.075	1.201
8	0.380	0.646	0.886	0.923	0.910	0.964	0.969	0.950	0.945	0.915
9	0.633	0.818	0.946	0.941	0.945	0.971	0.977	0.978	0.973	0.936
10	3.037	1.861	1.184	1.168	1.103	1.080	1.058	1.037	1.076	1.146
11	1.042	1.015	1.005	1.009	0.984	1.006	0.987	0.994	1.005	0.997
12	0.983	0.920	1.004	0.963	0.991	0.990	0.991	1.004	0.990	0.930
13	1.977	1.729	1.129	1.118	1.028	1.057	1.044	1.041	1.046	1.093

[26] Observing 0.366% of 55,000 lives die implies you saw 201.3 people die. Of course that really isn't possible, but the Beta distribution still works with those parameters.

14	1.084	0.908	0.985	1.009	1.009	0.999	1.002	1.015	0.988	0.983
15	0.469	0.708	0.921	0.898	0.937	0.973	0.991	0.965	0.974	0.942
16	0.330	0.548	0.918	0.903	0.910	0.934	0.958	0.944	0.940	0.901
17	2.224	1.476	1.130	1.075	1.065	1.035	1.044	1.043	1.018	1.085
18	0.521	0.834	0.948	0.978	0.984	0.972	0.970	0.974	0.970	0.924
19	1.377	1.094	1.065	1.049	1.011	1.044	1.016	1.016	1.022	1.045
20	1.310	0.977	0.987	1.024	0.999	0.997	1.009	0.991	0.997	0.986
21	0.734	0.910	0.968	0.971	0.983	0.990	0.983	0.983	1.005	1.024
22	3.434	2.127	1.247	1.183	1.102	1.085	1.056	1.062	1.046	1.199
23	0.288	0.471	0.856	0.872	0.923	0.942	0.945	0.960	0.933	0.882
24	0.965	0.976	1.031	1.011	1.022	0.994	0.992	0.995	1.022	0.996
25	0.774	0.778	0.921	0.957	0.950	0.995	0.985	0.979	0.986	0.927
26	1.702	1.170	1.044	1.029	1.027	1.023	1.020	1.027	1.023	1.095
27	0.274	0.392	0.825	0.861	0.908	0.919	0.951	0.933	0.925	0.829
28	1.483	1.212	1.087	0.999	1.011	1.021	0.999	1.015	1.003	1.041
29	0.537	0.911	0.925	0.955	0.983	0.996	0.979	0.987	0.959	0.963
30	2.453	1.547	1.100	1.121	1.092	1.054	1.033	1.046	1.045	1.084
31	1.115	1.108	0.985	1.031	0.977	1.014	1.021	1.018	1.027	1.006
32	0.474	0.694	0.912	0.952	0.944	0.977	0.976	0.976	0.973	0.919
33	0.505	0.736	0.926	0.921	0.943	0.961	0.969	0.982	0.965	0.941
34	0.316	0.367	0.831	0.879	0.890	0.948	0.936	0.945	0.927	0.869
35	1.073	1.545	1.010	1.032	1.037	1.032	1.004	1.015	1.005	0.978
36	1.498	1.436	1.060	1.008	1.023	1.045	1.012	1.003	1.022	1.039
37	0.677	0.833	0.937	0.978	0.962	0.975	0.970	0.997	0.986	0.927
38	0.454	0.842	0.913	0.949	0.961	0.959	0.981	0.977	0.984	0.918
39	0.374	0.539	0.864	0.851	0.910	0.932	0.954	0.949	0.940	0.889
40	1.192	1.036	1.040	1.043	1.025	1.011	1.012	1.008	1.017	1.017
41	0.260	0.570	0.863	0.863	0.901	0.930	0.949	0.940	0.932	0.896
42	0.526	0.516	0.879	0.918	0.938	0.935	0.946	0.944	0.944	0.920
43	0.880	0.976	1.062	0.996	0.972	1.006	0.999	0.999	0.976	1.029
44	0.617	0.778	0.917	0.919	0.974	0.957	0.965	0.974	0.959	0.911
45	0.954	0.830	0.976	0.977	0.985	0.988	0.986	0.987	0.984	1.005
46	0.692	0.729	0.954	0.977	1.005	0.983	0.981	0.988	0.976	0.948
47	0.913	1.197	0.981	0.994	0.967	1.000	0.985	1.006	0.995	1.009
48	0.646	0.692	0.905	0.951	0.960	0.995	0.955	0.971	0.966	0.902
49	0.910	0.891	0.970	0.989	0.995	0.981	0.994	0.982	0.991	1.000
50	1.357	1.334	1.039	1.054	1.013	0.997	1.027	1.002	1.022	1.044
51	1.168	0.921	1.051	1.038	0.987	1.020	0.997	1.014	1.013	1.005
52	0.384	0.658	0.907	0.896	0.962	0.963	0.963	0.951	0.937	0.877
53	0.284	0.470	0.843	0.854	0.890	0.911	0.939	0.944	0.912	0.840
54	0.564	0.809	0.972	0.933	0.967	0.980	0.974	0.981	0.976	0.923

