Pitfalls in the Probability of Ruin Type Risk Management

Jonathan Evans, FCAS, MAAA

Pitfalls in the Probability of Ruin Type Risk Management

Jonathan Evans, FCAS, MAAA

Abstract

Funding levels for many insurance and financial risk entities are often set to achieve a certain low probability of ruin. Specific real world examples which utilize the same essential methodology include: funding self insurance at a certain percentile of aggregate losses, Value at Risk (VAR) funding of investment banks, return period or PML funding of property catastrophe exposures, and probability of ruin through stochastic modeling commonly used in Europe as in Daykin (1994). We use the concepts of probability of ruin, return period, and percentile interchangeably in this paper. Butsic (1992) has pointed out that these analyses neglect to consider the severity of insolvency. This paper addresses a somewhat related issue. Probability of ruin may often be inconsistent with many other reasonable risk management criteria. For example, combining two independent risks may produce a required funding level at a 1% probability of ruin which is actually higher than the sum of the separate 1% probability of ruin funding levels for each of the risks. Use of this criterion for risk management may lead to the nonsensical result of discouraging risk sharing between independent risks. We examine several examples of this phenomenon and how it may lead to undesirable risk management strategies.

Homeowners Insurance, a Trivial Real World Example

A single house generally has a 90^{th} percentile loss in a given year of 0. However, a portfolio of 1,000,000 houses will invariably have a 90^{th} percentile loss in a given year much greater than 0. So the 90^{th} percentile of the combined risks is greater than the sum of the 90^{th} percentiles of the separate risks. If a homeowner wishes to minimize his 90^{th} percentile loss (or perhaps even 99^{th} percentile loss) he should buy no insurance at all, since the premium itself guarantees a 90^{th} percentile loss greater than 0. Equivalently, a large insurance group should form a separate member company for each policy, so as to keep the 90^{th} percentile losses at 0.

We can find trivial examples of this phenomenon at arbitrarily high percentiles less than 100%, or equivalently arbitrarily small probabilities of ruin greater than 0% (See Appendix – Theorem 1).

How Can This Happen ?

Many people are stunned by this result. They are properly taught to think of pooling or sharing of risk as a way of reducing or managing risk. This is always true if risk is measured by standard deviation. Two separate risks, whatever their correlation, will always have a total standard deviation less than or equal to the sum of their separate standard deviations. However, certain percentile type measurements may be greater for a combination than the sum of the separate parts, even for very high percentiles. It is important to note that the Normal distribution does not exhibit this phenomenon (See Appendix – Theorem 2).

A Symmetric Example

This phenomenon is not just a characteristic of skewed distributions. It can also happen for some symmetric distributions. Consider the sum or convolution, X1 + X2 of two identical and independent copies of the random variables X, as follows:

X	Probability	X1 + X2	Probability
1	20%	2	4%
0	60%	1	24%
-1	20%	0	44%
		-1	24%
		-2	4%

The 75th percentile of X1 and X2 separately is clearly 0, but the 75th percentile of X1 + X2 is 1.

Lognormal Example

It can also happen for smooth continuous distributions with only one local maximum. For an example using continuous loss distributions, consider two independent risks with simulated (65,000 iterations) lognormal distributions X1 and X2:

		X1	X2	X1 + X 2	Difference Between Percentile of Sum and Sum of Percentiles
	mean	100,000	300,000	400,000	
	CV	300%	200%	168%	
	sigma	1.51743	1.26864	NA	
	mu	10.36163	11.80682	NA	
Percentiles	99%	1,085,317	2,550,976	2,896,718	-739,575
	95%	389,469	1,066,725	1,313,245	-142,950
	90%	225,804	678,361	871,090	-33,076
	85%	154,968	499,047	664,510	10,495
	80%	113,881	388,787	537,483	34,815
	75%	88,751	315,419	445,164	40,994
	70%	70,837	260,696	379,135	47,602
	65%	57,266	218,513	326,772	50,994
	60%	46,875	185,408	284,290	52,007
	55%	38,724	157,213	247,084	51,147
	52%	34,419	143,000	228,476	51,057

Although at the 90th percentile we see a combined percentile less than the sum of the separate percentiles, as high as the 85th percentile the combined value is larger.

