Pricing Aggregate and Credit Risk for Risk Sharing Entities

John D. Deacon, FCAS

ABSTRACT

This paper recognizes that entering into a risk-sharing financial arrangement with another entity creates credit risk. One can use a distribution of outcomes to price both aggregate and credit risk. This paper presents a way to price aggregate and credit risk for deals in which another entity is contractually liable for losses.

INTRODUCTION

Can you say "balance"? When arranging any risk-sharing deal between an insurance company and another entity, the insurance company can be extending credit, whether it knows it or not. Pricing the aggregate risk in concert with credit risk must be considered to avoid exposing the insurance carrier to risks that were not realized, quantified, or priced for. A financial arrangement can quickly get out of balance when these two risks are not priced and evaluated together. I will introduce concepts that are fundamental to this study, discuss pricing of aggregate risk, and offer a method for pricing credit risk when aggregate cover is provided.

FUNDAMENTALS

The scope of this paper is small risk sharing entities (RSE), which generally do not have much equity to commit. The analysis in this paper can be applied to any size entity, however. A few typical entities that fall into the category of RSE are: group and single parent captives, risk retention groups, large insureds with large deductibles (or annual deductibles), and large insureds on retrospective rating plans. 'Small' is considered to have annual losses less than $10 million. Attempts can be made to price aggregate cover for entities with less than $1 million expected annual losses, but I argue that the losses for this entity will be too volatile to provide aggregate cover.

To price a deal with an RSE requires balance. Since these entities do not have much equity, they generally have an interest in having an annual cap on losses they will be responsible for. Entering into a risk-sharing deal with an RSE usually exposes the insurance carrier to another form of risk, credit risk.

Typically, an RSE will do business with a primary insurance carrier, taking advantage of the strength, flexibility, services, expertise, national presence, or product line rate/form/rule filings. The RSE needs the insurance carrier for one of these reasons, or a primary insurance carrier would not be involved. One of the risks for the primary carrier is that RSE losses will exceed the aggregate limit. When RSE losses exceed the aggregate limit, the RSE is relieved of responsibility for losses above that level which are now the responsibility of the insurance carrier. The insured and the claimant may not even be aware of the business arrangement, since the purpose of insurance is to indemnify those damaged as a result of doing their business. The primary company is ultimately responsible for paying claims to claimants, which makes it even more important that they structure deals to protect themselves and ensure their solvency.

Even if the all the premiums are sufficient to cover all the losses on a gross basis, it is certainly conceivable to mis-structure and mis-price a deal so that one entity or the other will certainly lose financially. This can happen quite easily. We propose keys in this paper to avoid financial disasters.

PRICING AGGREGATE - Simulation

What can actuaries do when we are pricing the aggregate charge for an RSE? We can use simulation. Two alternatives for a simulation approach to pricing aggregate cover:

1) simulation of frequency and severity and
2) simulation of loss ratios

Simulation of frequency and severity should be the more precise method as long as the parameters and loss distributions are accurate. Those are critical conditions, however. This method is not for discussion in this paper. The results from this first type of simulation should be theoretically the same as those from simulation of loss ratios.

Another possible way to use simulation in pricing aggregate cover is simulation of loss ratios. Simulation of loss ratios involves determining the distribution of expected outcomes based on the expected amount and variance of how losses relate to premium historically. One immediate caveat is that we need to on-level premiums and losses so that we are measuring true variance and not just variance from changing prices or trends.

Simulation of loss ratios is much quicker and more convenient than simulation of frequency and severity. The claim and severity distributions don't need to be determined explicitly. Loss ratio simulation should nonetheless provide a reasonable approximation of the variance of the outcomes on a prospective basis if on-leveling has been done. We can easily read loss ratios and their corresponding percentiles from the simulation output and relate them to key parameters of the deal. For instance, if the aggregate attachment is 90% of premium we can determine the likelihood that the RSE losses will exceed the agg in terms of percentile.

To help illustrate, the table below shows some assumptions for a possible deal and the relationships between simulated loss ratios and their likelihood of occurring, given the inputs. Let's assume that this is a homogeneous deal for General Liability coverage where the RSE retains $500k per occurrence and also has and aggregate cap of 90% of gross premium (%GP) on annual occurrence-limited losses. We have already simulated loss ratios given the assumptions.