55	1.452	0.907	1.028	1.043	1.028	1.015	1.024	1.001	1.012	1.075
56	1.101	1.216	0.999	1.025	1.063	1.043	0.997	1.006	1.024	1.019
57	0.743	0.749	0.999	0.971	1.007	0.977	0.998	1.002	0.985	0.959
58	0.683	0.841	0.948	0.950	0.973	0.973	0.965	0.962	0.972	0.909
59	1.007	1.134	0.999	1.007	1.003	0.995	1.017	1.001	1.007	1.031
60	0.763	0.870	0.958	0.979	0.983	0.960	0.978	0.980	0.973	0.964
61	1.747	1.228	1.101	1.077	1.030	1.023	1.035	1.038	1.033	1.118
62	1.404	1.237	1.009	1.026	1.012	1.008	1.011	1.008	1.018	1.015
63	1.238	1.390	1.034	1.039	1.040	1.038	1.004	1.027	1.021	1.052
64	1.476	1.165	1.087	1.078	1.045	1.032	1.035	1.019	1.029	1.088
65	0.976	0.925	0.974	1.004	1.029	1.008	1.013	0.996	1.015	0.994
66	0.951	1.215	1.019	1.038	1.008	1.041	1.014	1.017	1.017	0.990
67	1.104	0.942	0.993	1.015	1.014	1.019	0.999	1.023	1.000	1.057
68	0.534	0.482	0.902	0.917	0.938	0.958	0.946	0.976	0.946	0.885
69	0.517	0.903	0.932	0.947	0.966	0.979	0.994	0.985	0.970	0.945
70	1.340	1.012	1.064	1.039	1.014	1.007	1.031	0.996	1.008	1.036
71	1.779	1.389	1.110	1.091	1.071	1.047	1.042	1.046	1.045	1.073
72	1.469	1.196	1.129	1.115	1.058	1.044	1.038	1.032	1.049	1.104
73	1.010	1.265	0.994	1.016	0.994	1.000	0.994	0.996	0.994	1.006
74	0.912	0.897	1.017	0.982	0.954	1.015	0.992	0.998	1.012	0.950
75	1.445	0.933	1.024	1.009	0.983	1.012	1.005	1.008	1.021	0.991
76	0.645	0.711	0.890	0.952	0.961	0.962	0.979	0.952	0.978	0.945
77	1.085	0.995	1.028	0.989	0.990	1.024	1.013	1.004	0.998	1.058
78	0.547	0.756	0.917	0.894	0.942	0.963	0.960	0.983	0.956	0.926
79	1.583	1.195	1.006	1.080	1.037	1.034	1.030	1.030	1.017	1.098
80	1.416	1.417	1.050	0.990	0.991	0.991	1.006	1.013	1.005	1.017
81	0.636	0.769	0.947	0.984	0.991	0.982	0.993	0.980	0.952	0.940
82	1.028	1.033	0.990	1.017	0.995	1.010	0.990	0.998	1.001	0.962
83	1.137	0.929	0.995	1.017	1.032	0.994	1.015	1.011	0.998	1.072
84	5.624	2.732	1.280	1.180	1.160	1.106	1.082	1.099	1.085	1.254
85	0.765	0.748	0.970	1.011	0.996	0.991	0.981	0.983	0.983	0.972
86	0.786	0.962	0.937	0.960	0.980	0.984	0.995	0.978	0.990	0.990
87	0.552	0.656	0.895	0.918	0.951	0.960	0.985	0.967	0.982	0.909
88	0.740	0.672	0.930	0.966	0.952	0.967	0.983	0.982	0.985	0.934
89	0.785	0.909	0.930	0.981	0.960	0.976	0.977	0.995	0.976	0.981
90	0.647	0.799	0.944	0.926	0.955	0.970	0.980	0.974	0.963	0.943
91	1.392	1.316	1.085	1.035	1.053	1.003	1.027	1.027	1.038	1.043
92	1.988	1.481	1.066	1.094	1.033	1.043	1.041	1.050	1.042	1.101
93	1.389	1.255	1.076	1.091	1.035	1.035	1.017	1.020	1.031	1.041
94	0.783	0.901	1.015	0.957	0.953	1.001	0.989	0.981	0.992	0.994
95	1.101	1.285	1.000	1.050	1.028	1.026	1.028	1.017	1.035	1.035

