



Reflections on an Epidemic

By Paul Conlin

On New Year's Eve 2019, my family went out to dinner with another family. Our twin boys are friends with, but go to a different school than, their son, so we took advantage of the holiday to catch up with them and spend time at one of the Chicago area's most famous, and busiest, highest-variety seafood restaurants (Bob Chinn's in Wheeling, Ill., if you're familiar). We had to wait in line for quite some time for a table for seven, but I didn't even mind—I was so grateful for the prosperous decade just past, and so looking forward to ringing in the New Year and the new decade. The 2020s, wow! Where have the years gone? I don't know about you, but I remember New Year's 2000 almost like it was yesterday. And now we were heading into the twenties! What did the future hold? Finally, the rough edges would be smoothed off of technology, the gig economy, internet trolling, vaping by minors, opioid-based painkiller dispensing—and it would be onward and upward to our bright, shining future. Right? Right?

Well, as we all know, it took all of 60 days for that hopeful plan to go off course. A pandemic. Really, people, a pandemic? With 150,000 or more dead Americans? With social distancing? And 14 percent unemployment? And second quarter 2020 GDP growth headed to minus 30 percent or minus 40 percent? And untold intangible damage to the economy for an unknowable period of time? Really??

But ... although the timing was unknowable to all of us happy diners at Bob Chinn's, the event itself—at least the medical part—wasn't, or shouldn't have been. News coverage describes it as “unprecedented.” And viewed through the lens of “during our lifetime,” OK, I'll give you that one. This is the worst thing that's happened in my 54-year lifetime. Yet...

ACTUARIES SHOULD TAKE A LONGER VIEW

We're actuaries. We think in longer than 54-year time spans. And through that lens? Well, this has happened before, and you don't have to go back that far to find it.



Ancient Athens had a well-documented plague occur during the Peloponnesian War. About 500 years later, during the Roman Empire of the Antonines, returning soldiers brought back smallpox with them, probably from Mesopotamia, which became endemic in Europe and was brought by their descendants 1,400 years later across the Atlantic, exposing the nonimmune Native Americans (thus killing over half of them), and was only eradicated in the 20th century. Seventy years after being laid low by the Antonine plague, the even more destructive Plague of Cyprian (so named for a Carthaginian bishop who wrote about it) ran amok through the empire from A.D. 249 to 261 or so. Malaria had long been rampant in swampy areas of Rome itself (and other swampy areas of the world where mosquitoes and humans share living space). A mysteriously thinly documented Plague of Justinian traveled through the Byzantine Empire in 541 and 542, including the capital, Constantinople, itself; paving the way for, among other things, the Slavic peoples to migrate into central and eastern Europe. The bubonic plague made use of the roads that Genghis Khan's heirs opened and maintained through central Asia to travel on fleas, which traveled on rats, which traveled with merchants and sailors (it arrived in Europe in 1347 when a merchant ship made port at Messina, Sicily), and infected and killed possibly one-third of all humanity from 1347 to the mid-1350s.

Travel. Note that theme.

The otherwise-prosperous and progressive 19th century struck an unfortunate balance of intermingling peoples and nations nearly as much as the already lamented, dearly departed 2010s, but without having the medical understanding of the 20th and 21st centuries. And the end result was a toxic brew of pandemics: cholera, polio, typhus, diphtheria, yellow fever, and, yes, influenza. The 1918–19 Spanish Flu spared almost no one, sickening

even King George V of Britain and U.S. President Woodrow Wilson. Unlike COVID-19, it was particularly harsh on the young and healthy and contributed to a deep but brief depression in 1920–1921, which could be preview of the world economy in 2020–2021.

So, no, the COVID-19 pandemic is not unprecedented. Its likes have happened before, and will happen again. But enough history. You're here to read about modeling. Let's get to it.

MODELING AN EPIDEMIC USING DIFFERENCES-IN-DIFFERENCES

The cholera epidemic was a scary one. Even though cholera (which had been locally described by Portuguese and British merchants and likely was already endemic on the Indian subcontinent in the 1700s) broke out around the world in 1817, as late as 1855 its origin (was it bacteria? a fungus? “bad air”?) was still unknown. The man who solved it was one of the founders of epidemiology, John Snow.

Dr. Snow was a physician, but he was also a data modeler, as skilled in modeling as any human living in 2020. [Victorian Britain gave us a dream team of modelers and statisticians: Sir Francis Galton (1822–1911) invented the word “regression” as we understand it today in its regression-to-the-mean sense, studying the heights of successive generations of family members.] In an era without computing, Snow relentlessly collected data on deaths by subdistrict of Britain in different outbreaks (one in 1849, one in 1854) and compared the water companies of each district. This was complicated by two different companies sometimes having unique service areas, and then sometimes both serving the same towns.

Dr. Snow had a hunch, which turned out to be correct, that one of the companies was transmitting a germ (which turned out to be the bacteria that causes cholera, which reproduces and grows in just about any standing water it can embed itself into) via its water pump. He invented a statistical methodology, still used by economists today, called differences-in-differences. (The differences in Snow's model were the Deltas of the death counts in water districts using company A versus the death counts in water districts using company B versus the death counts in districts using companies A and B; the second dimension of differences was deaths in the 1849 outbreak versus in the 1854 outbreak.¹)

Differences-in-differences is a twist on multiple linear regression, in which the coefficients include the usual x -intercept (α), two coefficients (β and γ) of the independent variables, but then a differences-in-differences “dummy variable” with the coefficient (lowercase δ). Given enough data, the four regression coefficients can be solved for and divided by their standard deviation to get a t -statistic and statistical significance. If the delta coefficient is statistically significant, its value

can be assumed to be the quantification of the value of the interaction of the independent variables.