Assumptions

Loss Ratio Distribution	Lognormal
On_Level Limited Expected LR (MEAN)	60%
Standard Deviation of Lim LR (STD DEV)	21%
Aggregate Attachment (%GP)	90%

Percentile	LR	Agg Loss % GP	Probability Weighted Expected Agg Loss
5%	30%	0%	0.0%
10%	34%	0%	0.0%
15%	37%	0%	0.0%
20%	40%	0%	0.0%
25%	42%	0%	0.0%
30%	45%	0%	0.0%
35%	47%	0%	0.0%
40%	49%	0%	0.0%
45%	51%	0%	0.0%
50%	54%	0%	0.0%
55%	56%	0%	0.0%
60%	59%	0%	0.0%
65%	61%	0%	0.0%
70%	65%	0%	0.0%
75%	68%	0%	0.0%
80%	72%	0%	0.0%
85%	77%	0%	0.0%
90%	84%	0%	0.0%
95%	96%	6%	0.3%
99%	121%	31%	1.6%
TOT COST % Gr Prem			1.8%

For example, we would say that a 77% loss ratio is at the 85th percentile. Or, there is a 15% chance that the loss ratio will be higher than 77%.

The loss ratio for a given percentile can be seen as the average for a range. I equate the change in percentile of a given range with the probability of that (average) loss ratio occurring. These agg losses (loss ratio minus agg attachment) are weighted against these probabilities. For example, at a 96% loss ratio the expected agg losses are 6% but these are not certain. They need to be multiplied by the probability of that loss ratio occurring in that range. So multiply 6% by 5% (.95-.90 = the difference in the probability from 90%ile to 95%ile) to get 0.3%, expected agg loss for that outcome.

We would perform the same calculations (to determine expected agg losses) using the output from simulation of frequency and severity in pricing the agg as we do with simulation of loss ratios.

Using simulation to price aggregate cover is not the only method. Table M is a traditional way insurance companies price annual aggregate cover for workers' compensation (WC). The NCCI has developed Table M which uses 'size group' and 'entry ratio' to determine expected agg losses in pricing aggregate cover. Table M contains ranges of expected losses to determine the size group. The entry ratio refers to the relationship of the aggregate attachment point and the expected losses. When one looks up the size group and entry ratio, the table returns the expected aggregate loss as a percentage of RSE expected losses.

The table M charge was developed using all workers compensation business. There are adjustments to be made for differences in severity and whether or not the RSE has a cap on individual claims and others. So then, is Table M applicable to RSEs? This question needs to be evaluated. Table M is the benchmark and is the starting point for RSEs but may not be applicable for a few reasons: 1) homogeneity 2) risk sharing and 3) pricing agg for lines of business other than WC. I do not wish to destroy the credibility of Table M, as the theory is solid, but only to offer alternatives to pricing aggregate cover.

PRICING CREDIT RISK

Once we have determined the distribution of aggregate losses, we have another type of risk our hands, maybe without even realizing it. This risk is credit risk. As mentioned earlier, RSEs are generally not very well capitalized.

We now introduce the need to achieve balance in extending aggregate cover to an RSE. Earlier I spoke of determining ways to avoid financial disaster. Financial disaster can happen a few ways.

One example of financial disaster would be that the aggregate is set so high that the RSE is not able to pay for their losses (because the RSE does not have enough equity) if they were to reach the agg. In this case the primary carrier must pay the losses to the claimants anyway, and will encounter a credit loss. If the carrier does not charge appropriately for the credit risk, it will lose financially.

Another example of a financial disaster would be for the primary carrier to provide agg cover at a low level because the RSE can only afford to pay for losses up to a certain point above premiums. In this case the primary carrier can have very high agg losses since the likelihood of losses exceeding the agg is high. I would consider agg risk to be 'high' when either there exists a 20% chance or higher of an agg loss or if expected agg losses are larger than 10% of GP. These are very **rough** guidelines and can vary greatly from deal to deal. If the carrier does not charge appropriately for the agg risk, it will lose financially.

We would consider a deal to be balanced when we have both an aggregate attachment point that will be reached only infrequently (roughly less than 20% of the time and 10% of GP) but at the same time the RSE needs to be able to afford losses at that level.

