Risk Management Newsletter

August 2015
Issue 33

Fatness of Tails in Risk Models

By David Ingram

ALMOST EVERY BUSINESS DECISION MAKER IS FAMILIAR WITH THE MEANING OF AVERAGE AND STANDARD DEVIATION WHEN APPLIED TO BUSINESS STATISTICS. These commonly used and almost universally understood terms can be used as the basis for a new metric of "fatness of tails." This new metric, called "Coefficient of Riskiness," would complement the Coefficient of Variance measure that is commonly used by modelers to compare volatility of different models. This new metric would similarly allow for comparisons of tails.

David Ingram, FSA, CERA, FRM, PRM, MAAA, is EVP and head of ERM Advisory Services at Willis Re in New York, N.Y. He can be reached at *dave.ingram@willis.com*.

EXTRAPOLATING THE TAILS OF THE RISK MODEL

The statistical approach to building a model of risk involves collecting observations and then using the data along with a general understanding of the underlying phenomena to choose a Probability Distribution Function (PDF). The parameters of that PDF are then chosen to a best fit with both the data and the general expectations about the risk.

This process is often explained in those terms—fitting one of several common PDFs to the data. But an alternate view of the process would be to think of it as an extrapolation. The observed values generally fall near to the mean. Under the Normal PDF, we would expect the observations to fall within one standard deviation of the mean about two thirds of the time. Within two standard deviations almost 98 percent of the time. When modeling annual results, it is fairly unlikely that we will have even one observation to guide the "fit" at the 99 percentile.[1]

So, in most cases, we really are using the shape of the PDF to extrapolate to a 99 percentile or 99.5 percentile value. But our method of describing our models presents that fact in a fairly obtuse fashion. Sometimes model documentation mentions the PDF that we use for this extrapolation. Rarely does the documentation discuss why the PDF was chosen. In the cases where this selection process is discussed, it is almost never mentioned that it is judgment of the modeler that drives the exact selection of the parameters that will determine the extreme values via the extrapolation process.

After the 2001 dotcom stock market crash, many modelers of stock market risk adopted a regime switching model as a technique to create the fat tails that many realized were missing from stock market risk models.[2]

But how fat were the tails in these regime switching models? Would reporting the skew and kurtosis of the resulting model help with understanding of the model? Or is the regime-switching equity risk model now a black box that can only be understood by other modelers?

We use the idea of extrapolation to construct for this new proposed measure of fatness of tails. The central idea is that we will have a three point description of our risk model—mean, standard deviation and Coefficient of Riskiness. With these three terms we can describe the degree to which we can expect a risk to have common fluctuations that will drive variability in expected earnings (mean and standard deviation) as well as an indication of the degree to which this risk might produce extreme losses of the sort that we generally hold capital for.

COEFFICIENT OF RISKINESS

Many remember the words of David Viniar, CFO of Goldman Sacks, who famously observed during the financial crisis that "we are seeing things that were 25 standard deviation moves, several days in a row."[3]

As we will show shortly, for some models, moves of many multiples of standard deviations may be expected. The Coefficient of Riskiness (CoR) is defined to help with discussing this quality of risk models. The CoR is the number of standard deviations that the 99.9 percentile value is from the mean.[4]

$$\text{CoR} = (V_{.999} - \mu)/\sigma$$

The CoR can be quickly and easily calculated for almost all risk models. It can then be used to communicate the

"Many remember the words of David Viniar, CFO of Goldman Sacks, who famously observed during the financial crisis that '*we are seeing things that were 25 standard deviation moves, several days in a row*' "

way that the risk model predicts extreme losses, allowing for actual discussion of extreme loss expectations with non-modelers. We use the mean and standard deviation in defining the CoR, not because they are the mathematically optimal way to measure extreme value tendency, but because they are the two risk modeling terms that are already widely known to business leaders.

Potentially, the CoR could become a part of the process for the initial construction of risk models, taking the position of a Bayesian prior in the common situation where there are no observations of the extreme values. And, if CoR has been established as a common idea with non-modelers, they could have a voice in the process of determining how the model will approach that part of the risk modeling puzzle.

The CoR value will not be a reliable indicator for models where the standard deviation is not reliable. It is instructive to identify the characteristics of such models and the underlying risks that such models seek to capture.