This is the part of the article where I should start showing formulas. But I'm not going to do that. Because with differences-in-differences, there is no one hypothesis formula. It completely depends on the number of independent variables and the combinations of those independent variables that are assumed to interact versus those that are not. The hypothesis is completely at the discretion of the modeler.

Economists (most prominently Alan Krueger—sadly, just deceased in 2019) have used differences-in-differences for all kinds of experiments, often testing the economic impact on market participants from various forms of regulation or other underlying exogenous or endogenous factors. What the successful models do appear to have in common is that the underlying is trending; and various counterfactual trends can be examined via differences-in-differences models. The trending property of the underlying data makes it particularly interesting to actuaries.

Emanuel Derman, formerly of Goldman Sachs and Columbia University, once gave a good analogy for imagining the equation first, then testing it. One of the most important, and strangest, laws of physics is Kepler's Second Law, which states that “the line between the sun and a moving planet sweeps out equal areas in equal times” (e.g., a planet like Mercury, which is close to the sun, orbits very quickly).² How did Kepler, using 16th-century astronomical technology, discover it? Notice the first two words in the quote above: “the line” between the sun and a planet. Do you see the insight? The insight is this: There is no actual line between any two objects in outer space. Kepler completely imagined one, then tested out the results against his observations. Only upon seeing that the calculations worked did he conclude he had stumbled upon a law.

Hypothesis/model first, calculations/validation second.

Brilliant.

A DIFFERENCE IN PERSPECTIVE: FINANCE VS. MATHEMATICS

I conclude with an actuarial example. In 2019, *North American Actuarial Journal* published an article on property and casualty incurred but not reported (IBNR) runoff from 1994 to 2010, written by three finance professors.³ The authors downloaded Schedule P data from all of the orange blanks in the National Association of Insurance Commissioners database and did a multivariable linear regression against the 20 or so characteristics of the companies for all of the 26 or so lines of business that P&C companies write.

They tested the statistical significance of the coefficients and were fortunate to find a nice dispersion of a few, but just a few,

significant coefficients at the 90th, 95th and 99th percentiles. (My personal favorite “discovery” for health actuaries was that medical malpractice coverage has had favorably developing reserves for the past quarter century, which has contributed to favorable medical trend for health plans. The asterisk indicates 90 percent statistical significance.)

But note the difference in the economics or finance approach versus the mathematical approach. The finance professors viewed the coefficients as the repositories of data for analysis and conclusions, not the dependent variable.

Come up with your equation first. Then see if it works.

WHERE DO THE MATHEMATICIANS GO WRONG?

Mathematicians and probabilist statisticians have been plowing ahead during the past half century in this field. An iconoclastic Russian mathematician named Kolmogorov worked out the underlying field paradigm of probability events in the 1930s, which led to the development of stochastic processes. This finding led to ergodic theory, which is the study of processes that repeat in cycles, and what’s happening “inside the system” upon cyclical repetition of the process: Is the original starting state replicated (or nearly replicated), or is it permanently unrecoverable (i.e., has “mixing” occurred)?

A wonderful way to learn theoretical math for actuaries who don’t want to get bogged down in proofs that one plus one equals two, but who still respect the insights and structure offered by pure mathematics, is counterexamples. Counterexamples were originally developed to test the limits, literally and figuratively, of various mathematical theorems. But they work for applied mathematicians, like actuaries, because they force theoretical mathematician authors to step out of their equations and nota-

tion for five-minute chunks of time and give concrete (counter) examples of what they mean by a theory they are explaining.⁴

But if you’re a career actuary, you knew a long time ago you wanted to work in applied math, not pure math. Keep reading about statistical models, whether written by actuaries or economists or finance professors. Be author agnostic. And evangelize on what you learn. ■

Paul Conlin, FSA, MAAA, is an actuary at Aetna. He can be reached at conlinp@aetna.com.

ENDNOTES

- 1 Dr. Snow’s original differences-in-differences table can be found in *Mastering Metrics* by Joshua D. Angrist and Jorn-Steffen Pischke (Princeton University Press, 2015).
- 2 The reason this works mathematically is the Law of Conservation of Angular Momentum, one of the five great conservation laws of mechanical physics. Besides lacking modern observational and computational tools, Kepler lived in the pre-calculus era, so he was unable to express his observations in integral form. For a concise explanation, see *Continuum Mechanics* by A.J.M. Spencer (Dover Publications, 2004). Mechanical physics has lots of parallels to econometrics in that the equation (and type of equation) is sought first, then the calculations begin.
- 3 Barth, Michael M., Evan M. Eastman, and David L. Eckles. 2019. It’s About Time: An Examination of Loss Reserve Development Time Horizons. *North American Actuarial Journal* 23, no. 2:143–168.
- 4 The masterpiece of probability counterexamples is *Counterexamples in Probability* by Jordan M. Stoyanov (Dover Publications, 2013).