Basics

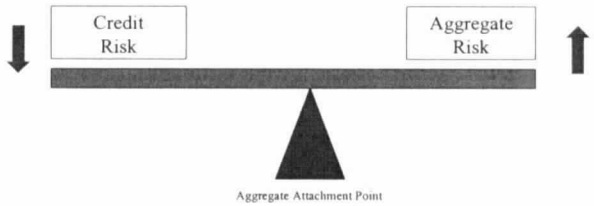

- Reducing risk of agg losses by increasing agg attachment increases credit risk,
- Reducing risk of credit losses by decreasing attachment increases risk of agg losses

The first illustration "Basics" shows how the balance between agg and credit risk works. Pushing on one makes the other pop up.

The second illustration "Possible Outcomes and How They are Covered" is a more detailed illustration of the risks at play. There are a few key sections of this graph to explain. The illustration shows a number line of all of the different outcomes of RSE losses. If losses are low, they are covered by the loss funds (portion of premium for losses, the non-expense piece only). If losses are high enough, they will exceed the agg and be the responsibility of the primary carrier or reinsurer providing the agg cover. The difference between the agg and loss fund is often referred to as the "gap". The gap is the maximum amount the RSE can actually lose, since it uses loss funds from premium to cover at least a portion of its liabilities.

Possible Outcomes and How They Are Covered

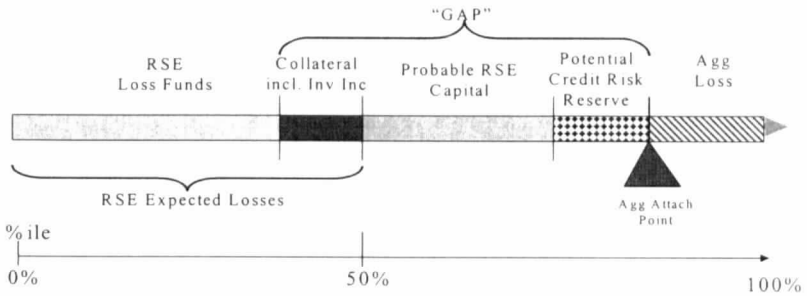

If losses exceed RSE loss funds, they will need to be covered by some other funds. Loss funds earn investment income until used to pay losses, and these investment gains could also be used to pay losses. The primary carrier should absolutely hold collateral (to protect against credit losses) at least up to expected losses of the RSE, preferably up to the agg if possible. In the illustration above, RSE loss funds are insufficient to cover RSE expected losses. This is natural for the same reason that it is common for insurance companies to write profitably above a 100% combined ratio. The reason they can do that is investment income.

In the above example, the RSE expected losses are at the 50th percentile of the distribution of outcomes. It is common when using a Lognormal distribution that the expected loss outcome is at 55th percentile or higher, due to the nature of a long-tailed distribution.

If losses are greater than expected, how will those losses get paid? The answer is the same way that they are paid when the premiums of a primary carrier are insufficient. They must be paid with equity or surplus. So, it becomes necessary to determine whether the RSE can even afford to pay for the "gap", if losses were to reach the agg. Even if the RSE can afford to pay for the "gap", there is some chance that they will not, for whatever reason. For the risk that the RSE will be unable or unwilling to pay, the insurance carrier will have to carry and charge for a credit risk reserve.

Now it is time to estimate the amount of credit risk of the primary carrier. While there are probably an unlimited number of ways to do this, I believe there is some value in having a simple model.

The steps for measuring credit risk for one contract period are as follows:

1. Determine the size of the gap
2. Determine the likelihood of the losses occurring in the gap
3. Determine how much capital the RSE has
4. Estimate the likelihood of receiving reimbursement for losses out of this surplus

I will now explain each of these steps in turn.

Determining the size of the gap should be the easiest step. All that is required is to subtract the loss funds from the aggregate attachment point (agg) which are both known values. The agg is often represented as a % of gross premium, which is 90% in the example. Specifying the agg attachment point in this manner is highly recommended since a set dollar amount can become inadequate quickly with any growth in exposures in the program from the time the pricing work is done to when final exposures are known. For that matter, increases/decreases in price adequacy throughout the contract year can also improve/deteriorate the level of agg protection. These issues should be addressed when setting up the deal. In the example, the "gap" is 35% of gross premium (GP).

Determining the likelihood of losses occurring in the gap can be accomplished by referring to the simulation of losses or loss ratios. The actuary can read percentiles of the distribution that relate to the endpoints of the gap (beginning = level of loss funds %GP, end = aggregate attachment %GP). Since the loss funds (55% GP) are at about the 50th percentile, the probability of losses occurring in the gap (or higher) is about 50% = 1-.50.

