

INTRODUCTION AND HISTORY

A catastrophe model represents the confluence of actuarial science, engineering, meteorology and seismology and has become intrinsically embedded in the risk assessment process of the Insurance Industry throughout all the major global economic regions. It is a complex mix of algorithms, mathematical functions, geodata and engineering assumptions, usually embedded in a software construct that allows for easy introduction of exposure data. Potential losses from any natural hazard which has been denoted a catastrophe can then be estimated. Catastrophe models are now considered to be essential tools for pricing, assessing capital adequacy and measuring profitability.

They made their first appearance 30 years ago with Applied Insurance Research and were then market-share based systems used only by treaty reinsurers and available only within the US market. As the usefulness of catastrophe models has increased, so has the reliance on them and now the majority of property insurers and reinsurers worldwide use them and indeed regulators *expect* that they use them. Over this time there has been a profound shift in the markets from widespread scepticism to blind acceptance. Karen Clark put it (Clark 2011):

“...what’s happened in the industry is we’ve gone from 20 years ago with no model, all underwriter judgements, and now we’ve gone to the other end of the spectrum, which is all models and no underwriter.”

Although models have matured significantly during the last 20+ years, barriers to entry have been increasing due to high software prices, costly technical infrastructure and, equally as costly, the IT experts required to manage the running of the models. Lack of openness has also restricted the ability to exploit available research in hazards and vulnerability and the weakness of understanding and quantification of uncertainties inherent in using models. The industry is now working with the Association for Cooperative Research and Development (ACORD) to set an industry standard for the collection and sharing of industry data and in 2012 a platform was set up by Oasis Loss modelling Framework to create and disseminate *open* multi-hazard cat risk modelling tools.

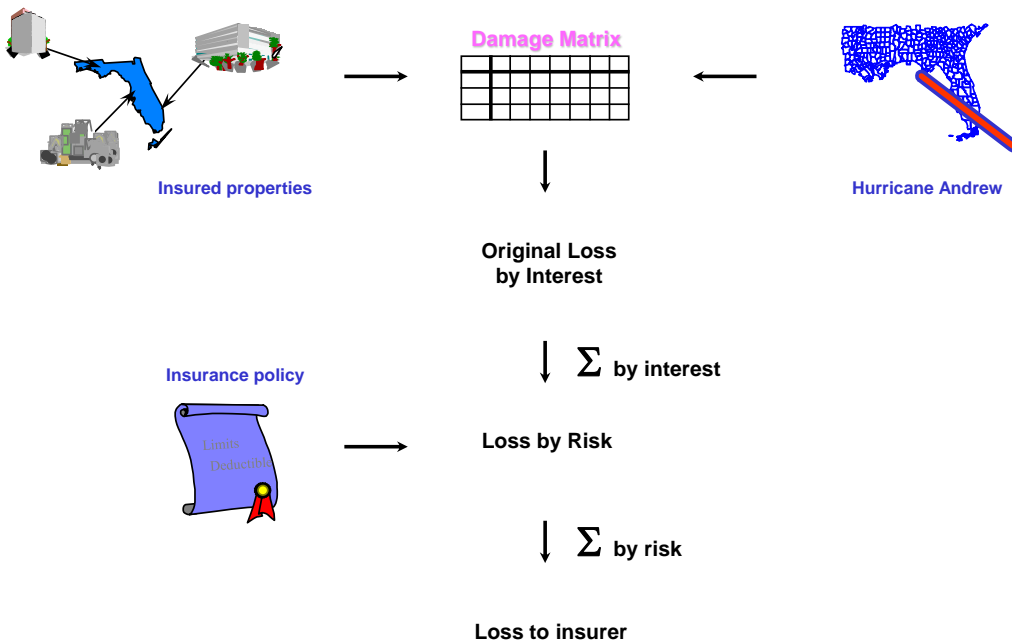
This chapter will outline how catastrophe models are constructed and address two aspects of these models crucial to their use – first, uncertainties arising from the assumptions that go into such models, and, second, decisions using the probabilistic outputs that these uncertainties generate. It is too easy to see models are providing “black box” definitive answers when in reality they are very crude approximations to reality and the inconvenient truth that we must face up to making decisions under uncertainty. The wise man or woman will want to know the uncertainties before deciding, whereas the fool will take a simplification and believe it to be the truth. As Richard P Feynman said (Feynman 1974):

“The first principle is that you must not fool yourself and you are the easiest person to fool.”

BUILDING CATASTROPHE MODELLING

The first step is to have a clear understanding of how catastrophe models are constructed.