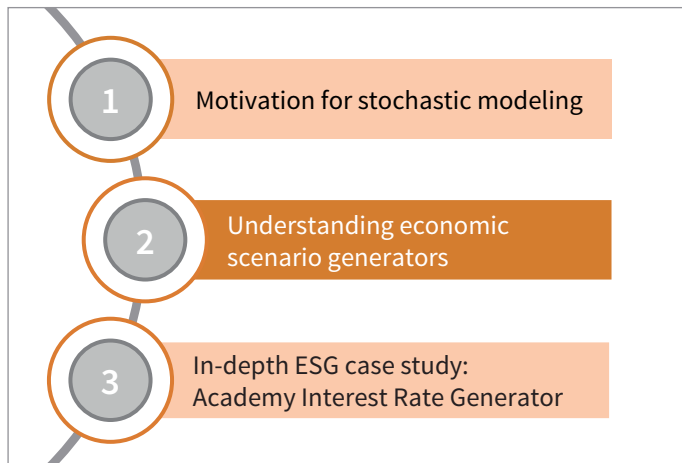
Economic Scenario Generators, Part II

Understanding Economic Scenario Generators

By Rahat Jain, Dean Kerr and Matthew Zhang

This article is the second installment of our three-part series on economic scenario generators (ESGs). Part I was published in the November 2019 issue of *The Modeling Platform*. Part II addresses critical considerations in selecting, building, using and validating ESGs (Figure 1).

Figure 1
ESG Three-Part Series Structure



ESGs, like other models, come in all shapes and sizes, with a wide array of applications. Decisions made in ESG selection and parameterization fundamentally impact the nature of results. It is therefore crucial to understand and challenge these decisions as an end-user. Only by asking the right questions, at a minimum those in Figure 2 (page 5), can one gauge whether a particular ESG is appropriate for their use case. This article walks through each of the decision-making steps shown in Figure 2.



1. RISK FACTORS

Risk factors are the stochastic variables the ESG is being used to generate. ESGs can model a specific individual risk factor or perform holistic simulations of multiple risk factors.

The following risk factors relevant to actuaries are often produced by ESGs:

- interest rates;
- equity or investment fund returns;
- credit spreads, corporate bond yields, default probabilities;
- exchange rates;
- nonmarket factors (e.g., mortality rates, lapse rates); and
- macroeconomic factors (e.g., GDP, inflation, unemployment).

As each scenario may represent a particular outcome, relationships across risk factors within that scenario must be internally consistent. These correlations are often modeled through a cascade structure, in which variables are established in sequence and preceding variables will affect succeeding variables.

An example of a simplified ESG cascade structure is presented in Figure 3, on page 5. In this model, risk-free interest rates are first generated independently, and subsequent risk factors are affected by those preestablished rates.

