**Risk Management**

April 2016
Issue 35

# What Is a CAT Model?

**By Theresa Krebs**

> Below, a vast database of economic exposure... in the middle, 10,000 hurricanes... and above, structural engineering curves of vulnerability to hazard.

I. Introduction: The Three-Decker Sandwich
II. Exposure
III. Hazard
IV. Vulnerability
V. The Financial Module
VI. Conclusion

## INTRODUCTION

An influential mentor at Risk Management Solutions once described a CAT model as a 'three-decker sandwich.' In the lowest layer, lies a vast database of economic exposures, including property values by locations and by coverage, along with attributes such as sector and occupancy. The middle decker is a powerful stochastic model of the natural hazard, which represents tens of thousands of individual geographic event footprints, each assigned a frequency of occurrence and each defined in terms of severity of the hazard at each location. The top decker is vulnerability, which begins with the physical fragility of a certain building construction to earthquakes, flood, and hurricanes, but also has to include secondary societal concerns such as business interruption and the cost inflation of repairs when there is a shortage of contract labor following a major catastrophe. In the end, once we've constructed our three-decker sandwich, comes the financial module, which overlays policy terms and conditions onto the 'ground-up' financial output.

So what is a CAT model? It may be best to describe it in terms of its inputs and outputs.

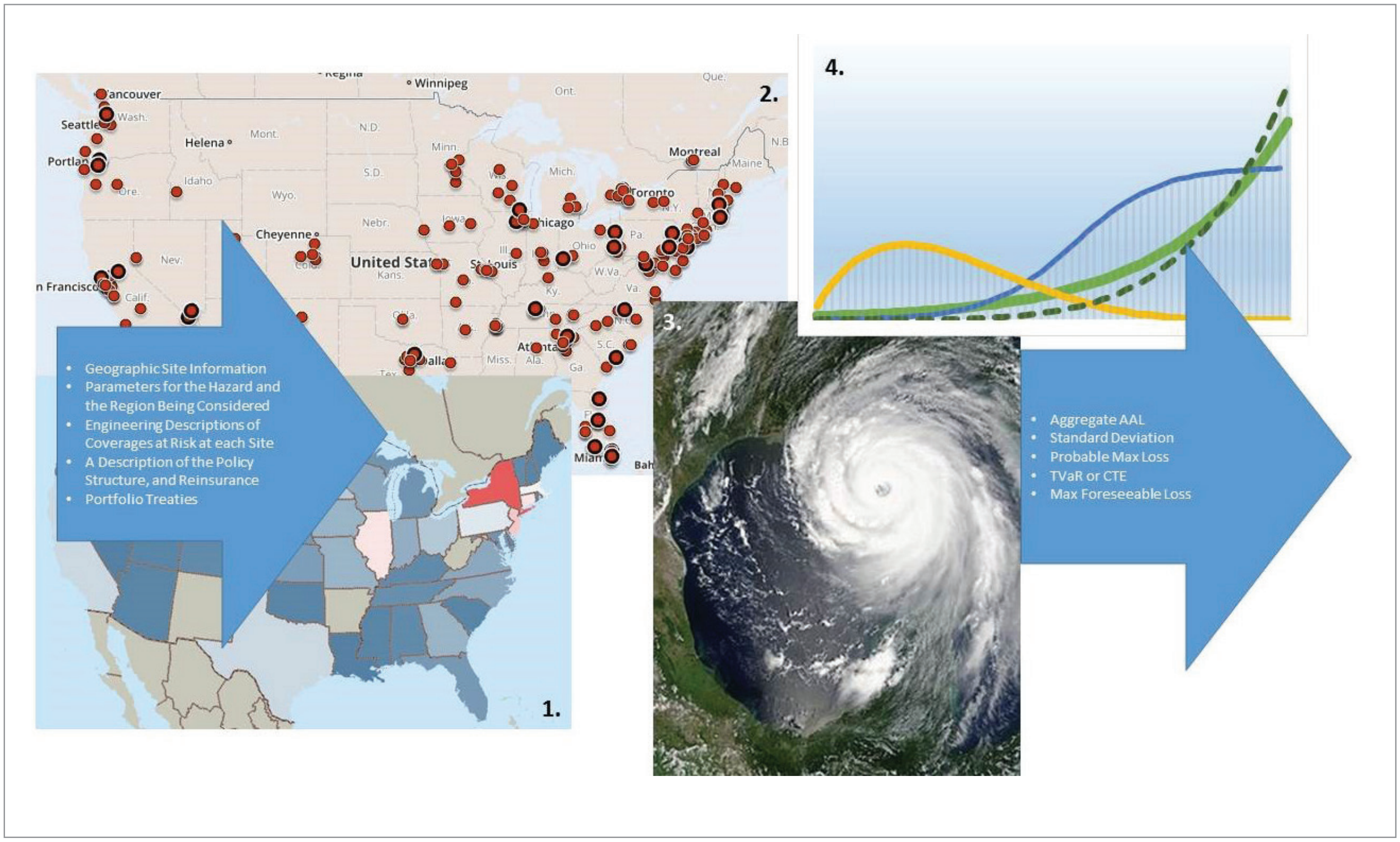

In the flow chart at left, one begins with the underlying economic exposure in image #1, to develop assumptions about geographic distributions of building characteristics. In image #2, we have geocoded the risks associated with the U.S. portion of a large multinational account, and we know the building characteristics and the economic values at each latitude and longitude, as well as any location-specific terms and conditions. In image #3, we evoke the hazard model with an image of Hurricane Katrina. And in image #4, we combine elements of the vulnerability model with examples of model output, both of which are further explained below.

The Probable Maximum Loss (PMLs) is a point along the Exceedance Probability Curve (EP Curve) in which individual scenario losses are ranked and plotted in terms of their cumulative frequency. The EP curve can be plotted both as a loss in an individual event and the aggregate loss over some time period, such as a year. In the aggregate loss there is the question as to what assumption to make about event clustering—whether to treat it statistically as a in a negative binomial assumption or whether to apply a model of spatiotemporal clustering.

PMLs correspond to 'average return periods.' If I talk about the 100-year PML, I do not mean that a given severity event will happen every 100 years. Rather, what I mean that this severity of loss has a 1 percent chance of happening this year, and next year it will again have another 1 percent chance of occurrence.

The Tail Value at Risk is a conditional integration over the EP curve past a certain return period. So the 100-year TCEAEP is the tail of the aggregate EP curve past the 100-year return period.

When CAT models are used to price accounts, the marginal TCEAEP can tell you what the account will do to the portfolio's tail, and this can be considered a risk load.