To illustrate how a catastrophe loss model works, we'll take a simple example of a deterministic scenario and then make it increasingly "probabilistic".



In the diagram above, we represent a severe windstorm scenario (hurricane Andrew) affecting properties in Florida. We know something of its characteristic damage in terms of insurance losses to properties from claims data. We can take those data and summarise them into "damage factors" representing the average damage to different types of property in different locations. For example, in Dade County, we might find residential property on average was damaged to an insurance recovery cost of 5% of insured values. Similarly we might find commercial property was damaged to a cost of 2% of insured values.

Putting these damages by property type and location together forms what is often termed a "damage matrix", being the potential damage of a similar windstorm was it to be repeated. This idea can be analysed at greater granularity in location and property type if the data are available. Following our simple example, then, suppose that an underwriter insured a commercial property in Dade County with estimated reinstatement value of \$1bn for a layer of \$50m excess \$50m, then a 2% loss would be \$20m and there would be no expected loss to the layer. On the other hand, if it had been insured for \$10m excess \$10m then on this simple average damage argument, the scenario windstorm would lead to a total loss to the layer.

It is often over simplistic to suppose that damage averages out, so a more realistic assumption would be to take a range of damage and sample from that distribution then average the results at the end. This is a "stochastic" model and represents the uncertainty of losses given that an event has occurred.

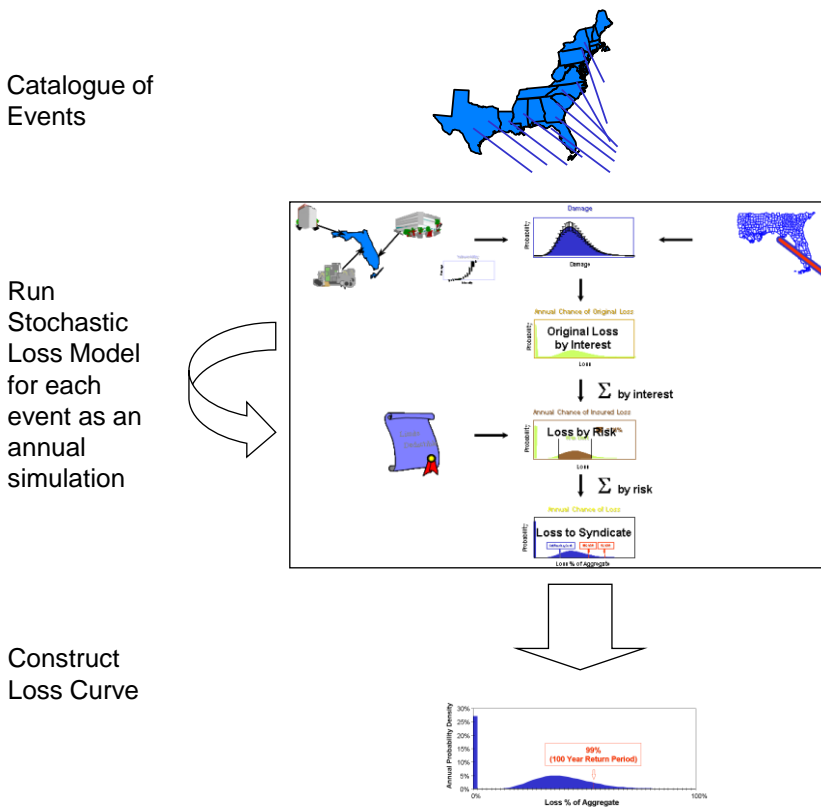
The crucial feature of this model is that it assumes that a distribution of loss can be associated to a given type of property for a given hazard characteristic. Whether parameterized in closed form or represented as an empirical histogram based on actual data, the crucial assumption is that this "damage matrix" (or "vulnerability function") represents real-life response to hazard severity.

CATASTROPHE MODELLING

Author: Dickie Whitaker – Oasis Loss Modelling Framework Ltd.

All catastrophe loss models depend on the justification of this association and how it can be defined in relation to property characteristics often expressed in terms of “primary modifiers” such as construction, occupancy, number of stories, age built (mostly because of Building Codes), together with “secondary modifiers” for particular perils such as roof construction for Windstorm where water damage is often an important factor in loss. We will return to this assumption later when analysing uncertainties in catastrophe models.

What happens next is that many scenarios are conjectured based on theories of the natural phenomenon and historic evidence in order to generate a catalogue of scenarios which can be sampled in order to generate simulations of a year’s events:



In this example the entirety of the sampled set of points from the damageability functions is passed all the way through to the output event/loss table.

