



**An Actuarial Approach to Stochastic Modeling of Casualty
Catastrophe (Cat) Risk**

Journal:	<i>CAS Publications</i>
Manuscript ID	CASPub-2022-EFor-0016.R1
Manuscript Type:	E-Forum
CAS Taxonomy:	Dynamic Risk Modeling < Actuarial Applications and Methodologies, ERM (Enterprise Risk Management) < Actuarial Applications and Methodologies, Extreme Event Modeling < Financial and Statistical Methods
Keywords:	Casualty Cat

SCHOLARONE™
Manuscripts

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

Neil M. Bodoff, FCAS, MAAA

Eric Dynda, FCAS

Krish Kamdar

Brandon Stevens, ACAS

Abstract

Motivation. Crucial to risk management of casualty insurance is the ability to quantify the likelihoods and severities of future extreme scenarios.

Method. We use foundational actuarial concepts to build a full stochastic model of casualty catastrophe (cat) risk.

Results. Casualty cat modeling is achievable within the actuarial community.

Conclusions. An actuarial approach can be used to build a forward-looking stochastic casualty cat model.

Availability. Some but not all of our proprietary work may be available to the reader.

Keywords. Casualty. Cat. Casualty Cat. Cat Modeling. Casualty Cat Modeling.

1. INTRODUCTION

In insurance and reinsurance risk management, understanding the likelihood and magnitude of future downside events is crucial. This need has been broadly acknowledged in the realm of natural catastrophe (nat cat) risk affecting property lines of business: cat models have been implemented throughout the insurance industry to address property cat. Yet an equivalent approach has yet to take hold in the casualty lines of business. One reason for this lack of progress has been the uncertainty and complexity relating to how to construct a cat model for casualty. In this paper, we show one approach for how to move forward on this crucial task.

1.1 Research Context

There is limited actuarial literature on this topic; primarily D'Arcy [1] has written on this issue.

1.2 Objective

This paper aims to describe how we built a casualty cat model and some of the things we

1
2
3
4
5 *An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk*
6

7 learned along the way.¹
8
9

10 **1.3 Outline**

11 The remainder of the paper proceeds as follows. Section 2 will discuss background and
12 current context and motivation for our new approach, Section 3 will discuss foundational
13 concepts of our model, Section 4 will discuss how to build the cat scenarios that comprise
14 the model, Section 5 will discuss the actuarial algorithms for generating detailed claims
15 output within the casualty cat scenarios, Section 6 will discuss parameter issues, Section 7
16 will discuss data issues, Section 8 will discuss numerical results, and Section 9 will discuss
17 software implementation.
18
19
20
21
22

23 **2. BACKGROUND, CONTEXT, AND MOTIVATION**

24 Measuring downside risk is important for insurance. Yet actuarial modeling has been
25 measuring casualty risk for quite a while; why do we need a novel casualty cat model to do
26 so?
27
28
29
30

31 **2.1 On the Need for a Casualty Cat Model**

32 Current methods typically use a frequency and severity modeling approach. The actuary
33 constructs probability distributions for the number of claims and the size-of-loss (“severity”)
34 of each claim. Using either simulation or analytical methods, the actuary obtains a probability
35 distribution for the aggregate loss distribution which shows the range of outcomes and the
36 downside risk of the insurance portfolio.
37
38
39
40
41

42 This approach is insufficient for the following reasons:
43

- 44 1. Statistically, for a granular individual line of business, the idea of using one overall
45 loss engine to represent the line of business is inaccurate and often does not produce
46 enough tail risk. The types of phenomena that drive the behavior of the tail of the
47 distribution are likely different than what drives the body of the distribution. As such,
48 a bi-modal or regime-switching loss process with more than one loss engine is likely
49 more descriptive; in other words, it’s more accurate to have separate loss engines for
50
51
52
53

54
55 ¹ These details will be generally applicable for any reader choosing to build a casualty cat model, despite
56 reflecting our specific experiences building Gallagher Re’s “eNTAIL” casualty cat model.
57

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

cat and non-cat rather than one overall loss engine.