Now, let's return to our three-decker sandwich.

## EXPOSURE

Each model vendor develops a proprietary database of the total and insured building stock, based on a combination of engineering expertise, economics expertise, and data that is purchased from the market. In some underdeveloped markets, without proper data collection capabilities, this exposure data may also be developed based on assumptions about GDP and insurance penetration that come into play. This database helps to calibrate the model to industry loss experience in the final stages of model development. Detailed information about building inventory and construction styles can also be used to supplement data when an underwriter doesn't have information about a risk. Given an occupancy such as a hospital, for example, within a specific territory, a certain percentage will be reinforced concrete, reinforced masonry, steel, or even wood frame. A weighted average of vulnerability curves for each construction type can be applied to the coverage losses.

This concept of conditional probabilities and weighted averages can apply to a number of elements that are not available, such as year built, number of stories, and even "secondary modifiers" such as the type of cladding on a building. This enables an insurer to obtain an overview assessment of an entire portfolio's losses in the face of uncertainty. When writing many risks, this can take advantage of averaging across the portfolio. However, a company that wants to carefully select the better risks will need to collect better data at the point of underwriting.

## HAZARD

The first step in modeling a risk begins with accurate geocoding. Given a complete street address, one can find a precise latitude and longitude. The location is assigned a grid cell, and that grid cell has an ID in the stochastic database. This allows for the retrieval of hazard intensity and frequency information for this specific, individual location.

Beyond capturing the properties of individual simulated events, there are questions around how the occurrence of one event affects the potential for another event in the same contract period. Where events are considered completely independent of one another, the Poisson distribution can be applied, in which each event has a unique frequency and is statistically independent of every other event. Where the occurrence of one event increases the potential for there being other events in that season, then the negative binomial distribution might be more appropriate. For earthquakes, the model of time-dependency can be applied where there is enough information about mean recurrence intervals and the time that has elapsed since the last rupture of a particular fault. Also stress transfer models may be appropriate for when the occurrence of one earthquake alters the probability of other earthquakes in the surrounding region. The choice of the model for event occurrence varies by model vendor.

In addition to ambient geophysical conditions such as wind speed, flood depth, and ground shaking, there is another type of hazard that is commonly referred as site hazard. In the case of earthquake, site hazard includes parameters such as the soil type, liquefaction potential and landslide potential at that location. In the case of windstorm, site hazard includes the effects of upwind surface roughness potential. Elevation will also be very critical for modeling the impact of storm surge.