This selection of events can allow for correlation – sometimes termed “clustering” - of events as appropriate and choices of different scenarios such as “near-term” event sets for windstorm to reflect beliefs in increased frequency of extreme events due to, say, global warming. Putting these together, then, we end up with a sampling large “event loss” dataset of ground-up losses to which insurance coverage and participations can be applied to estimate the losses to the insurer. In our simple example of one property and \$50m excess \$50m we would get, say, 10,000 years of events, small and large, and some would lead to losses and most to no loss. When expressed in terms of distributions of the largest loss in a year (so-called “Occurrence”) or total loss in a year (so-called “Aggregate”), then statistics relating to average losses, variances, or chances of losses exceeding given values can be calculated. These are generally nowadays expressed as “Exceedance Probability” or “Value at Risk” distributions.

COMPONENTS OF A CATASTROPHE MODEL

Putting these ideas together, we can characterise a catastrophe model as basically having three distinct sections:

1. *hazard module* (science module), which looks at the physical characteristics of potential disasters and their frequency
2. *vulnerability module* (engineering or damage module), which assesses the vulnerability, or damageability, of buildings and their contents and
3. *financial module* (insurance module), which calculates losses to insurance policies or similar financial instruments to protect against risk.

Hazard module

The purpose of this module is to produce simplified representations of the complex physical characteristics that are being modelled. For a US hurricane model, for instance, the first step would be to collate available historical hurricane data for the area under consideration. The USA holds data (of varying quality) for approximately the last 100 years. The coastline is divided into segments known for each historical event and these will range from 10 to 35 miles on average depending on the model. Input parameters will be collated for each hurricane modelled. The data are aligned within the model and probability distributions are then fitted to these parameters. Historical event data is statistically insufficient however to fully populate models so sampling from the probability distributions within the fitted parameters is used to produce 'synthetic' events. These synthetic events create the event database and are the backbone of the hazard component of catastrophe models. The size of the stochastic event sets required to construct a model varies by peril and by country but starts at a few thousand examples for a simple model in a small country up to half a million or more for earthquake or windstorm in the USA.

Input parameters vary widely from model to model but for a hurricane would generally include:

- forward speed (the speed over land)
- radius to maximum wind speed (distance between the eye of the hurricane and the maximum wind)
- central pressure
- size, usually measured on the Saffir-Simpson Intensity (SSI) scale: this integer-valued scale ranges from 1 to 5 and depends on the maximum wind speed and pressure
- landfall location (where offshore hurricanes hit the coast)
- peak wind speed
- direction of the hurricane track
- wind profile (asymmetric due to Coriolis force caused by the rotation of the Earth)
- surface roughness (used to estimate the slowing effect of the terrain on the forward speed of the hurricane)

The same collation principles apply to the modelling of other perils such as earthquakes, tornadoes, tsunami, flood etc. An earthquake catastrophe model might require the following:

- seismic fault geometry (location, depth, dipping angle)
- fault type
- frequency and severity of events on a fault (earthquake severity is measured by various types of magnitude; body-wave, surface-wave, moment etc.) The moment magnitude scale (MMS, M_w or M), which succeeded the 1930s Richter scale, is the most objective scale since it is based on the seismic energy released during an earthquake

Flood models utilise similar data sets but are considerably more complex owing to the nature of the risk. Tidal or fluvial event emphasis has given way to surface water flood models and hydrologic modelling in both 1D or 2D has been developed. Observable data is obtained from the likes of the U.S. Geological Survey, hydroelectric companies and automated weather sensors and combined with quantitative precipitation forecast data. Flood is not covered by the insurance industry in many countries and so there are relatively fewer stochastic flood than other catastrophe models available.

This hazard module within the catastrophe model gives the '*site hazard*' (the wind speed for a hurricane or other wind, or intensity of shaking for an earthquake for instance) at the sites under consideration.

There are a number of measures of site hazard *intensity*, especially for earthquake peril. In the case of an earthquake the simplest one, which uses personal reports and observations, is the **Modified Mercalli Intensity (MMI)**. This integer-valued scale (from I to XII) is therefore descriptive and not quantitative. MMI = I for example is described as '*Not felt except by a very few under especially favourable conditions*' and MMI = XII (the most severe) as '*Damage total. Lines of sight and level are distorted. Objects thrown into the air*'. A more efficient way of measuring site intensity is **Peak Ground Acceleration (PGA)**, measured by instruments such as accelerographs, which is not a measure of the total energy of an earthquake but rather of how hard the ground shakes in a specific geographic area. PGA measurements generally parallel the MMI scale.