Self Insured Workers Compensation, Frequency/Severity Example

We can extend the lognormal example to a real world frequency/severity process. Consider two large factories whose workers compensation risks are independent. The factories are considering pooling their self insured workers compensation. State law requires that self insured workers compensation be funded at the 75th percentile of gross loss before required per occurrence excess coverage. Let Factory 1 have a claim severity distribution equal to the first lognormal from the previous example and a Poisson claim frequency distribution with a mean of 2.1. Let Factory 2 have the second lognormal from the previous example for its severity and a Poisson claim frequency with mean of 1.2. A typical result from 65,000 simulations is:

		Factory 1	Factory 2	Factory 1 + Factory2	Difference Between Percentile of Sum and Sum of Percentiles
	Poisson Frequency	2.1	1.2	3.3	
	Severity mean	100,000	300,000	172,727	
	Severity CV	300%	200%	257.2%	
	sigma	1.51743	1.26864	NA	
	mu	10.36163	11.80682	NA	
Aggregate Loss Percentiles	99%	1,822,088	3,226,341	3,804,187	-1,244,242
	95%	786,388	1,465,536	1,906,977	-344,947
	90%	499,407	947,544	1,322,473	-124,478
	85%	368,077	693,410	1,035,558	-25,929
	80%	286,534	533,505	842,276	22,236
	75%	230,766	417,824	700,176	51,585
	70%	187,598	330,133	592,035	74,304
	65%	154,029	260,452	503,291	88,810
	60%	126,295	205,706	429,314	97,313
	55%	104,266	159,913	367,566	103,387
	52%	92,662	135,263	334,483	106,558

The factories choose not to pool their risk, since doing so would require a net additional contribution of $51,585 to their self insurance fund, even though the higher percentiles for the pooled risk are much less than the sum of the parts.

Property Catastrophe Example

The phenomenon can also happen with portfolios of property catastrophe exposures. Consider two such portfolios. One is for risks exposed to California earthquakes and the other is exposed to Atlantic Hurricanes. Catastrophe modelers typically calculate Poisson frequencies for loss events of different sizes. These events are sorted in descending order and frequencies are accumulated to give a Poisson frequency of an event of a given size or greater. The return period of these losses is defined as the inverse of this cumulative frequency. Portfolios are then evaluated by the size of loss for a given return period, or "PML". Since these two perils are independent and Poisson we can add the separate frequencies to get frequencies for a combined Poisson distributed portfolio.

	Incremental Frequency			Cumulative Frequency at Level and Above			Approximate Return Periods		
Size of Loss Event	Atlantic Hurricane	California Earthquake	Combined	Atlantic Hurricane	California Earthquake	Combined	Atlantic Hurricane	California Earthquake	Combined
100,000,000	0.0100	0.0100	0.0200	0.0100	0.0100	0.0200	100	100	50
60,000,000	0.0100	0.0100	0.0200	0.0200	0.0200	0.0400	50	50	25
20,000,000	0.0200	0.0200	0.0400	0.0400	0.0400	0.0800	25	25	13
10,000,000	0.0600	0.0600	0.1200	0.1000	0.1000	0.2000	10	10	5

Although there are differences between the meanings of frequency and the probability of one or more events in a year, for low frequencies these numbers are essentially the same. So a 100 year return period event has approximately a 1% probability of occuring one or more times in a year. By combining the portfolio we get a 25 year return period loss which is greater than the sum of the 2 separate 25 year loss events. However at the 50 year return period we get a combined loss less than the sum of the separate losses. If credit rating agencies, catastrophe reinsurers, and regulators evaluate companies based on the 25 year return period it does not make sense to combine these risks.

A Related Example: "The Reinsurance Broker's Gimmick"

A reinsurance salesman may propose the following scheme:

"Randomly select half of your property catastrophe policies. Cede 100% of these. You will be ceding half of your premiums and losses, but my assistant – a world renowned statistician and catastrophe management expert - will show you that your 100 year PML will decrease by 60% or more. This is an excellent, cost effective way to manage your Cat risk."