96	1.612	1.029	1.100	1.050	1.057	1.043	1.020	1.012	1.036	1.033
97	1.090	1.372	1.095	1.024	1.027	1.019	1.011	1.017	1.022	1.047
98	0.410	0.648	0.935	0.883	0.935	0.957	0.963	0.966	0.973	0.934
99	0.831	0.623	0.884	0.967	0.952	0.959	0.982	0.981	0.959	0.933
100	2.896	1.899	1.201	1.131	1.120	1.060	1.063	1.065	1.058	1.161
101	2.128	1.463	1.151	1.134	1.039	1.027	1.034	1.029	1.053	1.094
102	1.261	1.007	1.033	1.001	1.006	0.995	0.995	0.994	1.013	1.039
103	1.561	1.270	1.041	1.072	1.054	1.029	1.011	1.026	0.996	1.054
104	0.911	0.970	1.100	1.040	1.022	1.024	1.014	1.023	1.020	1.044
105	0.895	0.853	0.981	0.993	1.009	0.994	1.000	0.991	0.993	1.030
106	1.521	1.546	1.041	1.047	1.046	1.030	1.018	1.028	1.041	1.046
107	0.311	0.436	0.865	0.835	0.883	0.934	0.939	0.932	0.922	0.875
108	0.770	0.774	0.937	0.976	0.976	0.996	0.983	0.985	0.974	0.951
109	0.948	0.918	0.956	0.991	1.028	1.030	0.998	1.007	1.004	1.009
110	2.409	1.576	1.095	1.064	1.059	1.028	1.024	1.033	1.032	1.091
111	1.118	1.283	1.020	1.068	1.035	1.004	1.010	1.026	1.026	1.058
112	0.823	0.866	0.972	0.975	0.992	1.007	0.969	0.991	0.996	1.003
113	0.245	0.464	0.808	0.875	0.872	0.929	0.934	0.948	0.914	0.821
114	1.273	1.432	1.031	1.003	1.018	1.032	1.022	1.032	1.011	1.070
115	0.558	0.837	0.916	0.915	0.932	0.964	0.963	0.966	0.958	0.900
116	0.973	1.140	1.002	1.010	1.036	1.025	1.006	1.014	0.995	1.034
117	1.922	1.506	1.114	1.129	1.071	1.045	1.017	1.027	1.035	1.116
118	1.446	1.041	1.032	1.035	1.019	1.008	1.024	1.000	1.005	1.040
119	1.736	1.130	1.019	1.030	1.044	1.016	1.021	1.009	1.011	1.035
120	0.388	0.633	0.898	0.929	0.927	0.948	0.947	0.970	0.939	0.934
121	1.960	1.242	1.072	1.066	1.089	1.036	1.032	1.026	1.016	1.072
122	0.523	0.810	0.922	0.917	0.963	0.981	0.975	0.963	0.985	0.961
123	1.366	1.392	1.047	1.039	1.047	1.018	1.027	1.013	1.013	1.052
124	1.094	0.995	0.977	0.974	1.008	1.034	1.001	1.008	1.013	1.045
125	1.645	1.253	1.074	1.054	1.042	1.027	1.039	1.021	1.012	1.062
126	1.724	1.119	1.039	1.057	1.021	1.026	1.031	1.025	1.027	1.016
127	0.889	0.872	0.992	0.997	0.966	0.996	0.988	0.986	0.974	0.985
128	0.855	1.228	1.034	0.988	1.013	1.021	1.003	1.013	1.004	1.011
129	0.563	0.663	0.931	0.937	0.966	0.979	0.960	0.953	0.969	0.890
130	1.726	1.318	1.086	1.037	1.064	1.030	1.009	1.025	1.039	1.102
131	1.653	1.183	1.082	1.032	1.000	1.048	1.006	1.022	1.040	1.023
132	1.568	1.175	1.055	1.014	1.036	1.022	1.002	1.013	1.035	1.038
133	0.397	0.638	0.894	0.962	0.947	0.955	0.967	0.975	0.979	0.893
134	1.315	1.185	1.162	1.025	1.029	1.039	1.035	1.043	1.040	1.069
135	0.988	1.008	1.043	1.028	1.050	1.009	1.004	1.025	1.012	1.036
136	1.411	1.307	1.064	1.073	1.023	1.025	1.010	1.020	1.025	1.049