Figure 2
Key Factors When Making Decisions With ESGs

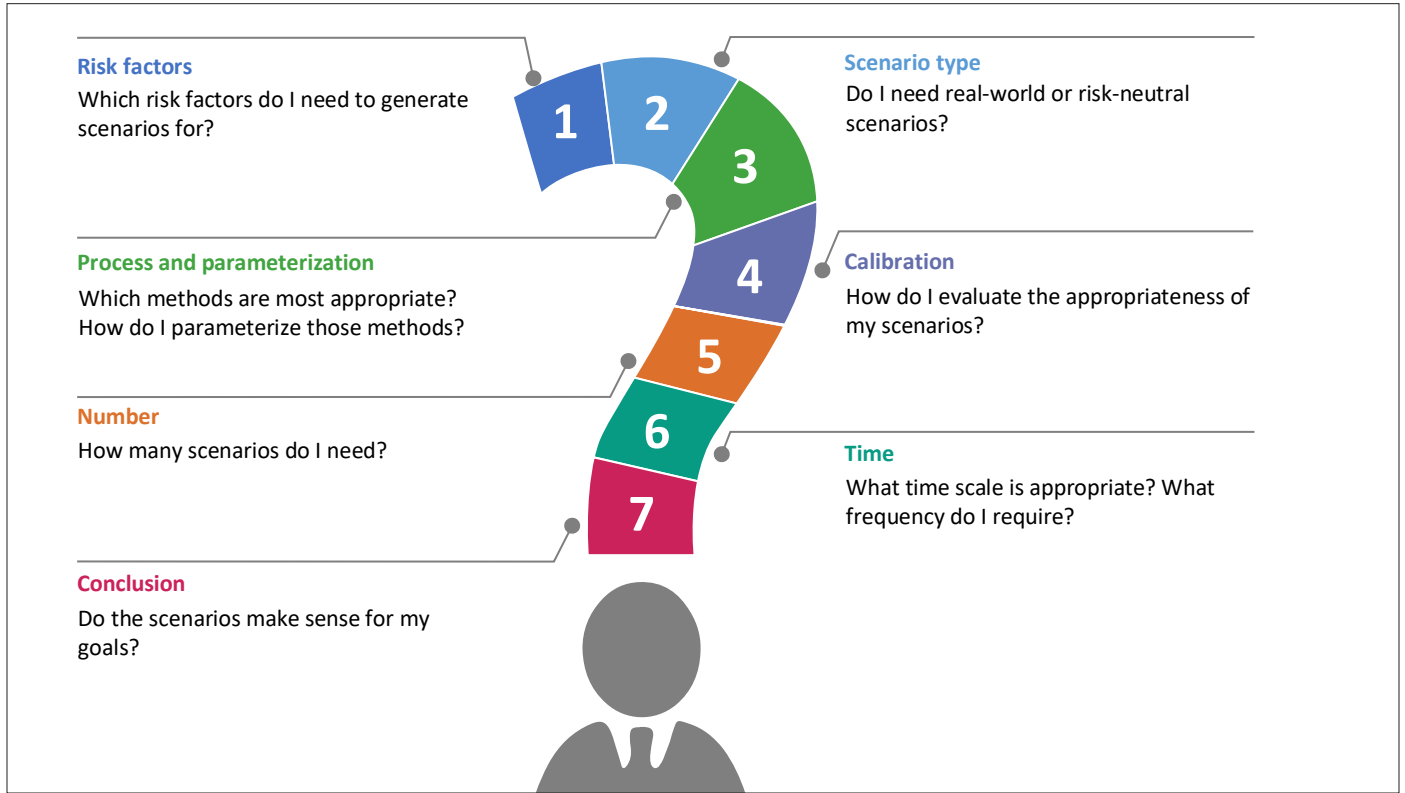
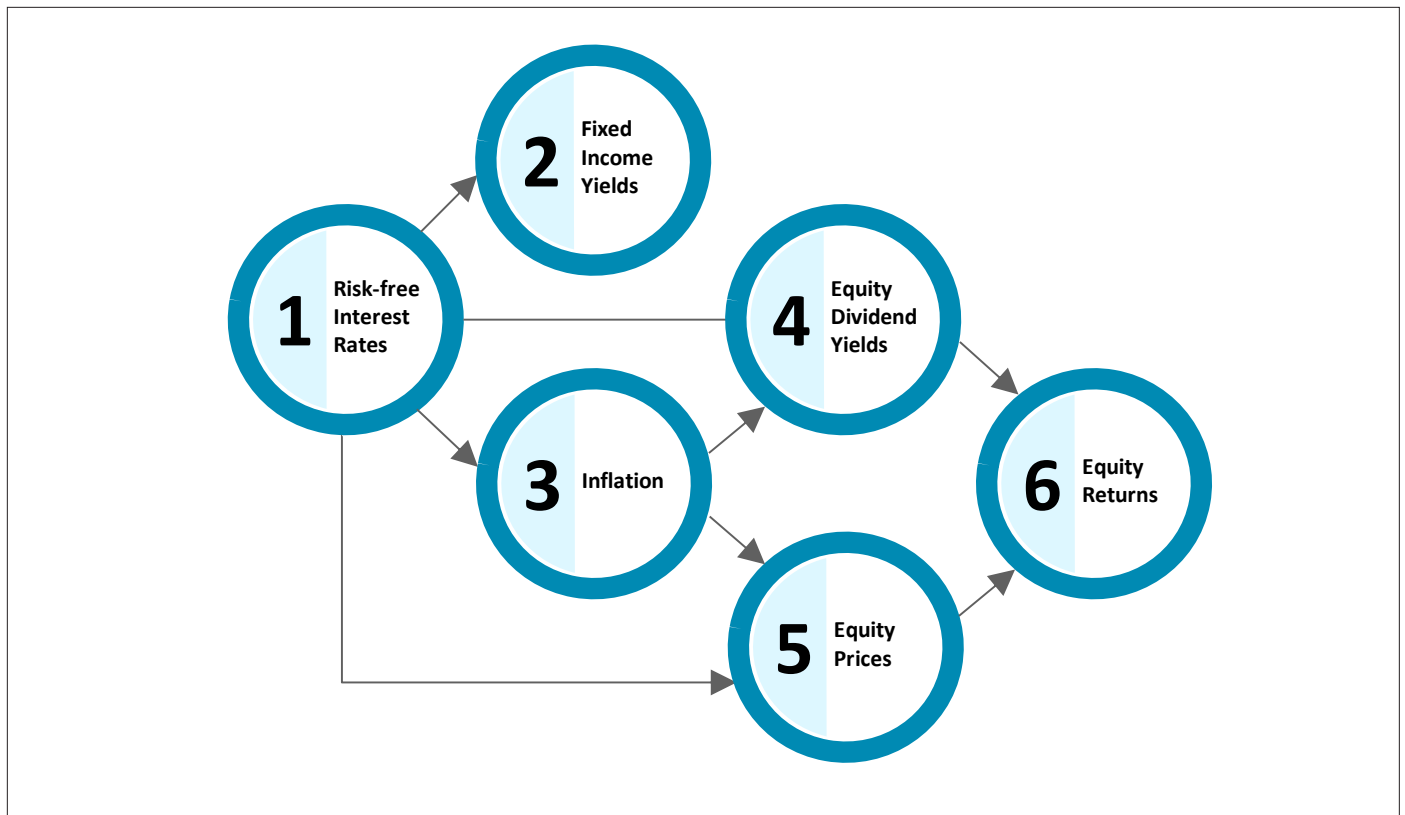


Figure 3
Example ESG Cascade Structure



2. SCENARIO TYPE

There are two primary types of scenarios: real-world and risk-neutral. Figure 4 summarizes key differences between real-world and risk-neutral scenarios.

