Measuring Value in Reinsurance

Excerpted from

Enterprise Risk Analysis

for Property & Liability Insurance Companies

Gary G. Venter

Spencer M. Gluck

Paul J. Brehm

Section 2.5: Measuring Value in Reinsurance[16]

By Gary G. Venter, FCAS, MAAA, Spencer M. Gluck, FCAS, MAAA and
Paul J. Brehm, FCAS, MAAA

Introduction

On one hand, investment portfolio managers have a wide variety of tools available to manage the risks in their portfolios, largely facilitated by fairly efficient and liquid markets. On the other hand, insurance portfolios, though risky, are fairly illiquid and have limited means to mitigate the assumed hazard risk. Reinsurance is, of course, the most readily available tool to transfer assumed hazard risk. Enterprise risk analysis can greatly facilitate effective use of reinsurance in hedging an insurance portfolio.

When asked to complete a cost-benefit analysis of their reinsurance purchases, cedents sometimes use the following calculation: First, they add up all the ceded premiums for the past several years, which they call the cost. Then, they add up all the recoveries and ceding commissions received, which they identify as the benefit. Subtracting cost from benefit gives the net benefit. Completion of this calculation is usually followed by 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 repriced soon after. Occasionally, some treaties return more than they cost over a long period but pay losses several years after the premium has been received, so that premium plus loss investment income exceeds recoveries. The cedent 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 cedents with such poor results that they envy the rest.

16 Venter [5] previously published a paper on this topic dealing primarily with measuring value in reinsurance based on cost versus stability achieved. His paper is reproduced in large part here, augmented with new sections.

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? We offer the following three related paradigms for measuring the value in a reinsurance structure.

1. Reinsurance provides stability. In the simplest terms, the benefit of reinsurance is that it provides stability of results. "Stability" includes protection of surplus against erosion from adverse fluctuations, improved predictability of earnings growth and customers' assured 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 cedent for gaining stability. In fact, this cost measure is what the naïve analyst receives as the net benefit.

2. Reinsurance frees up capital. Going one step further, the incremental stability gained by purchasing reinsurance frees up risk capital that would otherwise be required of the ceding company. That is, reinsurance can be a substitute for required capital. The value of reinsurance, then, could be gauged by the amount of income foregone to purchase the cover versus the amount of capital freed up. In accounting terms, both numbers would be negative for an insurer, so the ratio of the two will be positive. This ratio can be thought of as the ROE cost of the reinsurance purchase. If this ROE cost is less than the firm's target returns, the purchase is a good financial decision.

3. Reinsurance adds to the value of the firm. In the end, the activities undertaken by the firm in the course of business are meant to add value to the firm. If the company is publicly traded, we are speaking directly of adding market value. It would perhaps be ideal if we could measure the value in a reinsurance purchase as the incremental market value added to the company.

The next three sections cover the three paradigms, above, in turn.

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 – to name a few. Measures can be applied to surplus, earnings or related accounts. Some companies prefer to look at more than one measure.

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

FIGURE 2.5.1: ABCD COVERAGE CHART

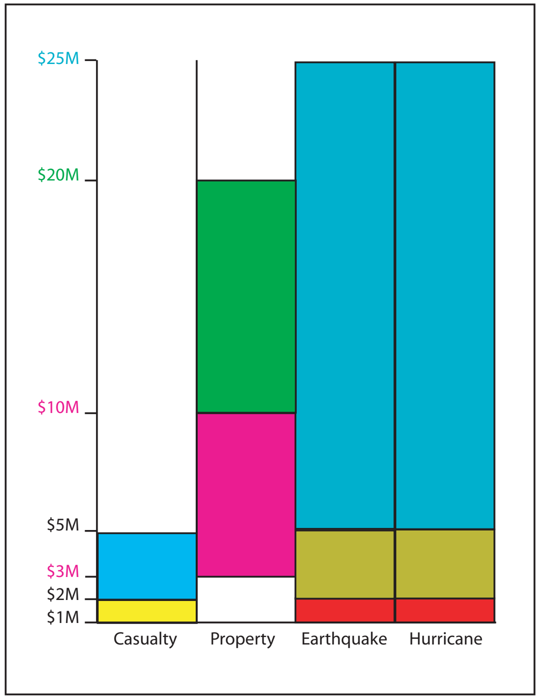

As shown in Figure 2.5.1, ABCD currently purchases a reinsurance program in several layers, providing 4 million x 1 million of casualty cover for $4.41 million, 17 million x 3 million of per-risk property cover for $2.36 million and a catastrophe program covering 95 percent of 24 million x 1 million for $1.53 million, with one reinstatement at 100 percent. This totals $8.3 million in ceded premiums prior to any reinstatement premiums. The catastrophe program is designed to cover at least up to the 1-in-250-year catastrophe event.

ABCD has been offered, as an alternative, a stop-loss program of 20 million x 30 million for a premium of $1.98 million. 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. Results 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 below.

Based on a simulation of 25,000 possible realizations of the underwriting results, average net recoveries after reinstatement premiums are $5.08 million for the current program and $0.98 million for the alternative. The ratio of these recoveries to ceded premium is 61 percent for the current program and 49 percent 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 $3.2 million for the current program and $1 million for the alternative. This difference is significant for ABCD, as its expected pre-tax income prior to ceded reinsurance is just $6 million ($2.5 million underwriting + $3.5 million 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? An analysis of the probability of adverse deviations from expected results is needed.

FIGURE: 2.5.2: ABCD SIMULATION OUTPUT

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

Figure 2.5.2, which is the ABCD simulation output, 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.

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 more than $5 million better in this very good year. However, the stop-loss has a higher standard deviation, and its worst result in 25,000 years is $9 million 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. Figure 2.5.3 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.