2. Statistically, on a total company level, when combining the probabilistic results of each line of business into the total company view, blending the lines of business will create too much diversification benefit and will show too light of a tail.
3. Even if one could adopt other strategies (correlation matrices, common shock, etc.) to overcome the statistical problems of the tail risk being too light, these measures fall short on ease of explanation. There is no coherent narrative about what is driving severe outcomes. A casualty cat model, however, would produce a coherent underlying narrative about what is driving the extreme downside results.

2.2 What is Casualty Cat?

Casualty cat is not always so simple to define or to describe. We intuit that it ought to refer to a large casualty event that creates a noticeable impact on the insurance industry. We also sense that it ought to be an event that is widespread and pervasive and thus affects more than just one insured business. For our model, we chose to describe an event as a casualty cat event if it satisfies the following two criteria:

1. It generated historical insured casualty claims losses of more than \$100 million in today's dollars
2. It generated claims on more than one "tower" of insurance coverage

2.3 Is Building a Casualty Cat Model Feasible?

We noted in the points above that there is significant motivation to develop and use a casualty cat model just like we use property cat models. But is the analogy to property cat sensible? If so, then why haven't casualty cat models had more adoption in the market? Isn't building a casualty cat model very different from building a property cat model?

On the one hand, adopting casualty cat modeling should be just as non-controversial as adopting property cat models. Yet, the unique attributes of casualty cause it to be quite different from property, to the point that many practitioners view any attempt at constructing a casualty cat model as impossible. In this paper, we address these unique attributes of casualty and demonstrate how we overcome them in our proposed framework for building a casualty cat model.

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

3. FOUNDATIONAL CONCEPTS OF PROPOSED MODEL

At a foundational level, property cat and casualty cat share some similarities. In both cases, we can break down a cat event into three components: the scenario, the exposure, and the loss. In property cat, the scenario is the occurrence of an event generating widespread property damage, the exposure comprises the insured properties vulnerable to damage from the event, and the loss comprises the individual claims frequencies and severities. In casualty cat, the scenario is the occurrence of an event generating widespread liability, the exposure comprises the insured entities vulnerable to liability from the event, and the loss comprises the individual claims frequencies and severities. From this high-level view, it appears possible to approach casualty cat modeling in a similar way to property cat modeling. However, complications quickly arise.

One of the key distinguishing attributes of casualty, in contradistinction to property cat, is that the risk is constantly shape-shifting. In property, the historical losses might arise out of Florida Hurricane events and future losses would arise from similar events. To this extent, in property we can observe some similarity, regularity, and stationarity regarding historical events and future events. In casualty, however, future events are often different than historical precedent events. Asbestos, which generated billions of dollars of claims and is one of the most notable casualty events in the history of insurance, is an illustration of this point. In the future, we would not expect such an event to repeat itself, for at least two reasons: first, insurance contracts have evolved to mostly exclude asbestos, and second, society has evolved to better avoid using this dangerous substance. More generally, we can say that in casualty business, after some item or substance or behavior is discovered to be injurious, the casualty event is unlikely to repeat itself because of changes in insurance contracts and changes in society's underlying economic activity. This diagnosis of one of the unique attributes of casualty leads us to the key concepts we've used in building our model:

1. Casualty risk is always changing, so it's crucial to focus on potential loss events in the future.
2. Although the future is unknowable, we can use Realistic Disaster Scenarios (RDS) to map out possible future situations; but because the future is unknowable, the RDS scenarios should not focus too narrowly on prescribed minutiae but rather focus on somewhat broad categories of future scenarios.

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

3. By using RDS scenarios that are defined broadly, we enable linking together new future scenarios to historical events that serve as “precedents” that “inspire” the future RDS scenarios. And by mapping historical losses to broad RDS categories, we bolster the realistic moniker of RDS. As a result, we obtain forward-looking casualty scenarios that are supported by historical events yet have enough freedom to not be straight jacketed by the past.
4. Construct RDS scenarios with stochasticity to obtain fully probabilistic output from a Monte Carlo simulation engine. These “Stochastic Realistic Disaster Scenarios” (SRDS) serve as the backbone of the model and generate output that allows us to answer questions such as “What is the 1-in-250-year downside risk for this casualty insurance portfolio?”