Figure 4
Comparison Between Real-World and Risk-Neutral Scenarios

	Real-World Scenarios	Risk-Neutral Scenarios
Objective	To generate subjectively “plausible” outcomes and represent real-world outcomes	To satisfy mathematical necessities of market completeness, risk neutrality and lack of arbitrage
Function	Realistic projection of assets and liabilities	Market-consistent valuation (i.e., what a liquid market would pay for these cash flows)
Discount rate	Riskier assets expected to earn a higher expected risk-adjusted return (risk-free rate plus premium)	Assets expected to earn the risk-free rate In rare situations will the discount deviate (i.e., illiquidity premiums for stress scenarios)
Usage	Assessing a range of possible outcomes [e.g., regulatory capital, rating agency requirements, “what-if” scenarios, tail results, value at risk (VaR) and conditional tail expectation (CTE)]	Replicating market prices using simulation, such as Monte Carlo (asset/liability pricing, duration/convexity, market-consistent embedded value, and fair value)
Output characteristics	Scenarios may exhibit “realistic behavior” (volatility clustering, rising yield curve, etc.)	All assets expect to earn the risk-free rate, regardless of risk profile

A given scenario is not inherently real-world or risk-neutral; this trait only becomes visible when considering all scenarios in a given scenario set.

Consider an example with binary outcomes, where prices for a given stock index can increase or decrease. The “best real-world estimate” is that there is a 60 percent chance of increasing. In this case, a real-world set of 1,000 generated scenarios should have 600 scenarios where the index increases and 400 scenarios where the index decreases. The “best estimate,” however, has no bearing on a risk-neutral scenario set. Instead, the scenarios themselves are mathematical fiction, derived to satisfy a specific calibration target (e.g., matching prices for market-observed assets).

3. PROCESS AND PARAMATERIZATION

Once it is decided what type of scenarios is needed for the selected risk factors, the next step is to determine how best to produce the scenarios. Common risk factors like interest rates and equity returns have a range of scenario generation methods backed by a wealth of academic research. The selection of an appropriate ESG model balances satisfaction of key properties against ease of use.

Simple models are generally intuitive and straightforward to parameterize but may not capture important properties. By comparison, more complex models are insightful but may be difficult to use, understand and parameterize.

Consider a simple generic continuous rate process, shown in Figure 5, which could represent anything from a risk-free rate or logarithm of equity return. A typical process consists of a drift component

that sets the general trend of that rate with respect to time (time is the deterministic component), and a volatility component that scales a Weiner process to introduce randomness.

Figure 5
Generic Rate Process

$$dr = \overbrace{\mu(r, t) dt}^{\text{Drift component}} + \underbrace{\sigma(r, t) dW}_{\text{Volatility component}}$$


A straightforward interpretation of this process can be to have the drift and volatility components set as constants.

$$\mu(r, t) = \mu$$

$$\sigma(r, t) = \sigma$$

This process would be clear to understand and parameterize. However, it may lack key properties. For example, if we are modeling interest rates, the use of constant drift will cause the trend of interest rates to monotonically increase with time, which is not realistic. The introduction of a mean-reverting drift process, transitioning us to the commonly used Vasicek model, can solve this problem:

Figure 6
Difference Between Risk-Neutral and Real-World Calibration



$$\mu(r, t) = a * (b - r_t)$$

$$\sigma(r, t) = \sigma$$

Although the results may be more intuitive, we must recognize that our model is no longer as easy to interpret and parameterize as before and must consider both a and b as additional parameters. More complex processes may see multiple factors being simultaneously projected or the setting of individual parameters as a stochastic process rather than constants. However, these complex models may be more capable of achieving specific goals, such as the production of fat tails, negative rates and so on.

4. CALIBRATION

Once a process is selected, one must ensure that parameters are chosen in a manner that satisfies the end need. The selection and cyclical rebalancing of parameters is called calibration.

Appropriate calibration is intrinsically tied to the intended use case, such as generating conservative scenarios for a risk management exercise or producing market-consistent scenarios that will replicate today's prices. The calibration process and priorities differ significantly between real-world and risk-neutral, as illustrated in Figure 6.

