Measuring Value in Reinsurance

Gary G. Venter, FCAS, MAAA

Measuring Value in Reinsurance

Gary G. Venter, Guy Carpenter Instrat

Abstract

Reinsurance produces value by producing stability. This can translate into higher earnings through reduced financing costs, improved access to markets, stronger product pricing and better employee job security. It can lead to a higher earnings multiple through reduction in the market price of possible bankruptcy and fewer misreadings of downwards earnings hits. Measuring these earnings and valuations effects is a developing science and for now we will only go as far as reviewing methods for measuring the stability that produces these benefits, and using these measures to compare alternative reinsurance programs.

Some of the conclusions are: standard deviation and variance can be misleading measures; using any measure with combined ratio can produce distortions in the analysis; the frequency of one reinsurance program producing a better result than another is not a very useful measure.

Measuring stability requires modeling. Making realistic models of insurers is a highly technical exercise with a number of possible pitfalls. Some of the technical details needed for modeling insurance financial risk are discussed.

Measuring Value in Reinsurance

Introduction

When asked to do a cost/benefit analysis of their reinsurance purchases, insurers' analysts sometimes do the following calculation. First they add up all the ceded premiums for the past several years. They call that the cost. Then they add up all the recoveries and ceding commissions received. That they identify as the benefit. Subtracting cost from benefit gives the net benefit. After this calculation usually follows a lament that the net benefit has been negative. Sometimes one or two treaties have had a positive net benefit, but these are usually canceled or re-priced soon after. Occasionally some treaties return more than they cost over a long period of time, but pay losses several years after the premium has been received, so that premium plus loss investment income exceeds recoveries. The analyst decides that reinsurance has been a losing proposition for the company for some time.

A moment's reflection, though, will reveal that this result was almost a foregone conclusion. Reinsurers are in business to make money, and some have succeeded at it. There are expenses involved. Thus over time total payouts by reinsurers have to be less than the premium they receive plus its related investment income. A given client can beat these odds in the short run, but probability eventually wins out – at least for the vast majority. And the exceptions usually are cedants with such poor results that they envy the rest.

So what's wrong with the analysis? Is reinsurance just a bad deal that should be shut down as soon as possible, or are there some other benefits that this calculation misses? In the broadest terms, we would assert that the benefit of reinsurance is in gaining stability of results. This includes protection of surplus against erosion from adverse fluctuations, improvement in the predictability of earnings growth, and the provision of customers with assurance of recovery of their insured losses. There is a cost to gaining this benefit, but the cost is not simply ceded premiums. Premiums less recoveries (including expense recoveries) would be a better measure of the cost to the cedant for gaining stability. In fact this cost measure is what the naïve analyst got as the net benefit.

So to sum up so far, the value of reinsurance is in the stability gained. The cost is the net of premiums and recoveries. For prospective analysis, the expected value of premiums less recoveries would be the comparable cost measure. The next step is quantifying this cost/benefit trade-off.

Quantifying Stability and Its Value

There are a few measures of stability that can be used – standard deviation and related quantities, percentiles or value at risk, and excess aggregates. Measures can be applied to surplus, earnings, or related accounts. Some companies prefer to look at more than one measure.

An additional step would be to translate the measure of stability into a measure of earnings gained through the increased stability. This could come in two directions:

1. reduction in capital costs for surplus not needed due to the reinsurance protection and from reduced costs of raising other capital because of improved debt ratings, etc.
2. other increases in earnings and firm value such as stock market valuation, risk elements in employee and management compensation, and improvements in access to customers and in the ability to get adequate premiums, perhaps related to improved claims paying ratings.

The former is perhaps more readily quantified, but the latter elements are also tangible financial benefits of a sound reinsurance program. The cost of the reinsurance program could well be offset by these effects on earnings and valuation. This is difficult to measure, however, and as it arises from the stability gained from the reinsurance program, we will be content here with just measures of that.

Perhaps the best way to illustrate these concepts is through an example. Consider ABCD, a small company or department that writes $33M of excess property and liability coverages. This consists of $14M in casualty coverages, with an expected loss ratio of 78%, and $19M in property covers with an expected loss ratio of 63%. Total expected losses are $22.9M and there is an expense ratio of 23% for a total expected combined ratio of 92%.

ABCD currently purchases a reinsurance program in several layers providing 4M x 1M of casualty cover for $4.41M, 17M x 3M of per-risk property cover for $2.36M, and a cat program covering 95% of 24M x 1M for $1.53M, with one reinstatement at 100%. This totals $8.3M in ceded premiums prior to any reinstatement premiums. The cat program is designed to cover at least up to the 1-in-250 year cat event.

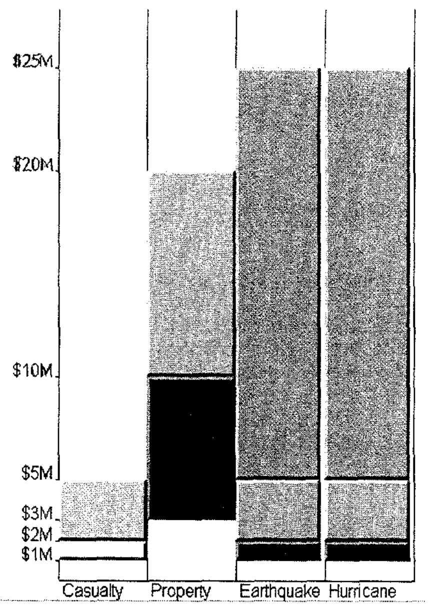

ABCD have been offered as an alternative a stop-loss program of 20M x 30M for a premium of $1.98M. Is this a better option? Cost/benefit analysis addresses such issues.

Doing a cost/benefit analysis requires first establishing cost and benefit measures. A reasonable cost measure, as discussed above, is the net excess of ceded premiums over expected recoveries. This can be estimated using a simulation study of financial results before and after reinsurance. Some of the technical issues of doing such a study are discussed later. Metarisk, Instrat's model building platform, was used to build a model of ABCD to simulate the financial results.