The actual procedures for generating all event footprints, ensuring they span the full range of possible occurrence and identifying the probability of each simulation are generally considered proprietary by the model vendors. One attribute that most models share, however, is called secondary uncertainty,

reflecting the uncertainty in loss estimation, whether related to missing knowledge about what is covered, or imperfect data around the vulnerability, or simply the inherent variability in the way a building has been constructed or the properties of the hazard at that location. Given one earthquake, and two adjacent contemporary buildings of identical construction and style, one may collapse while the other remains standing. In truth, this may come down to resolution and detail and our ability to capture differences in liquefaction potential across a distance measured in feet, but in practice this must be regarded as randomness in the model. As such, it is considered a form of uncertainty, a "distribution around the distribution," and it is used to load the results of the model, increasing the standard deviation and fattening the tail. A key differential across modeling companies comes from whether this secondary uncertainty has the potential to be correlated across the various simulated events in the model. Where the uncertainty is epistemic—i.e., related to our underlying knowledge—it can be expected to correlate across all the events. However some model vendors do not fully capture this correlation.

Each stochastic model is a hybrid of statistical and deterministic methods. At the heart of the model, there is almost always some form of parameterization. Even if we had enough historical data to initialize the model 10,000 times, that data would contain noise, and the noise in the initial conditions would cause the model to veer off. Academic scientists get around this using a method called "normal mode initialization." Modelers employ a range of techniques to explore bias in the output of climate models for example, testing the outputs against actual data—for example on the wind speeds of storms across Europe over the past fifty years. However for tropical cyclones, climate models have not achieved sufficient resolution, and parametric models may be applied, based on the copious information on past track behavior.

Typically, the models are at least somewhat parameterized, and they contain fewer degrees of freedom than the natural data. This does and does not carry implications for risk loading. For example, take the log-distribution of Rmax, the radius to maximum winds of a hurricane from the center. It is normally distributed. Each vendor can decide at what percentile to cut off the distribution. However, this has less implications for the tail than you would think—the further you get out into the normal tail, the lower the event rate and the less it contributes to the whole, so it doesn't necessarily fatten the loss tail nearly as much as you would think. So while companies do load the models for uncertainty in data quality, they generally see no need to compensate for the bounded scatter in the distribution.

Hazard modules vary greatly in their sophistication between vendors. For example, storm surge models can range from fully time-stepping, numerically discrete solutions of the high-viscosity (and therefore highly nonlinear) Navier-Stokes equations, or they may be a simple lookup table that relates the angle of attack of the hurricane, to the continental shoreline, to its minimum pressure and maximum offshore wind speeds. Both actuaries and scientists need to exercise astute caution in interpreting the output of the various CAT models, and understand their strengths and weaknesses.

## VULNERABILITY

While every module of a CAT model is uncertain, the vulnerability module can be at once the most uncertain and the most influential. At the heart of the vulnerability model is the "vulnerability curve," a classically S-shaped curve that is bounded by a standard deviation, and which relates hazard on the X-axis to a ground-up "damage ratio" on the Y-axis.

The standard deviation around the vulnerability curve is the essence of secondary uncertainty. Each vendor assigns a proprietary distribution around that standard deviation and integrates over the distribution. The vulnerability team carries the greater burden of uncertainty and expert judgment.

Vulnerability requires a strong understanding of the performance of different building types under a range of loads. Wood frame, for example, performs extremely well in earthquakes, while reinforced concrete, built to code can also be very resistant to earthquake shaking, although not when built without proper reinforcing, or the attentions of an engineer. Steel performs relatively well, while masonry, especially unreinforced masonry, performs terribly.

In a hurricane, wood frame performs terribly. Masonry is prone to water damage, and the stiffness of the walls can cause tension between the walls and the roof. Concrete, including reinforced concrete, performs extremely well. Steel performs well, but glass surfaces and certain forms of cladding do not. The behavior of the roof—its configuration, its attachment to the frame, the materials of which it is made, and whether the frame forgives the flexing of the roof—all can be strong determinants of the performance of a building in a hurricane.

## THE FINANCIAL MODULE

Once the ground-up loss is calculated for a single location, location-level policy terms such as site limits and site deductibles can be applied to achieve a gross location-level loss that is only net of the location-level terms. This in turn can be aggregated over the entire location schedule underneath the layer for each stochastic event. Lastly, policy terms and conditions are applied. The gross and net-of-fac losses can be re-allocated to the location level using deconvolution, a task that is made much easier by the assumption of a Poisson distribution.

## CONCLUSION

CAT models are complex, and it often can be difficult to determine what is driving a large modeled loss or a change upon renewal. The levers are many. However, modelers strive to be rational, a-political, neutral arbiters of the true financial loss, indifferent to hard or soft markets. To this end, each model is steadily becoming more state-of-the-art. Over time, each geophysical model becomes less and less parameterized—in fact, some storm surge models are fully dynamic, time-stepping, academic models that are run on Linux clusters to develop the stochastic database. CAT models are getting better every day, and as they improve, uncertainty decreases. ■

Theresa Krebs is CAT risk manager, North American Property & Specialty Lines at Chubb Global Property and Specialty in Philadelphia, Penn. She can be reached at *theresa.krebs@chubb.com.*