Risk-Neutral Calibration

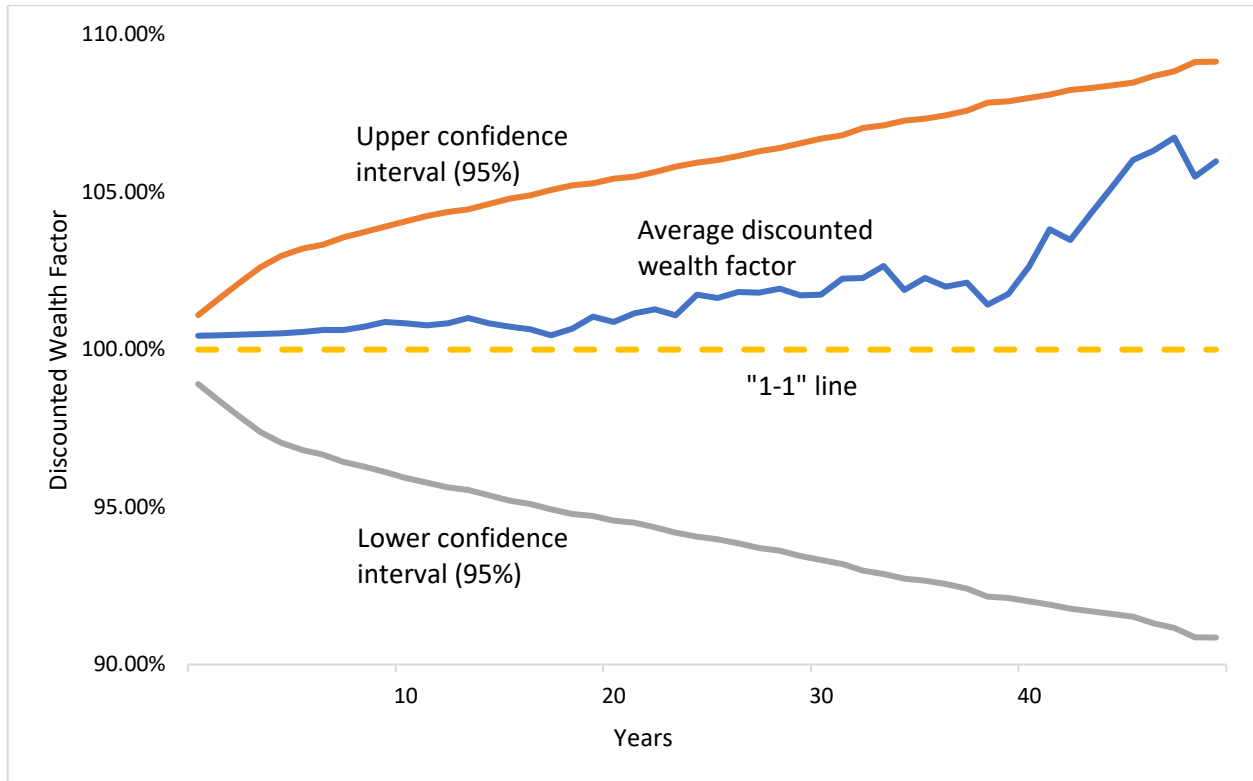
The goal of risk-neutral calibration is market consistency. With the underlying assumption being that markets do not permit arbitrage, projected scenarios must be consistent with today's pricing of market-observed data. Any other outcome suggests the existence of arbitrage. Whether the resultant risk-neutral scenarios are "realistic" has no bearing. Figure 7 summarizes the key tenets of risk-neutral calibration.

Figure 8 demonstrates an example of a Martingale test performed against a risk-neutral set of scenarios. This test is conducted by discounting the total equity fund accumulation back to time zero at the risk-free rate. Although volatile, we see that the average dis-

Figure 7
Key Tenets of Risk-Neutral Calibration

Arbitrage Free—"Fit observable market prices"	Risk Neutral —"No risk premium for riskier investments"
Ensure that scenarios do not imply arbitrage against the observed market	Do not allow risk premium for investing in risky assets
Assessment of fit can be subjective, as one cannot match the entire market	Standard test is the Martingale test (sometimes called the "1=1 test"), which aims to check that assets have the same expected return as the risk-free rate
Goodness of fit should be aligned with level of risk exposure (e.g., if pricing optionlike assets, matching option prices is a priority)	

Figure 8
Martingale Tests for Risk Neutrality



counted wealth factor for the risk-neutral set is reasonably close to 1 and therefore passes the test for risk neutrality.

Real-World Calibration

Real-world calibration has no defined calibration requirements; two different people rarely have the same outlook for the future. As prediction of the future is inherently subjective, different assumptions, approaches and biases will naturally result in different outcomes.

In reviewing real-world scenarios for reasonableness, an actuary should undertake both quantitative and qualitative assessments. A qualitative assessment typically relies on “stylized facts.” Common stylized facts include:

- yields for longer-term bonds being higher than yields for shorter-term bonds;
- equity returns exhibiting higher volatility than interest rates; and
- riskier assets exhibiting higher volatility than less risky assets.