COEFFICIENT OF RISKINESS FOR VARIOUS PROBABILITY DISTRIBUTION FUNCTIONS

The CoR for the Normal PDF is 3.09. This is true for all models that use the Normal PDF, because all values of a Normal PDF are uniquely determined by the mean and standard deviation.

Another commonly used PDF is the Lognormal. The lognormal model has two characteristics that make it popular for risk models—it does not allow negative outcomes and it has a limited positive skew.

As it turns out, the CoR is a function of the Coefficient of Variance for the Lognormal PDF.

Table 2 suggests that very large CoR values are possible for models of risks with standard deviation that are very small compared to the mean (CV close to zero above).

The Poisson PDF is also widely used because of its relationship to the binomial distribution. Since the Poisson PDF is fully determined by a single parameter, the CoR is always approximately 3.5.

CONTINUED ON PAGE 8

Table 1: Lognormal PDF: CoR

Coefficient of Riskiness for Various Means/ Std Dev Combinations

		Mean				
		100%	80%	40%	20%	10%
Standard Deviation	7%	17.7	14.9	9.6	7.7	8.2
	10%	13.5	11.7	8.3	7.7	9.3
	15%	10.5	9.3	7.7	8.4	10.6
	20%	9.0	8.3	7.7	9.3	11.3
	25%	8.3	7.8	8.0	10.0	11.5
	30%	7.9	7.7	8.4	10.6	11.5
	40%	7.6	7.7	9.3	11.3	11.0
	50%	7.7	8.0	10.0	11.5	10.4
	60%	7.9	8.4	10.6	11.5	9.7
	70%	8.2	8.8	11.0	11.3	9.1
	80%	8.6	9.3	11.3	11.0	8.6
	90%	8.9	9.7	11.4	10.7	8.1
	100%	9.3	10.0	11.5	10.4	7.6
	120%	9.9	10.6	11.5	9.7	6.8

Table 2: Lognormal: CoR vs. CoV

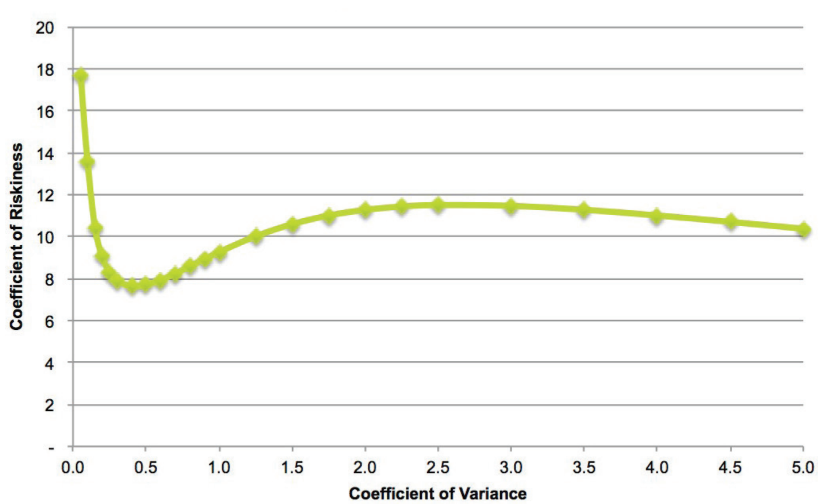

The Exponential PDF and its close cousin, the Pareto PDF, are used for a variety of types of risks. These risks all have the characteristic that they are usually fairly benign but in rare instances, they produce extremely adverse outcomes. Operational risks are sometimes modeled with an Exponential PDF. Risks from extreme windstorms and earthquakes are also modeled with Exponential PDFs as is pandemic risk.

The Exponential PDF models can produce a wide range of CoR values. Standard deviation, the Normal PDF concept, does not always work well for an Exponential PDF. In theory, the standard deviation (as well as the 99.9 percentile value) can actually be infinite. This may be an insurmountable problem with using the CoR on Exponential PDF risk models.

To solve that problem, some models use truncated exponential models. Truncated exponential models will have finite variance but might still have unstable sample values at the 99.9 percentile and therefore unstable CoR.