FIGURE: 2.5.3: ABCD COVERAGE OPTION PROBABILITY DENSITIES

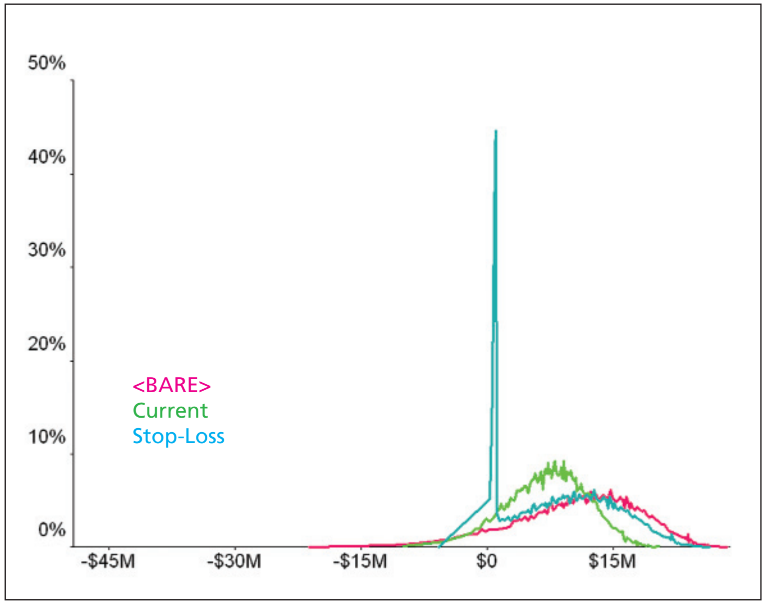

This is also a problem with using standard deviation as a measure of volatility: Standard deviation measures upward and downward deviations and can be reduced by eliminating the favorable deviations. Measures that capture only unfavorable deviations are more useful and will be discussed below.

Also apparent in Figure 2.5.3 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 in Figure 2.5.4 (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.

FIGURE 2.5.4: ABCD COVERAGE OPTION CUMULATIVE DISTRIBUTION

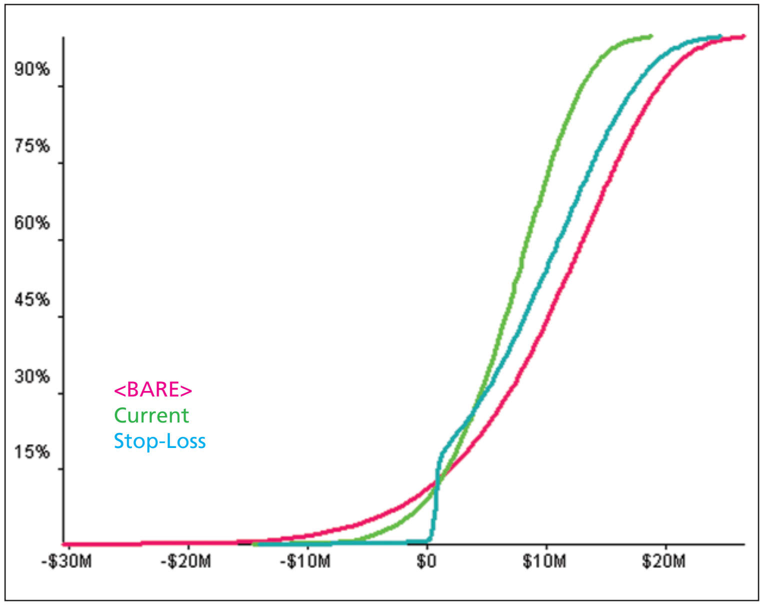

These distributions are shown in Figure 2.5.5. The current program better protects the worst-case event; but by the 0.25 percent level (worst case in 400 trials), the stop-loss is better. From the 12 percent to 26 percent levels, the current program is better, by as much as $1,100,000. But in the worst years, the stop-loss could be more than $6,000,000 better than the current program, 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 Figure 2.5.5 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).

FIGURE 2.5.5: ABCD CDF OUTPUT

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

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. Thus, the necessary decision is which probability distribution it desires. The decision is facilitated by comparing the ending probability distributions of the various programs, not by looking at the distribution of differences between programs. The company may gain a psychological benefit from thinking that its program is better more often than others. However, 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 figure 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.

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 methods more intuitive. Some of these are illustrated using the data from the ABCD example. Figure 2.5.6, known as the box or space needle 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, that is, from 25 percent to 75 percent. 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 percent to 95 percent. The next range is from 1 percent to 99 percent, and the outer boxes get to the 1-in-500 levels: favorable and unfavorable.

FIGURE 2.5.6: ABCD COVERAGE OPTION BOX OR SPACE NEEDLE VIEW

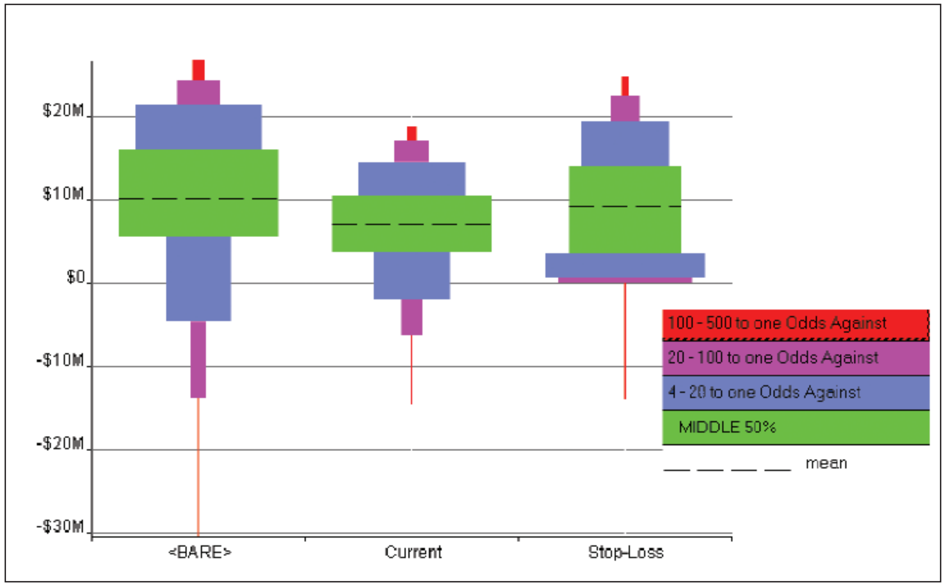

The current program can be seen at a glance to be most compressed, but achieves this compression by sacrificing profitability in the good years. The stop-loss program shows more protection in the 1-in-20 and 1-in-100 years, but is about the same as current at 1-in-500.

Figure 2.5.7 is a cost-benefit diagram at selected probability levels. Each point shows the cost of a program versus 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 – that is, the more expensive programs provided more benefit at most probability levels.

