RMS WHITE PAPER

HIGH-DEFINITION SIMULATION

Empowering Models and Their Users

RMS™

Executive Summary

Current approaches to catastrophe modeling have served the industry well. It is becoming clear, however, that the industry needs a more sophisticated loss modeling approach that more accurately captures loss uncertainty and correlation, temporal event dependencies, and a wider variety of insurance contracts. In the RMS(one)™ environment, this is now possible through high-definition (HD) simulation modeling, the latest in RMS' 25-year history of market-leading catastrophe modeling innovations. The HD-simulation framework enables the explicit modeling of clustering, seasonality, uncertainty, correlation, and time-dependent contract terms, and provides the flexibility and transparency to make adjustments reflective of a company's own view of risk.

HD simulation improves the modeling of uncertainty and loss correlation among locations, because these factors are taken into account at the ground-up level, when losses are first sampled. Contract terms are subsequently applied directly to the sampled losses, instead of to a mean loss with an assumed beta distribution. This eliminates the need for approximations introduced by the RiskLink approach while increasing the ease and transparency with which the RMS(one) financial model operates. Simulation also allows for a more sophisticated correlation framework, where ground-up correlations are determined and then allowed to propagate naturally through the financial model. Finally, the approach explicitly models zero- and total-damage states.

The new RMS(one) financial model, enabled by the RMS Cloud, can process several hundred thousand simulation periods at once, with the period length dynamically adjusting from months to up to six years based on the treaty terms. This computing power makes simulation modeling possible; if too few simulations are performed, the results may not converge to those of a traditional approach that considers all events simultaneously. HD simulation leads to accurate modeling of a wider variety of policy and treaty terms, including multi-year, second-event, and aggregate limits and deductibles, as well as the incorporation of event clustering and seasonality—whether RMS' view or a company's own view. Results are stored as discrete losses, simplifying grouping, roll-up, and the calculation of loss metrics using RMS Analytics.

Two loss sampling methods will be offered to allow for a broader understanding of uncertainty in a variety of use cases. Period loss table (PLT) mode represents the RMS Reference View for stochastic model runs and is used when the timing of an event is significant (e.g., contracts with hours clauses or aggregate limits, or peril regions where clustering is a significant factor). PLT mode samples from the ground-up event loss distribution each time an event is encountered on a pre-simulated timeline. In contrast, sub-event loss table (SELT) mode samples several points per event, and is used in scenario analyses to understand uncertainty in individual event losses.

HD simulation will be implemented in RMS models with a phased approach. The initial RMS(one) release will utilize simulation to model catastrophe treaties, and users will benefit from the incorporation of a simulated loss output, clustering, seasonality, and the straightforward application of a wider variety of catastrophe treaty terms. The RiskLink financial model methodology is maintained to calculate ground-up, gross, and per-risk treaty loss. By keeping results identical from the ground-up to net pre-cat level, this implementation starts clients on a path to improved risk modeling, but does not introduce inconsistencies in the market between clients on RMS(one) and those still transitioning. HD simulation will first appear in full—at the ground-up level—with the late 2014 release of the RMS® Probabilistic Terrorism Model. All new RMS models, such as Europe flood, U.S. flood, and West Pacific typhoon, will use the HD-simulation approach. Ultimately all RMS models will be upgraded to full simulation.

Catastrophe Modeling: Where We Stand Today

For years, catastrophe models, including RiskLink, have operated using an event-based approach. Every event in a precompiled event set occurs with a given annual frequency and causes a loss with some degree of uncertainty, from which we derive metrics such as the average annual loss and exceedance probability (EP) curves. This has been a useful working approach for the industry; the assumptions it made were necessary given constraints on computing power and storage. However, it is not without limitations.

For example, an event-based approach does not consider the date on which an event occurs; but the event date is critical for hours clauses as well as aggregate and *n*th event terms in policies and treaties. This approach also does not consider event seasonality, which can be a major issue when contracts are in place for more or less than a year. It also assumes that events occur independently—one event's occurrence does not affect the probability of a subsequent event—ignoring event clustering, now known to be a factor for certain perils. Modeled losses for such events are thus not as accurate as they could be.

An event-based approach also introduces mathematical issues related to assumptions for modeling uncertainty and loss correlation. At each financial perspective, a particular correlation weight and uncertainty distribution for event losses is assumed. While appropriate at the ground-up level, continuing to assume this after the application of policies and treaties introduces analytical issues. It also requires the use of calculus and complex analytics in the "back end," giving the process a black-box feel to most users.

