Gallagher Re

Do Actuaries Dream of Cyber Cats?

Developing a practical framework for actuaries to communicate cyber tail risk model results

By: Ryan Wilkins
Special thanks: Justyna Pikinska, Kiana Lashgari and Albert Choi

TABLE OF CONTENTS

This paper will outline a practical framework for actuaries to integrate and communicate the results of cyber catastrophe modeling, to aid the formation of their carrier's view of cyber risk exposure. By leveraging the existing structures arising from natural catastrophe modeling, actuaries can translate cyber cat modelled results into a familiar and intuitive presentation. This objective type of analysis lends credibility to the modelled results and allows stakeholders with less cyber familiarity to make informed and consistent decisions.

The paper will also provide actuaries with a base understanding of the implied cyber industry cat loss curves arising from the proposed framework. These resources will allow actuaries building out cyber tail risk analyses from scratch to benchmark their own assumptions. This is caveated by noting that each major cyber cat modeling vendor releases major model updates at a yearly cadence, materially altering modelled results.

Finally, the paper will explore predictions about the future of cyber catastrophe modeling and comment on potential non-modelled risks.

What is Tail Risk?

At the highest level, tail risk refers to the chance of a loss occurring due to a rare, extreme event. Some sources attempt to apply a statistical definition to tail events, for example, three standard deviations above the mean or more,[1] or use less precise language, such as the often-quoted 'black swan' terminology. However, by talking with risk management professionals, one will find varying opinions on what constitutes a tail event.

One of the most significant insurance industry innovations of the late 20th and 21st Centuries comes in the form of natural disaster modeling tools, otherwise known as **catastrophe** or **cat** models. These models have become the standard for property insurance risk aggregation quantification, covering major perils such as earthquake, hurricanes and tornados, along with evolving to model additional perils such as flood or wildfire. Nat cat-modelled results are used by actuaries, underwriters, risk managers and others to set risk guidelines and to price (re)insurance or exotic insurance products, such as Insurance Linked Securities (ILS).

The success of Nat cat models has inspired existing model vendors and new market entrants to consider the potential of expanding this technology outside of the property insurance context, including casualty, pandemic and terrorism risk, with the potential for cyber cat insurance now being explored.

[1] https://www.investopedia.com/terms/t/tailrisk.asp

Cyber as an Aggregation Risk

How actuaries communicate tail risk

Tail risk is a major area of concern for insurance companies. The theoretical maximum limit of liability that insurance companies hold on their books is many multiples of the premium that is collected by issuing policies. Insurance carriers invest heavily in monitoring their risk aggregations and estimating worst case scenario events to ensure the survival of the business even in the face of extreme losses.

As such, modeling these remote events plays a significant role in driving product development and risk tolerances at (re)insurance companies. Senior management relies on actuaries to understand their tail exposures and communicate the risks in a manner that allows informed decisions to be made. The vital statistics used by actuaries include:

- **AAL** (Average Annual Loss): The statistical expectation of catastrophe losses in the immediate next 12 months, typically across 10,000 or 50,000 simulation years.
  - Usage example: Primary policy pricing
  - The main contribution to the premium is typically the expected policy loss. For a given portfolio of homogeneous risks on a cat-exposed line of business, the actuary must determine how to allocate the AAL across the book. This may be achieved by simply adding an equal cat load on each policy or allocating based on exposure.
- **OEP** (Occurrence Exceedance Probability): For a loss sized X, the probability that a single event or occurrence produces a loss of X or more in the next 12 months.
  - Usage example: Reinsurance pricing
  - The definition of an "occurrence" in cyber is an active area of discussion in the market, as the boundaries around what constitutes a single event are less clear.
  - For event Excess of Loss (XoL) reinsurance, actuaries use the OEP curve to estimate the likelihood of cedants making a recovery on their purchase. This guides decisions on where to set attachment and limit exhaustion points and, subsequently, the reinsurance price.
- **AEP** (Aggregate Exceedance Probability): For a loss of sized X, the probability that all events or occurrences produce combined losses of X or more in the next 12 months.
  - Usage example: Setting capital requirements
  - Many regulatory requirements boil down to having to hold enough capital for the insurance company to financially withstand a tail event at a given probability, often in the vicinity of 0.5% or 1-in-200. Furthermore, internal risk controls may set acceptable aggregate loss exposure guidelines with further granularity to preserve enterprise value.
- **RP** (Return Period loss): The return period is simply the inverse of the O/AEP, expressed as a number of years. In fact, the RP is typically the quoted statistic when communicating cat results. A reinsurance underwriter might ask their actuary, "what is the 1-in-250-year loss?" when deciding to write a piece of business.

Exploring the Existing Tail Risk Modeling Landscape

A key observation of Nat cat modeling is that an outside observer can typically approach the subject with some pre-existing knowledge. Most can intuit at a high level what damage may arise from a California Earthquake or a Florida Hurricane and why insurers are concerned about aggregation risk.

Consider the natural catastrophe risk landscape as a plane, each **peril** can roughly be represented as a mutually exclusive bubble, covering some section of the plane. Within each peril are the **modelled perils**, which have familiar names from common cat models, usually corresponding to a geographical area. Below we see this graphically:

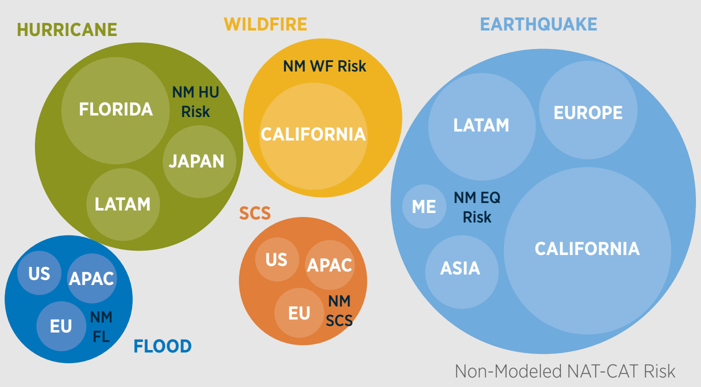

What are some of the key observations here?

- **Each Nat cat peril is well-defined:** When setting desired risk levels, there are clear demarcations between each subset of Nat cat perils. Each risk (earthquake, hurricane, etc.) can be analyzed separately, or disregarded if clear exclusions are in place.
- **Perils can be subdivided based on a portfolio's actual exposure:** Nat cat exposure has the advantage of being bound by geography. An insurer who writes only Florida-based homeowners policies can, for all practical purposes, ignore the California earthquake modules and focus analytics resources on Florida hurricanes.
- **Credit for portfolio diversification can be empirically tested:** By having distinct components of an overall cat load, the loss curves for each peril can be modelled individually and their sum compared to an aggregate modeling approach, with the differences being attributable to risk diversification within the portfolio.
- **The market accepts there is non-modelled risk:** All stakeholders are aware that these models have limitations. The area within each peril currently not modelled can be accounted for with an explicit load if said peril is covered by the book of business.