The results of the analysis, based on a simulation of 25,000 possible realizations of the underwriting results, is that average net recoveries after reinstatement premiums are $5.08M for the current program and $0.98M for the alternative. The ratio of these recoveries to ceded premium is 61% for the current program and 49% for the alternative, which makes the current program sound more favorable. The proposed cost measure, however, is not ceded loss ratio but premium less expected recoveries. This is $3.2M for the current program and $1.0M for the alternative. This difference is significant for ABCD, as its expected pre-tax income prior to ceded reinsurance is just $6M ($2.5M underwriting + $3.5M investment).

The stop-loss program thus has a lower ceded loss ratio but costs less than the current program. Can it possibly provide enough protection? For this an analysis of the probability of adverse deviations from expected results is needed.

The table at the right from the simulation run shows some summary statistics for net premiums minus losses (gross less ceded) prior to any expenses or investment income. The difference in the means is the relative net cost differential between two programs.

Statistic	<BARE>	Current	Stop Loss
Mean	$10.1M	$6.9M	$9.12M
Standard Deviation	$8.09M	$5M	$6.24M
Skewness	-0.8619	-0.4235	0.0945
Safety Level, Percent	99.0%	99.0%	99.0%
Safety Level, Value	$24.3M	$17M	$22.3M
Smallest Simulated	-$49.3M	-$23.2M	-$32.2M
Largest Simulated	$30.9M	$22.7M	$29M
Number of Simulations	25,000	25,000	25,000

The safety level shown in this case is the best result at the 1-in-100 level. It shows that the stop loss program is over $5M better in this very good year. However, the stop-loss has a higher standard deviation, and its worst result in 25,000 years is $9M more adverse than the current program. Thus under some measures the current program provides more protection than the stop-loss.

Most companies do not manage to a 25,000-year event, so a comparison is needed at more realistic probability levels. The graph below shows the simulated probability densities for the net premium less net losses. It shows that the current program does produce a compression of results, but much of this compression comes by cutting off the profitability of the good years.

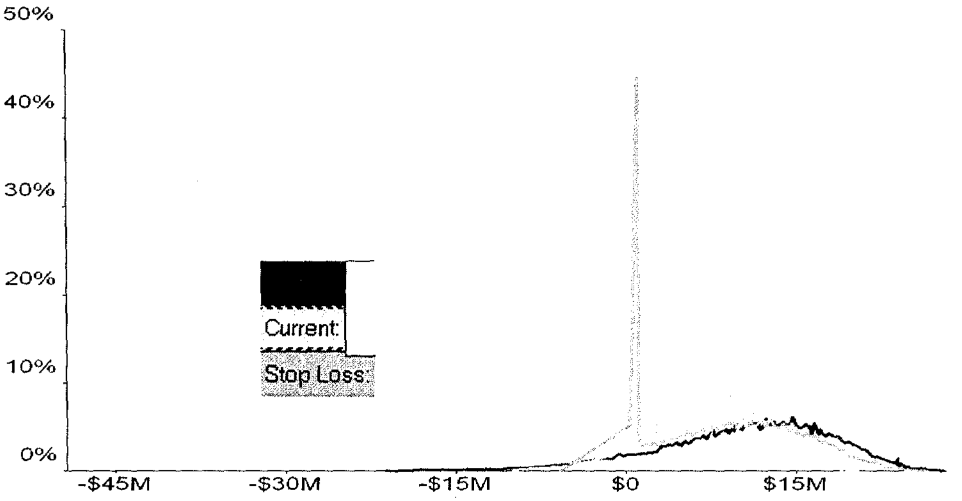

This is also a problem with using standard deviation as a measure of volatility: Standard deviation measures upward and downward deviations, and can be reduced by eliminating the favorable deviations. Measures that capture only unfavorable deviations are more useful, and will be discussed below.

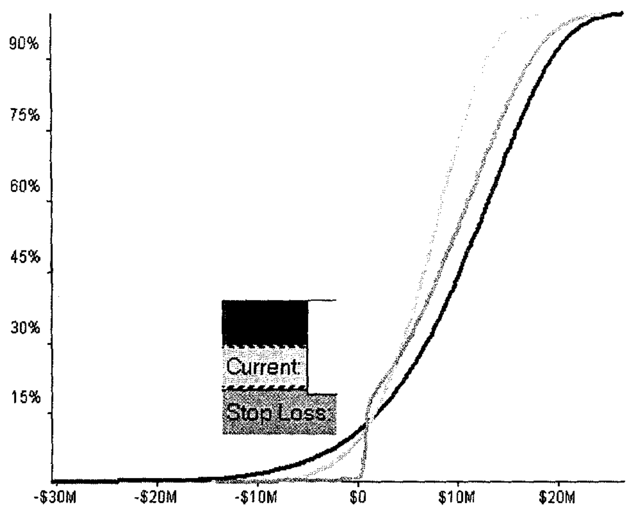

Also apparent in the graph is the concentration of events at the retention of the stop-loss program, and the similarity of the stop-loss and the gross or bare positions in good years.

The cumulative probability distributions (here truncated at the 1-in-500 levels good and bad) give another perspective on the relative performance of the alternative programs. The upper right part shows that the stop loss is indeed more profitable in the good years. But in the 1-in-10 to 1-in-4 range, the current program provides more protection. For the years beyond 1-in-10, the stop loss gives a considerably more favorable result.

These distributions are shown in table form below. The current program better protects the worst case event, but by the 0.25% level (worst case in 400 trials) the stop loss is better. From the 12% to 26% levels, the current program is better, by as much as $1,100,000. But in the worse years the stop-loss could be over