Finally, the approach lacks flexibility, as RMS' built-in assumptions on uncertainty distributions and correlation cannot be altered. While these were put in place to reflect RMS' best understanding of the risk, clients currently cannot test these assumptions and cannot implement their own views.

Moving to A Simulation Approach

A new approach, based on sampling events and their losses over simulated time periods rather than precompiled event sets, can solve these issues. Instead of representing uncertainty as a standard deviation around a mean loss, and continuing to recalculate these metrics after each contract is applied, uncertainty is considered once: when a single value—a "loss quantile"—is randomly chosen from the event's loss distribution each time that event is encountered on a simulated timeline of events. Losses to contracts are calculated arithmetically, without the need for complex mathematics. Initiating simulation at the ground-up level eliminates the analytical issues introduced at subsequent financial perspectives, bringing greater transparency to financial model calculations.

Laying events along a pre-simulated event timeline achieves additional benefits, such as the incorporation of new science on event clustering and seasonality. Instead of assuming a constant event frequency, simulation can model periods of low activity, as well as periods of high activity from storms of similar track and intensity. This better captures clustering events like European windstorms Lothar and Martin (1999). Finally, the timeline lets users explicitly model time-dependent contracts without workarounds. This more accurate, flexible, and comprehensible approach has positive implications for pricing risk, managing capacity, and capital modeling.

Given the benefits of full simulation and the limitations of the current event-based approach, RMS will migrate all models to simulation over the coming years. This cannot be done in a single, rapidly executed step because of the impact on certain modeled losses, including the RMS Reference View in peril regions with clustering, and users' modeled losses for portfolios with time-dependent contracts. A gradual migration is intended to make change management easier for RMS clients. RMS is committed to releasing responsible changes to its models, and will only introduce new methodologies after undergoing a thorough change management process, which allows model users to test model results themselves before accepting them as best practice. The phased approach also facilitates clients' varying needs with respect to RMS(one) adoption. As this migration takes place, RMS will ensure that RMS(one) loss results are consistent with those of RiskLink to reasonable levels and timeframes, as described below.

All RMS models will undergo two distinct migration steps. First, at the initial release of RMS(one), all models will use the RiskLink financial model methodology up to the net pre-cat financial perspective; only catastrophe treaties will use simulation. This will allow for consistent results between RiskLink and RMS(one) through the net pre-cat level. Results may differ at the post-cat level, based on the degree to which the peril and treaties are dependent on event timing (e.g., clustering or seasonality of the peril, aggregate treaty terms). However, because incorporating simulation for catastrophe treaties is widespread in the market, and there are many RMS and non-RMS tools available to do this today, RMS does not anticipate that this will lead to issues in the market.

The second step is a full conversion to the HD-simulation approach, where simulation begins at the ground-up level, instead of at the net pre-cat level. Substantial benefits to the financial model will be realized, particularly in the modeling of uncertainty and correlation. The breadth of sensitivity tests available will also expand. The model conversion process will be spread over multiple years, beginning with the late 2014 release of the RMS Probabilistic Terrorism Model and continuing over the following 2–3 years with the release of RMS models for Japan typhoon, Europe flood, U.S. flood, and basin-wide West Pacific typhoon; and earthquake models for North America, Japan, and New Zealand. For a deeper look at how these two steps are implemented and their mathematical benefits, see the Appendix.

Working in a Simulation World

What implications will simulation have on business workflows and results exchange? RMS' goal is to move much of today's Excel- and SQL-intensive EDM and RDM analysis onto RMS(one) as part of a powerful, integrated catastrophe modeling workflow. RMS Analytics, the platform's risk intelligence and visualization capabilities, can be customized to give users the analytics they need as soon as an analysis is run, eliminating manual data interrogation steps.

When users need to perform off-platform data interrogation and exchange results with counterparties, loss metrics will be exported as "results objects" in a variety of formats. Over the coming years, RMS anticipates that the majority of users will have moved away from RDMs and to these new forms of market currency by the time full simulation rolls out. The move to results objects will be accomplished in two phases: when simulation is introduced for catastrophe treaties in translated models, and when HD simulation is fully implemented in new and translated models.