What Does the Cyber Tail Risk Model Landscape Look Like?

At the time of writing, Gallagher Re licenses three of the commercially available cyber aggregation models that estimate tail risk on client cyber insurance portfolios, also licensed by many (re)insurers for their own analyses. The past few years have seen immense investment being made in the development of these tools, with major updates made annually. Curious readers can also review the CyberIQ paper "Evaluating Cyber Models" for more information on how Gallagher Re scrutinizes the underlying framework of these models.

These models are recognized as being in an early stage of build out, with significant improvements to come as the cyber line of business evolves. Consequently, when envisioning the **Cyber Catastrophe Risk Landscape**, it is important to note there exists significant non-modelled risk. The models share many characteristics in their parameterizations, yet each have unique areas of focus, as demonstrated below.

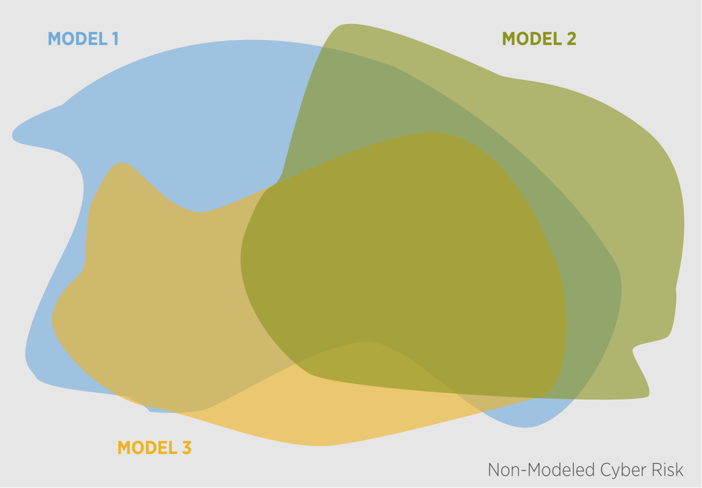

This representation provides some high-level observations on current cyber aggregation modeling limitations:

- **Modelled results are often presented as a single output, or at the vendor-defined level of granularity:** In the Nat cat world, actuaries may blend vendor model results with their in-house property models to produce a credibility-weighted result. By using the language presented by the developer of the model, actuaries lose an amount of flexibility in integrating the results.

  Example: A major writer of Florida property insurance may have their own internal hurricane model, which is combined with a commercially available earthquake model to form their property cat view of risk. With an in-house cyber model and no additional granularity, one must choose between the internal or commercial model results.

- **Comparing models based solely on raw output faces interpretation challenges:** Each model is parameterized using different modeling philosophies and assumptions. While many similarities can be seen between them, there is usually no obvious 1:1 mapping of scenarios. However, all models produce the same set of tail-risk statistics outlined above, which will form the basis for our framework.

  Example: Assume we've run 10,000 simulations of two models and extracted the 9,940–9,960th highest scenarios from each, representing the 'neighborhood' around the 1-in-200 loss. If the magnitude of losses from each model are similar, how best can we compare the neighborhood output from each model?

- **Overlap between risks modelled within scenarios/modules can be non-obvious:** When scrutinizing the assumptions backing each modelled scenario, we often find overlap. An argument can be made that this unfairly increases modelled cat risk through double counting.

  Example: Two data breach scenarios could begin with the same single point of failure (SPoF) but disseminate outwards differently.

- **Differing definitions of coverages and exclusions can cause modeling to not represent true exposure:** Each insurer has its own coverage wording which might not exactly fit to the vendor model's definition of coverage, causing losses to be covered that, in reality, should not be, or vice versa.

  Example: Breach and containment coverage is a significant driver of modelled losses. Some insurers may sublimit components of this coverage, such as ransomware payments, which some models don't provide the flexibility to reflect.

- **Underlying technographic risk data may differ between models:** Most components of simulated loss within the cyber models depend significantly on the technology being used by their IT infrastructure. This information is not always available and thus is often estimated or derived from external sources. Two models can therefore have differing views on SPoF concentration within the same portfolio.

  Example: If a simulated scenario results in Amazon Web Services US-East-1 going down, one model may assume a distribution of affected risks based on AWS market share, whereas another may rely on data from an "outside-in" scanning vendor.

Framing Cyber Tail Events in a Familiar Context: a Peril-Based Approach

To allow actuaries to credibly and succinctly communicate the results of cyber aggregation modeling to a wider audience, the industry needs to move away from vaguer concepts such as a "1 in 100 cyber event" and towards more distinct nomenclature and scenario classification.

It is common knowledge by purveyors of cyber modeling that there usually exists significant divergence in estimated insurance losses between model vendors. These models are nascent, with major version changes happening on a yearly basis which often significantly change results despite the same input.

Gallagher Re has adopted a **peril-based approach** to modeling systemic cyber risk, paralleling the common techniques used by Nat cat modelers. This approach has also seen significant adoption throughout the cyber insurance industry.

As of mid 2023, Gallagher Re considers three perils:

- **Ransomware:** Typically, these losses involve malicious actors gaining access to the internal file systems of an insured and encrypting the contents, demanding payment in return for decryption. Historically, the most influential driver of loss experience, ransomware losses have heavily influenced the development of the cyber insurance market, especially since 2019. Consequently, it has been an area of focus for all of the major vendor models.
  - » Common coverages: extortion, data recovery, business interruption (BI)
- **Data breach:** A loss of company or customer data has wide-reaching financial implications for the affected firm. Costs associated with reputational harm, loss of business and the subsequent fallout can quickly outstrip the immediate breach response amounts. Prior to the wave of ransomware activity, these incidents represented the bulk of insured cyber losses. As such the models have parameterized their scenarios with substantial incident history in mind.
  - » Common coverages: Breach response, liability

- **Service provider outage:** Whereas the other two perils are typically considered "human caused" losses (similar to terrorism or arson risks), the service provider outage peril deals with the knock-on effects of the failure of a software or service that a firm depends on to run its business. In these scenarios the insured is not usually the actual target of the attack, even if the original source of the issues is considered man-made. These are largely focused on cloud outage potential, due to the ubiquity and interconnectedness of these services, but can also include other vendors, such as payment providers.
  - » Common coverages: liability, BI, contingent BI

We note that this paper focusses on affirmative cyber losses for standalone cyber policies. Other types of cyber losses may include affirmative losses on blended lines of business such as Tech E&O, or non-affirmative ("silent cyber") losses on non-cyber lines of business. Recent tightening of underwriting and the introduction of cyber exclusions in non-cyber policy forms has significantly reduced the perceived risk of silent cyber, though it is still an area of concern for corporate level aggregation studies.