Probability	BARE	Current	Stop Loss
0.00%	-49,263,333	-23,198,963	-32,243,333
0.25%	-25,817,548	-12,416,243	-9,439,234
0.50%	-21,827,529	-10,377,108	-6,311,695
0.75%	-17,837,510	-8,337,973	-3,184,156
1.00%	-13,847,491	-6,298,838	-56,618
1.25%	-12,641,527	-5,703,459	237,924
1.50%	-11,677,654	-5,290,176	286,117
1.75%	-10,713,781	-4,876,893	334,311
2.00%	-9,749,908	-4,463,610	382,505
4.00%	-5,892,701	-2,551,287	575,365
6.00%	-3,602,653	-1,315,561	689,867
8.00%	-2,008,347	-409,204	769,583
10.00%	-686,845	284,986	835,658
12.00%	416,042	951,819	890,802
14.00%	1,448,699	1,464,523	942,435
16.00%	2,415,661	1,919,933	990,783
18.00%	3,226,822	2,388,329	1,251,605
20.00%	3,905,868	2,802,539	1,925,868
22.00%	4,554,807	3,190,684	2,574,807
24.00%	5,209,039	3,549,185	3,229,039
25.00%	5,513,974	3,713,920	3,533,974
26.00%	5,832,081	3,880,394	3,852,081
28.00%	6,371,517	4,205,322	4,391,517
30.00%	6,891,421	4,514,526	4,911,421
32.00%	7,401,904	4,827,688	5,421,904
34.00%	7,856,716	5,146,708	5,876,716
36.00%	8,321,687	5,428,461	6,341,687
38.00%	8,761,854	5,694,960	6,781,854
40.00%	9,208,534	5,962,559	7,228,534
42.00%	9,639,097	6,244,632	7,659,097
44.00%	10,021,333	6,495,969	8,041,333
46.00%	10,439,457	6,780,995	8,459,457
48.00%	10,823,625	7,026,301	8,843,625
50.00%	11,191,515	7,269,232	9,211,515

$6,000,000 better than current, and the median result is almost $2,000,000 better. As the stop-loss is less costly and usually provides a better result, sometimes dramatically so, it would have to be considered a more useful program for ABCD.

A more careful use of vocabulary is actually appropriate here. Even though we would use the above table to say that the stop-loss is $6,350,000 better at the 1-in-100 level, the 99th percentile loss event is unlikely to be the same event for the two programs. Thus the difference between the programs in the 1-in-100 year gross loss event could be more or less than $6,350,000, as could the 99th percentile of the distribution of the difference between the programs. What the table actually allows us to calculate is the difference in the 99th percentiles of the net result under the two programs (or in this example the 1st percentile, since we are looking at earnings.)

This would, however, seem to be the most meaningful comparison. In the end the company is going to select a single program, and it will end up with the probability distribution produced by that program. So the decision to be made is which probability distribution it wants. This is measured by comparing the ending probability distributions of the various programs, not by looking at the distribution of differences between programs. There might be psychological benefits to thinking you got a program that was better more often, but if that program does not produce a better final distribution of net results, that psychological benefit will not translate into a better financial position for the company.

The above shows the general features of a cost/benefit comparison of alternative reinsurance programs. The cost is the expected income foregone by buying the program, and the benefit is the protection against adverse deviation. One way to specifically quantify the value of the protection is to look at the capital that would be absorbed by a loss at a selected adverse level – say the 1-in-100 level. In this example the gross loss at this level is $13.85M compared to $0.06M for the stop loss, which is a difference of $13.8M in capital needed. If capital costs 15%, $2.1M would be needed to raise this much, compared to the $1M net cost of the reinsurance. The comparison is similar if pre-tax income is used as the basis rather than premium less losses. An even more dramatic savings would be shown if the capital requirement were set so that only 25% or some other limited part of surplus would be eroded by the worst year in 100.

Other Comparisons

Once financial risk can be simulated, a variety of methods are available to compare reinsurance programs. Different analysts and decision-makers will find different ones more intuitive. Some of these are illustrated below, using the data from the ABCD example.

The next graph, known as the box view, shows probability in ranges. The area of each box is proportional to the probability of being in the range from the bottom to the top of the box. The middle box shows the interquartile range, i.e., from 25% to 75%. The two boxes on either side show the range from 1-in-4 to 1-in-20.

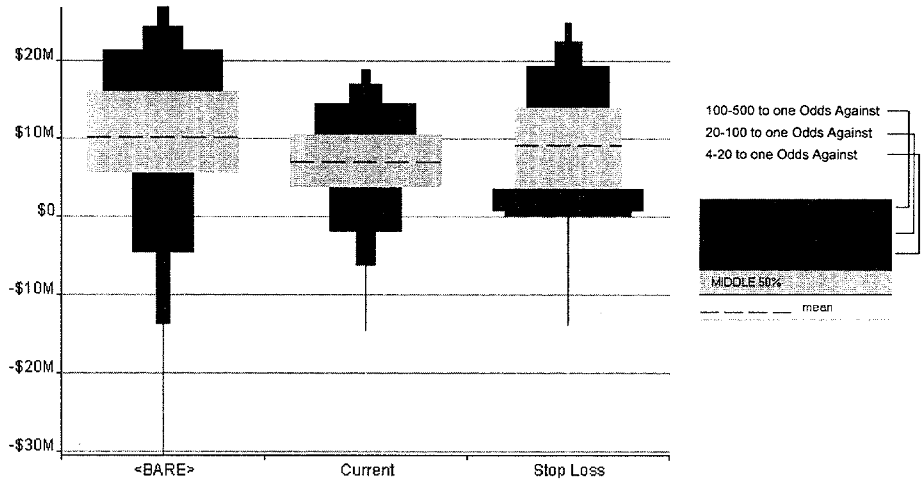

Thus the outside of the middle three boxes is the range from 5% to 95%. The next range is from 1% to 99%, and the outer boxes get to the 1-in-500 levels favorable and unfavorable.

The current program can be seen at a glance to be most compressed, but achieves this by sacrificing profitability in the good years. The stop-loss shows more protection in the 1-in-20 and 1-in-100 years, but is about the same as current at 1-in-500.