Working with Results: RMS(one) Initial Release

Because simulation will only be utilized for catastrophe treaties in the initial RMS(one) release, present-day workflows will require minimal change: net pre-cat event loss tables and exceedance probability (EP) data can be exported to their familiar tables in RDMs. These metrics, as well as post-cat period loss tables (PLTs) and others can be exported as results objects and stored in a variety of data formats including flat files, Excel files, and proprietary formats such as ACORD XML. Much of this process can be automated with RMS-built APIs. In addition, secure tools will allow users to grant counterparties access to results stored in the RMS Cloud, avoiding the need for downloading and sending results over FTP sites. Those who still use RiskLink can import pre-cat RDMs from counterparties and apply catastrophe treaties manually in RiskLink.

Working with Results: Models with Full HD Simulation

Models run with full simulation in RMS(one) will produce period loss tables (PLTs) or sub-event loss tables (SELTs) at all financial perspectives (see Appendix for details). Users can continue to export PLTs, SELTs, and other metrics as results objects in a variety of formats. PLTs will represent the preferred RMS risk currency because they use a simulation timeline.

HD Simulation: Empowering Models and Their Users

Powered by the RMS Cloud and anchored in simulation, the RMS(one) financial model realizes a revolutionary improvement in the precision and capabilities of catastrophe modeling. A wider variety of contract terms will be handled with greater speed, transparency, flexibility, and accuracy. Risk intelligence and visualization capabilities, provided by RMS Analytics, will allow data interrogation to largely move to the RMS Cloud, eliminating the need for post-analysis querying and manipulation of RDMs to derive certain metrics. Instead, results objects will provide the needed model output for efficient transfer in the market.

Bringing HD simulation to market in a phased approach provides a practical solution for implementing change at a measured, responsible pace. With the initial release of RMS(one), model users can seize the opportunities simulation modeling provides to better understand the intricacies of model assumptions and uncertainty, and ultimately take control of discovering, implementing, and owning their view of risk.

Appendix

Step 1: Simulating Catastrophe Treaties on RMS(one): Initial Release

At its initial release, RMS(one) will use simulation technology to calculate catastrophe treaty loss and the remaining net loss. While RMS(one) will use rebuilt hazard and vulnerability modules, achieving order-of-magnitude speed increases, in order to keep results consistent with RiskLink, the RMS(one) financial model largely retains the methodology present in the RiskLink financial model up to the net pre-cat perspective.

The Process

Using mean damage ratios and their standard deviations from the vulnerability module, the RMS(one) financial model applies policy and per-risk treaty terms to generate an event loss table (ELT) for the net pre-cat financial perspective. Simulation occurs when an ELT is merged with a period event quantile table (PEQT). These tables, listing the period, event ID, event occurrence date, and loss quantile, are pre-generated by RMS based on each event's frequency in the event set. Each time an event is encountered on the PEQT, a sampled loss is drawn from the specified quantile of the event's loss distribution (defined by its mean loss and standard deviation) stored in the ELT. The result is a period loss table (PLT), which lays out the events and losses occurring over a number of simulated years. Once generated for the net pre-cat perspective, catastrophe treaty terms are applied arithmetically directly to the losses. The result is a net post-cat PLT. This process is outlined in Figure 1.