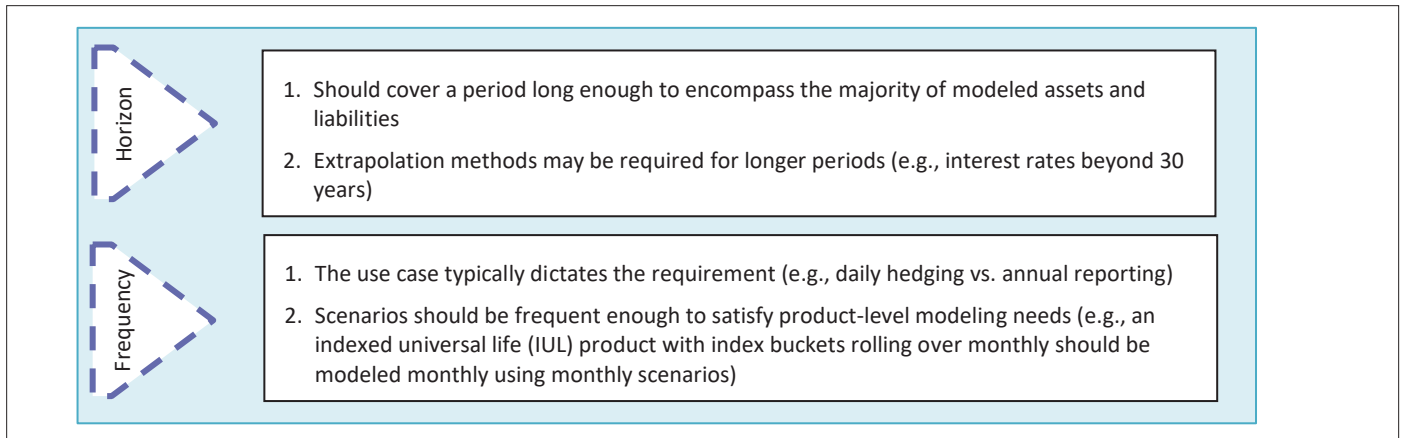
For both aggregate statistics and stylized facts, fidelity is benchmarked against historical experience. If scenarios are significantly divergent from past observations, the scenarios may be less defensible as “realistic.” The historical period used for comparison is important, as is accommodating future expectations. For example, few historical data sets have featured negative interest rates; however, negative interest rates exist today and can be reasonably expected to exist in the future.

5. NUMBER

The number of scenarios within a scenario set is an important assumption as each scenario has unique potential to expose different outcomes and risks. Using more scenarios is always more informative than using fewer scenarios; however, the relative usefulness of increasing the scenario count is very much tied to the objectives and the nature of the underlying model. A variable annuity product with exposure to multiple investment funds and dynamic policyholder behavior will exhibit different (and more volatile) interest-rate scenario-driven outcomes than a payout annuity.

Using a smaller number of scenarios may pose a risk to convergence of resultant metrics and statistics and reduce the consistency of outcomes. At the same time, a drawback of too many scenarios is model runtime and data constraints. Hence, the

Figure 9
ESG Actuarial Timing Considerations



key is to strike a balance between the number of scenarios that achieve robust results without becoming unmanageable.

6. TIME

Time horizon and scenario frequency are both considerations when building an ESG. Time horizon is the period over which the scenario forecasts (e.g., projecting 100 years into the future). Frequency is how often a scenario’s values change (e.g., every month). Figure 9 describes some common considerations.


7. CONCLUSION

The seventh—and most important—question is whether the scenarios being produced by an ESG make sense for a particular use case. Actuaries often outsource the production of economic scenarios and may view scenario generation as being beyond their core skill set. By failing to challenge fundamental assumptions underpinning ESG results, the risk of misapplication increases.

Indeed, Actuarial Standard of Practice (ASOP) 56 states clearly in section 3.1.6 that “for models that use assumptions as input, the actuary should use, or confirm use of, assumptions that are **appropriate given the model’s intended purpose**” (emphasis added). The seven-step ESG decision-making process outlined in this article can serve as a starting point for actuaries to understand and challenge the scenario sets upon which they so often rely.

Stay tuned for the final installment of our three-part series as we take a closer look at the Academy Interest Rate Generator (AIRG), the most commonly used real-world ESG for U.S.-based actuaries. As a case study, we will be applying the seven-step framework outlined in this article to the AIRG. ■

The views or opinions expressed in this article are those of the authors and do not necessarily reflect the official policy or position of Oliver Wyman.



Rahat Jain, FSA, CERA, MAAA, is a consultant at the actuarial practice of Oliver Wyman. He can be reached at Rahat.Jain@OliverWyman.com.



Dean Kerr, FSA, MAAA, ACIA, is a partner at the actuarial practice of Oliver Wyman. He can be reached at Dean.Kerr@OliverWyman.com.



Matthew Zhang, FSA, CERA, MAAA, is a consultant at the actuarial practice of Oliver Wyman. He can be reached at Matthew.Zhang@OliverWyman.com.



Tidy Data Formats, Part 2 Applications of Tidy Mortality Tables

By Matthew Caseres

In this article we continue our work on tidy mortality tables by creating a format that can join mortality rates to data sets in a single line of code. For purposes of this demonstration, I have combined these mortality rates with lapse rates and demographics from a 2014 RGA study in order to create a semirealistic data set. This data set is then used for a Monte Carlo simulation to generate a distribution of claim amounts. I then discuss the creation of open-source implementations of regulations as well as trade-offs between the freedom to use actuarial judgment and consistency of results across companies.

R code snippets are included for the discussion of simulating claim amounts. Any programming language can be used for the tasks described here, but a language meant for data work, like R or Python, will provide constructs that make tasks easier.

Figure 2
Joining the Tidy Mortality Table to Experience Data



TIDY COMBINED SELECT AND ULTIMATE MORTALITY RATE TABLES

In Part 1 of this series, published in the April 2020 issue of *The Modeling Platform*, I wrote a script that allowed me to reformat hundreds of select mortality tables into a tidy format. I now create a representation of a mortality basis that can be attached to a data set in a single line of code. See that Figure 1 is a tidy mortality table that contains select and ultimate mortality rates for each combination of issue age and duration up to the end of the mortality table.