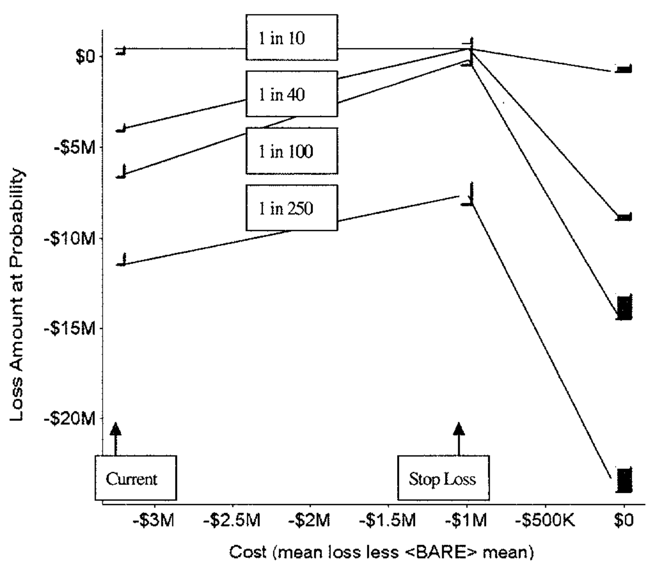

This is a cost/benefit diagram at selected probability levels. Each point shows the cost of a program vs. its loss amount (net premium less net loss) at a given probability level. To be efficient at a selected probability, a more expensive program has to have a lower loss level at that probability. In this example, the current program is not efficient at any of the levels shown, although it is at a few other levels, as discussed above. The choice of programs becomes more difficult when programs of different costs are all efficient – i.e., the more expensive programs provided more benefit at most probability levels.

Other financial measures can also be compared. The table below shows the probability distribution for pre-tax income net of each reinsurance structure. The comparison and the decision process is very similar to that for premium less loss, but the monetary values include expenses and investment income. For ABCD this shows a 20% probability of a loss with no reinsurance, 28% with the current program, and 26% for the stop-loss. Besides giving a reinsurance comparison, these figures give ABCD management perspective on their prospects of overall profitability.

Pre-tax Income

Probability	BARE	Current	Stop Loss
0.00%	-$55,178,595	-$28,306,230	-$37,630,975
0.25%	-$31,005,991	-$17,189,245	-$14,119,948
0.50%	-$26,892,281	-$15,086,897	-$10,895,456
0.75%	-$22,778,572	-$12,984,549	-$7,670,964
1.00%	-$18,664,862	-$10,882,200	-$4,446,471
1.25%	-$17,421,513	-$10,268,365	-$4,142,799
1.50%	-$16,427,760	-$9,842,270	-$4,093,112
1.75%	-$15,434,007	-$9,416,175	-$4,043,424
2.00%	-$14,440,254	-$8,990,080	-$3,993,736
4.00%	-$10,463,474	-$7,018,476	-$3,794,897
6.00%	-$8,102,434	-$5,744,442	-$3,676,845
8.00%	-$6,458,705	-$4,809,988	-$3,594,659
10.00%	-$5,096,235	-$4,094,278	-$3,526,536
12.00%	-$3,959,159	-$3,406,773	-$3,469,682
14.00%	-$2,894,490	-$2,878,175	-$3,416,448
16.00%	-$1,897,552	-$2,408,648	-$3,366,601
18.00%	-$1,061,245	-$1,925,731	-$3,097,694
20.00%	-$361,149	-$1,498,681	-$2,402,529
22.00%	$307,908	-$1,098,503	-$1,733,472
24.00%	$982,421	-$728,889	-$1,058,959
25.00%	$1,296,808	-$559,048	-$744,572
26.00%	$1,624,777	-$387,412	-$416,603
28.00%	$2,180,935	-$52,412	$139,555
30.00%	$2,716,957	$266,377	$675,577
32.00%	$3,243,264	$589,248	$1,201,884
34.00%	$3,712,176	$918,157	$1,670,796
36.00%	$4,191,560	$1,208,645	$2,150,180
38.00%	$4,645,373	$1,483,405	$2,603,993
40.00%	$5,105,900	$1,759,300	$3,064,520
42.00%	$5,549,810	$2,050,117	$3,508,430
44.00%	$5,943,896	$2,309,246	$3,902,516
46.00%	$6,374,982	$2,603,107	$4,333,602
48.00%	$6,771,059	$2,856,018	$4,729,679
50.00%	$7,150,354	$3,106,480	$5,108,974

Financial ratios, on the other hand, may give considerably different comparisons of net results. The combined ratio, for example, combines premium, loss, and expense in a fairly different way than does net underwriting income. Underwriting income subtracts direct losses and expenses and ceded premium from direct premium, and adds in loss and expense recoveries. The combined ratio subtracts loss and expense recoveries from direct loss and expense, and divides by direct less ceded premium. This can give a misleading result, especially if there are minimal ceded expenses, as part of the ratio is direct expense divided by net premium.

The graph below illustrates this for ABCD's reinsurance alternatives.

Combined Ratio

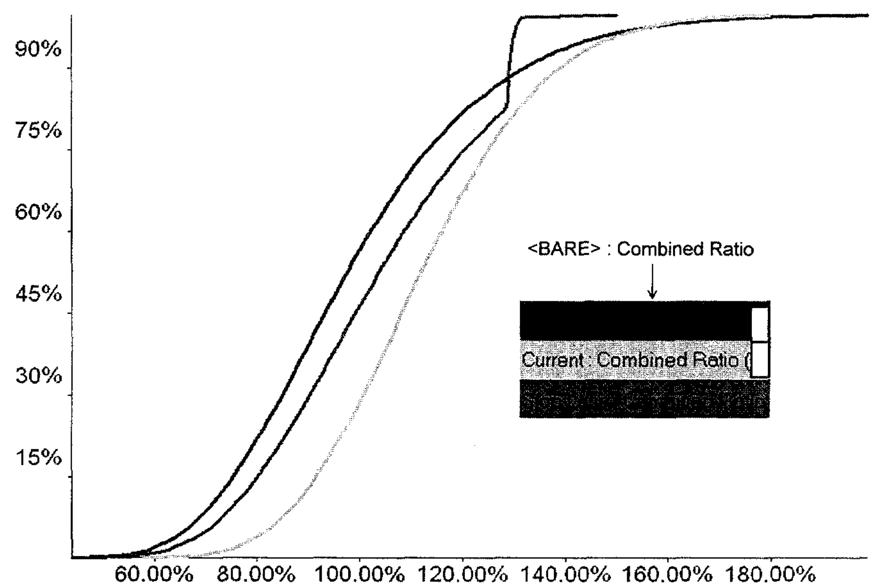

Here the current program shows up as not better than the stop-loss at any probability level, and barely ever better than no reinsurance, even though in many adverse cases it provides considerable income benefit over the direct position and is sometimes better than the stop-loss. This distortion is due to this program's relatively high ceded cost impacting the expense ratio.

