LLOYD'S MARKET ASSOCIATION

CATASTROPHE MODELLING

GUIDANCE FOR NON-CATASTROPHE MODELLERS

JUNE 2013

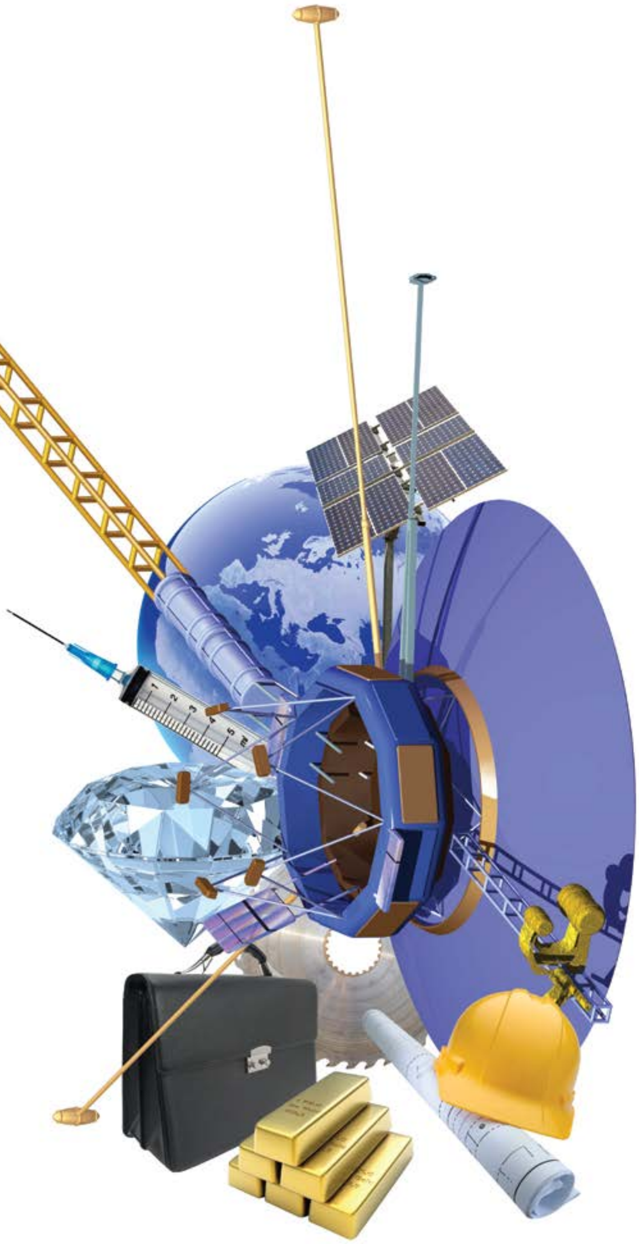

Lloyd's Market Association (LMA) provides professional technical support to the Lloyd's underwriting community. For more detailed or current information on a range of topics please visit the LMA website.

If you have feedback on this guide, or would like to suggest a question for future inclusion, please contact: gary.budinger@lmalloyds.com

---

LLOYD'S MARKET ASSOCIATION

GUIDANCE FOR NON-CATASTROPHE MODELLERS

Contents

Introduction

The LMA Exposure Management Working Group (EMWG) was formed to look after the interests of catastrophe ("cat") modellers working in the Lloyd's market. Given the increasing use of cat models in the market in recent years and their perceived complexity and sophistication, the EMWG agreed that it would be useful to put together a simple guide for colleagues not working on cat modelling.

The aim of this guide is therefore to provide some of the background information as to what a cat model is, explain the terminology frequently used by modellers, and answer many of the most frequently asked questions on cat models. This document is not intended to be a comprehensive guide to cat modelling or to answer all questions on this subject. It will hopefully provide non-cat modellers with a basic understanding of the work of cat modellers, and also act as a quick reference document on cat models.