Policies spread throughout Florida or California overall may have a low average correlation for a given hurricane or earthquake event. This is because a given event in either state is relatively localized inside of the state. When viewed from the perspective of two randomly split portfolios recombined this situation may exhibit a similar pattern to the previous example which used a Florida portfolio and a California portfolio. So in exchange for 50% of the premium the 100 year loss may come down by 60% or more, but what the salesman and his brilliant assistant neglect to mention is that the 200 year loss may come down by only 40% or less.

Stochastic Simulation Example

A European investor spends 200 million German Marks to capitalize an insurance company to underwrite maintenance, warranty, and recall insurance for a large European auto manufacturer over a 5 year period. Expected annual losses for routine claims are 1 billion German Marks, with a coefficient of variation of 10%. Investment income exactly offsets underwriting expenses, the risk load is 5% (reduced to 4% after the first year) of routine expected losses, premium is collected and losses are paid annually, and only autos sold and owned in Europe are covered. The investor runs into difficulty after an actuary working for European Union officials models 10,000 stochastic simulations of the company with a Gamma distribution (Billions of Marks are Gamma distributed with Alpha =100, Beta = 0.01) for routine annual claims and a 1% chance in any year that there will be a large model recall costing 2 billion Marks. The actuary discovers that the company has a 12.1% chance of bankruptcy over the course of its 5 year operation. European Union officials state that the absolute maximum probability of bankruptcy they will accept is 10%.

Fortunately, the European investor has a cousin who works as an investment banker on Wall Street in New York and is quite expert at engineering financial derivatives. The cousin proposes to offer annual aggregate loss reinsurance coverage for 400 million Marks xs 1,100 million Marks. In exchange the investor will cede 2.4% of premium and agree to assume the costs for North American owned autos also in the event of a recall, which his cousin had previously agreed to insure. The cost of the North American autos covered in the event of a recall will be another 2 billion Marks. When the actuary adjusts his model for the new reinsurance derivative, he generates a ruin probability of 9.4% The officials concede and the deal is finalized. Some key simulation results are:

Percentiles	Liquidation Loss Before Financial Engineering Deal	Net Liquidation Loss After Financial Engineering Deal
99%	-1.779	-3.828
95%	-0.162	-0.053
90%	-0.020	0.000
80%	0.000	0.000
70%	0.000	0.000

Sample Simulation (Billions of Marks):

	Year 1	Year 2	Year 3	Year 4	Year 5
Beginning Surplus	0.200	0.107	0.059	0.135	0.203
Premium	1.050	1.040	1.040	1.040	1.040
Losses	1.143	1.088	0.964	0.972	0.999
Cat Loss	0	0	0	0	0
Ending Surplus	0.107	0.059	0.135	0.203	0.244
Liquidation Loss Before Financial Engineering Deal	0.000				
Net Beginning Surplus	0.200	0.125	0.052	0.103	0.146
Ceded Premium	0.025	0.025	0.025	0.025	0.025
Ceded Losses	0.043	0.000	0.000	0.000	0.000
Assumed Cat Loss	0	0	0	0	0
Net Ending Surplus	0.125	0.052	0.103	0.146	0.162
Net Liquidation Loss After Financial Engineering Deal	0.000				

What the officials did not consider was that the expected policyholder deficit or expected value of insolvency, which the actuary's model generated, was 84 million Marks before the reinsurance derivative and 167 million Marks after the reinsurance derivative. This is the expected cost to the auto manufacturer (or government guarantor) due to the insurer's default. The default cost has doubled because even though the probability of default has decreased modestly the average cost of default has risen dramatically.

Probability of ruin simulations and analyses, which do not include other risk measurements, are particularly likely to miss the dangers of exotic reinsurance agreements or financial derivatives. With the growing use of Value at Risk (VAR) by investment bankers to analyze derivatives this danger may also be present in banking. A somewhat mitigating factor is that many VAR calculations estimate a variance and then use a Normal distribution to get a percentile. The Normal distribution is not in itself vulnerable to this inconsistency with regard to percentiles versus standard deviations (See Appendix – Theorem 2). The Normal distribution is also generally not vulnerable to inconsistencies between percentile type measures and expected policyholder deficit type measures, see Butsic (1992).

Two Possible Defenses of Probability of Ruin Type Methods

There is a strong case for a minimum probability threshold for risk management. A reasonable value judgement may be that events which have less than a 1 in 1,000,000 chance of happening should simply be ignored. It may be ridiculous for routine decisions to be based on worst possible outcomes. Similarly, perhaps some people would say that we can allow the 1 in 50 chance event to be worse in exchange for lowering the 1 in 25 chance event.