137	0.830	0.851	0.964	0.939	0.969	0.978	0.986	0.970	0.974	0.928
138	0.483	0.480	0.930	0.916	0.928	0.938	0.946	0.970	0.947	0.901
139	2.501	1.755	1.149	1.133	1.074	1.067	1.051	1.057	1.071	1.105
140	1.824	1.049	1.042	1.066	1.010	1.007	0.992	1.006	1.034	1.103
141	0.779	0.680	0.902	0.890	0.924	0.957	0.965	0.971	0.948	0.942
142	1.240	1.085	1.022	1.068	1.042	1.033	0.997	1.003	1.006	1.017
143	1.317	0.909	1.053	1.022	1.025	1.014	0.990	1.020	1.005	1.016
144	1.739	1.266	1.081	1.112	1.059	1.024	1.015	1.022	1.038	1.083
145	0.539	0.614	0.876	0.887	0.886	0.936	0.931	0.942	0.929	0.906
146	0.382	0.677	0.902	0.916	0.916	0.941	0.944	0.969	0.963	0.871
147	0.774	1.243	1.008	1.004	1.003	1.012	0.997	1.006	1.008	0.968
148	0.732	0.751	0.925	0.950	0.929	0.961	0.967	0.959	0.955	0.921
149	1.350	1.150	1.051	1.007	1.029	0.993	0.993	1.001	1.017	0.998
150	1.671	1.439	1.075	1.025	1.034	1.042	1.018	1.015	1.029	1.042
151	2.844	1.927	1.170	1.210	1.104	1.078	1.059	1.056	1.066	1.154
152	1.297	1.162	1.037	1.023	1.033	1.028	1.014	1.022	0.989	1.037
153	0.460	0.709	0.918	0.920	0.960	0.968	0.947	0.959	0.955	0.947
154	0.259	0.512	0.800	0.866	0.896	0.935	0.925	0.947	0.929	0.828
155	0.978	1.338	1.008	0.982	1.016	1.009	1.008	1.012	1.033	1.026
156	0.688	0.849	0.944	0.958	0.948	0.974	0.975	0.993	1.000	0.990
157	1.990	1.889	1.134	1.129	1.056	1.066	1.048	1.052	1.052	1.134
158	3.173	2.289	1.168	1.160	1.097	1.069	1.064	1.060	1.076	1.128
159	3.508	1.701	1.183	1.110	1.099	1.053	1.062	1.059	1.055	1.113
160	4.691	2.362	1.252	1.233	1.150	1.107	1.084	1.084	1.123	1.243
161	1.013	1.204	0.976	0.961	1.007	1.002	1.004	1.010	0.978	1.009
162	1.769	1.213	1.061	1.059	1.056	1.034	1.031	1.030	1.037	1.042
163	0.488	0.645	0.891	0.895	0.921	0.956	0.959	0.979	0.966	0.913
164	0.516	0.724	0.924	0.932	0.924	0.951	0.972	0.975	0.969	0.952
165	1.212	1.170	0.997	1.037	1.020	1.023	1.028	1.006	1.013	1.023
166	0.519	0.801	0.894	0.903	0.926	0.943	0.966	0.969	0.956	0.919
167	2.266	1.165	1.078	1.071	1.044	1.033	1.028	1.028	1.026	1.073
168	3.999	2.044	1.229	1.159	1.153	1.096	1.064	1.060	1.087	1.182
169	0.599	0.738	0.993	0.932	0.932	0.969	0.991	0.981	0.972	0.938
170	1.008	1.159	0.997	1.011	1.013	0.995	0.998	1.006	1.003	1.019
171	1.116	0.850	0.991	1.011	1.005	0.997	0.997	0.990	0.992	0.970
172	1.092	1.086	1.037	1.014	0.986	0.997	1.005	1.008	1.017	1.023
173	1.260	1.101	1.020	0.954	0.999	0.980	1.007	0.989	1.010	0.966
174	0.787	0.907	0.988	1.055	1.018	0.999	0.989	0.993	1.000	0.940
175	1.967	1.569	1.141	1.131	1.057	1.012	1.044	1.033	1.032	1.059
176	0.824	0.863	1.043	0.987	1.018	1.004	1.004	1.000	0.997	1.001
177	1.861	1.484	1.075	1.082	1.068	1.053	1.024	1.040	1.019	1.050