Uses of Cat Modelling

A cat model is a computerised system that generates a robust set of simulated events and estimates the magnitude, intensity, and location of the event to determine the amount of damage and calculate the insured loss as a result of a catastrophic event such as a hurricane or an earthquake. As with any emerging industry the boundaries are not defined, and many myriad of uses of cat modelling continue to be established. At the core, within the insurance industry, cat modelling supports:

- Risk pricing – using local software a quick repeatable risk assessment over the known locations of a risk being offered can be run. In addition to supporting the calculation of a robust internal technical price, the cat modelling process can provide a wealth of additional information regarding the potential hazard exposure. This includes; susceptibility to hurricanes or earthquakes, proximity to liquefaction (the process of something moving from a solid to liquid mass) storm surge risks such as flooding, or vulnerability assessments (for example, which building standard code was in place when the locations were built and what relative impact this could have). While there is still a lot of uncertainty and complexity to assessing such risks, the benchmark figures produced allow relative comparisons between risks, and over time. All of this is intended to supplement an Underwriter's wider knowledge about a risk and lead to optimal decision making over the long term, including calculating the correct price.

- Portfolio management – As above for individual risks, so for an entire portfolio of risks, cat modelling is used to rapidly accumulate across a portfolio to communicate the combined profile. For example, acting as a common currency, cat modelling can put a high value industrial facility in the US in direct comparison to a warehouse in Belgium. At this level cat modelling supports business strategy, both identifying areas of concern (such as with too great an accumulation of correlating risks) or identifying opportunities (where diversifying risks could be added to the portfolio with marginal impact).

- Capital requirements – The robust standardised approach to assessing cat risk that cat modelling provides, can benefit other processes undertaken by insurers. The main usage is in calculating solvency and other regulatory or economic capital requirements, where the output from a cat modelling process will provide a risk profile that can be combined with other forms of business risk to inform capital requirements.

Cat Modelling Practice

Operating the software is only a small part of what it takes to effectively utilise cat modelling within a business. As with any model that attempts to simplify and represent real world phenomena it is vital that there is a strong understanding of the appropriate usage and limitations of the model.

Different components of a cat model

A cat model is built up of a number of modules that must all operate in coordination to produce the desired risk assessment. It is important to note that two of these (the hazard and vulnerability modules) could be considered individual models in their own right, and the combination of one feeding the other brings with it challenges that need to be understood.

The vendors of cat models create and fix a set of events. While these are a small subset of the range of potential outcomes they provide a sensible number of scenarios that will represent the underlying hazard, while remaining at a practical level to make quick decisions. These events are run consistently each time a model is operated, so there is no random element involved, and they can be compared between risks and between different companies using the same model. As they only represent a subset, increased uncertainty should be applied to events at the extreme tail.

The components of a cat model are:

- Exposure data module – every cat model needs an input of risks against which an assessment is to be made. This usually consists of capturing multiple details about a risk, along with recording insurance policy terms. The two essential features of a risk that need to be known are the geo-location and an insured value. After this, depending on the type of risk, there might be options to enter primary characteristics, such as construction type or year built for a property risk, and even secondary characteristics such as roof type. Some models provide approaches for geocoding locations based on addresses, and for estimating characteristics if unknown.

- Hazard module – each generated event is tagged with the core components relevant to the hazard. For a Hurricane this might be landfall location and direction of travel, peak wind-speeds and central pressure; for an earthquake this would normally be the epicentre and magnitude. The hazard module must combine this information with the exposure data being provided and any information the model has on salient features such as surface roughness (for windstorm hazard) or soil type (for earthquake hazard) at each location. For each event an assessment of the hazard impact at each location being assessed must be established.

- Vulnerability module – the resulting output of the hazard module is then passed to the vulnerability module. The hazard at any one location is independent of the risk that is actually there, but what we are interested in is how the risk at that location will respond to the predicted hazard conditions. The vulnerability module contains a number of vulnerability curves, with the appropriate one chosen depending on the primary characteristics of the risk. These are often derived from engineering studies or past experience, and represent how a risk will respond under different conditions. For an earthquake, for example, peak ground acceleration (PGA) is often the most important factor when considering how badly damaged a building will be. As the PGA increases, so does the expected damage. The relationship between the two is described in the vulnerability curve. Secondary characteristics, if provided, will often be used to make minor modifications to the vulnerability curves. The result of these calculations is a damage ratio to be applied to the risk at the given location.