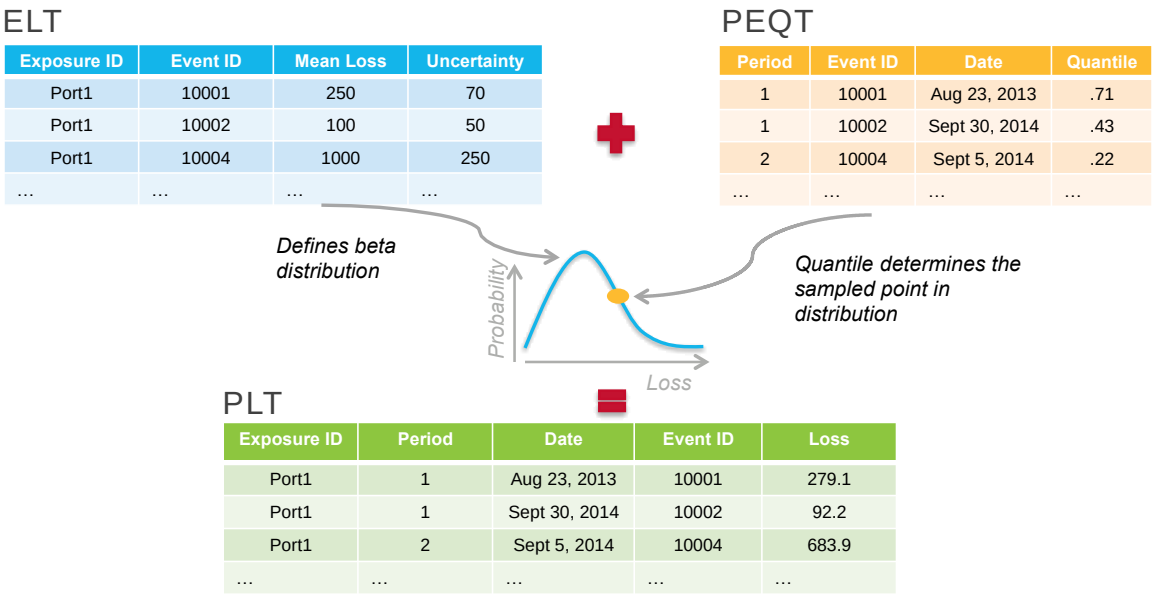

Figure 1: Net pre-cat ELT merged with pre-generated PEQT to generate a net pre-cat PLT. Cat treaty terms are then applied directly to the losses.

Benefits

This methodology in RMS(one) will provide several improvements over the RiskLink financial model:

1. Uncertainty is taken into consideration when losses are first sampled from the severity distribution. Treaty terms are then applied arithmetically directly to the losses. The process is similarly easy when rolling up and grouping losses. In the RiskLink financial model, a mean loss and standard deviation are propagated, and complex analytics are used to determine the financial situation. The new process is quicker to implement and easier for users to understand.
2. The analysis window for each period ranges in length from months to six years; the length is automatically chosen at runtime based on the treaty terms. This flexibility means that treaties with aggregate terms can be accurately modeled, including catastrophe bonds with multi-year terms.
3. The RMS Reference View will consider the necessary number of periods to achieve convergence of results by peril region balancing the need for accuracy with storage and run-time issues. However, the ultimate number chosen is up to the user—model runs can be "fit for purpose." Fewer periods can be used for more frequent perils such as severe convective storms or when "ballpark" figures are sufficient. More periods can be used for infrequent perils or when a high degree of accuracy is necessary, such as with concentrated portfolios or for facultative risk.
4. Clients may specify their own simulation timelines, thus incorporating their own knowledge on event frequency, clustering, and seasonality.

Step 2: Full Simulation

Though catastrophe treaty simulation in the initial RMS(one) release has many benefits, it is only the first step in RMS' transition to a simulation framework—one that uses simulation as the very first step in the financial model. The first model translated will be the RMS Probabilistic Terrorism Model in late 2014; all RMS models will eventually be moved to HD simulation, including any new peril regions such as U.S. flood.

The Process

In HD simulation, sampling occurs at the ground-up level: the input is ground-up loss to a location coverage (e.g., building, contents) by an event. From here, two sampling modes can be used: PLT mode or SELT mode.

PLT mode is very similar to the cat treaty simulation detailed in Step 1 above, except that it operates on ground-up mean damage ratios and standard deviations (instead of net pre-cat ELTs) and the quantile selected is chosen at runtime (instead of being predetermined along with the PEQT). RMS' new correlation methodology is used for quantile selection: the degree to which two locations are correlated affects the degree to which similar quantiles are chosen for them. The result is a ground-up PLT by location coverage. Policies and treaty terms are applied directly to losses as they are rolled up to, ultimately, a net post-cat PLT.

In SELT mode, no simulation timeline is employed. “Sub-event” losses are generated from the scenario’s event loss by sampling from its loss distribution dozens of times. The results are then output to the SELT. Once again, policy and treaty terms are applied directly to losses (Figure 2).