FIGURE 2.5.7: ABCD COVERAGE OPTION COST BENEFIT

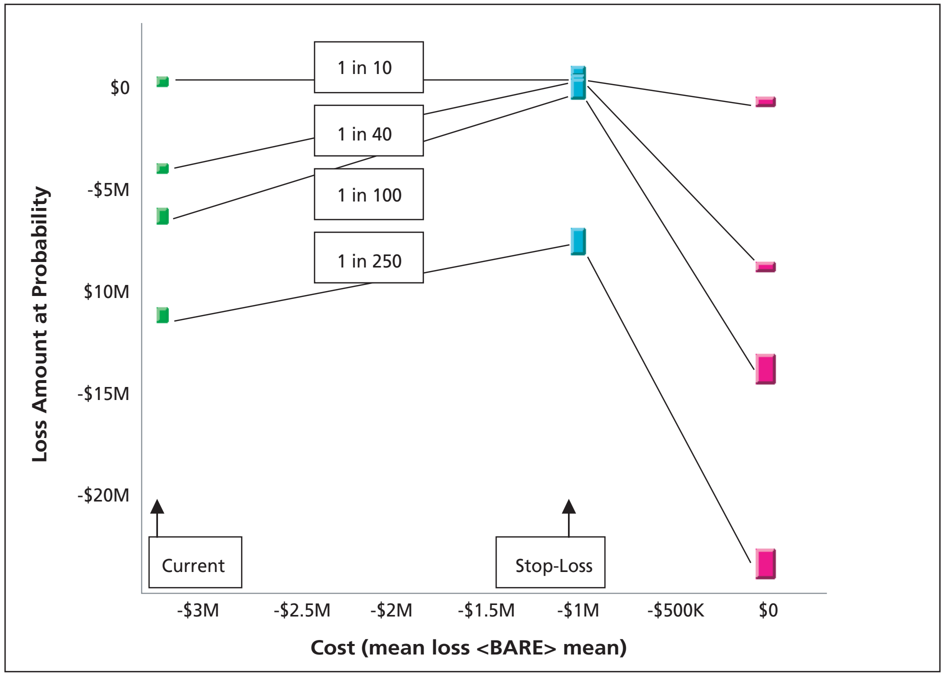

Other financial measures can also be compared. Figure 2.5.8 shows the probability distribution for pre-tax income net of each reinsurance structure. The comparison and the decision processes are very similar to those for premium less loss, but the monetary values include expenses and investment income. For ABCD, this shows a 20 percent probability of a loss with no reinsurance, 28 percent with the current program and 26 percent for the stop-loss. Besides giving a reinsurance comparison, these figures give ABCD management perspective on their prospects of overall profitability.

FIGURE 2.5.8: ABCD PRE-TAX NET 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. Figure 2.5.9 illustrates this for ABCD's reinsurance alternatives.

FIGURE 2.5.9: ABCD COVERAGE OPTION COST BENEFIT

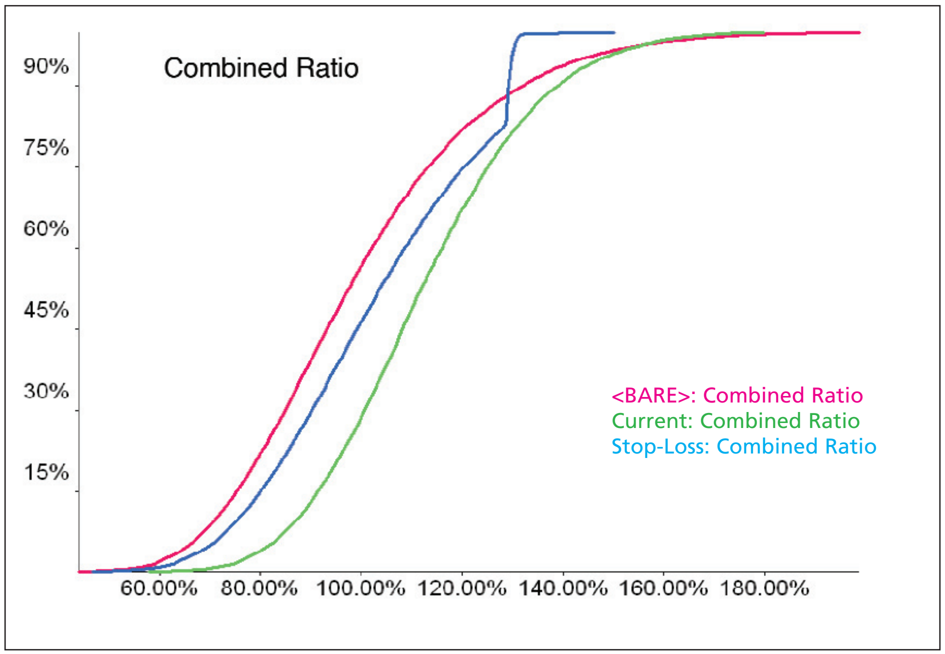

Here, the current program shows up as not better than the stop-loss at any probability level, and rarely better than the option of no reinsurance, even though in many adverse cases it provides considerable income benefit over the direct position and is sometimes better than the stop-loss. This distortion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

Efficient frontier charts, like the one shown in Figure 2.5.10, that graph a scattering of alternative scenarios according to their risk (by some measure) and return (by some measure) are a common way of assimilating and summarizing the vast amount of data and information from the modeling of reinsurance structures. Efficient frontier analyses can explicitly show the tradeoff between cost and benefit. Such analyses can help identify programs that are clearly inefficient and point the analyst in the direction of a structure that could be more advantageous. The efficient frontier can also illustrate that competing programs may not be inherently better or worse but rather alternative points on an optimal curve that can be distinguished only based on company preferences or budgetary constraints.