4. BUILDING THE CASUALTY CAT SCENARIOS

What steps did we take to construct the casualty cat scenarios in our model?

1. Collect historical data: recalling our description of casualty cat as an event that causes more than \$100M of loss and affects more than one tower of insurance coverage, we sifted through publicly available historical data to find candidates in the historical record.
2. Create forward-looking “scenarios” by identifying broad categories of events: after collecting data on historical events, we noticed that several of them could be slotted into broad, general categories. For example, after collecting data on the MGM Grand Hotel fire and other similar events, we realized we could create a broad forward-looking category of “Building Fire”. This forward-looking scenario acknowledges that Building Fire is an ongoing risk and yet is broadly described because we do not really know where the locus of the next fire will be.
3. Organize scenarios into broad-based “perils” based upon the underlying “nexus of aggregation”: after constructing all the various casualty cat scenarios, we can see that some of them have areas of commonality despite being quite different. Recall that a casualty cat scenario is something that creates losses to multiple different insureds arising out of one common underlying cause. What are these underlying causes that serve as the nexus for aggregations of claims into a cat scenario? We used four types

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

of “nexuses of aggregation” for casualty cat scenarios, as follows²:

- a. “Single Physical”: multiple casualty claims in which the nexus of aggregation arises from a specific single physical location in a scenario such as Building Fire (e.g., MGM Grand Hotel fire, etc.)
- b. “Single Professional”: multiple casualty claims, especially professional liability claims, in which the nexus of aggregation arises from a specific single company in a scenario such as Firm Financial Collapse (e.g., WorldCom, etc.)
- c. “Systemic Products”: multiple casualty claims, especially products liability claims, in which the nexus of aggregation arises from a widely used product in a scenario such as Construction Materials (e.g., Chinese Drywall, etc.)
- d. “Systemic Professional”: multiple casualty claims, especially professional liability claims, in which the nexus of aggregation arises from a widely adopted behavior in a scenario such as Financial Advisors (e.g., Lloyd’s agents mis-selling, etc.)

5. ACTUARIAL ALGORITHM FOR SIMULATING CLAIMS

We wrote down the following Monte Carlo simulation algorithm to simulate thousands of years of casualty cat outcomes for a given portfolio of policies, starting with simulation year #1:

1. Given simulation year #1, simulate whether scenario #1 occurs or does not occur
2. Given simulation year #1 and scenario #1, simulate whether there is an insurance claim generated to policy #1 in the insurer’s portfolio
3. Given simulation year #1 and scenario #1 and policy #1 and claim #1, simulate the size-of-loss of the amount of the claim from ground up, apply policy limits and attachment points, and obtain the insured claim loss amount
4. Loop through all simulation years; loop through all casualty cat scenarios; loop through all policies in the insurer’s bordereau; loop through all claim counts; generate all claims amounts

² This classification was developed by Andrew Newman, Head of Global Casualty, Gallagher Re

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

The output of the simulation is an array of claim amounts simulated across policies, cat scenarios, and simulation years. This fully granular output set allows the actuary to conduct diagnostic checks and understand drivers of loss, down to the individual policy level. It also allows for the construction of exceedance probability (EP) curves for various return periods for individual scenarios, a select group of scenarios, or a total combined view. These curves can be constructed in a manner identical to property cat modeling.

6. ACTUARIAL PARAMETERS

Given the algorithm described above, we needed to estimate several sets of modeling parameters.

First, for each casualty cat scenario in the model described as a stochastic RDS, we needed to estimate annual scenario frequency: the likelihood of the scenario occurring in any given year. We collected empirical historical data to estimate these values. We chose a Poisson distribution to model annual cat scenario frequencies, using annual scenario likelihoods estimated from the empirical historical data as the mean. Although at first we used a Bernoulli distribution, we switched to a Poisson distribution in order to incorporate the possibility (though small) of multiple cat occurrences in the same year from the same type of cat scenario (e.g., multiple Building Fire events separately occurring within the same year).