Overlaying this framework on our previous visualization of the cyber modeling landscape, some of the advantages of this approach are detailed below:

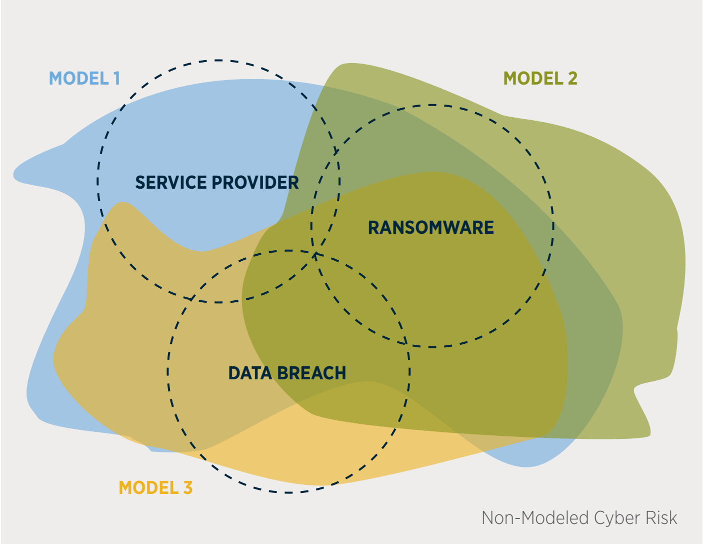

A peril-based framework allows actuaries to:

1. **Compare models on a granular, like for like basis:** Instead of reviewing the modelled output as a whole, each result set can be broken into its component parts. Where models have materially different aggregate results, what drives the differences can be seen. Alternatively, two models with similar overall loss estimates may reveal stark differences in their expectation of the most significant events.
2. **Adjust models independently:** As modeling best practices continuously improve, each cyber analytics team will develop their own preferences for the settings and assumptions within each model in accordance with their book. This framework allows expert judgment to be applied within each model while still allowing the flexibility of a generic end-product.
3. **Increase interpretability by using a common language:** There are many cases where abstracting away the specifics of the scenarios can be useful for communicating results. A non-cyber expert audience concerned with the "big picture" of the tail risk may be better suited worrying about which peril their book of business is more exposed to and the various types of potential causes, rather than homing in on one particular case.
4. **Explicitly determine a separate load for non-modelled risk:** If the in-house view of an insurer's cyber modeling team feels strongly about one area of risk their portfolio may be exposed to, this model allows for their creation of a "fourth peril" to account for it explicitly. Alternatively, a generic "non-modelled" risk parameter can be established based on a view of where the cat loads generated from this method differ from expectation.
5. **Address perceived model deficiencies:** In recognition of the nascent state of current cyber vendor models, many carriers have developed internal guidelines as to where their analytics teams disagree with the assumptions or outputs from each model. By applying the same logic to each model to derive peril curves, an actuary can produce multiple views of risk for the same peril each of which are consistent with internal views.

Actuaries using this approach must be aware of and accept some limitations. Firstly, the peril-based approach does not use the entire model. By trying to capture the broadest categories of tail-driving losses, we lose bespoke insights from the areas that model vendors have individually decided to focus on. In cases where these outside scenarios are mutually exclusive from our three perils, we're now lumping said exposure in with other "non modelled" risk.

Secondly, there are known overlaps between the defined perils. Lines of cause and effect are usually less clear in the cyber world than in property. A 'data breach' scenario might be the result of failure to pay a ransom from a 'ransomware' originating incident. When reviewing results, actuaries must determine if any intersection credits need to be accounted for, or if when comparing two models on the same peril if there is sufficient delineation within the underling scenarios to be on a like for like basis.

In addition, there are non-modelled risks within each defined peril. Much like how an earthquake in Florida could potentially lead to insured losses but is rarely a focus for modeling a national portfolio, there are significant gaps in modeling the landscape of each cyber peril. The potential misrepresentation of true exposure is amplified by the large number of 'unknown' root causes of these loss events as the line of business continues to mature.

Finally, granularity within each peril is omitted through this approach. Ransomware losses are sometimes identified as resulting from 'targeted' or 'non-targeted' attacks. Similarly, one paricular data breach scenario may result in a heavier proportion of third-party losses vs. first-party losses when compared to another data breach scenario. This is another form of basis risk resulting from a generalized peril approach where coverage terms may be applied more broadly than their true exposure.

One solution is to team up with those that have dedicated major resources towards learning the intricacies of each cyber vendor model and objectively applying this framework. When combined with strong modeling capabilities, it presents a powerful tool for tailoring views of tail risk towards each insurer and reinsurer's risk perspective. Later, we will discuss where we expect this framework to develop further. (Using our proprietary vendor blender tool, we produce multi-model, peril-derived aggregate cyber cat curve bespoke for each analysis based on what our research determines is most appropriate for that client portfolio.)

Applying the cyber peril-based approach

We examine two practical applications of cyber cat modeling using the **peril-based approach**, which parallel the common approaches used by property cat actuaries. Note that we assume the underlying portfolio consists of only **standalone cyber** insurance policies. Actuaries can adapt methodologies to extend this framework to other lines of business with cyber exposure, most often Tech E&O.

Projected Expected Loss Ratios (ELRs)

When tasked with estimating the future loss ratios of an insurer's portfolio, it is desirable to break down the component parts of the mean loss ratio. For cyber, the first level of granularity to be considered is exposure to single risk losses (non-cat) versus cat events. The non-cat component can be estimated using traditional actuarial methods of triangle development and on-levelling experience after stripping out any internally defined cat experience. Using our framework, the cat component can be divided into the contributions from each of our perils, along with a separately calculated non-modelled cat load to adjust for internal views on cat risk.

$$ELR_{\text{Cyber}} = \mu_{\text{Non-CAT}} + \mu_{\text{CAT}}$$

$$ELR_{\text{Cyber}} = \mu_{\text{Non-CAT}} + (\mu_{\text{Modelled CAT}} + \mu_{\text{Non-Modelled CAT}})$$

$$ELR_{\text{Cyber}} = (\mu_{\text{Attritional}} + \mu_{\text{Large Loss}}) + (\mu_{\text{SP}} + \mu_{\text{RW}} + \mu_{\text{DB}}) + \mu_{\text{Non-Modelled CAT}}$$

$$ELR_{\text{Cyber}} = (\text{Experience}) + (\text{Exposure}) + \text{Actuarial Judgement}$$

Value at Risk (VaR) Metrics

Determining reinsurance needs or capital requirements relies heavily on calculating the Value at Risk (VaR) metrics of the insurer's portfolio. For a cat-exposed line of business such as cyber, these losses are likely to be the largest contributor of losses to extreme events. Using a peril-based approach allows the insurer to target their risk mitigation efforts towards the areas of risk concentration specific to their mix of business. Having a logical and familiar cat framework can assist in justifying to reinsurers and regulators that proper credit should be given for the risk mitigation efforts taken.