FIGURE 2.5.10: ABCD COVERAGE OPTION

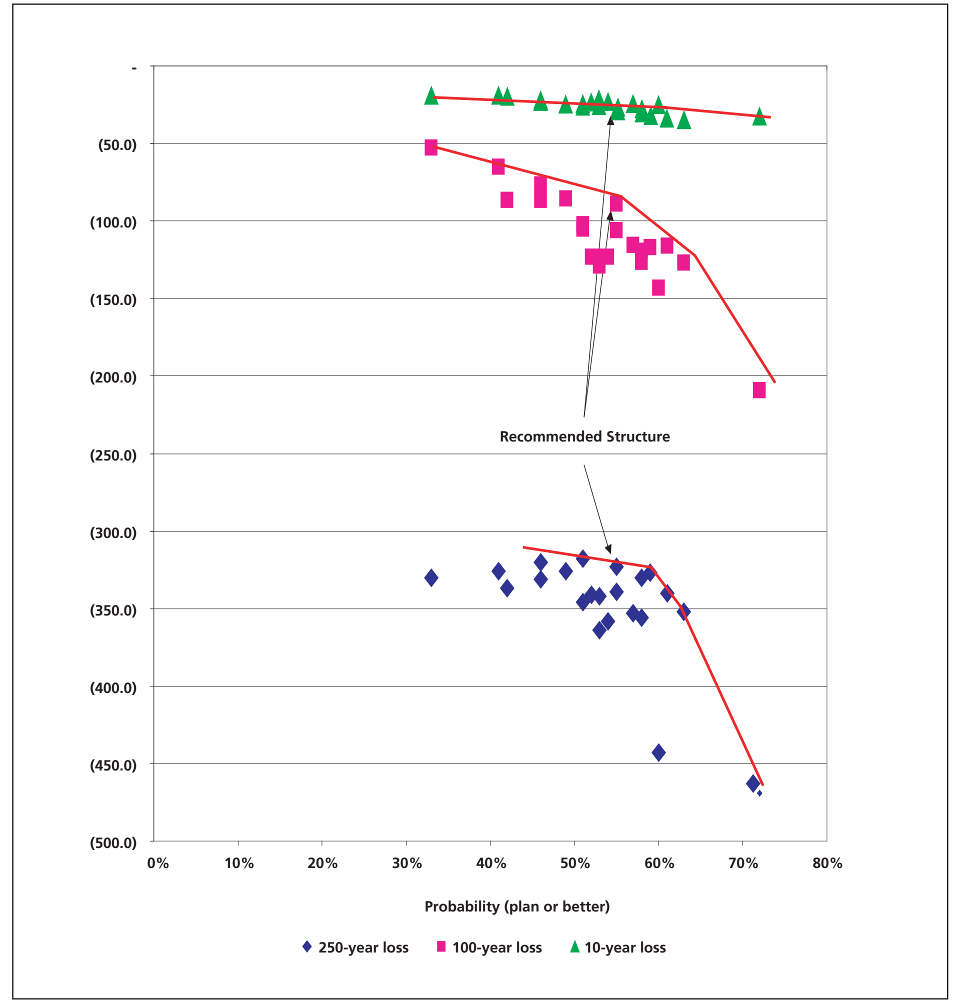

Efficient frontiers are often reviewed at several probability levels, as shown above. Here, the 1-in-10, 1-in-100 and 1-in-250 levels are shown for a number of alternative programs. In this case, the comparison of risk and return is between loss at the probability level and the probability of the combined ratio being at plan or better (scale at top). Each possible program is at the same point on the horizontal scale but is shown three times to represent its loss at each probability level on the vertical scale.

The more expensive programs make meeting plan less likely, due to the cost of the reinsurance. Thus, even though they have less possibility for adverse loss, they are not recommended. In fact, the most expensive programs are not efficient at the 1-in-250 level, as they are low-attaching programs that also run out of limit too soon.

There are a variety of ways stability can be measured and portrayed, in efficient frontier graphs or otherwise, depending on management criteria. Emphasis on rating agency or regulatory requirements and constraints (see also Section 2.3) or market expectations suggests some common, real-world options that might include the probability:

- Of surplus dropping below 2x RBC,
- Of surplus dropping below a BCAR score supporting a target rating,
- That an expected loss in a 10-year return period exceeds a threshold level of surplus or
- Of an x percent drop in quarterly earnings per share.

The paradigm of measuring value in reinsurance in the tradeoff between the net costs of the reinsurance cover versus the stability gained in the purchase is superior to a more simple dollars-out-against-dollars-in analysis. In the end, however, significant judgment is still required to evaluate the efficacy of the cost-benefit tradeoff. The next sections take this analysis further by essentially trying to quantify the value of stability.

Reinsurance as Capital

Insurers hold capital or surplus in part as a contingency fund to pay claims and expenses even in those scenarios where actual losses and expenses exceed available revenues. It stands to reason that the more volatile a company's results, the more surplus ought to be held. Stated in the context of the previous section, the more stable the company's results, the less surplus is required. The relationship between stability and surplus (or equity or required capital) forms the basis for the second paradigm for measuring value in reinsurance.

The second paradigm adds the step of translating the measure of stability into a measure of required capital or surplus. Since reinsurance is meant to stabilize results, the directional change in required capital should be negative. Since capital – either debt capital or equity capital – carries a cost, the reduction in required capital translates into a reduction in capital costs. In the cost-benefit considerations of the previous section, the reduced cost of capital is the benefit of the reinsurance. The cost of the reinsurance is still the net amount foregone in the reinsurance transaction on a present-value basis.

In this second paradigm, then, an explicit dollar benefit can be compared to an explicit dollar cost. This cost-benefit comparison allows for direct and unambiguous comparisons of reinsurance structures. Presumably, the only role for non-numeric preferences would come in those instances where the net monetary benefits of competing structures were the same.

Alternatively, the cost of reinsurance (a negative value, as it is outflow) can be divided by the change in required capital (also a negative amount, since capital is released) to calculate what amounts to a marginal ROE measure. Reinsurance structures with the better ROEs would be preferred. Furthermore, reinsurance structures with marginal ROEs above the company's cost of capital would be preferable to going bare.

While the numerator in our marginal ROE is straightforward, the denominator is more of a challenge. In general, we divide models of required capital into two classes:

1. "Theoretical models" – those that derive required capital and changes in it based on the calculated risk metrics from an enterprise risk model and
2. "Practical models" – those that derive required capital for the company by concession to the reality of various rating agency formulas (e.g., BCAR, S&P CAR), regulatory requirements (e.g., RBC, ICAR) or actual capital.

The more theoretical models would establish a level of required capital based on a risk metric consistent with management's views towards risk. Some of the usual suspects include V@R, TV@R, XTV@R, WXTV@R and more. Each has its advantages and its disadvantages and is further discussed in Section 2.2. Having selected the appropriate risk measure and threshold value that define the required capital, the methodology is simple. Required capital is calculated as the threshold value of the selected risk metric using the risk distributions produced for the company for each of the competing reinsurance structures. Using the current structure (or perhaps the scenario where the company is bare) as the base case, the marginal changes in required capital can be easily calculated.