A second related argument arises if real world entities such as regulators, reinsurance markets, or credit rating agencies are fixed on a certain percentile level for things like pricing reinsurance, setting capital requirements, and assigning credit ratings. If this is the case, then a risk manager or insurance executive may still find the optimal strategy to be based on percentile type measures. That is to say, a single player in the market place may not be wise to ignore existing standards.

Possible Solution

Fixing either a certain tolerable probability of ruin or minimizing the probability of a loss of a certain magnitude allows for undesirable results, primarily because it ignores other levels of probability or time horizon. The optimal strategy may change dramatically for different levels of probability or time horizon. Standard Deviation considers all levels of probability but may give unreasonable weight to large rare events. A possible compromise is to introduce another measure which covers many or all levels of probability/loss size. For example, a utility function with decreasing weight for less probable levels of loss could be used to weight the magnitude of ruin at various levels of probability.

Conclusion

Probability of ruin type calculations are pervasive throughout insurance and finance. However, their use as a standard for setting risk based capital requirements or as a selection criterion for comparing different risk management strategies may lead to nonsensical and undesirable consequences. In some cases this is obvious, such as when it implies that homeowners should not buy any insurance since doing so would increase their 90^{th} percentile losses. Other cases are more subtle, such as the case where randomly ceding half of a portfolio of catastrophe exposed property risks reduces a 100 year PML by more than half, even though this reduces the 250 year and higher PMLs by less than half. Any application of probability of ruin type methods to risk management should be accompanied by consideration of alternative measurements of risk.

References

Butsic, Robert P. (1992) Solvency Measurement for Property-Liability Risk-Based Capital Applications, *Casualty Actuarial Society Discussion Paper Program* May 52-83 , Casualty Actuarial Society, Arlington, Virgina

C.D. Daykin, T. Pentikainen, M. Pesonen (1994) *Practical Risk Theory for Actuaries (Monographs on Statistics and Applied Probability, Vol 53)*, CRC Press

Appendix

Theorem 1

For any "percentile" f such that $0 < f < 1$ there exist 2 independent nonnegative random variables such that the f percentile of their sum is greater than the sum of the f percentiles of each of the random variables.

<u>Proof ("Flipping 2 weighted Coins"):</u>

Let X1 be a random variable with a probability of being 0 equal to f and a probability of being 1 equal to $1 - f$. Let X2 be a random variable identical to and independent of X1.

The f percentiles of X1 and X2 are both equal to 0.

Prob(X1 + X2 > 0) =
Prob (X1 > 0 OR X2 > 0) =
Prob(X1 > 0) + Prob(X2 > 0) – Prob(X1 > 0 AND X2 >0).

Independence implies

Prob(X1 > 0 AND X2 > 0) = Prob(X1 > 0) *Prob(X2 > 0).

So, Prob(X1 + X2 > 0) = 2 * (1 - f) – (1 - f)^2

Since $0 < f < 1$ we also know $0 < 1- f < 1$

Therefore (1 – f) ^ 2 < 1 - f and 2 * (1-f) – (1-f)^2 > 1 - f

So the Prob(X1 + X2 > 0) > 1 – f

QED

Theorem 2

The Normal distribution does not demonstrate the phenomenon in Theorem 1.

<u>Proof:</u>

Consider two independent normal distributions:

X1 = Normal(Mean1, Sigma1) and
X2 = Normal(Mean2, Sigma2)

It immediately follows that

X1 + X2 = Normal(Mean1 + Mean2, SigmaTotal).

For any percentile there exists a unique constant k, such that for any normal distribution the value of that percentile is equal to mean + k Sigma. So we have the following percentile values:

Risk	Value at Percentile
X1	Mean1 + k Sigma1
X2	Mean2 + k Sigma2
X1 + X2	Mean1 + Mean2 + k SigmaTotal

SigmaTotal is always less than or equal to Sigma1 + Sigma2. For percentiles greater than the 50th, $k > 0$. So for percentiles greater than the 50th percentile the value of X1 + X2 is always less than or equal to the sum of the corresponding percentile values for X1 and X2.