A standard Monte Carlo simulation approach to integrating cat modeling into portfolio level VaR calculations is as follows:

1. Parameterize a non-cat (optionally separate attritional and large loss) random variable, X , for the given portfolio using traditional actuarial techniques on known experience.
2. For each of the three cat perils, P1, P2, P3, determine the most appropriate (or blended) source model and produce the corresponding AEPs for each of the N cat simulation years.
   a. Optional: Adjust the AEP results for perceived non-modelled risk, or introduce a fourth cat AEP curve representing non-modelled risk from an alternative source.
3. For some appropriately high number of portfolio simulations M ,
   a. Loss_m (noncat) ~ X
   b. Loss_m (cat) = $P1_{m \bmod N} + P2_{m \bmod N} + P3_{m \bmod N}$
   c. Loss_m = Loss_m (noncat) + Loss_m (cat)
4. To determine portfolio VaR at probability p , order each Loss_m into ascending order array R, then:
   a. $\text{VaR(Portfolio)}_p = R_{M*p}$

Note, the above assumes independence between the non-cat and cat contributions to overall losses. An extension of this methodology may introduce correlations. Correlations between cat perils themselves will depend on the parameterization and weights used by each original model.

Case Study: Proxy Industry Cat Loss Curves

Goals

As the number of insurers offering cyber products grows, many actuaries building cyber models from the ground up face data availability challenges. Actuaries often seek alternative views of risk to benchmark and validate in-house analysis. This case study aims to provide an objective and adaptable view of the current cyber cat landscape as seen by the market leading cyber vendor aggregation models.

Methodology

Defining the cyber insurance landscape

Modeling a cyber insurance policy portfolio requires two types of data: policy term information, and firmographic[2] exposure data on the underlying risks. Insurers often organize this in an **in-force policy bordereau**, which contains data for all active, unexpired insurance policies written by the insurer at a given evaluation date. The input schema will vary between models, but common fields include:

- **Example policy terms**
  - » Limit (per occurrence/aggregate)
    - Cyber specific coverage limits (ransomware, breach response, etc.)
  - » Attachment
  - » Deductible, self insured retention
  - » Premium
- **Example firmographic data (underlying insured)**
  - » Yearly revenue
  - » Geographic location
  - » Industry category
  - » Employee count
  - » Record count

To best represent the cyber risk landscape, **synthetic in-force portfolios** that represent mutually exclusive 'sections' of the existing cyber insurance market are created. At this point in the methodology, a trade-off is required between granularity and practicality. For each additional variable parameter, the number of portfolios to be created exponentially increases, reducing their broad applicability to the reader's specific case.

Through sensitivity testing, **firm size** has been identified to be the most influential piece of firmographic data on modelled cyber losses. Second, a split between **primary vs excess** insurance policies provides a level of insight into what size of losses drive each peril. Readers who often speak of insurance business at the portfolio level might feel this makes intuitive sense — for example, a portfolio manager may describe their cyber business in a single sentence as focusing on "primary limits on small to medium sized risks." Thus, for purposes of this study we will examine **eight aggregate market segments**:[3]

RISK SIZE	ATTACHMENT TYPE		
	PRIMARY	EXCESS	HIGH EXCESS
Micro	✓		
Small	✓	✓	
Medium	✓	✓	
Large	✓	✓	✓

[2] Firmographic data describes characteristics of the firm, such as geographical location, in contrast to technographic data, which describes the technologies that make up the firm's IT systems.

[3] Definitions and parameterizations available in **Appendix A**.

Building synthetic industry portfolios

Gallagher Re has built a **cyber insurance Industry Exposure Database (IED)**, providing the most comprehensive dataset within the industry on the types of policyholders purchasing cyber protection. Through this, informed assumptions can be made about the policy characteristics within the synthetic portfolios and then augmented with actual firmographic risk data. The algorithm for creating these portfolios is as follows:

For each of the eight portfolio permutations

1. Determine the appropriate number of risks, N, to simulate by breaking the problem into component parts that can be estimated from the IED:

$$\text{Policy Count} = \frac{\text{(Total Portfolio Limit)}}{\text{(Average Policy Limit)}} = \frac{\text{(Average Policy Rate on Line)}}{\text{(Average Policy Limit} \bullet \text{Total Premium)}}$$

Where,

- **Average rate-on-line:** Estimated based on known IED premiums on policies fitting the criteria for this portfolio, on leveled to 2023 rate levels.
- **Average policy limit:** The median policy limit seen for applicable IED policies.
- **Total premium:** $100 million gross written premium for all portfolios.

2. Create a policy limits and attachment profile, allocating the **N** risks in discrete buckets corresponding to our target portfolio makeup. The following provides an example of our **primary, large risks** synthetic portfolio makeup:

	LOWER ATTACHMENT	$0M	$1M	$2M	$5M	TOTAL
WHOLE LIMIT	UPPER ATTACHMENT LIMIT AT PARTICIPATION	$1M	$2M	$5M	$10M	
500,000	500,000	0	0	0	0	0
1,000,000	1,000,000	0	0	0	0	0
3,000,000	3,000,000	0	0	0	0	0
5,000,000	5,000,000	0	0	0	0	0
10,000,000	5,000,000	30	3	0	0	33
10,000,000	10,000,000	66	12	0	0	78
25,000,000	10,000,000	174	13	0	0	187
25,000,000	15,000,000	27	7	0	0	33
25,000,000	25,000,000	0	0	0	0	0
50,000,000	50,000,000	0	0	0	0	0
	TOTAL	297	35	0	0	331

3. Create **N** 'synthetic policies' with the permutations of limits and attachments from above. This process is automated via R script that reads in the profile.

4. Randomly select **N** insured risks and their associated firmographic data from the IED, within the firm size constraints.

   a. We restrict ourselves to US-based risks only.

   b. We implicitly assume that by randomly sampling, the representation of firmographic data within our synthetic portfolio will be close to the market average. This is an important assumption to note when benchmarking, for instance an insurer's specific book may be heavier in certain **industry classes** or completely exclude others.

5. Randomly assign the **N** risk to the **N** policies. Here, we are making the assumption that any risk within the portfolio can potentially purchase any of the policy layers within the portfolio. We note this assumption is less of an issue for large risks that are likely to buy large towers but may cause distortions for smaller risks.

6. The target $100 million premium is allocated to policies based on their share of total portfolio limit. While this is a weak assumption, we will not be examining results on an individual risk level and rather are most concerned with the total portfolio exposure being proportional to the premium, and that portfolios are comparable to each other. This is achieved through using actual average rate-on-line and policy limit data when calculating policy counts above.