A more useful hazard intensity measure still is **Spectral Acceleration (SA)**, measured in g (the acceleration due to Earth's gravity, equivalent to g-force). This takes into account the amplification or weakening of vibrations caused by both the soil structure and the building characteristics at the site, giving a closer approximation to the *actual* motion of a building in an earthquake. This measure is the hardest to estimate and of course there is little historical data but catastrophe modellers increasingly use it in site hazard analysis.

A combination of these measures is then used to estimate earthquake site hazard intensity using '*attenuation functions*' which define the relationships between the earthquake size (magnitude), distance from rupture to the site and the soil structure found at the site.

Modellers will also consider the various earthquake-induced phenomena that cause losses and which are part of seismic hazard:

liquefaction: the liquid-like behaviour of water-saturated unconsolidated soils leading to loss of their bearing strength. Buildings situated on such soils can topple or sink in a liquefaction event

landslide/landslip: a deep failure of slopes, shallow debris flows and rockfalls caused by a loss of soil strength along the sliding surface

fire following: fire triggered by broken gas lines during an earthquake that spreads quickly because of the earthquake-caused disruption to the emergency response system (water supply, transportation etc.)

Vulnerability module

The likely amount of catastrophic damage at a specific site can be estimated if we are in possession of certain facts. In the case of buildings, construction information such as design, materials, quality of construction and inspection etc. will give approximate figures depending on the severity of the catastrophe. This information is often limited however and therefore assumptions become a necessity. There are also many other random variables to consider, such as whether the building has had a change of use. An office or industrial property may have been retrofitted for use as a residential building. Load-bearing walls may have been interfered with and the contents will certainly be different, all of which change the behaviour of a building.

There are practical limitations however on the resolution of data required by each model to ensure a reasonable assessment of the risk. In practice, CRESTA zone data (an internationally recognised method of accumulating exposures - generally along recognised political boundaries) is a practical minimum ex-USA and county or zip code a minimum within the USA. Data that is coarser than this will cause the uncertainty of the results to be considerably increased, possibly to a level that might be considered unusable. Industry accord that location data should be kept at zip/postcode level has led to commensurate improvements in results.

It is not only the construction value of a risk however that insurers and reinsurers need to know in the event of a catastrophe, but also the replacement value. The rebuild costs can, and often do, increase following a disaster. Some policies include business interruption alongside contents and the structure itself and it is possible that, even if a building has no actual structural damage, the loss due to business interruption or contingency business interruption could be a total loss. This vulnerability - the 'unknown unknowns' - has to be taken into account by the modeller and assumptions made to cover them as comprehensively as possible.

Three types of information in combination are considered:

1. *Empirical information.* For a considerable number of years now the scientific and catastrophe modelling community has collected detailed data following catastrophic events. To be of use, information about factors such as wind load or intensity, and detailed information on damage and building type must be collated. This level of data gathering has considerable value as reduces some of the uncertainty shown in the more theoretical approaches. Empirical information however often unfortunately consists of claims data 'contaminated' by the application of limits and deductibles and it can be difficult therefore to adequately extract the damage values.
2. *Engineering consensus.* Expert opinion is a mainstay in the production of vulnerability functions. One main source for the derivation of earthquake vulnerability curves is the non-profit research organisation the Applied Technology Council, established by the Structural Engineers Association, in California which is one of the few organisations to provide generic information on this type of damage.
3. *Engineering reliability analysis.* Computer-aided design (CAD) models analyse the exact components of buildings in order to ascertain the probabilities of various failure modes under specific circumstances. From the theoretical perspective this analysis is the most sound; however, the resulting fragility curves still have to be calibrated using *actual* loss data, which in many cases is insufficient. Another drawback is that it estimates the damage of an 'average' structure and the construction of the buildings in question may be better or worse than average. This is a relatively expensive method and so is mostly used for complex commercial or industrial buildings.

Financial module

This last component translates physical damage costs into real-time, actual financial loss figures and is therefore the module of most interest to the insurance market. Insured losses are assessed and estimated by applying local policy conditions such as deductibles, attachment points and limits to the loss figures. Some models have the additional ability to calculate the effects of a number of types of reinsurance, including excess of loss, facultative and some types of proportional.

MODELLING WITH UNCERTAINTY

Probably the single most important factor in the use of catastrophe models is an understanding and recognition of the uncertainty inherent in the modelling and how this uncertainty is handled.