- Financial module – armed with the expected damage ratios for each location that we are assessing, the cat model can then begin to accumulate upwards through the financial and insurance terms. Starting with a calculation of the Ground Up loss to the individual location the financial module will typically accumulate through location-level terms, to policy and then programme level conditions, at each stage applying limits, deductibles, and special conditions that have been coded into the model. The resulting output is an Event Loss Table that provides an assessment of the financial risk exposure to individual events. This can then be combined to an exceedance probability (EP) curve to give further measures for the entire risk.

CAT MODELLING TERMINOLOGY

The cat modelling industry is full of terminology and acronyms, many of which have been borrowed from mathematics or actuarial modelling. This section will attempt to explain some of the most common ones used by cat modellers.

EP Curve

An EP curve communicates the probability of any given financial loss being exceeded. It can be used in one of two ways: provided with a financial loss the EP curve could be read to give you the probability of this loss (or a greater loss) occurring; or alternatively provided with a probability level the EP curve could be read to show you the financial loss level that this corresponds to.

It is important to note that this refers to a loss being exceeded, and not the exact loss itself. This approach is used for cat modelling, as it is beneficial to identify attachment or exhaustion probabilities, calculate expected losses within a given range, or to provide benchmarks for comparisons between risks or over time. Calculating the probability of an exact financial loss is of little value.

OEP and an AEP curve

OEP stands for Occurrence Exceedance Probability; AEP stands for Aggregate Exceedance Probability.

The OEP represents the probability of seeing any single event within a defined period (typically one year) with a particular loss size or greater; the AEP represents the probability of seeing total annual losses of a particular amount or greater.

They can be used in tandem to assist with managing exposure both to single large events, as well as accumulations of multiple events across a period.

VaR and TVaR

VaR stands for Value at Risk; TVaR stands for Tail Value at Risk.

They are both mathematical measures used in cat modelling to represent a risk profile, or range of potential outcomes, in a single value.

Value at Risk is equivalent to Return Period, and measures a single point of a range of potential outcomes corresponding to a given confidence or fixed position. When used to compare two risks, in conjunction with the mean loss, it communicates a measure of uncertainty in the loss assessment.

Tail Value at Risk (or Tail Conditional Expectation) measures the mean loss of all potential outcomes with losses greater than a fixed point. It helps to communicate 'how bad things could get'. When used to compare two risks, along with mean loss and Value at Risk, it helps communicate how quickly potential losses tail off.

With current modelling techniques any EP curve is limited by the number of theoretical events or simulation years used to make it up. In the tail of a distribution there can be large jumps between individual points. Value at Risk points read at high return period / confidence levels can perform strangely as the limited number of sample points makes figures jump back and forth between assessments. The TVaR measure provides a small amount of protection against this effect. By considering the average of all points in the tail it is less sensitive to such effects and can provide a more stable measure.

However the TVaR is necessarily reliant on the quality of modelling in the tail of the distribution, where models will always be fairly weak.

Event Loss Table (ELT)

An ELT is a collection of theoretical cats (hurricanes, earthquakes etc.) along with the modelled losses estimated to occur from each event. This forms the raw data that is used to build up EP Curves and calculate other measures of risk.

Coefficient of Variation (CoV)

The CoV is the standard deviation divided by the mean (annual average loss). The wider the variation on the distribution of data, the higher the COV.

Difference between Near Term, Long Term and Historical rates

Models for North Atlantic Hurricane need to take into account the strong influence that global climate and oceanic conditions have on them, potentially affecting everything from frequency and strength to landfall location.

Long term or Historical analyses use all available information on past Hurricane activity (stretching back to around 1850) to advise on likely frequencies to be seen in the coming year.

Near Term analyses by AIR (which are referred to as Medium Term analyses by RMS) attempt to better represent current conditions.