7. Run the portfolio through each model as normal.

8. Extract each peril, P, AEP results from each model, M, using the internally defined Gallagher Re View of Risk for adjusting the models.

Curve blending

For purposes of creating industry loss curves, it must be decided how to determine an average curve for each peril from the results of each model. Fortunately, methods for blending curves are well documented in traditional property cat modeling. Here, the methods used by Homer and Li[4] are adapted to produce the following algorithm:

Assuming three models, for peril P,

1. Order **peril**, P, AEP results M_{1i}, M_{2i}, M_{3i} of each model in ascending **simulation year**, i, in a joined table. Note each model must have been run on the same number of simulations, in this case 50,000.
2. Simulate a random number U_i between 0 and 1 for each i and append as a new column.
3. Append a **blend** column with entries B_i.
    a. For each row in the table,
        i. if $U_i < 0.333$, $B_i = M_{1i}$
        ii. else if $U_i < 0.666$, $B_i = M_{2i}$
        iii. else $B_i = M_{3i}$
4. The resulting column B contains our blended AEP curve for Peril P.
5. Repeat for the remaining perils.

To determine a total blended AEP curve T, we have the option to blend the T_1, T_2, T_3 combined AEP curves from each individual model, or add together each B_i. For our purposes we've taken the first option to provide a "average model" aggregate view.

The resulting output for each of our synthetic portfolios is provided graphically below. Numeric output is available in the appendix.

[4] "Notes on Using Property Catastrophe Model Results" https://www.casact.org/sites/default/files/2021-02/2017_most-practical-paper_homer-li.pdf

Modelled results (individual perils)

Consistent Y-Axes

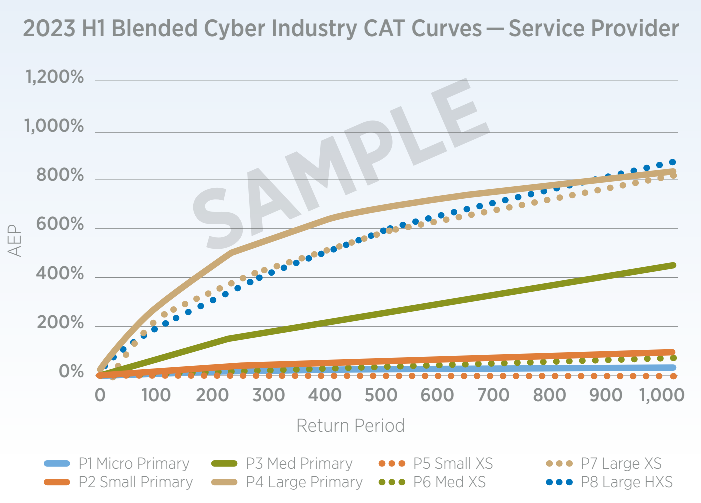
2023 H1 Blended Cyber Industry CAT Curves — Service Provider

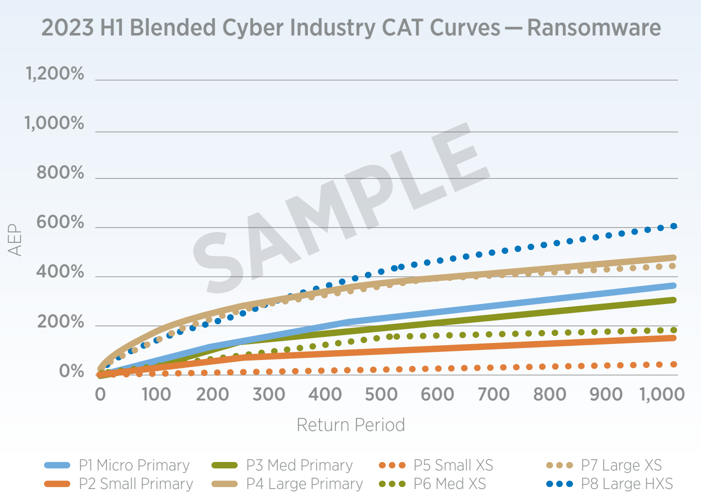
2023 H1 Blended Cyber Industry CAT Curves — Ransomware

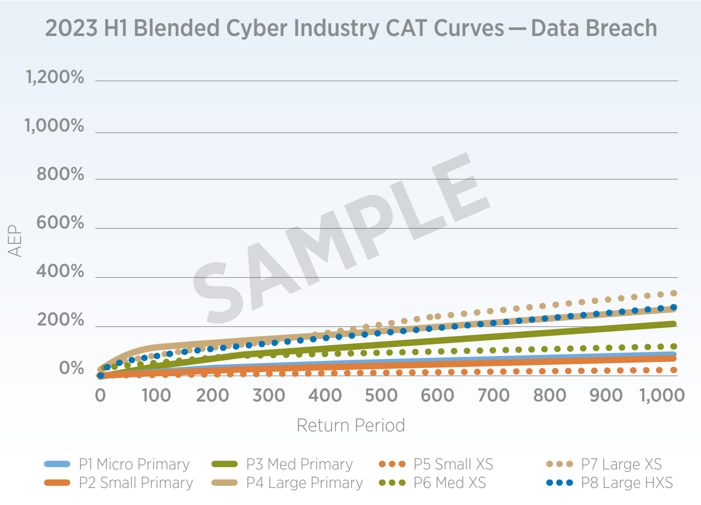
2023 H1 Blended Cyber Industry CAT Curves — Data Breach

Freeform Y-Axes

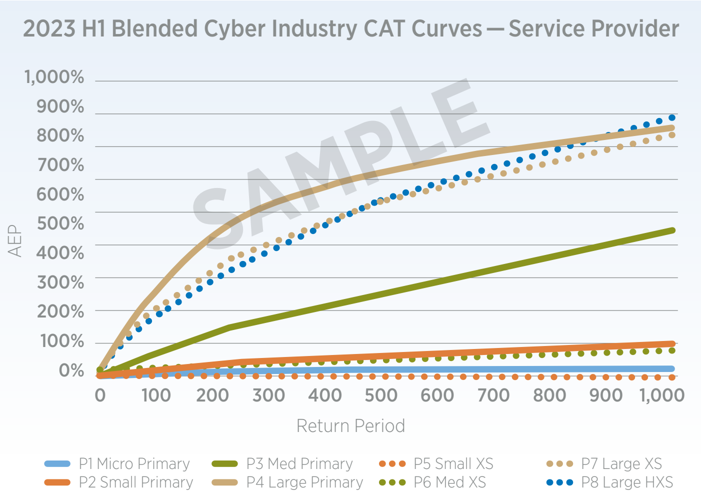
2023 H1 Blended Cyber Industry CAT Curves — Service Provider