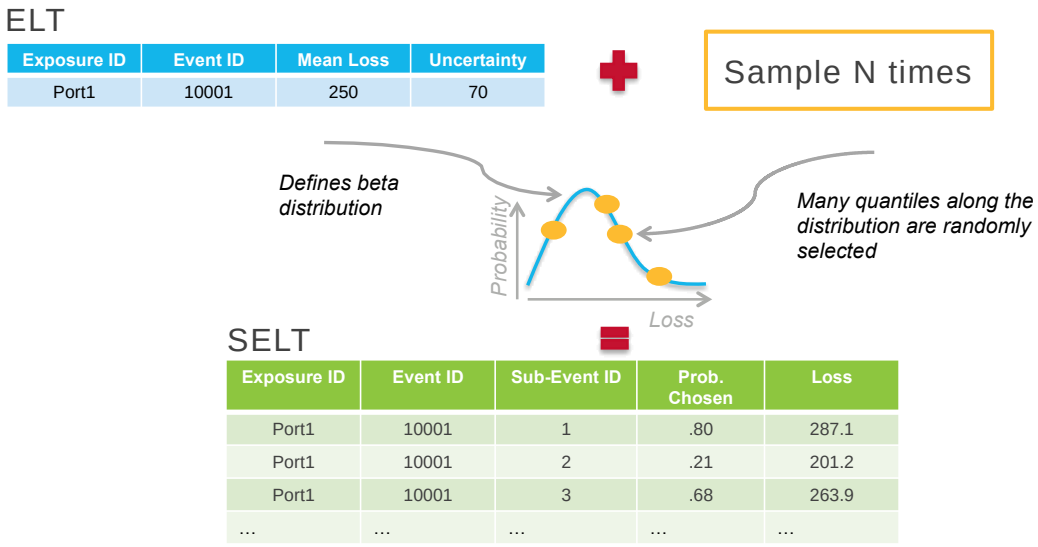

Figure 2: Several points are sampled from the ground-up event loss distribution to generate the SELT.

Using Two Sampling Methodologies

Why release two methodologies? Each mode has specific use cases.

PLT mode is the RMS Reference View. The simulation timeline allows it to appropriately deal with the full range of contract terms modeled by RMS, including contracts spanning multiple years, aggregate limits or deductibles, and hours clauses. Seasonality and clustering—whether RMS’ view or an adjusted view—can also be incorporated. Finally, PLTs can be easily rolled up with other losses in capital models and other frameworks.

SELT mode is used for scenario analyses and cat response. Because it samples single events dozens of times—users decide exactly how many times—it reveals the uncertainty in individual event losses.

Benefits

With HD simulation, RMS(one) enables the full benefits of a flexible, intelligent, and easy-to-use financial model:

1. The model is more transparent and easier to follow. In the ELTs output by the RiskLink financial model, mean losses and standard deviations are given. The model assumed that these parameters define a beta distribution, and calculus is required to work with the curve and determine the probability that a particular policy or treaty takes a loss. In HD simulation, contract terms and conditions are applied directly to the sampled losses. This approach can be used because uncertainty was taken into account when the ground-up loss was first sampled

from a severity distribution. The new approach is easier to follow and in keeping with the RMS goal of facilitating greater transparency into financial modeling.

The simulation framework makes it easier to calculate loss metrics. OEP, AEP, TCE-OEP, and TCE-AEP values are arithmetically derived directly from examining losses per period in the PLT (e.g., the aggregate exceedance probability of a loss threshold is the percentage of simulated years with an aggregate loss exceeding that threshold). Average annual loss, excess average annual loss, and expected shortfall are similarly straightforward calculations. This is an intuitive framework for understanding the financial model and following losses from ground-up through post-cat.

2. Uncertainty and correlation are modeled more accurately. The RiskLink financial model assumes that the loss distribution at each financial perspective follows a beta distribution, even when contract layers are applied to curves. It also assumes fixed proportions of independent and correlated standard deviation when rolling up losses. While these were necessary shortcuts in world without distributed processing, they are not needed for HD simulation. Specific distributions (e.g., beta) are only assumed to model ground-up loss uncertainty for locations; once sampled, the losses can follow any distribution. Correlation is similarly assumed only at the ground-up level when samples are first drawn; at higher levels it is an observed output. This more accurately reflects reality: damage ratios are partially correlated for nearby locations for a given event, but the degree of correlation may vary when ground-up losses propagate to gross and net levels.
3. The new framework is more flexible for users. Reworking the financial module allows for more contracts to be modeled via the contract definition language (CDL). This means that hours clauses, aggregate policy limits and deductibles, aggregate sub-limits, and franchise and inuring deductibles will be fully included and correctly applied. Users can also explicitly adjust the damage ratio at which a total loss is assumed. Finally, the flexible correlation framework allows RMS to significantly upgrade its modeling of loss correlations. Instead of assuming constant correlation weights by peril region, these weights are allowed to vary by locations' proximity and the peril region.