Mathematically, we can say as follows:

Let X = random variable representing the number of occurrences of a casualty cat scenario

Let X follow a Poisson distribution with parameter λ , for all k equal to non-negative integer values:

$$f_X(k; \lambda) = \lambda^k / k! * e^{-\lambda} \quad (6.1)$$

Further, the model is made up of several independent cat scenarios, so we can say as follows:

X = matrix or “vector” of cat scenario random variables

$X = \{X_1, X_2, \dots, X_n\} = \mathbf{X}$

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

Each element in \mathbf{X} follows its own independent Poisson distribution with parameter λ such that λ is also a vector:

$$\lambda = \{\lambda_1, \lambda_2, \dots, \lambda_n\} = \boldsymbol{\lambda}$$

Then we can say in vector space:

$$f_{\mathbf{X}}(\mathbf{k}; \boldsymbol{\lambda}) = \lambda^k / k! * e^{-\lambda} \tag{6.2}$$

Second, given a particular scenario, we needed to estimate the likelihood that the scenario could lead to a claim on any given policy in the insurer's bordereau of policies. This requires two steps:

1. In the first step, we need to describe which exposure categories of policies might even be "eligible" to record a claim on a given scenario. For example, in a Pharmaceuticals scenario, we only want Products Liability insurance policies covering companies in the Pharmaceuticals industry to potentially record a claim, whereas some unrelated policy such as a Trucking Liability policy for a fleet of long-haul truckers would have a zero chance of manifesting a claim in this scenario. This task can be accomplished so long as all the policies in the insurer's portfolio are mapped to suitable categories. Conceptually this is straightforward, but operationally it can be challenging, especially if the insurer has not historically mapped its casualty policies to key "cat exposure categories".
2. In the second step, we need to quantify the likelihood that any given policy within the eligible exposure categories would accrue a casualty claim against it. We can estimate this claim likelihood parameter from the empirical data on historical events and then model the claim likelihood as a Bernoulli random variable. Within an eligible exposure category, we have identified two additional factors that can play a role in predicting claim likelihood. One of these factors is the size of the insured. Historical data shows that very large companies have elevated claim likelihood in casualty cat scenarios. These companies are often specifically targeted in liability lawsuits due to the "deep pockets" effect. The other factor is the global region of the insured. Regions with particularly litigious environments, the United States chief among them, have a higher frequency of casualty cat occurrences, and when they do occur, they tend to be more

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

widespread.

Therefore, by collecting the exposure category, size, and region of each insured policy, we can obtain an informed idea of the claim likelihood for each policy in a given casualty cat scenario.

Mathematically, we can say as follows:

Let Y = random variable representing the claim count for a particular policy, given that a particular casualty cat scenario X has occurred.

Let Y follow a Bernoulli distribution with parameter p for $k = 0$ or $k = 1$:

$$f_Y(k; p | X) = p \text{ for } k = 1; \quad (6.3)$$

$$f_Y(k; p | X) = 1-p \text{ for } k = 0;$$

Further, we have various types of policies and cat scenarios, so we can say as follows:

Y = matrix of policies and cat scenarios = \mathbf{Y}

Each element in \mathbf{Y} follows its own independent Bernoulli distribution with parameter p such that p is also a matrix:

$$p = \mathbf{p}$$

Then we can say in matrix space:

$$f_Y(k; \mathbf{p} | \mathbf{X}) = p, \text{ for } k = 1; \quad (6.4)$$

$$f_Y(k; \mathbf{p} | \mathbf{X}) = 1-p, \text{ for } k = 0;$$

Third, we needed to quantify the size-of-loss dollar amount of the claim severity if a claim were to occur to any given policy within any given cat scenario; essentially, we needed to construct a claim severity loss distribution conditional on the specific cat scenario in play and conditional on the specific insured policy getting hit with a claim. The parameters of this random loss distribution can be estimated based upon historical data; we recommend modeling the mean and variability of the size-of-loss on a pure “from ground up” loss basis and then overlaying the specific attachment point and policy limit explicitly as part of the

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

simulation algorithm. We collected historical data and fit a lognormal distribution to empirical historical mean and coefficient of variation to obtain the parameters of the severity curve.³

Mathematically, we can say as follows:

Let Z = random variable representing the claim size-of-loss amount i.e., “severity” for a particular claim, given that a particular casualty cat scenario X and claim Y has occurred.