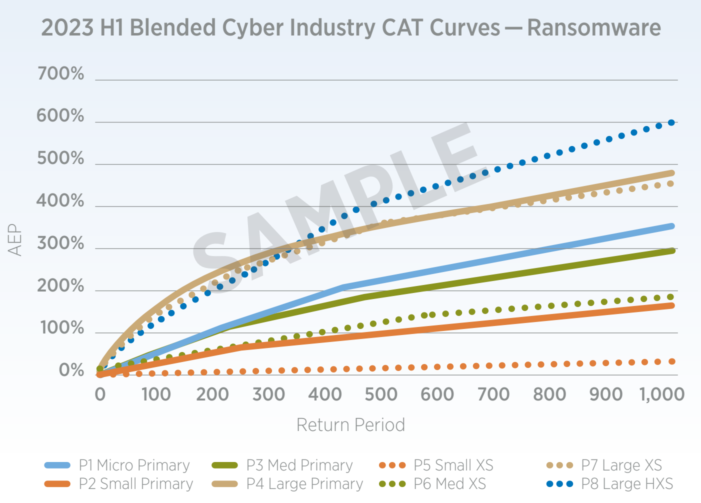
2023 H1 Blended Cyber Industry CAT Curves — Ransomware

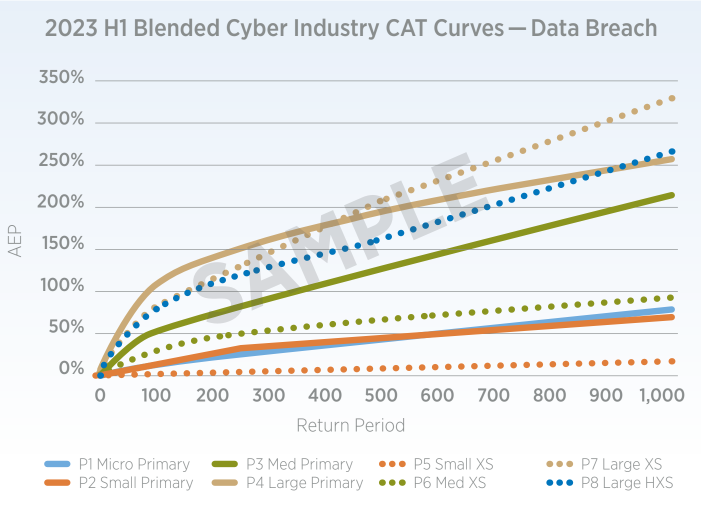
2023 H1 Blended Cyber Industry CAT Curves — Data Breach

Reviewing synthetic portfolio results

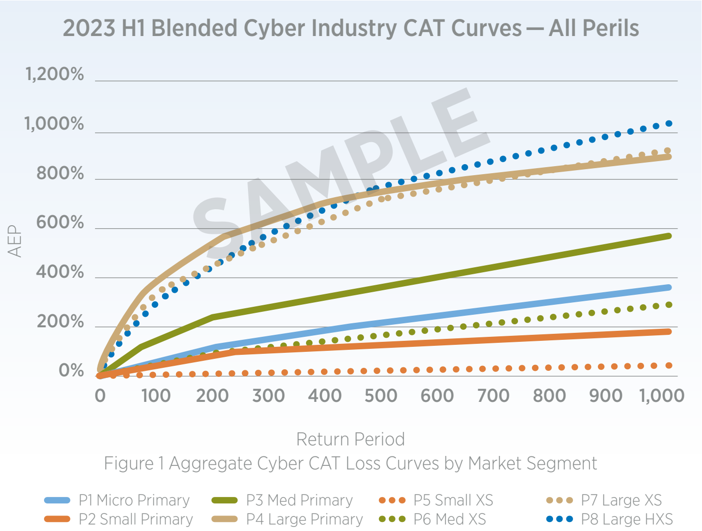

Figure 1 Aggregate Cyber CAT Loss Curves by Market Segment

The interpretation of these results is as follows: for a given sector of the standalone cyber market, where mutually exclusive sectors are defined by a combination of insured risk size and policy attachment level, the graphed line depicts the blended cyber model expectation for a 1-in-X aggregate year loss ratio for **catastrophic** cyber events. As discussed, these results are based on the Gallagher Re interpretation of each model and are subject to the assumptions listed in the appendix. An actuary wishing to find **full portfolio level** tail losses will need to overlay an internal view of non-cat loss distributions over a blended cat curve representing their specific mix of business.

Using the **medium sized, primary attachment** portfolio as an example, we can make some initial observations. We can see from the aggregate graph that medium sized risks are resulting in approximately 2x–3x the loss ratio in the tail as compared to small primary risks, diverging the further gone into the tail. In reviewing individual peril breakdown, we see this effect is driven by both the service provider outage and data breach perils, and less so by differences in ransomware exposure. To expand, here are some high-level observations we've made with the current iteration of cyber models:

1. **Extreme aggregation events are driven by larger risks:** Modeling suggests that the worst losses suffered by the industry when an event happens will result from aggregations of large losses to large risks. This makes some intuitive sense — the low number of risks within high limit, large enterprise portfolios means each incremental limit loss results in significant loss ratio increases.
2. **Tail service provider outage events are expected to be far more severe for larger risks:** When theorizing potential major cyber events such as a large cloud provider outage, it is logical to assume significant impact to large companies with many interconnected systems.
3. **All sizes of risks are exposed to significant ransomware losses in the tail:** Ransomware has historically been the driver of major losses within cyber. Recent discussions have suggested that attackers are turning their attention more towards small and medium sized businesses as laws around paying ransoms have changed. Not to be outdone, large risks are still considered valuable target due to their deeper pockets, despite stronger cybersecurity. All the models agree that ransomware drives losses across the tail.
4. **Excess business tends to model low in the tail:** Many actuaries (author included) unscientifically describe excess business loss experience as "spikey". This usually refers to high volatility in the results, with some years running near loss-free, and some with high losses. Our modeling here shows some bad experience in the tail for excess business, but potentially lower than one would expect, especially for SME business. This might suggest that extreme scenario **footprints** are spread too thin — in reality we'd expect more concentrated, big tower losses.
5. **Exposure of smaller risks to large aggregation events is potentially undermodelled in the current cyber model landscape:** Past loss experience would suggest that SME cyber business has not performed as well as the current cyber models claim it will. This is partially explained by rates increasing cumulatively more than 100% since 2020, halving loss ratios on an all else equal basis, however the curves seen in these charts are still flatter than one would expect. For example, our aggregate curve suggests an 83% and 132% cat loss ratio at 1-200 and 1-500 respectively for small, primary risks, much tighter than what is implied by current ASL pricing. Some vendors have focused in this area of the modeling by introducing more granular "micro micro" risk categories to further differentiate the sectors.