In order to evaluate the surplus of the RSE, or whether the RSE will even be able to pay if losses occur in the "gap", we can analyze the financial information of the RSE. Generally speaking, the equity or surplus must be able to take a financial hit in the amount of the gap (and then some since equity needs to cover multiple years of exposure). The equity must be greater than the gap. If it isn't the deal is not really financially viable since the RSE will not be able to pay their liabilities and have some surplus left over to stay in business. This is a straightforward high-level check to make sure the RSE is capable of covering its risk. In the example, the RSE is barely capable of paying for losses if they reach the agg since dollars of gap are close to equity. The step of determining the amount of RSE equity can be fairly difficult since:

1. Accurate financial information may be tough to get (but should be required)
2. Assets and liabilities are constantly changing so this information can become quickly out of date, especially with a small RSE
3. The financial information may not have been audited
4. The RSE may have business with other insurance carriers, which makes any individual carriers "stake" of the RSEs equity extremely difficult to determine.

Since the RSE equity and surplus can bounce around, it behooves the actuary of the primary company to be in tune with the way that the RSE books liabilities and budget for any under reserving through adjusted surplus or conservative estimates, we would recommend the former if possible. It is possible to structure a deal so that the RSE will be unable to pay, so the onus is on the actuary of the primary company to recognize when this is happening and make alternative recommendations.

The final element to evaluate is the likelihood that the RSE will pay (or be able to pay) its liabilities. I name this creditworthiness. Creditworthiness can be determined from financial analysis of the RSE balance sheet and operations. Evaluation of this element is not for the scope of this paper. The number is critical and must be estimated, since ignoring it implies 100% creditworthiness, which is an optimistic and impossible assumption. Every entity has some likelihood of being unable or unwilling to pay its liabilities.

Resulting from this exercise is a dollar and %GP estimate of credit risk that needs to be priced for and managed. For the example, this is $116,361 or 1.7% of GP. This risk also needs to be funded. It should ideally be charged back to the RSE on an expected basis since it is a real risk that the primary carrier takes on. There exist circularity issues with this method since if the price of credit risk is added in with the expenses, the loss funds will be lower. So it is an iterative process. I will not go through the iteration here.

Structure of the deal, key parameters, and Credit Pricing

Program		Value	Formula
Est Annual Gross Premium (GP)	A	$7,000,000	given
Per Occurrence limit	B	$500,000	given
Annual Aggregate Attachment point	C	90%	given
Program Expenses incl. reins.	D	45%	given
"GAP"	E	35%	C-H
"GAP"	F	$2,482,721	AXE
RSE Information			
RSE Expected Loss	G	60%	given
RSE Loss Fund	H	55%	1-D
RSE Equity	I	$3,000,000	given
Creditworthiness	J	0.375	given
Probability of Losses above loss funds	K	50%	from sim table
Credit Risk Pricing			
Probability Expected Credit Loss	L	$116,371	A*E*K-I*J
Probability Expected Credit Loss	M	1.7%	L/A
Discount factor for time value of money	N	0.86	3 years at 5%
Indicated Price for Credit Risk	O	1.4%	M*N

Assumptions with this method of evaluating credit risk:

1. Measuring credit risk can be done with a simple model
2. Reliant on estimates of probabilities from the simulation
3. Liabilities on the RSE financial statements are adequately stated (mentioned earlier); i.e. all equity is allocated to prospective contract period.

This exercise could be applied to older contract periods in much the same manner. The distribution of outcomes (i.e. the variance) should decrease as the period matures.

CONCLUSION

When pricing a financial arrangement with an RSE, the actuary must be cognizant of the different nuances of the deal. The actuary must recognize that RSEs do not usually have a lot of equity to commit, so that entering into a deal with an RSE can create credit risk. Knowing where the attachment point is with respect to the expected losses and quantifying the relative position of these key elements is critical to a well-structured arrangement. If the credit risk and agg risk are not priced in concert, the arrangement can quickly get out of balance, and the primary carrier could have great financial losses.

In most cases, entering into a deal with an RSE also means extending credit. For small RSEs without much surplus, this can be true. In order to price for credit risk, we should consider the amount the RSE has at risk, the likelihood of losses exceeding loss funds, the amount of equity the RSE has, and the likelihood the RSE will pay losses when they exceed loss funds. Using this information, the actuary can recognize, quantify, and price for credit risk and aggregate risk using the same distribution of outcomes.