In many areas of knowledge the certainties of the 20th century have given way to a 21st century recognition that we may *not* be able to know everything – it is the ‘age of recognised uncertainty’. David Spiegelhalter (Spiegelhalter and Riesch 2011) stated when commenting on risk:

- *‘Seek transparency and ease of interrogation of any model, with clear expression of the provenance of assumptions*
- *Communicate the estimates with humility, communicate the uncertainty with confidence*
- *Fully acknowledge the role of judgement’*

Models are only representations of a problem for which there are inputs and rules with which to drive useful output metrics. For catastrophe risk assessment these metrics include average annual losses (the ‘technical rate’) for pricing and a percentile of annual losses needed for capital reserves as determined by, say, a regulator. European regulatory legislation enshrined in solvency II for instance requires 99.5% of predicted annual losses to be held as capital. Any given model takes inputs of data and assumptions and this gives the potential for uncertainty in three areas: *the data, the model itself and what the model doesn’t cover*. It is wise therefore to be cautious about expecting models to give reliable answers. A particularly harmful instance of computer model use was for rating and correlating financial instruments based on subprime mortgages, which promoted a massive bubble of debt and which led to the financial crisis (see Montross 2010).

Sources of uncertainty in catastrophe loss models

It is worth bearing in mind that a catastrophe model is a primitive fixed representation of a peril. Maybe in thirty years’ time we will have high-resolution dynamic representations of windstorms and earthquakes and floods modelled in real time, but the present state of the art is much cruder, it is a static representation of the largest intensity of a peril which is assumed to drive damage. Events are represented by a “footprint” of some measure(s) of intensity such as maximum gust speed for windstorm, peak acceleration for earthquake, or maximum depth for flood. Catastrophe models represent the many uncertainties in the peril intensity footprints and responding vulnerabilities through probability distributions.

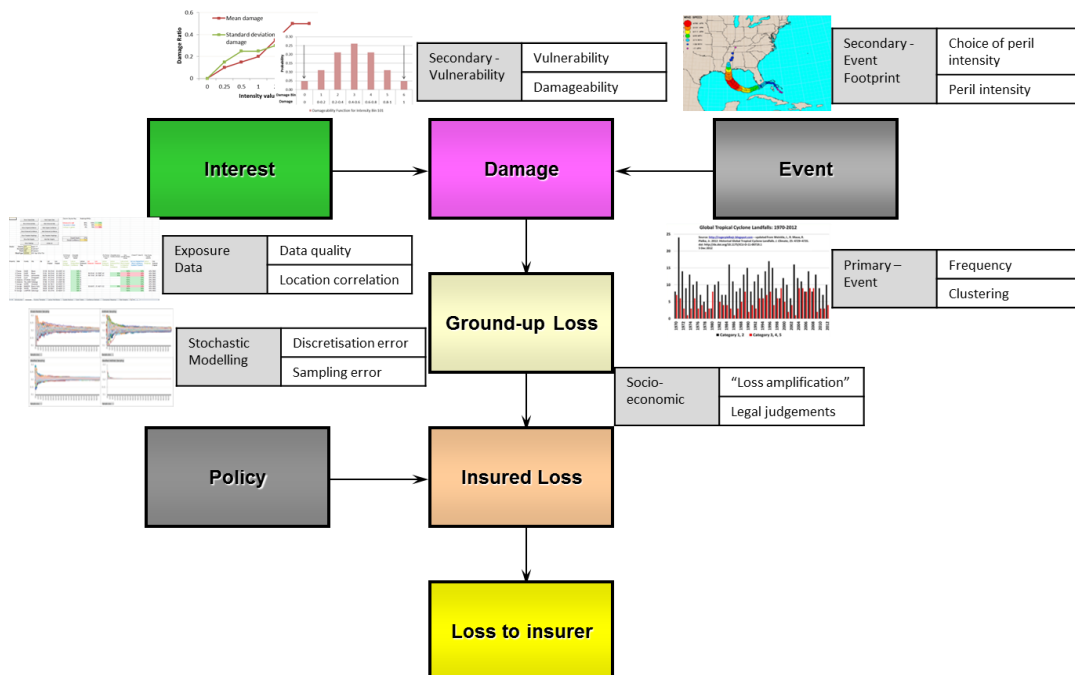
Driving of damage from peril intensity maxima is a causal presumption which may not be true as the damage may not have been due to the metric being used (for example if the damage was caused by water speed or pollution rather than depth). In practice many models use correlations of intensity to loss (e.g. measured by damage ratios), which may not be causally originated or at least not directly.

To varying extents these probability distributions reflect intrinsic uncertainties in the science or engineering (so-called aleatory uncertainty) and our knowledge of the facts that are used in the model (so-called epistemic uncertainty). To some extent this distinction is unhelpful as these two roots of uncertainty usually overlap as we have limited data with which to characterise the intensity or damage probability distributions which are themselves intrinsically uncertain.