Let Z follow a lognormal distribution with parameters μ and σ .

Further, we have various types of cat scenarios, policies, and claims, and so we can say as follows:

Z = matrix of claim severities = \mathbf{Z}

Each element in \mathbf{Z} follows its own independent lognormal distribution with parameters μ and σ such that μ and σ are conditional on \mathbf{X} , \mathbf{Y} .

To recap, we have developed parameters for annual cat scenario frequency, conditional claim likelihood (given a scenario has occurred), and conditional claim severity (given a claim has occurred in a given scenario). These parameters can all be derived from empirical historical data and fit to common statistical distributions. With these parameters and distributions, we now have all the required information to execute the Monte Carlo simulation algorithm described in Section 5.

6.1 Correlation

In the framework described above, we eschew correlations rooted in abstract statistical methods such as correlation matrices. Rather, correlation emerges organically: when a simulated casualty cat scenario occurs, the probability \mathbf{p} of an associated claim rises for all policies in the class of exposures that could be hit by the “storm path” of such a casualty cat scenario. This approach reaps the benefits of generating correlation via an underlying common shock model⁴ together with the benefits of the clear narrative coherence that identifies a specific cat scenario as the “epicenter” of the underlying shock.

³ Here, too, we found a differential behavior for larger companies; note that when the “deep pockets effect” applies, in which larger companies with larger limits accrue larger losses, the “consistency tests” articulated by Miccolis [3] do not necessarily apply.

⁴ See Meyers [2].

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

6.2 Scenario Heterogeneity and Parameter Variance

In keeping with the foundational concepts discussed in Section 3, specifically the idea of creating scenarios supported by and derived from historical events but broad enough to not be “straight jacketed” by the past, we found it appropriate to build in an additional layer of uncertainty into our parameter set. For any cat scenario parameterized in the model, when a future event occurs, that instance of future event might be more severe or more benign than typical; this variation relates to the underlying nature of the event itself but also in how it intersects with the unique attributes of unanticipatedly salient exposures in the insurer’s portfolio. For example, sometimes an insurer’s underwriters may focus on a particular “niche” of an industry when writing policies, creating an exposure accumulation at a level more granular than that captured by the relatively broad “exposure categories” within our model. When any casualty cat event occurs, it is possible to hit a particular “niche” within the broader exposure category particularly hard depending on the specific cat scenario. If this “niche” is the one that the insurer has built up a large exposure in, they will suffer an outsized number of claims compared to the rest of the insurance market.

Mathematically, we are saying that the matrix \mathbf{p} of conditional claim likelihood probability parameters are not fixed parameters but rather are random variables \mathbf{P} with means \mathbf{p} and associated variability. Rather than introduce independent variability for each element in \mathbf{P} , we use a common shock model to introduce correlated variation across all the claim likelihood parameters conditional on the cat scenario.⁵

7. DATA

Compared to property cat, collecting data on historical casualty cat events involves some additional complications. The first complication is the shape-shifting, protean nature of casualty risk. There is much more idiosyncrasy in the historical record of casualty cat events compared to property cat events which makes it more difficult to derive future scenarios. We solve this issue by using broad-based future scenarios, as described earlier in this paper. Another data-related complication unique to casualty cat involves determining accurate historical loss amounts – data around court judgements, settlements, and appeals, along with information on whether these losses were insurable, can be difficult to find and often

⁵ See Meyers [2].

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

requires multiple sources of data. A related issue is that these court-related delays can create very long periods between the event occurrence and the determination of final loss amount for the cat, creating lingering uncertainty long after the event has occurred.

7.1 Updating the Model with New Vintages of Data

The model is refreshed with updated parameters as new data arrives. Sometimes the new data arises from new loss development on already known historical events (e.g., Asbestos); other times, the new data arises from a completely new casualty cat event that has occurred (e.g., Surfside building collapse). Whenever new information arrives, it can lead to changes in empirically measured attributes such as the frequency of casualty cat events, the conditional likelihood of claims given an event, and the conditional claim severity given a claim within a cat event; these changes in the empirical data then cascade into updated mathematical parameters for the updated go-forward stochastic casualty cat model.