Calculation with the "practical" models proceeds in a similar fashion. However, the analyst need not be concerned with risk metrics. A threshold value is still required, such as BCAR = 175 percent, or RBC = 4x authorized control level. The company's current or prospective model score is calculated as the basis. The model score is then recalculated for each reinsurance program under consideration and compared to the basis case. The models in question are all relatively similar – capital factors are applied to premiums, reserves, assets, etc. Capital requirements are reduced by reinsurance because premiums and possibly current or expected reserves are reduced. There is, however, typically a small corresponding increase in required capital due to a factor applied to reinsurance receivables.

While the practical method is, in some sense, easier to implement, it suffers from the disadvantage that it measures capital based on risk-producing proxies (e.g., premiums), rather than explicitly modeling the risks themselves. So, in the previous example of the stop-loss contract versus the current structure, the stop-loss would likely show little effect on rating agency or regulatory required capital, as it would have little impact on premiums and no prospective impact on reserves.

One way to compensate for the above disadvantage is to build the various rating agency and regulatory required capital models into the enterprise risk model. A capital score can then be calculated for each scenario and with each iteration. Rather than rely on the marginal differences as described above, required capital can be set at predetermined levels based on the probability distributions of the regulatory scores. For example, a company may define required capital at that level where there is less than a 10 percent probability that BCAR drops below 130 percent (the level typically associated with an A-).

Consider an example of a company deliberating changing the retention and/or the limits of an excess of loss reinsurance contract. The sample loss distributions of the alternatives are shown in Figure 2.5.11.

FIGURE 2.5.11: RETENTION/LIMIT COVERAGE OPTION CDFs

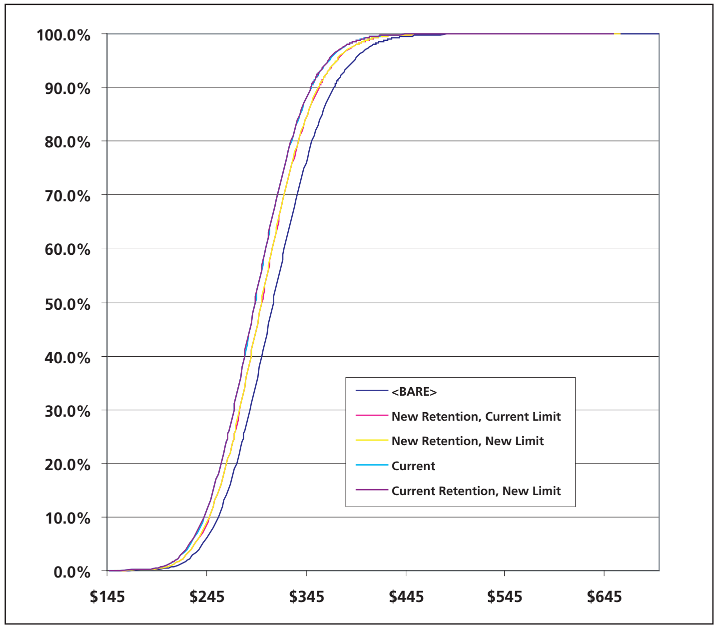

In this example, the cost or benefit was measured as the net present value of the alternative's ceded premiums less the net present value of ceded losses as compared to the current program. Current capital was based on management's view of minimum capital required to retain existing ratings, which coincided with the V@R at the 99.98th percentile. This V@R value was then used as the proxy for required capital for each of the options, and capital was read off of the curves above.

FIGURE 2.5.12: COMPARISON OF ALTERNATIVE REINSURANCE STRUCTURES WITH MARGINAL ROE

	<BARE>	Option 1	Option 2	Current	Option 3
Retention	0	New	New	Current	Current
Limit	Unlimited	Current	New	Current	New
△ NPV Ceded Premium - (less)/more	-39.6	-8.9	-10.1	0.0	-0.2
△ NPV Ceded Loss - (less)/more	-23.3	-9.1	-9.2	0.0	-1.5
△ NPV Net Benefit/(cost)	16.2	-0.2	1.0	0.0	-1.4
Capital Consumed/(released)	51.6	1.5	1.8	0.0	0.3
After Tax Marginal ROE	20%	-8%	35%		-339%

Option 2, above, was deemed superior, as it had the highest marginal ROE. Alternatively, one could have compared the increased capital required (1.8) with its associated cost of capital (say, 10 percent) to the benefit achieved (1.8 x 0.1 = 0.18 vs. 1.0).

The above example is illustrative, but the specific algorithm is not necessarily recommended. The use of higher percentile V@R estimates to gauge required capital can be very volatile. And V@R itself, while intuitive, restricts the definition of risk to a single point on the loss distribution. We generally recommend using a risk metric such as XTV@R at a lower percentile. Figure 2.5.13 shows the XTV@R calculation, by peril, for a variety of thresholds for three different reinsurance options. The capital consideration would be expressed as a multiple of these XTV@Rs, for example, six times the 90th percentile (1,182 for Option A and 1,218 for Option B) or 4.5 times the 95th percentile (1,193 for Option A and 1,220 for Option B).