AIR does this by marking each historic year as either having the Atlantic in a “warm phase” (where sea surface temperatures in the Atlantic are warmer than the long term average) or a “cold phase”. At present we are assessed to be in a “warm phase”, so AIR uses only historic years in a similar phase to advise on likely frequencies for the model.

RMS takes a different approach, instead eliciting a number of academic “models” designed to forecast the next 5 years of events. They then apply a weight to each model according to how accurately it is able to represent the previous 5 years, to form a blended assessment of future frequencies.

Difference between Ground Up, Gross, Net and Final Net losses

A Ground-up loss is the loss to the policyholder or risk insured; gross loss typically refers to claim made to insurer; net loss typically refers to gross loss net of reinsurance; final net loss typically refers to the gross loss net of reinsurance and reinstatements.

FREQUENTLY ASKED QUESTIONS

Why is it that every time an event occurs I hear that it was not covered properly by the cat models?

A model is only a representation of reality. Depending on the questions being asked a model could be highly complex or extremely simple, and it is in understanding the limits of a model that its value can be properly achieved.

First and foremost it must be understood what a model is attempting to represent in the first place. More recently model vendors have begun to explicitly state the elements of loss that their model is intended to represent, and more importantly they have started to identify known elements of loss that they explicitly do not cover. Cat models do not pretend to cover all elements of all cat risks worldwide, and it is therefore the responsibility of individuals to ensure that they clearly understand both of these. Vendors do continue to work to add to the suite of risks covered by their models, but this is a continual work in progress and is driven largely by market demands.

However, even within risks that are covered we would still expect to see elements that are not perfectly represented. Producing a model of a real world phenomenon is only as good as the information that is available and the investment spent in studying it. Loss Amplification (price increases following a major event caused by a scarcity of resources and an increased demand) is a known impact, but relatively little recorded information about it is available historically worldwide, and how it varies between events that occur once every 10 years to events that occur once every 100 years is almost non-existent. An attempt to allow for this is included in a number of models, but it is highly likely that this will need to develop over time.

Models must be considered in the context of the purpose for which they were designed. For most cat models this is to assess the overall risk profile of a set of locations to particular hazards. To achieve this practically, certain assumptions and approximations are required. When used for its intended purpose these reductions should produce negligible impact, however drilling down too far into any model will reach a point below which the model is no longer appropriate. The climate simulation models, used by the IPCC (Inter-governmental Panel on Climate Change) to estimate the impact of climate change on the planet, would do an appalling job of telling you what the weather will be like at your house on your birthday, but still remain valid approaches for predicting worldwide temperature changes over decades.

When an individual event occurs and the resulting profile is compared against the cat models it is important to identify when an outcome casts doubt on a key assumption relevant to the overall value of the cat model, or whether the particular features of the event simply fall outside of the subset of generated events, but within the consideration of the overall model.

What is the impact of poor quality data on results?

A model is only as good as the data that feeds it. Even if we had perfect exposure data the challenge of cat modelling is still huge, and the results that are produced will contain numerous uncertainties, however if the input data is of poor quality then no amount of modelling will produce correct output. Poor quality data can be of two forms, inaccurate or incomplete.

Models are unable to identify inaccurate data, so will continue to assess the risk based on the information being correct. This means that output will be presented back to the user with no indication that the results being analysed are inappropriate, and if this information continues to feed further down the chain incorrect decision making will follow.

An incorrect location could put the risk further into a hazard zone, or further away. Incorrect primary characteristics could imply the location was more or less vulnerable than reality. If inaccuracies are minor and random and spread through a large enough portfolio of risks then are unlikely to cause too many problems, however if the inaccuracies are systematic, or if they occur on peak risks they have the potential to significantly mislead.

Incomplete data causes problems for a different reason. Models need certainty to proceed, so missing information is usually replaced by estimates. This is beneficial in that it allows us to proceed with a modelled analysis even with information missing, but what is not always clearly communicated is the additional uncertainty that this brings.