8. NUMERICAL OUTPUT

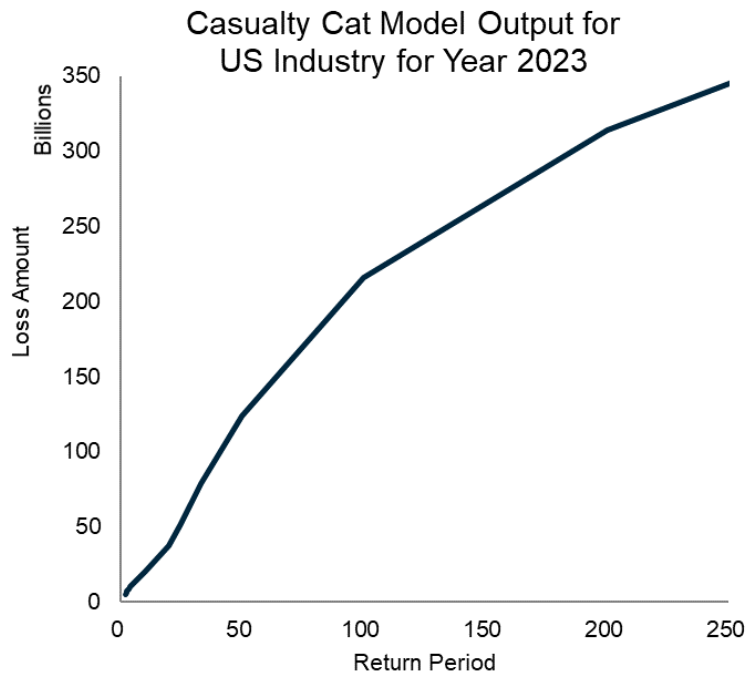
Our model is designed to apply to a specific insurer's unique bordereau of policies and unique exposures. As a complement, we also built a synthetic bordereau designed to approximate the US industry portfolio. After applying our model to the estimated US industry portfolio, we obtained an estimated "1-in-250 year aggregate loss" of \$345 billion in current dollars for year 2023.⁶

As previously mentioned in section 5, the numerical output from the model would consist of all simulated claims. Similar to property catastrophe modeling, the numerical outputs can be viewed in granular detail by claim by policy and can be rolled up to any level of aggregation or such as by peril, by scenario, by business unit, and so forth. The graph in the table below illustrates the aggregate exceedance probability (AEP) curve for the estimated US industry portfolio.

Exhibit 1

⁶ Compare and contrast to D'Arcy [1] who obtained an estimate of approximately \$141 billion in 2016

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk



8.1 the Reasonableness of Numerical Output

Validating

There are two primary ways that we validate the reasonableness of the numerical output. The first step is to check that the software is generating loss amounts at the industry level that align with industry's loss data for historical events, adjusted from the past to the present. The second step is to check that the software is generating loss amounts at the company portfolio level that align with the company's loss data for historical events, adjusted from the past to the present. Although the software is designed to be used for go-forward modeling of new casualty cat scenarios, it is reasonable to validate the model's numerical output by reviewing it against historical casualty cat events.

9. SOFTWARE IMPLEMENTATION

Our desire to implement our actuarial modeling techniques led us to three "journeys" in actuarial software implementation. Our three journeys were:

1. Software: we initially wrote down many of our actuarial algorithms in Excel, which allowed for ease of use and visual inspection. Once we were ready to overlay basic

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

simulations, we migrated from pure Excel and began to incorporate Visual Basic for Applications (VBA), which allowed us to further develop our software. Ultimately, we implemented the eNTAIL model in the R programming language to reap the benefits of a modern, industrial-strength software engine.

2. Platform: we initially developed the software for the desktop, which had its limitations. We migrated to a fully cloud-based platform for ease of use and for consistency of user experience.
3. Algorithm: we initially developed our software algorithms using loops, as described in Section 5. As we progressed in our “journey,” we learned that although traditional loops are possible, it’s dramatically more efficient to use “vectorized” approaches.