There are many sources of uncertainty within catastrophe loss models, and categorising them is somewhat arbitrary, but a suggested taxonomy is tabled below:

Area	Uncertainty	Description	Type of Uncertainty
Primary – Event	Frequency	How often does this event occur?	Model – Parameter Risk
	Clustering	How correlated (or independent) are events in time?	Model – Parameter Risk
Secondary - Event Footprint	Choice of peril intensity	Which measure of intensity is appropriate (if any) as a driver of damage?	Model – Parameter Risk
	Peril intensity	Given the measure, how uncertain is the intensity?	Model – Parameter Risk
Secondary – Vulnerability	Vulnerability	How does vulnerability vary by type of property?	Model – Parameter Risk
	Damageability	Given a property suffers the peril intensity, how uncertain is the damage?	Model – Parameter Risk
Exposure Data	Data quality	How accurate and precise are the property data?	Data (general)
	Location correlation	How does damage uncertainty correlate due to location?	Data – Wrong location
Stochastic Modelling	Discretisation error	What errors are introduced due to discretisation (choice of probability bins)?	Model - Calculation Error
	Sampling error	What confidence level is associated to random sampling?	Model – Calculation Error
Socio-economic	“Loss amplification”	How uncertain are loss cost estimates as a function of size of catastrophe?	Model – Model Inadequacy
	Legal judgements	How uncertain are rulings of liability?	Unmodelled – Wordings

The diagram below shows where each of these sources of uncertainty applies in the loss modelling process:



It is of particular interest that many sources of uncertainty lie outside the model:

Model Uncertainty	
Model Inadequacy	
Model Risk	
Parameter Risk	
Calculation Error	

Data Uncertainty	
Wrong location	
Understated Values	
Building Risk profile	
Schedule out of date	

Unmodelled Uncertainty	
Secondary Perils	
Contingent BI	
Policy Wordings	
Loss Adjustment Expenses	

Model Uncertainty covers the uncertainties that exist within the scope of the model. *Model inadequacy* represents the unsuitability of the underlying physical hazard and vulnerability models, as well as the effect of ignoring secondary perils such as demand surge (sometimes termed Loss Amplification) and fire following earthquake, or secondary coverage such as business interruption (BI). *Model risk* is the risk that the particular model is wrong or provides an incomplete and misleading picture. *Parameter risk* is the risk that the model’s parameters (often expressing assumptions or calibration factors) are wrong, and *calculation error* the risk that the model has been insufficiently discretised or sampled.

Data Uncertainty in this case applies to information describing the insured exposures which are typically properties and structure. This can cover *location* which is generally the most important attribute of an interest (such as a building), and *understated values* where, for various reasons (including *out of date* schedules), the insured values do not reflect the indemnity that would be incurred. Data uncertainty can also include the *risk profile* of a building, i.e. the attributes that characterise propensity to damage - “primary modifiers” such as construction and occupancy, plus “secondary modifiers” which are other attributes that affect damage such as roof geometry. And of course, any data may be *out of date*, as is so often discovered when a catastrophe occurs and a building is damaged which is not on the schedule used for underwriting.

Unmodelled Uncertainty is set apart from the data and the model as it relates to the part of the representation of the problem that is outside the domain of the model. It covers a range of factors, which include *secondary perils* such as business interruption, demand surge, and fire following (although these can sometimes be included in the model and any weaknesses would then fall under “Model Inadequacy”). A particular cause of unmodelled loss in recent years has been *contingent business interruption* losses caused by supplier failure. The interpretation of *policy wordings* by the relevant jurisdiction can also materially affect losses, and a component of insurance cost generally ignored by modellers is the expenses and fees from adjusters and lawyers and other third parties involved in a claim, termed *loss adjustment expenses*. In some cases these can be a material overhead on top of the pure indemnity cost.

All of these factors go into the risk assessment process and contribute to cost and cost uncertainty, so it is worth checking how many of the uncertainty factors get addressed when looking at a particular model.


The principal concerns are model risk and the unmodelled uncertainties. Model risk is a problem to which solutions are being developed. Unmodelled uncertainties need to be allowed for with judgement and subjective allowances, and some may in due course be brought under the modelled umbrella. When handling model uncertainty safety margins are added to compensate for the

"unknown unknowns". The New Zealand and Japan earthquakes – in which the impact of un-modelled perils such as liquefaction, aftershocks and tsunami and backwash proved so devastating - models often miss sometimes large portions of the overall loss.