Technical Background for Modeling

This kind of modeling of insurer financial risk needs to have access to advanced statistical methodology. Some of these elements built into the Metarisk modeling platform are discussed in the following.

Information Risk

Insurers expect random fluctuation of number of claims and cost of claims to affect their results for an accounting period. Insurers attempt to quantify this risk by trying to get a handle on the frequency and severity distributions they face. Large companies should display more stability than small companies. However, an initially surprising finding is that large companies in many lines of business, while they may be somewhat more stable than the smaller companies, have less stability than you would think from scaling up the frequency and severity distributions. This is similar to a classic result of Charles Hewitt.[1] He found that large insureds have more variability than would be expected from treating them as independent combinations of small insureds.

This can be reconciled by considering information risk, which is the risk that the frequency and severity distributions might be different than they are thought to be. Information risk has two key components: estimation risk, which comes from limits on data and imperfection of models in matching the data; and projection risk, which is the possibility that the underlying parameters might change between estimation and realization. Projection risk in particular does not go down with the size of the company, as there are industry and economic forces that affect the results of all companies.

Estimation risk can be quantified within the model estimation process. The technical key is that when using maximum likelihood estimation, the covariance matrix of the parameters can be derived directly as the inverse of the matrix of partial derivatives of the log-likelihood function. What this means in practice is that estimation risk can be quantified.

Instead of getting a severity distribution, however, when you do this you get a joint distribution of possible parameters of severity distributions. If you are generating scenarios from a simulation model, first the parameters of a severity distribution have to be generated from the distribution of possible parameters. Then the losses for the scenario are generated from the parameters selected.

[1] *Loss Ratio Distribution – A Model*, C Hewitt PCAS LIV, 1967

This is illustrated at the right. Each point in the plane represents a pair of parameters that together define a Pareto distribution. The surface represents the probabilities of those points, and so shows the relative probability of each pair.

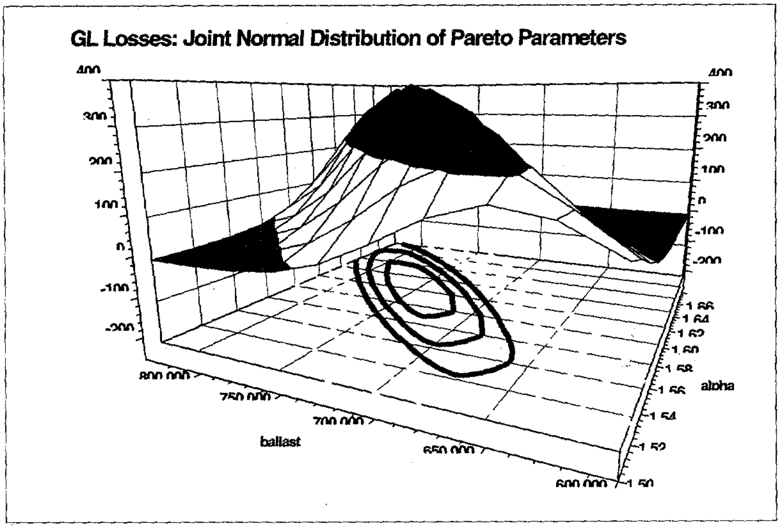

Projection risk includes any changes that could occur in the process that generates claims. Changes in the economy could affect driving, for example, either towards more or less, or even toward faster or slower. Inflation could affect claims costs, as could jury attitudes and new theories of risk.

Although not as statistically established as estimation risk, it is still possible to get some handle on projection risk. One method is to look at fluctuations in parameters over time, and extrapolate that. A related approach is to look at total fluctuation of losses over history, and see how much of that can be explained by the other types of risk. The remainder can be considered projection risk.

The graph at the right illustrates the first approach. Each vertical line represents a point in time. The graph shows the parameter value at that point and its associated uncertainty. Here the parameter is the scale parameter of a Pareto distribution fit to the losses at each point in time. A linear projection of the parameter is shown, with the last four periods being the projection interval. The projection risk is quantified using regression theory, which projects greater uncertainty for the later projected periods.

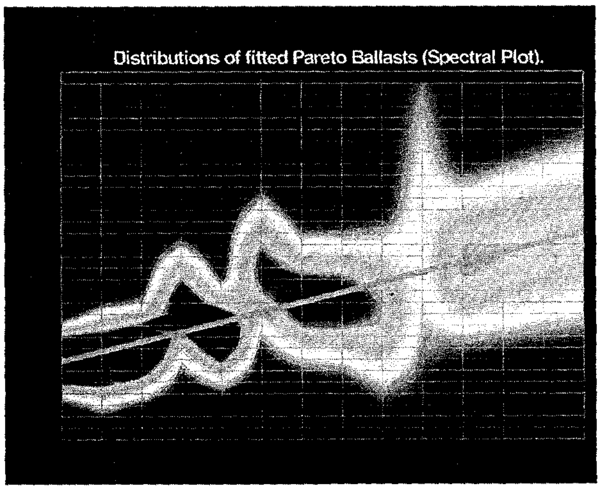

Correlation

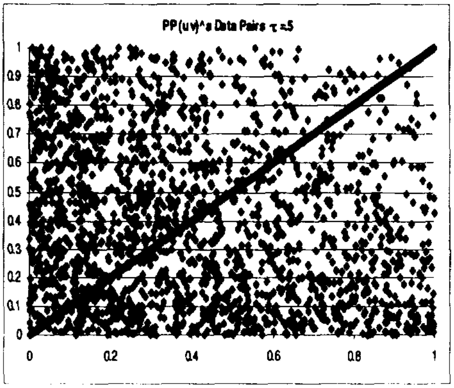

In many cases it is not reasonable to assume that losses from different business centers are uncorrelated. For instance, different profit centers could be affected by the same events, or losses could be affected in the same way by economic conditions. However, some of the standard statistical methods used to handle correlation are not always applicable to insurer losses. For example, in some cases the correlation could arise only from large events, with the usual run-of-the-mill claims not correlated.