To demonstrate why vectorization is more efficient than loops, we’ll compare two functions to simulate the number of occurrences of a scenario that follows a Poisson distribution with mean **lambda** across **n** simulation years. First we’ll look at loops and then we’ll look at vectorization:

```
simulate_scenarios_with_loop <- function(n, lambda) {
  # initialize object to store values
  scenarios <- c()
  # loop through each simulation year
  for(i in 1:n) {
    # simulate 1 value from Poisson distribution
    v <- rpois(1, lambda)
    # add to vector of results
    scenarios <- c(scenarios, v)
  }
  # return the n values
  scenarios
}
```


An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

Each time the loop iterates, R replaces the existing vector in memory with a copy of itself appended with the new value. While this works, it's not the best way to aggregate data in R and we can avoid the overhead of loops while obtaining the same result with vectorization.

```
simulate_scenarios_with_vectorization <- function(n, lambda) {
  # simulate n values from Poisson distribution
  rpois(n, lambda)
}
```

Because R is an interpreted language, it doesn't benefit from the optimization that occurs in compiled code. Vectorization, however, allows you to perform an operation across an entire vector of values. In the function above, all **n** random numbers are generated in one go.⁷

```
library(bench)
# define parameters
n <- 10000
lambda = 0.5
# measure the time it takes to run each function
mb <- suppressWarnings(
  bench::mark(
    loop = simulate_scenarios_with_loop(n, lambda),
    vec = simulate_scenarios_with_vectorization(n, lambda),
    check = FALSE
  )
)
mb[c("expression", "min", "median", "itr/sec", "n_gc")]
```

⁷ Under the hood, these values are ultimately iterated through compiled functions written in C.

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

```
## # A tibble: 2 x 4
##   expression      min   median `itr/sec`
##   <bch:expr> <bch:tm> <bch:tm>    <dbl>
## 1 loop          172ms   186ms     4.96
## 2 vec           264µs   337µs    2818.
```

Comparing the two functions, we see the vectorized approach is over 300 times faster. Applying the principle of vectorization throughout our model creates two key benefits: more concise code that's easier for the developers to manage and faster runtimes for end users.⁸

10. CONCLUSIONS

Casualty cat modeling is an important component of the overall insurance risk management landscape. Though presenting some unique challenges in comparison to property cat, it is feasible to build a casualty cat model that provides meaningful insight into tail risk for casualty insurer portfolios.

⁸ If the reader is set on using loops, Morgan-Wall [4] discusses ways to implement them more efficiently in R.

An Actuarial Approach to Stochastic Modeling of Casualty Catastrophe (Cat) Risk

11. REFERENCES

- [1] D'Arcy, Stephen P., "Casualty Catastrophe Analytics: Where We Are Now and Where We Should Be on This Critical Risk," *Variance* 10:2, 2016, pp. 292-311. <https://www.casact.org/sites/default/files/2021-07/Casualty-Catastrophe-Analytics-Darcy.pdf>
- [2] Meyers, Glenn G., "The Common Shock Model for Correlated Insurance Losses," *Variance* 1:1, 2007, pp. 40-52.
- [3] Miccolis, R., "On the Theory of Increased Limits and Excess of Loss Pricing," *Proceedings of the CAS*, 1977, <https://www.casact.org/pubs/proceed/proceed77/77027.pdf>
- [4] Morgan-Wall, Tyler (2017, September 5) "There's no need to apply yourself" <https://www.tylermw.com/theres-no-need-to-apply-yourself/>

Biographies of the Authors

Neil Bodoff is an EVP at Gallagher Re and can be contacted at neil_bodoff@ajgre.com

Eric Dynda is a VP at Gallagher Re and can be contacted at eric_dynda@ajgre.com

Brandon Stevens is a VP at Gallagher Re and can be contacted at brandon_stevens@ajgre.com

Krish Kamdar is an AVP at Gallagher Re and can be contacted at krish_kamdar@ajgre.com