The problem common to all modelling, though, is the quality of the source data. “Garbage In Garbage Out” data uncertainties are crucial, though they do depend on the model. Incomplete geocoding or geospatial data, or incorrect data – where there have been errors in collection or where there are inherent errors in the data – are contributory causes to uncertainty. The event catalogues upon which each model is based will generally be derived from published data available in the public domain. For this reason there is usually a reasonable consistency between different models. New research that refines these catalogues is continually ongoing and is incorporated on a regular basis. However, for some geographical areas, in particular those that are relatively inactive, such as the New Madrid area in the USA, expert judgement must play a significant role. In another instance, for an aggregate model at US County level using occupancy and high-level construction type, it is hardly necessary to know the longitude/latitude or the roof geometry as the model is insensitive to these attributes of a property. Arguably the issue with data is the sensitivity of the model output to the granularity (precision) and accuracy of the data. We currently lack methodologies to deal with these uncertainties.

Second-order Risk

Probability, the great deceiver! Many models are probabilistic, which is to say that they express their output in the form of chances of outcomes. For instance, catastrophe loss models provide the probability distribution of annual losses. This all sounds most confidence-building as the use of probability seems to express uncertainty in a complete and mathematically rigorous way. Yet it is something of a magician’s illusion, as a weakness of probabilities is that they can be combined to wash out important underlying uncertainties. Guy Carpenter illustrated this problem in a booklet on uncertainties in catastrophe loss modelling (Guy Carpenter 2011):



Source: RMS and Guy Carpenter

The feature this chart exhibits - the probability of probabilities - has a long history in the literature on risk, dating back to 1981 when Stanley Kaplan and John Garrick published “On the Quantification of Risk” (Kaplan 1981).

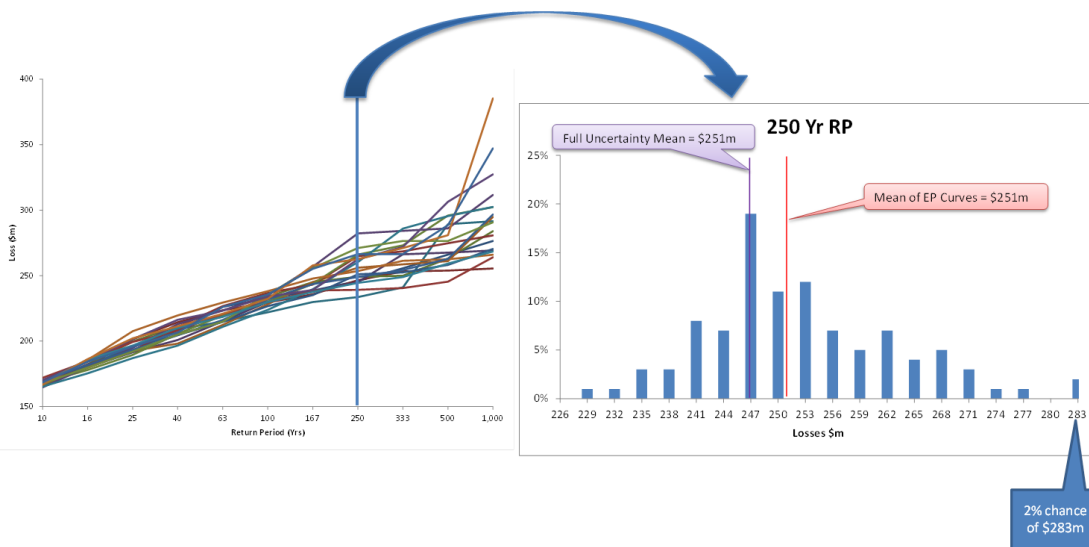
Garrick has more recently produced a book on catastrophe risk (Garrick 2008), of particular interest to loss modellers as it comes from outside insurance, where he works through the scenario-based and frequency methodology, all in the light of the second-order uncertainties he and Kaplan identified nearly 30 years before.

CATASTROPHE MODELLING

Author: Dickie Whitaker – Oasis Loss Modelling Framework Ltd.

As an illustration, suppose a broker presents two risks with the same loss probability (EP) curves and statistics to an underwriter. One may have a very small spread of fuzziness around the average annual loss and the 99.5% annual loss percentile, and few “unmodelled” elements. The other might have large spreads and be highly sensitive to data accuracy and model parameters, or have many “unmodelled” elements of risk. The underwriter who tracks down these uncertainties can make a better decision and use judgement more effectively than by using “black box” statistics.