The situations where the CoR cannot be reliably applied to an exponential PDF are those that are characterized as "Wild Randomness" and "Extreme Randomness" by Mandelbrot[5] on his seven point scale of varieties of randomness. If you want to use CoR to compare your risk models, you can just mark these models with infinite variance as WR or ER. Hopefully, your WR and ER risks will be a small part of your overall risk profile and there will be a finite variance for the entire company risk model.

Extreme value analysis (EVT) does not, by design, permit a generalized look at a statistic like CoR because it is fundamentally an approach that divorces the tail risk analysis from the data regarding the middle of the distribution that make up the mean and standard deviation. However, individual risk models that blend a model of expected variation around the mean with a specific model of the extremes based upon the generalized extreme value distribution can produce values that would lead to a CoR calculation and CoR could help to provide a metric for comparing risk models that incorporate EVT with other risk models that do not.

EXAMPLES FROM INSURANCE RISK MODELS

The author has obtained summary information from approximately 3400 models of gross (before reinsurance) property and casualty insurance risks that were performed over the 2009 to 2013 time frame by actuaries at Willis Re.

In addition, we have obtained summary output from stand-alone natural catastrophe model runs for property insurance.

It is interesting to note that none of these models showed a 99.9 percentile result that was 25 standard deviations. But, as you see, the natural catastrophe models did produce CoR values as high as 18.

What you can see from three examples is that CoR does seem to be bounded for these actual models into the range of 3–18 and that existing processes for modeling insurance risks do already produce a range of CoR values.

COMMUNICATING RISKINESS WITH COR

Non-technical managers are usually familiar with the ideas of mean and standard deviation as the defining terms for statistical models. The CoR described here is proposed as a substitute for a discussion of the characteristics and implications of the selection of PDF that in general, is needed but is not taking place.

The CoR, if adopted widely, could come to be used similarly to the Richter scale for earthquakes or the Saffir-Simpson Hurricane Wind Scale. If you were presenting a model of hurricanes or earthquakes and mentioned that you had modeled a 2 as the most severe event, everyone in the room would have a sense of what