Instrat has done a fair amount of research in correlation of losses in the case where it is important to quantify not only the degree of correlation but also where in the distribution the correlation takes place. For instance, Rodney Kreps of Instrat's Seattle center has pioneered a methodology to combine independent and perfectly correlated scenarios to fine-tune exactly where the correlation occurs[2]. The diagram above shows a sample generated from one of the joint distributions based on this methodology. Here the correlated and independent pieces clearly stand out, with less independence with larger losses.

This research also includes copula methods. A copula $C(u,v)$ is a function that expresses a joint distribution function $F(x,y)$ in terms of $F_X(x)$ and $F_Y(y)$, the individual (or marginal) distribution functions for X and Y, so $F(x,y) = C(F_X(x),F_Y(y))$. By selecting among the many available copulas, a good deal of control can be exercised over where the correlation takes place. For instance, a copula with a great deal of concentration in the right tail is discussed in Venter (2001)[3]. This Heavy Right Tail copula is graphed below. The strong correlation in the joint right tail is evident.

The same paper also describes methods for identifying characteristics of copulas by calculating certain functions of the copula. Looking at these functions for the data and comparing to those for several copulas, the copula that is most applicable to the data can be determined.

---

[2] *A Partially Comonotonic Algorithm For Loss Generation*, R. Kreps, ASTIN Colloquium, 2000

[3] *Copula Tails*, G. Venter, ASTIN Colloquium 2001

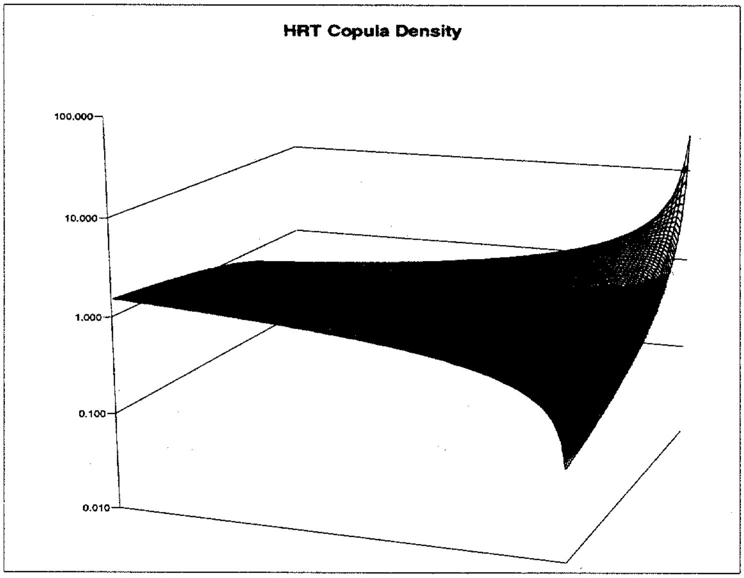

For instance, tail concentration functions L(z) and R(z) can be defined for the left and right tail.. These are just the total probability in the square from the point (z,z) to either (0,0) or (1,1), respectively. These functions are graphed below for two popular copulas, the Frank and the Gumbel.

The Gumbel copula has greater concentration in the tails, especially the right tail. The HRT copula (not

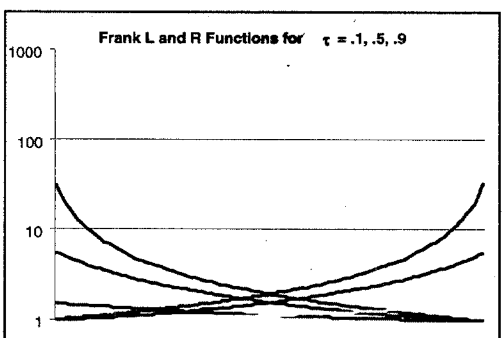

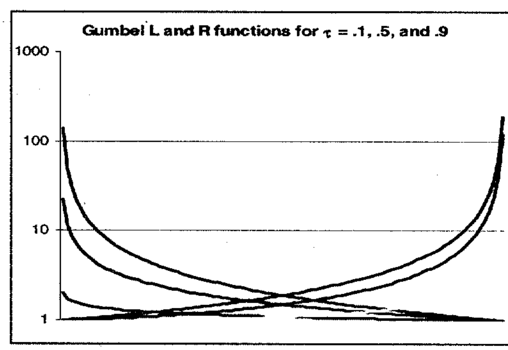

shown) is similar to the Gumbel in the right tail, but more like the Frank in the left tail.

Another function is the cumulative tau, which shows how Kendall's tau changes for different parts of the distribution. At z this is defined as the value of tau restricted to the square from (0,0) to (z,z). This differs in shape among copulas, and so provides a way to compare copulas among themselves and to data. It is illustrated below for the normal and HRT copulas.

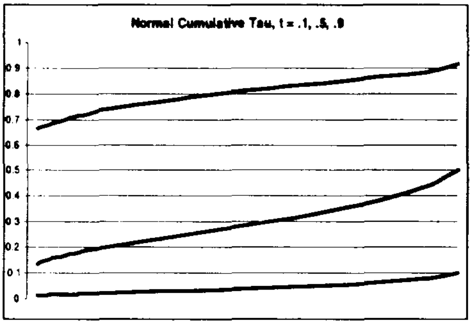

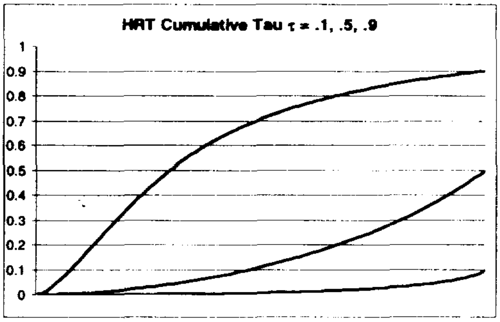

Assets and Liabilities

Reinsurance and other hedging methods are used to control the risks in the asset and liability portfolios. In addition, the asset and liability risk provides a context for the current year underwriting risk that traditional reinsurance programs address. Thus even companies that do not want to review hedging strategy for assets and liabilities may want to have these risks incorporated in a model of reinsurance impact on the total risk of the company. Instrat's models for asset and liability risk are discussed below.