FIGURE 2.5.13: COMPARISON USING XTV@R RISK METRIC

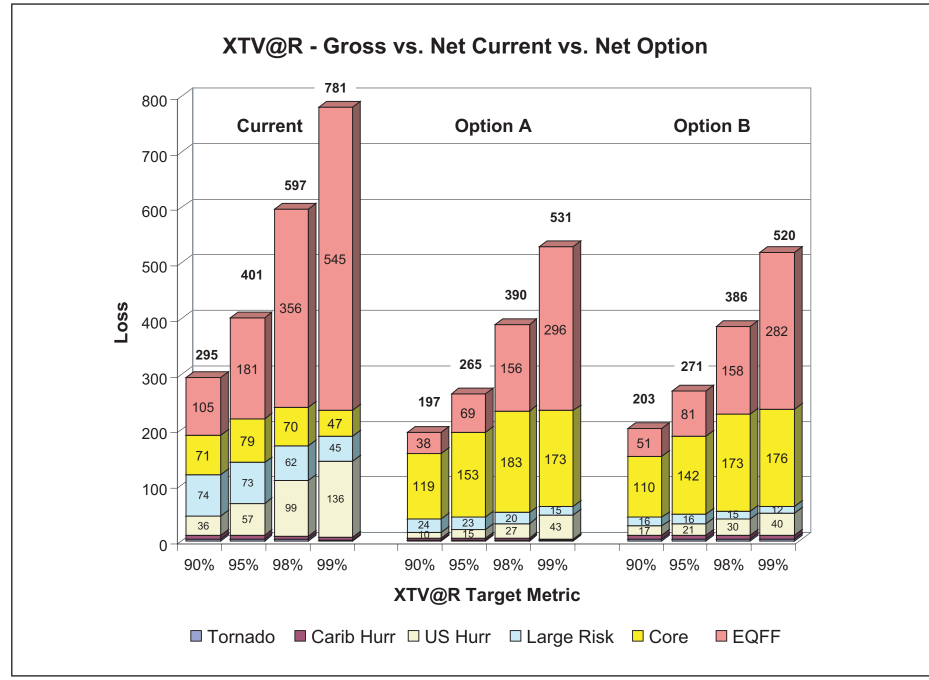

In summary, the second paradigm attempts to judge the value in the reinsurance program based on a comparison of the marginal cost of capital implied for each program to the respective marginal cost of the reinsurance. The company can use the results of the analysis to select the best reinsurance program from among alternatives or choose to retain the risk, perhaps with additional capital.

Reinsurance, Capital and Accumulated Risk

Whether capital requirements are defined by enterprise risk models, rating agency formulas or regulatory ratios, loss reserve risk requires capital. In the prior section, the ROEs are based on capital for a single year, but business that contributes loss reserves absorbs capital for more than one year. For long-tailed business, the difference is highly significant and should be considered when analyzing the value of reinsurance.

Accumulated loss reserves create accumulated risks. The accumulation of risk is exacerbated by aspects of risk that are correlated across accident years. As a convenient mechanism for modeling capital absorbed for many future years, we introduce the notion of *as-if* loss reserves. For an accident (or underwriting) year of new business, the as-if loss reserves are the loss reserves that would exist at the beginning of the accident year, if that business had been written in a steady state (except for trend) in all prior years. The capital absorbed in the current year by the combination of the accident year and the as-if loss reserves is a surrogate for the present value of the capital absorbed by the accident year over time.

The as-if reserves mechanism provides two practical advantages:

1. It can measure the impact of accumulated risk caused by correlated risk factors.
2. The reinsurance being analyzed or considered can be applied to the accident year and the as-if reserves, providing a more valid measure of the impact of the reinsurance on accumulated risk and on capital absorbed over the full life of the accident year.

The results of an analysis of accumulated risk will be highly dependent on the form of the underlying risk model with respect to time-related projection risk. The illustrative examples below incorporate several such features that are further discussed in subsequent chapters; specifically:

- Severity trend (and its associated uncertainty) is modeled as applying through the date of loss payment. All unpaid losses, therefore, continue to be exposed to this trend risk. Loss reserve risk models that incorporate calendar-year trend are discussed in Section 5.2.

- Severity trend risk is modeled according to the AR-1 (first order autocorrelated) process that is introduced in Section 3.2.

The example is intended to be typical of direct middle-market commercial liability insurance written in the United States, with a maximum policy limit of $2 million.

Before considering the effect of reinsurance, let us examine the impact of accumulated risk on the direct (i.e., BARE) results. Figure 2.5.14 compares probability density functions (PDFs) and cumulative distribution functions (CDFs) for the accident year alone versus the accident year plus as-if reserves.

FIGURE 2.5.14: DISTRIBUTIONS OF UNDERWRITING LOSS — CURRENT YEAR AND ACCUMULATED RESULTS

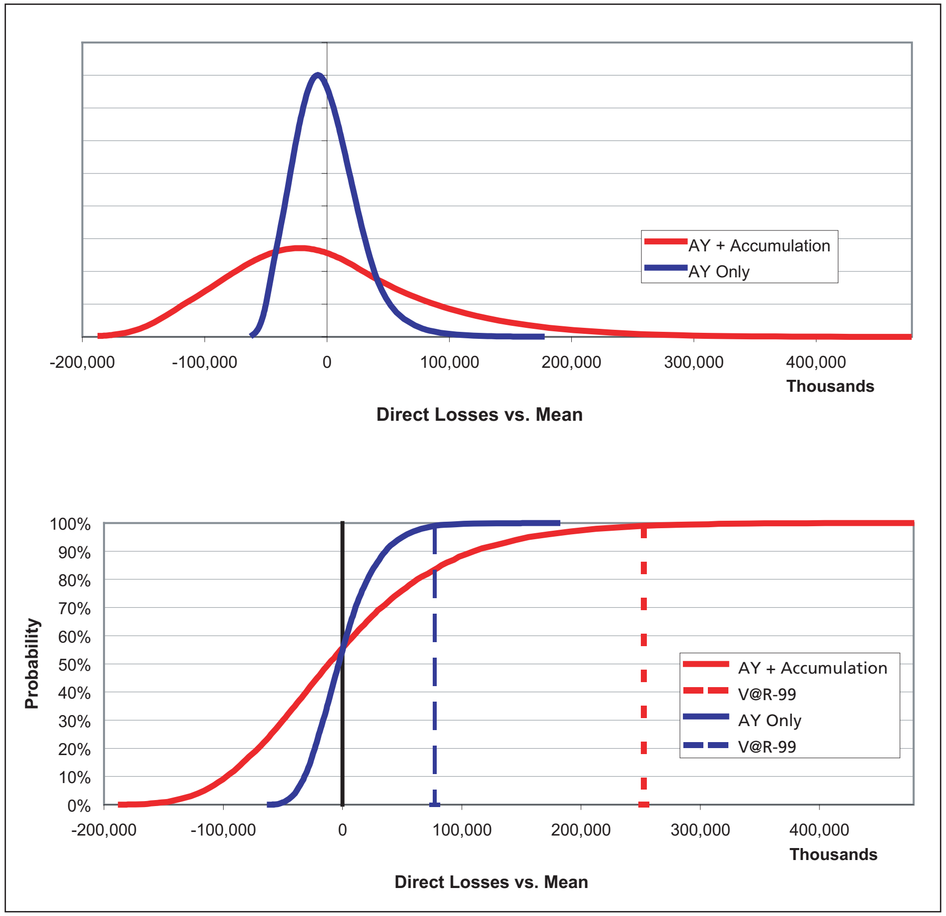

For a given outcome, we define capital consumed as the present value of losses and expenses minus the present value of underwriting funds available. For the accident year, funds available are premiums, and we have assumed breakeven underwriting. For the as-if reserves, funds available are the nominal value of the reserves. Present values are at 4 percent per annum. Figure 2.5.15 compares PDFs and CDFs for capital consumed.