The diagram below, taken from a test run of the open source software Oasis, shows the “wheatsheaf” of EP curves from the loss model and how this distribution gives a spread of VaR at 250 years (99.6% VaR) where the upper 2% confidence bound is \$30m greater than the mean:



Please note, though that this is very much only the stochastic sampling error element of uncertainty as per the taxonomy of uncertainty.

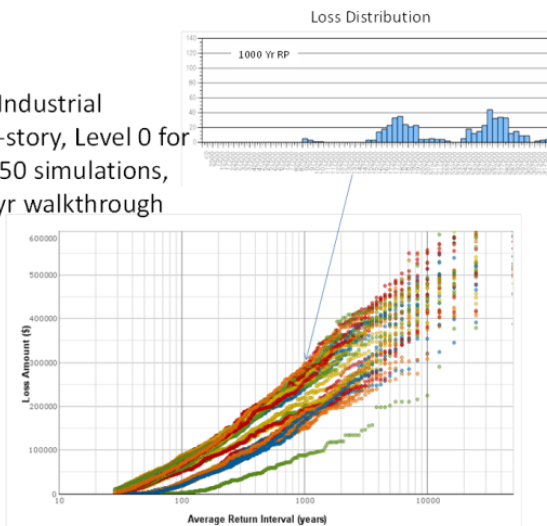
Example – the effect of Data Quality

Lack of knowledge is well-represented using the spread of EP Curves. As a simple example, ImageCat’s Seismicat California earthquake model gives multi-modal large spreads when there is very limited information on construction available:



Occupancy - Industrial
L.A., 1991, 2-story, Level 0 for all subtypes; 50 simulations, with 50,000-yr walkthrough each

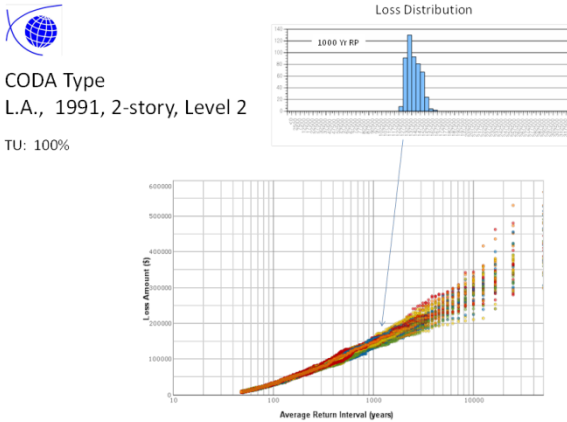
- TU: 40%
- BSW_C: 20%
- BSW_M: 20%
- W_E: 10%
- CBF: 7%
- LS: 3%



CATASTROPHE MODELLING

Author: Dickie Whitaker – Oasis Loss Modelling Framework Ltd.

But when the construction type is known (e.g. by an engineering inspection) as “Tilt Up” we get a very different result (Huyck 2013):



CODA Type
L.A., 1991, 2-story, Level 2
TU: 100%

The underwriter with the better data would thereby be able to make a more confident and different judgement on price and capital required.

THE FUTURE

We are entering a new era in the use of models, one in which the convenient certainties of yesteryear have been seen in many cases to not only be delusional but also hide crucial information on the uncertainty of our knowledge. Information which can help us appreciate what we know about the risk and inform our pricing and setting of adequate capital.

This is all being made possible by technology which enables us to see these “fingerprints of risk”. Coupled with a commercial and regulatory environment which requires “your own view of risk” then the future, whilst demanding more of us, will give a deeper and more profitable view of catastrophe risk.

Acknowledgements

Grateful thanks to Peter Taylor for his thoughts on uncertainty and the Oasis Loss Modelling Framework.

References

Clark K (2011) *P/C Industry Depends Too Much on Catastrophe Models, Says Pioneer Clark*, Insurance Journal April 14th, 2011

Feynman, Richard P (1974) *Cargo Cult Science* (e.g. http://neurotheory.columbia.edu/~ken/cargo_cult.html)

Garrick B J (2008) *Quantifying and Controlling Catastrophic Risks*, Academic Press

Guy Carpenter (2011) *Managing Catastrophe Model Uncertainty issues and challenges*, Guy Carpenter December 2011

Huyck (2013) *Private Communication*.

Kaplan S Garrick B J (1981) *On the Quantitative Definition of Risk*, Risk Analysis Vol 1 No 1 1981

Kay J (2010) *Obliquity*, Profile Books

Montross F (2010) *Model Madness*, GenRe

Spiegelhalter D J and Riesch H (2011) *Don't know, can't know: embracing deeper uncertainties when analysing risks*, Phil. Trans. R. Soc. A (2011) 369, 4730–4750