Assets

In the US market, most asset classes are correlated with interest rate movements, so these make a good starting point for asset modeling. Instrat's model starts off with short-term interest rates, and builds up the yield curve from there. A three-factor model of short-term rates is the starting point. This model, from the academic literature[4], has been shown to have the minimal level of complexity needed to capture the dynamics of short-term rate movements identified to date. The factors used are the short-term rate itself, the volatility of that rate, and the locus of mean reversion.

---

[4] *Stochastic Volatility and Mean Drift in the Short Rate Diffusion: Sources of Steepness, Level and Curvature in the Yield Curve*, T. Andersen and J. Lund, Northwestern University Department of Finance Working Paper 214

Bonds of varying maturities are priced by discounting their cash flows along a number of stochastically generated paths, then taking the average. The paths are generated by the three-factor model with an adjustment for market-price-of-risk. This is the only way to generate arbitrage-free yield curve evolution. However it is not enough that the yield curves be arbitrage-free. They also need to have appropriate distributional properties so that they adequately represent the probabilities of various outcomes.

A yield curve is a vector of interest rates for various terms, and so the probability distribution of yield curves can be represented as a multi-variate joint distribution. This can get unwieldy, however, so simpler ways of representing the multi-variate probabilities are needed. One method we have used[5] is to look at yield spreads – i.e., the difference in yield between bonds of different terms, as a function of the short-term rate.

Inverted curves are often associated with high short-term rates, and in fact longer yield spreads tend to decline with increasing short-term rates. This is not a strict deterministic rule, however, and there is historically some degree of deviation from this norm. Both the trend and the deviation from it can be used to assess the reasonability of a set of yield curve scenarios. Short-range scenarios would of course be expected to match current conditions, but longer-range forecasts should start to resemble historical distributions.

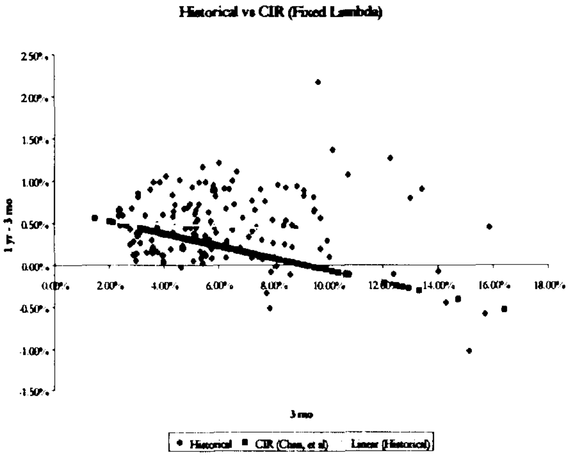

The graphs above and below show the three month to one year historical yield spreads as a function of the short-term rate, a linear fit to that, and the distribution of the conditional spreads from two models – the Cox, Ingersoll Ross[6] model and the Andersen and Lund model.

The CIR model in the graph above can be seen to have no deviation from the declining relationship, which appears to decrease too rapidly in any case. On the other hand, the Andersen-Lund model shown below seems to capture the historical spread distribution fairly reasonably.

---

[5] *Asset Modeling – Empirical Tests of Yield Curve Generators*, G. Venter, General Insurance Convention and ASTIN Colloquium Papers vol. 2, 1998, p. 175

[6] *A Theory of the Term Structure of Interest Rates*, J. Cox, J. Ingersoll and S. Ross, Econometrica 53, 1985

For longer spreads, however, we found that an adjustment to the market-price-of-risk process for the Andersen-Lund model improved the match to historical yield distributions.

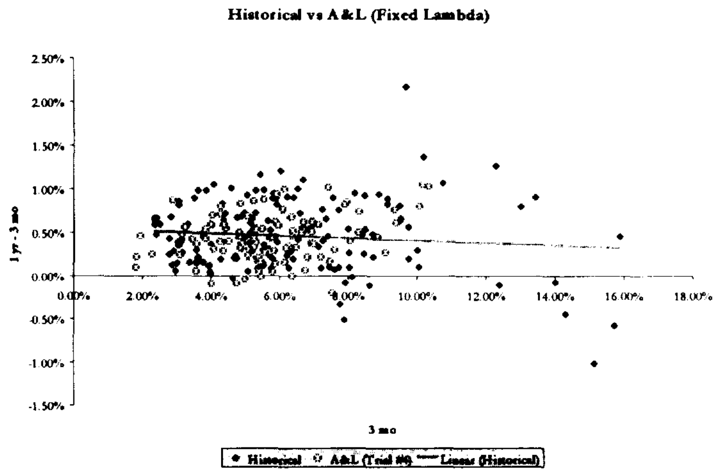

The graph below shows the three year to ten year spread as a function of the short-term rate historically and as derived from the Andersen-Lund model with fixed market-price-of-risk. Although there is some dispersion in the model-generated scenarios, they are tighter around their trend than is the historical data. Also the trend is steeper than historical.

By adding a stochastic process for the market-price-of-risk, the generated spreads were distributed more like historical, as is shown in the next graph.

This effort in modeling interest rates is needed to get a set of realistic interest rate scenarios. It is not enough to generate a lot of scenarios. Unrealistic scenarios should not be generated, but the full range of reasonably possible scenarios should be produced.