FIGURE 2.5.15: DISTRIBUTIONS OF CAPITAL CONSUMED — CURRENT YEAR AND ACCUMULATED RESULTS

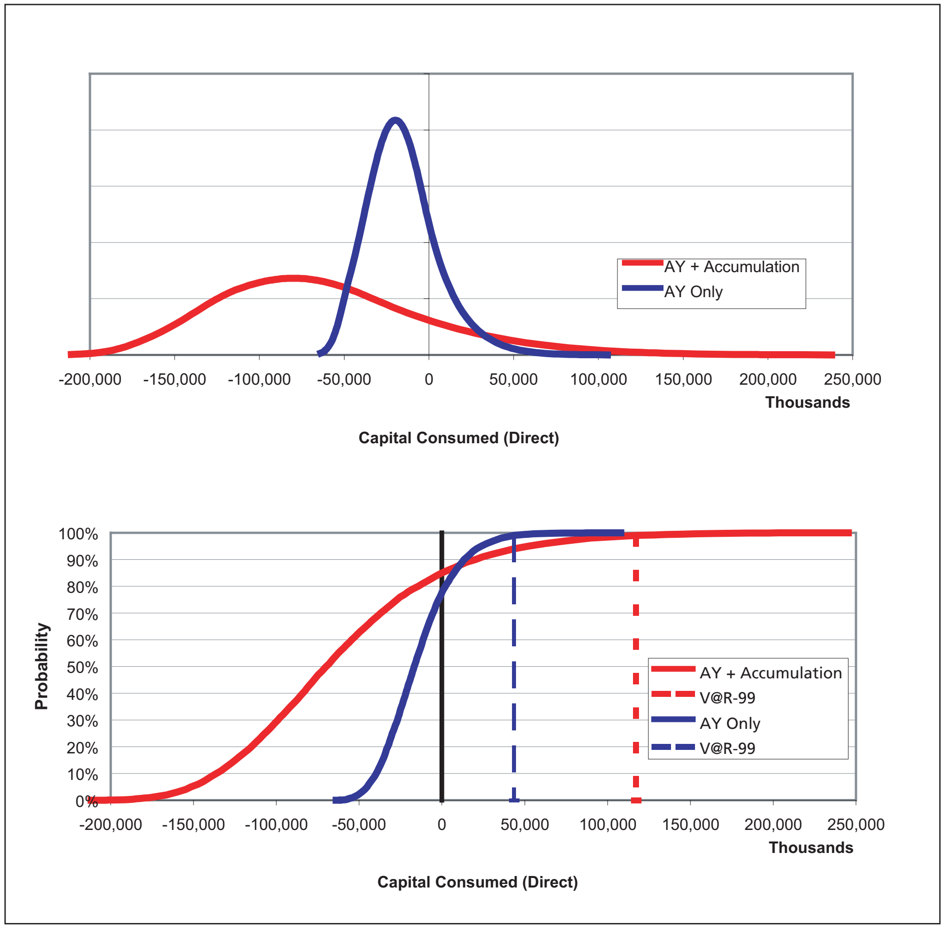

Clearly, the analysis of accumulated risk provides a dramatically different picture of the capital consumed by writing a long-tailed line of business.

Next, we add reinsurance to the picture. We illustrate a simple per-occurrence XOL program covering 1.5M x 0.5M. In order to properly reflect the impact of the trend-risk model, it is essential that the model of the ceded losses reflect not only the XOL process risk but also the XOL payment pattern and the leveraged effect of changes in severity trend on the XOL layer.

FIGURE 2.5.16: TAILS OF CAPITAL DISTRIBUTIONS

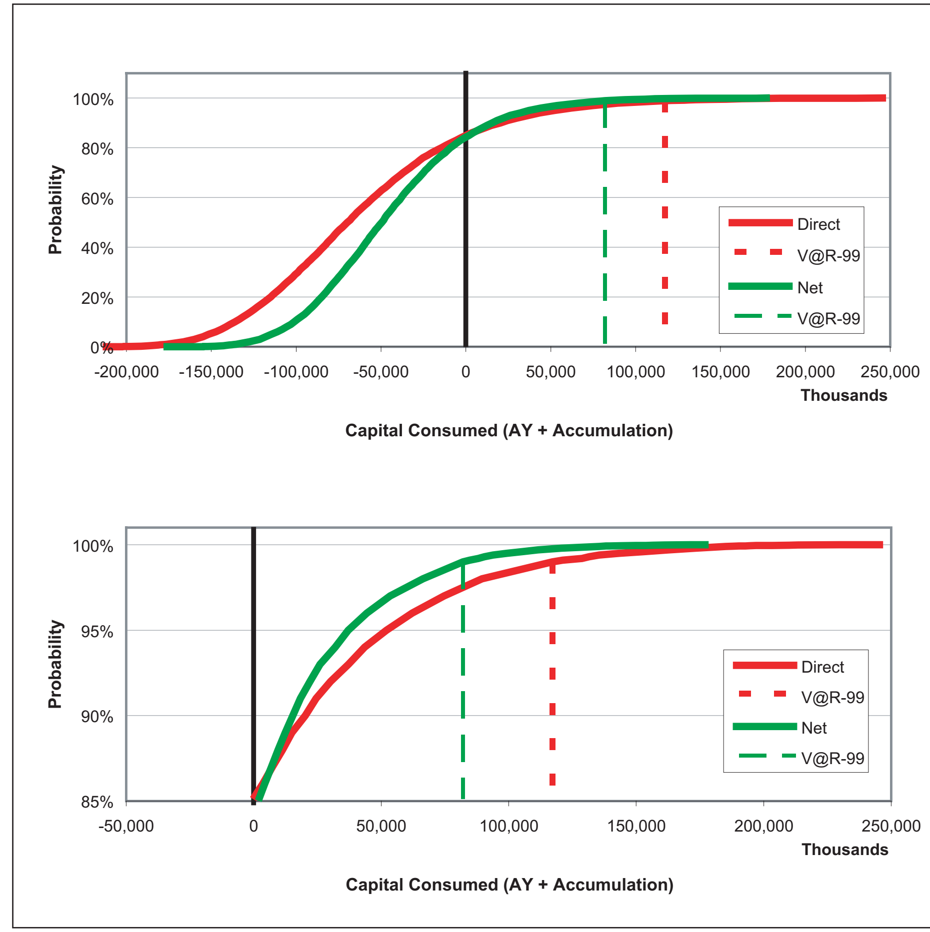

The reinsurance in the example is priced to a 5 percent underwriting loss, with no ceding commission and a 15 percent reinsurer expense ratio. Figure 2.5.16 displays the CDFs of the capital-consumed distributions, direct and net.