In cat modelling communicating and understanding the uncertainty is vital, however in the case of incomplete information no additional uncertainty is added to the results. If there were sufficient time to reprocess the analysis with the complete range of potential inputs it would be more obvious that the missing information will have introduced a far wider range of potential outcomes than is otherwise suggested. When dealing with natural cats the difference between building codes, or distances from the coast or a fault line can make the difference between a risk having no loss or being a total loss.

Why do I need aggregates if I have a model?

Cat modelling is just one of many tools in an arsenal for understanding and managing cat risks. As noted there are many elements of cat risk that cannot currently be modelled, or are in the early stages of being developed into a cat model. Additionally, while modelling helps to push the boundaries of loss forecasting, the limitations and uncertainties are unlikely to go away any time soon, and one must never lose sight of common sense approaches to managing risk.

The recording and monitoring of aggregate positions can provide a useful fall-back and sense check against which the complex output of cat models can be reviewed and challenged.

What is a 1 in 250 return period?

Future losses from cat events cannot be accurately predicted. Instead the purpose of any form of modelling is to use what knowledge we do have about the likelihood of events occurring, along with estimates of the potential impacts that each event could have, to build up a picture of the range of potential outcomes.

To translate this range of outcomes into something meaningful it is common practice to select a fixed confidence level to report against. Asking for the 1 in 250 return period is, like gambling odds, simply an easier way to represent asking for the monetary loss in the range of outcomes where only 1/250 = 0.4% of potential outcomes are worse. In mathematical terms this is the 1 – 0.4% = 99.6% confidence point, and you are stating that you are '99.6% confident' that losses will not be larger than this value.

'Return Period' figures must therefore always be considered within the context of the analysis. For example: Which regions and perils have been included in the assessment, and are there additional potential losses not included? It is important to note that this is simply a way of representing how confident you are about potential loss outcomes being reviewed and is not directly intended to be translated into a multi-year assessment of event frequency, where other considerations would be required.

Why do 1 in 100 year losses happen every few years?

'1 in 100' relates to the probability of a loss in a particular region to a particular peril. Imagine you have a 100 sided dice. With just one dice your chance of rolling a 100 would be 1 in 100 or 1%. However, if you had 10 dice, your chance of rolling a 100 would be 10 times greater - so 10%. The different dice represent the different perils and regions that are insured around the world so, unfortunately, '1 in 100 year events' should be expected every 10 years, if not more frequently.

What is an n-year event?

The "events" used by the models are theoretical and used as vehicles to support the calculations, and should be used with caution. EP curves are then built up considering all events and scenarios, and how they interact with each other and the final resulting curve should be considered as separate from the individual events that went to make it up. This EP curve can now be used to reach your n-year loss, but there is no such thing as an n-year event.

The trouble with converting the purely financial EP curve to a real world comment on events can be seen when you consider that a rare Category 5 hurricane that skims the coast can cause far less financial loss than a more frequent Category 3 that drives onshore.

Additional useful information can be gained from looking at the events or scenarios that cause losses at the n-year level; however it is important to remember that there are a large number of different combinations that could achieve the same result and it will not always be possible to determine this from the model alone. For example, careful examination of the EP curve may lead you to find that your n-year loss is being strongly driven by one country/peril or another; or that you have more or less exposure to single large events than multiple small events.

Why can you not add up return period losses?

A 'Return Period Loss' is the monetary amount, given a range of potential outcomes, where a given fixed percentage of outcomes result in worse monetary losses (see also 'What is a 1 in 250 return period?).

Combining two analyses means combining two sets of potential outcomes. In some cases the two sets may be independent, leaving you simply with a single larger set of outcomes. In other cases the two may interact – perhaps a large loss outcome from the first analysis is linked to a large loss outcome in the second, such as if both have been caused by the same theoretical Hurricane.

The new 'Return Period Loss' for the combined analysis now depends heavily on how the two different sets of outcomes interacted, which can't be seen by looking at the individual analyses alone, and must be recalculated once the grouping has been performed.