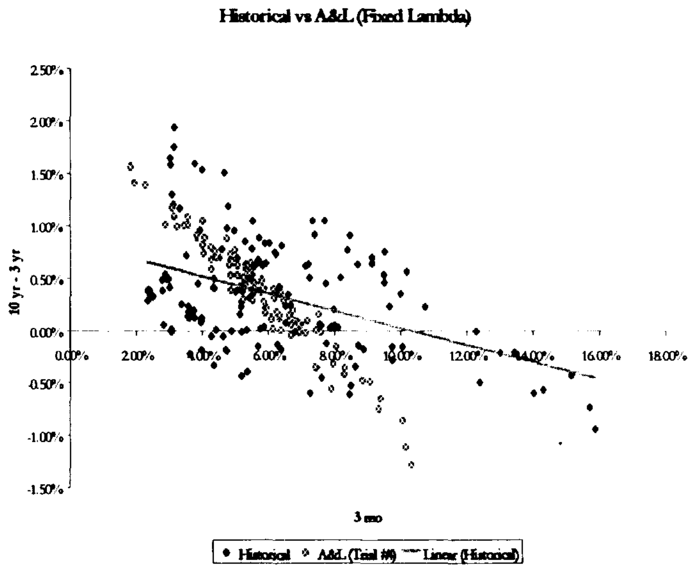

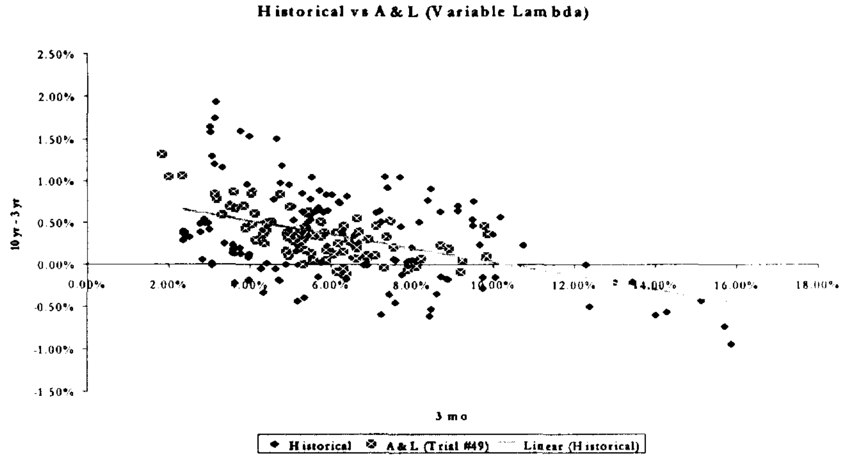

This requires firstly an arbitrage-free generator. Even though some published yield curve histories might appear to have small arbitrage opportunities, these are hard to find and take advantage of in practice. There is in fact not a precise yield curve at any point in time. There are bid-ask spreads, and trades taking place at slightly different times that get into the published curves. Also, doing large-volume trades will change the yields and eliminate apparent arbitrage possibilities. Having arbitrage possibilities in the generated scenarios, however, will make strategies that take advantage of these opportunities appear preferred. Even if strategies are restricted to types that could not take advantage of many such opportunities, the optimality of some of the forbidden strategies could distort the modeled performance of the allowed strategies.

Just having an arbitrage-free generator, however, is not enough to guarantee that all and only the reasonably likely scenarios will be produced. For instance, the CIR generator produces yield curves having the property that any yield-spread from any curve produced is exactly predicted by a linear function of the short-term rate. Actual curves will arise that vary from this relationship, so too limited a universe of curves is produced. The conditional distribution of yield spreads given the short-term rate is one statistical measure of the variety of yield curves produced, and the Anderson-Lund model with stochastic market-price-of-risk is able to produce a realistic set of curves by this measure. However there could be other measures of the statistical distribution of yield curves that would require even more advanced models.

Liabilities

Modeling asset and current underwriting risk is still incomplete. Insurers may have a considerable liability risk. Sometimes reinsurance is purchased to handle this risk. Even if not, quantifying liability risk helps give context to the risk transfer that reinsurance provides.

The principal liabilities of most insurance companies are loss and expense reserves. There are numerous stochastic models in the actuarial literature for such reserves. We have come up with a classification method for these models[7], with 64 classes, which are based on answering six binary questions about the model:

Six Questions

Q1. Development depends on emerged?

Q2. Purely multiplicative development?

Q3. Independent of diagonal effects?

Q4. Stable parameters - e.g. factors?

Q5. Normally distributed disturbances?

Q6. Variance of disturbances constant?

If all answers are yes, the assumptions of the age-to-age factor method are satisfied, so that model can be used to generate runoff scenarios. If some of the answers are no, other models may be indicated:

Alternatives

Development ~ ultimate (Q1=No) – BF new incrmnt = f*ultimate + e

Additive not multiplicative (Q2=No) - Cape Cod new incrmnt = a + e

Diagonals inflation sensitive (Q3=No) – Separation new incrmnt = f*[prev emerged]*(1+i) + e

Factors change over time (Q4=No) – smoothing Use just last 3 diagonals, exponential smoothing . . .

Lognormal disturbances (Q5=No) - take logs

Disturbance proportional (Q6=No) - weighted MLE

Combined models - e.g.: new incrmnt = a + f*ultimate*(1+i) + e

Reduced parameters - e.g.: new incrmnt = $(1+i)$*ultimate/$(1+j)^k$ + e

Finding the model that is most consistent with a company's runoff history is key to generating development scenarios that are likely to represent the company's loss emergence process. Regression procedures can be used to answer the six questions above, which can then narrow down the modeling choices.

[7] *Liability Modeling – Empirical Tests of Loss Emergence Generators*, G. Venter, General Insurance Convention and ASTIN Colloquium Papers vol. 2, 1998, p.421

Conclusion

Cost/benefit analysis provides a useful methodology for insurers to quantify the value in their reinsurance transactions, and to compare among alternative structures. A good cost measure is the net decrease in underwriting earnings expected from the program. Benefit measures would ideally show the increased earnings from reduced financing costs, better claims paying ratings, etc., but these are difficult to quantify. A reasonable substitute is to measure the increased stability that arises from the reinsurance program. The earnings improvements follow from the increased stability in any case. Several risk measures based on various financial accounts give similar comparisons. However, using the combined ratio can give a distorted picture of the effects of reinsurance on earnings. Also the variance and standard deviation can give misleading results, as they can be lowered by eliminating the chance for favorable deviations. Looking at the distribution of differences in programs is also not as useful as looking at the differences in the distributions.

To be captured appropriately, several of the risk components themselves require fairly sophisticated financial models, including the measurement of information risk and the correlation among lines. For a more complete quantification of the net risk of the insurer, asset and liability modeling is required as well. Asset models need to be arbitrage-free, but also need to generate scenarios with reasonable statistical properties; doing both is difficult and a number of models fail in one or the other aspect. Liability models should identify the process that generates the liabilities, which can usually be done within a regression framework when working with development triangles. Having a model of the underwriting, asset, and liability risks can at least quantify some of the major stochastic elements driving insurers' variability of financial results, and allow testing of the degree of variability left after various reinsurance proposals.