FIGURE 2.5.17: STRUCTURE COMPARISON WITH ROEs

		DIRECT	NET
Mean Profit (Discounted Basis)		14,555,893	10,865,516
Capital	V@R-99	117,190,899	82,032,910
	TV@R-99	139,381,010	97,578,009
Spread ROE	V@R-99	12.42%	13.25%
	TV@R-99	10.44%	11.14%

Figure 2.5.17 compares expected profit (cost) – direct, net and ceded – versus required capital according to several tail measures. This type of comparison can measure the value of reinsurance and compare competing reinsurance options, as has been described in the prior section.

Note that the required capital illustrated above is relative to the distribution for a stand-alone line of business, rather than to an allocation of company capital or capital cost. As such, the realistic capital levels are not as deep in the tail of the distribution as they would be for the company as a whole.

Reinsurance and Market Value

In an ideal world of business, every competitive action of a firm would be performed with the intent of increasing the value of the company. Where constraints force a decision between competing activities, the firm would presumably prefer those that increase value the most on a risk-adjusted basis. These statements are certainly true for publicly traded companies, where market value is a ready report card on the company's performance, but they likely hold for mutual companies as well.

The third paradigm for measuring value in reinsurance extends the notion in the previous section – that the stability garnered from reinsurance is a substitute for capital and can therefore be judged accordingly – to relate the concepts of capital consumption and stability to the ultimate value of the firm. The third paradigm is perhaps the holy grail of cost-benefit analysis.

We can all agree that the value of the firm is favorably impacted by the effective use of capital, stability of earnings and steady growth. Several academic studies [2] have chipped away at the relationships between capital, earnings, growth and value. Recent studies have found:

- Insureds demand price discounts of 10 to 20 times the expected cost of the chance of an insurer default (see Phillips, Cummins and Allen [3]).
- A 1 percent decrease in capital gives a 1 percent loss in pricing, and a 1 percent increase in the standard deviation in earnings leads to a 0.33 percent decrease in pricing (see Sommer [4]).
- A ratings upgrade is worth 3 percent in business growth, and a ratings downgrade can produce a 5 to 20 percent drop in business (see Epermanis and Harrington [1]).

The third paradigm is an area of ongoing research, but our research into this class of models is encouraging. Section 2.6 introduces FLAVORED models, which seek to measure the market value of risk management.

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 the net present value of earnings expected from the program. Conversely, we find using combined ratios can give a distorted picture of the effects of reinsurance on earnings.

The simplest measure of benefit – our first paradigm – is the increased stability gained from the reinsurance transaction. Measures of stability, 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. Efficient frontier analysis is often a useful tool.

Benefit measures would ideally show the increased value of the firm (third paradigm) from the increased earnings from reduced financing costs, better claims paying ratings, etc. A reasonable substitute is to relate the increased stability that arises from the reinsurance program to capital requirements. Value in reinsurance can be measured versus a cost of capital or in terms of marginal ROEs. Several risk measures based on various financial accounts give similar comparisons.

References

[1] Epermanis, Karen, and Scott E. Harrington. "Market Discipline and Reaction to Ratings Changes in U.S. Property-Liability Insurance Markets." University of South Carolina (March 2001).

[2] Froot, Kenneth, Gary G. Venter, and John Major. "Capital and Value of Risk Transfer." AFIR Colloquium, Boston (November 2004). Available online at http://www.actuaries.org/AFIR/colloquia/Boston/Froot_Venter_Major.pdf.

[3] Phillips, Richard D., J. David Cummins, and Franklin Allen. "Financial Pricing of Insurance in the Multiple-Line Insurance Company." *The Journal of Risk and Insurance,* vol. 65, no. 4 (1998) pp. 597-636.

[4] Sommer, David W. "The Impact of Firm Risk on Property-Liability Insurance Prices." *The Journal of Risk and Insurance,* vol. 63, no. 3 (1996) pp. 501-636.

About the Authors

Paul J. Brehm

Paul Brehm is a Managing Director of the Instrat® unit at Guy Carpenter & Company, LLC. Paul has more than 20 years of experience in product pricing, forecasting, reserving, reinsurance, catastrophe modeling and portfolio management.

Paul is a Fellow of the Casualty Actuarial Society (CAS) and a member of the American Academy of Actuaries. He has chaired the CAS Valuation, Finance and Investment Committee, and was an original member of its ERM Committee.

Paul has authored a number of articles relating to pricing, reserving, risk transfer and financial models. He earned a B.S. in economics from the University of Minnesota.

Spencer M. Gluck

Spencer Gluck is a Senior Vice President in the Instrat® unit at Guy Carpenter & Company, LLC. Spencer has 30 years of industry experience and has worked with many lines of business as well as complex alternative risk transfer and capital markets transactions.

Spencer is a Fellow of the Casualty Actuarial Society (CAS) and a member of the American Academy of Actuaries. He has authored articles on reserving and risk transfer in the *Proceedings of the Casualty Actuarial Society* and the *CAS Forum*. He earned a B.A. in mathematics and an M.A. in education from Cornell University.

Gary G. Venter

Gary Venter is a Managing Director in the Instrat® unit at Guy Carpenter & Company, LLC. Gary has more than 30 years of experience as a practitioner in the insurance and reinsurance industry, is the author of a number of published research articles and frequently lectures on actuarial topics worldwide.

Gary is a Fellow of the Casualty Actuarial Society and a member of the American Academy of Actuaries and has served on numerous committees. He earned a B.S. in mathematics and philosophy from the University of California, Berkeley, and a Master's degree in mathematics from Stanford University.