Initial testing of the mid-2023 model version releases have shown significant shifts in modelled micro & small losses, indicating that vendors have heard and reacted to this commentary.

Non-modelled risk consideration, or a fourth peril?

Looking ahead, some current topics of discussion in the cyber industry hint at where modeling might go next. Ransomware continues to be an economic issue. However, some insurers are reporting that Fund Transfer Fraud (FTF) is increasing significantly as a share of overall paid losses. FTF is accounted for with current models, but potential systemic events may be underrepresented.

Separating out nation-state sponsored events is increasingly important. With markets such as Lloyd's taking a hard stance on excluding these types of losses, and others taking a less strict approach, there will be increasing demand for clear segmentation within modeling. In the simplest form, this could be integrated as a pure fourth state-sponsored cat peril. More sophisticated approaches might split the existing perils into each component, i.e state-sponsored ransomware vs all other ransomware, much like how California earthquake and Middle Eastern earthquake are segmented from the overall earthquake peril.

Final remarks

The stated goal of this study was to provide Actuaries with a practical, objective approach to communicating cyber cat modeling results. This paper argued that modelled output taken directly from a black-box cyber model lacks carrier-specific considerations. By adapting the familiar nomenclature from natural catastrophe modeling, we've shown that the intricacies of existing vendor cat models can be presented in a context that is intuitive for non-cyber expert audiences, without sacrificing significant credibility. The limitations of this approach are clear — in house cyber modeling experts will still be required to have deep understanding of the assumptions driving their analysis when tasks require more granular results.

The appendices of this study outline each of the assumptions used for our case study, providing the enterprising actuary with a starting point for their own cyber modeling development, or a data point to benchmark against existing work. Due to the rapidly evolving nature of cyber risk, we note that these benchmarks will quickly become dated — model updates in 2023 are expected to materially change results. Actuaries wishing to stay up to date on cutting edge cyber analytics are encouraged to reach out directly to the authors or any Gallagher Re Cyber Analytics colleagues with feedback or questions.

Appendix A: Modeling Assumptions

Portfolio parameter definitions

We aimed to build synthetic portfolios that represent the true exposure of each market segment as seen in actual cyber insurance portfolios. The Gallagher Re Cyber IED allows us to observe actual limits and attachment deployments, and mimic the definitions of revenue size used by insurers and vendor models. Below are the high level parameters used to create the limits and attachment profiles, corresponding to roughly the 25th to 75th percentile of the market:

Company Size by Revenue

REVENUE	REVENUE BAND
Micro	0–10M
Small	10M–250M
Medium	250M–1B
Large	1B+

Policy Limit Bands (ex. Participation)

	100% LIMIT	
REVENUE	PRIMARY	EXCESS
Micro	200k–1M	N/A
Small	1M–5M	5M–10M
Medium	3M–10M	5M–10M
Large	5M–20M	10M–50M

Policy Attachments Definition

ATTACHMENT	ATTACHMENT BAND
Primary	0–1M
Excess	1M–25M
High Excess	25M+

Rate on line assumptions

The chart below outlines the relative premium assumptions used in our study as part of the **Policy Count** calculation. These rates were selected based on information within Gallagher Re's Cyber IED and adjusted to reflect 2023 rate levels. Since each portfolio was calibrated to a theoretical $100M of premium, the modelled results for each portfolio can be compared directly. Actuaries using these results as a benchmarking tool may require adjustment of the loss ratios should their internal views of rate levels vary from our estimates here:

Rate on Line Relativities (Indexed to Medium, Primary Rates)

REVENUE	PRIMARY	EXCESS	HIGH EXCESS
Micro	0.26	N/A	N/A
Small	0.62	0.47	N/A
Medium	1.00	0.76	N/A
Large	1.27	0.96	0.65

Appendix B: Numeric Modelled Results

The tables below contain the data feeding each of the graphs in the modelled results section. Recall that these can be adjusted for any differentiation in view on the above assumptions with reasonable justifications. Actuaries may also wish to interpolate between market segment curves to represent their in-house book of business. If you are an Actuary looking for more detailed benchmarking and guidance on cyber modeling, please don't hesitate to reach out to the experts at the Gallagher Re Cyber Analytics team.

Blended, All Perils, CAT Only (2023H1 Modeling)

PERCENTILE	RP	P1 MICRO PRIMARY	P2 SMALL PRIMARY	P3 MED PRIMARY	P4 LARGE PRIMARY	P5 SMALL XS	P6 MED XS	P7 LARGE XS	P8 LARGE HXS
10.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
20.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
30.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
40.0%	2	0%	0%	0%	0%	0%	0%	0%	0%
50.0%	2	0%	0%	0%	0%	0%	0%	0%	0%
60.0%	3	0%	0%	0%	0%	0%	0%	0%	0%
70.0%	3	0%	0%	0%	1%	0%	0%	0%	0%
80.0%	5	1%	1%	5%	20%	0%	0%	13%	6%
85.0%	7	1%	3%	10%	37%	0%	1%	28%	25%
90.0%	10	3%	5%	17%	64%	0%	5%	52%	42%
95.0%	20	8%	11%	35%	125%	0%	14%	97%	82%
96.0%	25	11%	14%	43%	145%	1%	18%	115%	100%
98.0%	50	24%	28%	78%	235%	3%	31%	198%	174%
99.0%	100	54%	51%	141%	368%	6%	51%	319%	290%
99.5%	200	100%	83%	229%	532%	11%	90%	452%	453%
99.6%	250	122%	95%	251%	585%	13%	107%	504%	508%
99.8%	500	222%	132%	363%	752%	23%	173%	724%	768%
99.9%	1,000	349%	188%	570%	891%	36%	283%	917%	1,036%

Blended, Service Provider Outage, CAT Only (2023H1 Modeling)

PERCENTILE	RP	P1	P2	P3	P4	P5	P6	P7	P8
		MICRO	SMALL	MED	LARGE	SMALL	MED	LARGE	LARGE
		PRIMARY	PRIMARY	PRIMARY	PRIMARY	XS	XS	XS	HXS
10.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
20.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
30.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
40.0%	2	0%	0%	0%	0%	0%	0%	0%	0%
50.0%	2	0%	0%	0%	0%	0%	0%	0%	0%
60.0%	3	0%	0%	0%	0%	0%	0%	0%	0%
70.0%	3	0%	0%	0%	0%	0%	0%	0%	0%
80.0%	5	0%	0%	0%	0%	0%	0%	0%	0%
85.0%	7	0%	0%	0%	0%	0%	0%	0%	0%
90.0%	10	0%	0%	0%	5%	0%	0%	3%	0%
95.0%	20	2%	3%	8%	57%	0%	0%	42%	27%
96.0%	25	3%	4%	11%	86%	0%	0%	59%	41%
98.0%	50	7%	10%	26%	158%	0%	3%	122%	95%
99.0%	100	15%	20%	66%	278%	1%	9%	222%	194%
99.5%	200	19%	41%	140%	476%	2%	18%	355%	308%
99.6%	250	19%	47%	171%	535%	2%	22%	415%	374%
99.8%	500	29%	74%	256%	707%	5%	46%	602%	610%
99.9%	1,000	37%	107%	452%	863%	11%	81%	844%	892%