Example to illustrate combining two event sets

EVENT SET 1			
Event	**Annual Rate**	**Loss**	**EP**
CA EQ 101	0.10%	15,000	0.10%
CA EQ 202	0.30%	14,500	0.40%
CA EQ 103	0.10%	13,000	0.50%
CA EQ 304	0.20%	12,500	0.70%
CA EQ 105	0.15%	12,000	0.85%
CA EQ 206	0.21%	11,000	1.06%
CA EQ 407	0.10%	10,500	1.15%
100 Year Loss is about 11,000			

EVENT SET 2			
Event	**Annual Rate**	**Loss**	**EP**
US HU 101	0.12%	20,000	0.12%
US HU 202	0.11%	16,000	0.23%
US HU 103	0.15%	14,000	0.38%
US HU 304	0.22%	13,500	0.60%
US HU 105	0.25%	12,750	0.85%
US HU 206	0.12%	12,600	0.97%
US HU 407	0.10%	11,500	1.07%
100 year Loss is about 12,500			

COMBINED EVENT SET			
Event	**Annual Rate**	**Loss**	**EP**
US HU 101	0.12%	20,000	0.12%
US HU 202	0.10%	16,000	0.22%
CA EQ 101	0.10%	15,000	0.32%
CA EQ 202	0.30%	14,500	0.62%
US HU 103	0.15%	14,000	0.77%
US HU 304	0.22%	13,500	0.99%
CA EQ 103	0.10%	13,000	1.09%
US HU 105	0.25%	12,750	1.33%
US HU 206	0.11%	12,600	1.44%
CA EQ 304	0.20%	12,500	1.64%
CA EQ 105	0.15%	12,000	1.79%
US HU 407	0.10%	11,500	1.88%
CA EQ 206	0.21%	11,000	2.09%
CA EQ 407	0.10%	10,500	2.19%
100 Year Loss is about 13,500			

What is Pure Premium?

The Pure Premium represents the average of all potential outcomes considered in the analysis, and could be considered to be the break-even point if such a policy was to be written a very large number of times.

The nature of cat risk means that the profit made when actual losses are lower than this assessed average is heavily outweighed by how much of a loss you could have when actual losses are higher than this assessed average; and that real experience will be very 'spikey' i.e. several years of no loss, followed by a large loss.

Because of this underwriters usually add an "uncertainty" load to reach a technical premium which the models can assist with calculating.

In addition the actual premium charged by underwriters should include consideration for potential losses not included in the modelled assessment, these can include claims handling capabilities, moral hazard, loss record, Loss Adjustment Expenses and other perils (fire, flood, theft etc).

How can I still get a loss to a layer when the mean loss is less than the attachment point?

Future losses from cat events cannot be accurately predicted. Instead the purpose of any form of modelling is to use what knowledge we do have about the likelihood of events occurring, along with estimates of the potential impacts that each event could have, to build up a picture of the range of potential outcomes.

The mean loss given by a model is then just the average of this range of outcomes – the break-even point if this scenario were to be repeated a large number of times – however when applying financial structures the models retain the full range of potential outcomes, and use these when considering losses to insurance policies.

While the average loss may be below the attachment point, the uncertainty involved in predicting exact losses may mean that there are some potential scenarios that do in fact exceed the attachment. It is therefore important that we consider these when calculating possible losses to the written policy.