CONTINUED ON PAGE 10

Chart 1

3400 Insurance Risk Models[6]

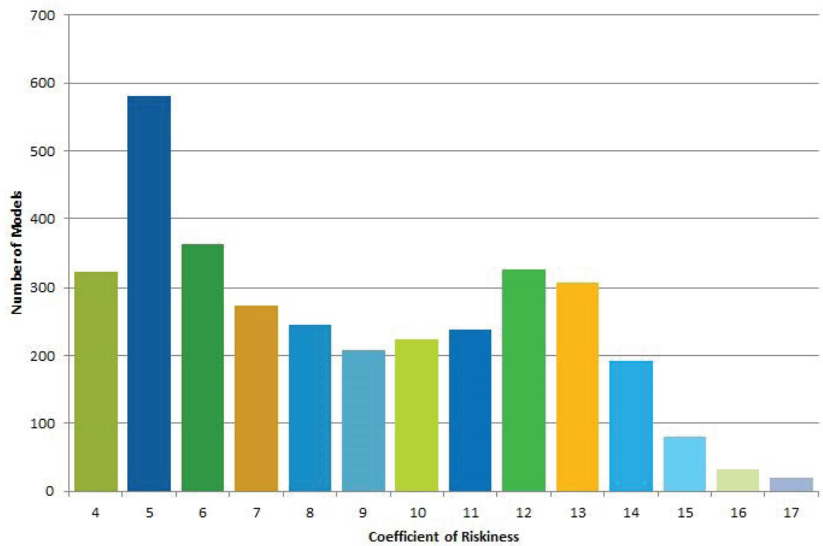

Chart 2

400 Natural Catastrophe Models

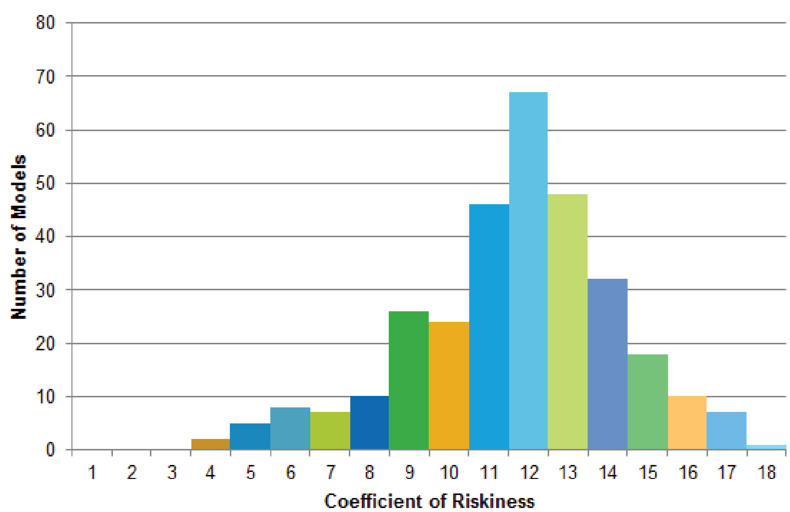

that meant, even if they do not know anything about the details of the modeling approach. They will have an opinion about whether a 2 is the appropriate value for the most severe possible hurricane or earthquake. They can easily participate in a discussion of the assumptions of the model on that basis.

The CoR does not really add any information about the model for PDFs such as the normal, lognormal and Poisson. However, the adequacy of those models to produce appropriate extrapolations of the fat tails actually experienced should by now be highly suspect. But it then does allow for quick comparison of fatness of tails of those models that use a single PDF with those models where different PDFs are used for frequency and severity of risk, for example.

The CoR could become a familiar tool for broad communication of model severity. If you believe that Vineir's comment about 25 standard deviations was actually based upon a measurement (rather than a round number exaggeration to make a point), then you would doubtless reject the validity of the model with a CoR of 3 or 4. If non-technical users of a risk model gained an appreciation of which of the company's risks have CoR of 3 and which were 12's that may be a large leap of understanding of a very important characteristic of the risks.

The hope is that by turning away from the technical, statistical discussion about choice of PDF and parameterization, the discussion can actually tap into the extensive knowledge and experience and gut feel of the non-technical management and board members. Perhaps the CoR can become like the Richter scale of risk models. Few people understand the science or math behind the Richter scale, but everyone in an earthquake zone can experience a shake and come pretty close to nailing the Richter score of that event without any fancy equipment. And they know how to prepare for a 4, a 5 or a 6 quake. The same goes for the Safir-Simpson Hurricane Wind Scale.

CONCLUSION

"If you don't know where you are going,
any road will take you there"
– Lewis Carroll

People naturally observe risk in the form of the range of experienced gains and losses. In statistical terms, those observations are represented by standard deviation. Statistical techniques that have long been applied to insurance company risks to develop central estimates are being used to calculate values in the extreme tails of the distribution of gains and losses. These processes are essentially an extrapolation from the "known" risk of volatility near the mean to "unknown" risk of extreme losses.

To date, there is no established language to talk about the nature of that extrapolation. The CoR described here is an attempt to bridge that gap. The CoR can be used to differentiate risk models according the fatness of the tails and could become a standard part of our discussion of risk models. With the use of a metric like the CoR, we believe that the knowledge and experience of non-technical management and board members can be brought into the discussions of risk model parameterization. The end result of such discussions will both ultimately improve the models and increase the degree to which they are actually relied upon for informing important decisions within a risk taking enterprise.

This article is a summary of the paper that won the Best Practical Paper awarded by the JRMS at the 2015 ERM Symposium.

ENDNOTES

[1] Some one-year loss calculations are performed by calculating a value for a much shorter period and extending that calculation to the full year by making an heroic assumption about the relationship between that short period and the full year. That substitutes a problem from that time period assumption for the lack of actual data about full year risk. And whether practitioners realize it or not, that process is an extrapolation into the unknown.

[2] Mary R. Hardy, "A Regime-Switching Model of Long-Term Stock Returns." *North American Actuarial Journal* Volume 5, Issue 2 (2001).

[3] *Financial Times*, August 13, 2007.

[4] The 99.9 percentile is chosen to be beyond the values most often used from the model. All of the ideas presented here about CoR would apply with a different chosen reference point.

[5] Benoit B. Mandelbrot, *Fractals and Scaling in Finance*, Springer, 1997.

[6] For this chart and the following, the CoR of 4, for example, indicates a value between 3 and 4.