Blended, Ransomware, CAT Only (2023H1 Modeling)

PERCENTILE	RP	P1	P2	P3	P4	P5	P6	P7	P8
		MICRO	SMALL	MED	LARGE	SMALL	MED	LARGE	LARGE
		PRIMARY	PRIMARY	PRIMARY	PRIMARY	XS	XS	XS	HXS
10.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
20.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
30.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
40.0%	2	0%	0%	0%	0%	0%	0%	0%	0%
50.0%	2	0%	0%	0%	0%	0%	0%	0%	0%
60.0%	3	0%	0%	0%	0%	0%	0%	0%	0%
70.0%	3	0%	0%	0%	0%	0%	0%	0%	0%
80.0%	5	0%	0%	0%	0%	0%	0%	0%	0%
85.0%	7	0%	0%	0%	0%	0%	0%	0%	0%
90.0%	10	0%	0%	0%	0%	0%	0%	0%	0%
95.0%	20	0%	2%	10%	45%	0%	2%	36%	25%
96.0%	25	1%	3%	13%	56%	0%	4%	48%	40%
98.0%	50	3%	8%	26%	94%	0%	11%	80%	75%
99.0%	100	38%	22%	49%	146%	0%	24%	127%	115%
99.5%	200	94%	48%	97%	231%	3%	50%	204%	192%
99.6%	250	117%	59%	112%	256%	5%	64%	240%	226%
99.8%	500	218%	99%	190%	357%	15%	125%	350%	412%
99.9%	1,000	349%	164%	299%	479%	30%	185%	459%	594%

Blended, Data Breach, CAT Only (2023H1 Modeling)

PERCENTILE	RP	P1 MICRO PRIMARY	P2 SMALL PRIMARY	P3 MED PRIMARY	P4 LARGE PRIMARY	P5 SMALL XS	P6 MED XS	P7 LARGE XS	P8 LARGE HXS
10.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
20.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
30.0%	1	0%	0%	0%	0%	0%	0%	0%	0%
40.0%	2	0%	0%	0%	0%	0%	0%	0%	0%
50.0%	2	0%	0%	0%	0%	0%	0%	0%	0%
60.0%	3	0%	0%	0%	0%	0%	0%	0%	0%
70.0%	3	0%	0%	0%	0%	0%	0%	0%	0%
80.0%	5	0%	0%	0%	0%	0%	0%	0%	0%
85.0%	7	0%	0%	0%	0%	0%	0%	0%	0%
90.0%	10	1%	1%	3%	7%	0%	0%	0%	0%
95.0%	20	3%	4%	13%	24%	0%	5%	19%	18%
96.0%	25	3%	5%	16%	31%	0%	9%	25%	25%
98.0%	50	6%	9%	32%	57%	1%	18%	50%	50%
99.0%	100	11%	14%	50%	105%	3%	28%	79%	75%
99.5%	200	21%	23%	68%	137%	5%	40%	111%	108%
99.6%	250	26%	28%	79%	146%	6%	45%	125%	120%
99.8%	500	39%	40%	119%	193%	10%	61%	208%	160%
99.9%	1,000	73%	67%	214%	257%	15%	88%	330%	263%

Global and Local Reinsurance

Drawing on our network of reinsurance and market specialists worldwide, and as part of the wider Gallagher company, Gallagher Re offers the benefits of a top-tier reinsurance broker, one that has comprehensive analytics and transactional capabilities, with on-the-ground presence and local understanding. Whether your operations are global, national or local, Gallagher Re can help you make better reinsurance and capital decisions, access worldwide markets, negotiate optimum terms and boost your business performance.

Author

Ryan Wilkins, FCAS
Cyber Pricing Actuary
Ryan_Wilkins@GallagherRe.com

Contributors

Justyna Pikinska
Partner, Head of Analytics for Specialty Lines
Justyna_Pikinska@GallagherRe.com

Kiana Lashgari
Senior Actuarial Analyst
Kiana_Lashgari@GallagherRe.com

Albert Choi
Vice President — Actuarial
Albert_Choi@GallagherRe.com

It's the *way* we do it.

For more information, visit **GallagherRe.com.**

Gallagher Re

Gallagher Re is a business unit that includes a number of subsidiaries and affiliates of Arthur J. Gallagher & Co. which are engaged in the reinsurance intermediary and advisory business. All references to Gallagher Re below, to the extent relevant, include the parent and applicable affiliate companies of Gallagher Re. Some information contained in this document may be compiled from third party sources and Gallagher Re does not guarantee and is not responsible for the accuracy of such. This document is for general information only and is not intended to be relied upon. Any action based on or in connection with anything contained herein should be taken only after obtaining specific advice from independent professional advisors of your choice. The views expressed in this document are not necessarily those of Gallagher Re. Gallagher Re is not responsible for the accuracy or completeness of the contents herein and expressly disclaims any responsibility or liability, based on any legal theory, for damages in any form or amount, based upon, arising from or in connection with for the reader's application of any of the contents herein to any analysis or other matter, or for any results or conclusions based upon, arising from or in connection with the contents herein, nor do the contents herein guarantee, and should not be construed to guarantee, any particular result or outcome. Gallagher Re accepts no responsibility for the content or quality of any third-party websites that are referenced.

The contents herein are provided for informational purposes only and do not constitute and should not be construed as professional advice. Any and all examples used herein are for illustrative purposes only, are purely hypothetical in nature, and offered merely to describe concepts or ideas. They are not offered as solutions for actual issues or to produce specific results and are not to be relied upon. The reader is cautioned to consult independent professional advisors of his/her choice and formulate independent conclusions and opinions regarding the subject matter discussed herein. Gallagher Re is not responsible for the accuracy or completeness of the contents herein and expressly disclaims any responsibility or liability based on any legal theory or in any form or amount, based upon, arising from or in connection with for the reader's application of any of the contents herein to any analysis or other matter, nor do the contents herein guarantee, and should not be construed to guarantee any particular result or outcome. Gallagher Re is a trading name of Arthur J. Gallagher (UK) Limited, which is authorised and regulated by the Financial Conduct Authority. Registered Office: The Walbrook Building, 25 Walbrook, London EC4N 8AW. Registered in England and Wales. Company Number: 1193013. AJG.com/uk.