178	1.321	1.307	1.035	1.060	1.027	1.039	1.004	1.016	1.012	1.023
179	1.263	1.041	1.099	1.067	1.054	1.029	1.016	0.999	0.997	1.033
180	0.977	1.322	1.030	1.008	1.004	1.000	1.004	1.008	1.002	1.017
181	1.274	1.319	1.010	1.063	1.062	1.045	1.025	1.029	1.026	1.054
182	0.790	0.903	1.018	1.009	1.005	0.990	1.000	0.997	1.003	0.988
183	1.865	1.004	1.126	1.060	1.021	1.024	1.015	1.019	1.041	1.033
184	0.204	0.549	0.789	0.842	0.882	0.896	0.930	0.940	0.918	0.840
185	1.162	1.000	0.997	0.948	0.994	1.001	0.994	0.975	0.972	0.977
186	2.277	1.504	1.112	1.109	1.061	1.070	1.041	1.048	1.044	1.095
187	2.218	1.576	1.169	1.112	1.094	1.082	1.058	1.029	1.024	1.094
188	0.688	0.782	0.890	0.953	0.955	0.977	0.974	0.984	0.974	0.948
189	3.051	2.427	1.274	1.203	1.150	1.097	1.081	1.072	1.094	1.230
190	1.716	1.342	1.084	1.055	1.075	1.039	1.022	1.028	1.039	1.104
191	1.120	1.010	0.998	1.024	0.995	0.998	0.994	1.000	0.973	0.980
192	1.937	1.168	1.028	1.060	1.032	1.014	1.026	1.042	1.026	1.044
193	1.342	1.289	0.977	1.020	1.006	1.021	1.010	1.020	1.003	1.034
194	1.441	1.436	1.084	1.089	1.057	1.015	1.032	1.023	1.023	1.097
195	2.950	1.905	1.281	1.145	1.158	1.062	1.075	1.073	1.079	1.153
196	1.334	1.377	1.027	1.031	1.040	1.013	1.022	1.012	1.031	1.083
197	1.253	0.891	0.992	0.980	0.972	0.986	0.998	0.991	0.999	1.030
198	0.603	0.724	0.918	0.923	0.944	0.960	0.974	0.985	0.976	0.938
199	0.511	0.808	0.914	0.936	0.950	0.984	0.973	0.985	0.967	0.943
200	1.141	1.101	1.087	1.018	1.027	1.031	1.008	1.021	1.032	1.049

References

Chow, L., Friedrich, C., and Helwig, D., "Quantification of the Natural Hedge Characteristics of Combination Life or Annuity Products Linked to Long-Term Care Insurance" Society of Actuaries, 2012.

Loomis, R., Facing Reality: *LTCi Is Broken – Let's Fix It*, Society of Actuaries - Long-Term Care News, Issue 34, p. 9-11, June 2013

Loomis, R., Hauser, P., Klugman, S. and Pahl, A., *Future Actuarial Tools: Understanding and Refining Simulation Models*, presented at the 12th Annual Intercompany Long-Term Care Insurance Conference, March 20th 2012, [ONLINE] Available at: http://iltciconf.org/2012/index_htm_files/30-FutureActuarialTools-2.pdf. [Accessed 11 April 2014].

Loomis, R., Hauser, P., Klugman, S. and Pahl, A., *Future Actuarial Tools: Defining Moderately Adverse Experience with Simulation*, presented in Session 9 of the 12th Annual Intercompany Long-Term Care Insurance Conference, March 19th 2012, [ONLINE] Available at: http://iltciconf.org/2012/index_htm_files/09-FutureActuarialTools.pdf. [Accessed 11 April 2014].

Loomis, R., *Measuring Subsidization in LTC*, Society of Actuaries - Long-Term Care News, Issue 23, p. 7-9, May 2009

Knight, F. H., *Risk, Uncertainty, and Profit,* Houghton Mifflin Company, 1921

Silver, N., *The Signal and the Noise, Why So Many Predictions Fail, But Some Don't,* Penguin Press, 2012

Stroup, W. W., *Generalized Linear Mixed Models*, CRC Press, 2013

Venter, G. and Sahasrabuddhe, R., *A Note on Parameter Risk,* Casualty Actuarial Society E-Forum, Summer 2012

Stochastic Modeling: Theory and Reality from an Actuarial Perspective, International Actuarial Association, 2010