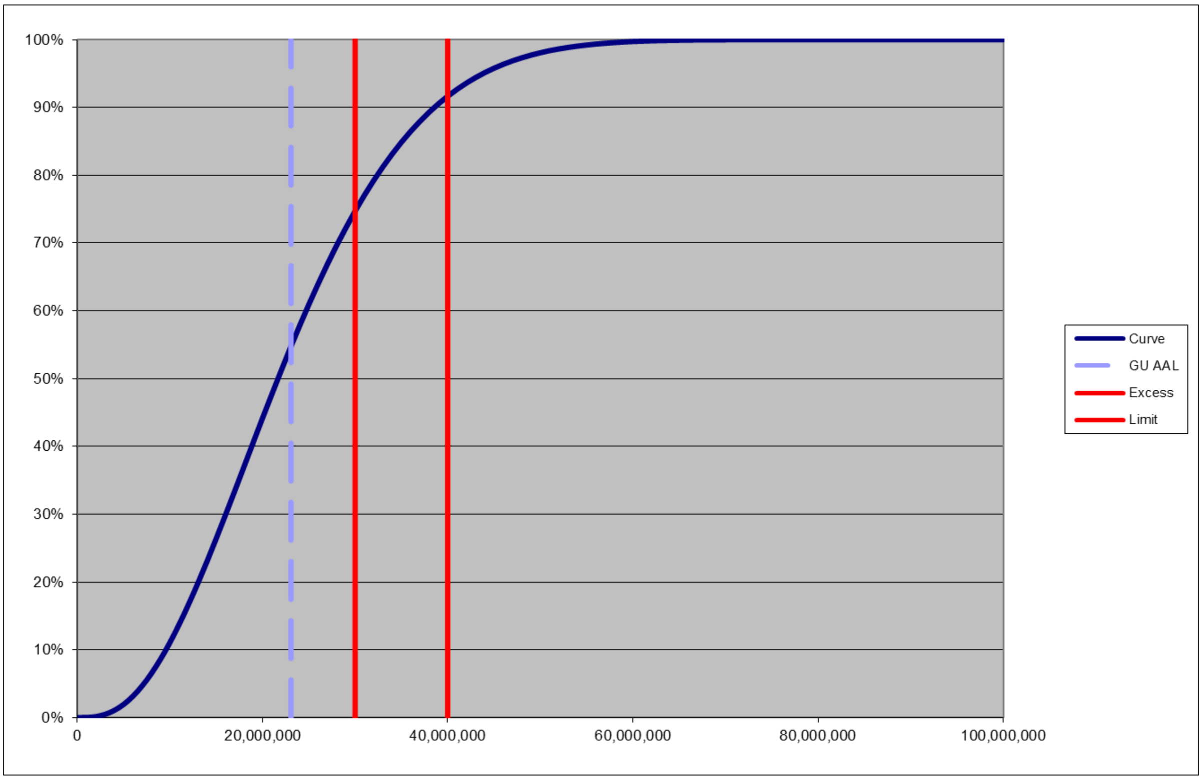

Diagram shows how losses can enter a layer despite the mean loss being less than the attachment point. Red bar = layer; Blue dash line = Ground up Average Annual Loss (GUAAL); Solid blue line = range of potential losses.

To give an example, a model might suggest that out of 10 potential future years they will see 9 clean years, and have one year with a single $100m loss. If you were to consider writing a $20m xs $20m policy on this risk the mean loss is $100m/10 = $10m, which is below the attachment point;

however in reality you would have a 9-in-10 chance of a zero loss, but a 1-in-10 chance of a total loss of $20m, giving you an average loss to the policy of $2m.

Why is the 10,000 year loss in RMS not the worst case loss for this account or portfolio?

This question confuses the AIR / simulation approach to modelling, with RMS's approach.

AIR uses a simulation methodology prior to setting up their model, whereby they run their model to create a potential year of cats 10,000 times. Each time the model is run you get different combinations of events, selected according to pre-coded frequencies. When we run the model in-house we get the resulting losses from these 10,000 potential simulated years. The EP curve that AIR builds up is created by ranking losses in descending order, and assigning each simulated year an equal likelihood of occurring. In this case the 1 in 10,000 year loss is the largest in the set.

RMS takes an entirely different approach. Each event in their model represents a scenario with a range of uncertainty, and each scenario is given an "event rate" that represents a likelihood of occurrence (a weighting). EP curves are built up mathematically from all events in the catalogue, resulting in a final distribution of potential loss outcomes that stretches out as far as they are willing to calculate. In practice this will result in the model being able to give figures for return periods in excess of 1 in 1million, although very little confidence should be given to the modelling anywhere near this part of the curve.

The reality is that neither model can tell you what the "worst case" loss for the account actually is, because our knowledge of cat is still developing. The only sensible answer to this is "total loss". Both models are simply stopping calculations at an extreme point.

END