Figure 1
Tidy Mortality Table With Select and Ultimate Mortality Rates

issue_age	duration	attained_age	q_sel	q_ult	q
60	60	119	NA	0.50000	0.50000
60	61	120	NA	0.50000	0.50000
61	1	61	0.00152	0.00772	0.00152
61	2	62	0.00310	0.00860	0.00310
61	3	63	0.00525	0.00957	0.00525

Use in Experience Studies

Once this table has been created, I am able to easily join mortality rates using a multi-key join, as demonstrated in Figure 2.

Storage

We currently use mortality tables from a web page maintained by the Society of Actuaries (SOA). However, I believe storing common tables as CSV files in cloud file storage would be easier to



maintain and use than the existing solution. I have created a **One-Drive** directory as an example but do not plan on actively maintaining the reformatted tables. I hope that others working in this area can extend this work into something that is well maintained.

I investigated Google BigQuery (a serverless analytics data warehouse) as a storage option but found that the platform is meant to store large data sets for analytics, rather than collections of tables. Tables are not discoverable in BigQuery, since there is no graphical representation of a directory of files.

Simulating Data

As I am unaware of any data sets that can be used to demonstrate methods related to lapses or premium patterns, I created a simulated data set that contains decrements from both lapse and mortality to aid in the demonstration of methods.

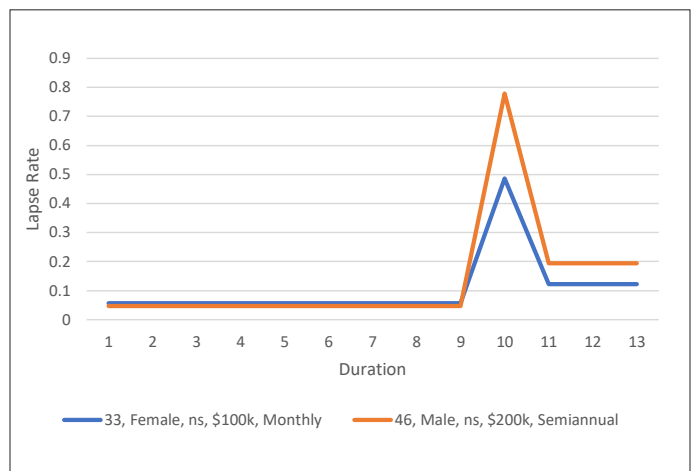
Methods

The data created is for a 10-year term life product with decrements due to lapse and death. The approach is to simulate a random number indicating a death when the simulated number is less than the mortality rate. I simulated another random number to determine if the policy has lapsed within a given year. It is pos-

sible that a policy would both lapse and die within a given year, in which case I used a tie-breaker random number that assigns the policy as lapsing or dying, each with a probability of 0.5. If the policy neither lapsed nor died, I moved to the next time step and ran the simulation again with rates for the new attained age. The simulation returns the year and cause of decrement.

My previous work on tidy mortality rates allowed for extraction of mortality rates up to the end of the mortality table for a given issue age, gender and underwriting classification. For lapse rates in duration 10, I used the general linear models factors from page 108 of the 2014 **Report on the Lapse and Mortality Experience of Post-Level Premium Period Term Plans** by RGA. Durations past the 10th duration were assigned one-fourth the duration 10 rate, and durations before the 10th were determined using a set of factors that I thought would give the data a reasonable shape. Figure 3 shows the lapse rates for two simulated policies.

Figure 3
Lapse Rate by Issue Age, Gender, Underwriting, Face Amount, Premium Mode



The gender, underwriting, face amount and payment mode were simulated using the data proportions given in the previously referenced RGA study. Issue ages were randomly assigned to ages from 30 to 55.

Figure 4 shows an example of what each simulated record looks like at this point.

To make this a VM-51 format, I randomly selected an issue date between Jan. 1, 1990, and Jan. 1, 2020, and randomly selected a

Figure 4
Example Simulated Records

Issue Age	Gender	Underwriting	Face Amount	Premium Mode	Decrement Year	Decrement Cause
35	Male	Nonsmoker	\$100,000	Monthly	9	Lapse

termination date from the year of decrement. When the simulated termination date was past the assumed current date of Jan. 1, 2020, I changed the decrement cause to “Truncated” to indicate that at the date of analysis I did not know the cause or time of decrement. A date of Jan 1., 2020, was assigned to the termination date field for truncated data. Now our data has an issue date and termination date as required.

You can access the 100,000 simulated VM-51 format policy records in BigQuery:

```
SELECT * FROM `soa-mortality-demo.tidy_mortality.VM_51`
```

MONTE CARLO MORTALITY DISTRIBUTION IMPLEMENTATION

I believe that the best way to communicate methods is to produce open-source code. Communicating methods in natural language can lead to differences in results between companies.

Method for Simulating Deaths

I used the simulated data and the expstudies R package to create a data set with rows for each policy year. Taking the 100,000

simulated VM-51 records from our previous discussion, I generated 752,990 intervals representing policy years for each policyholder.

```
expstudies::addExposures(records)
```

Some manipulations are performed to put the data in the mortality study format described in the SOA publication on [Experience Study Calculations](#). In the SOA publication, exposures depend on whether the year is a leap year or not. This has the consequence of inflating mortality in leap years by a factor of 366/365 compared to non-leap years. In my calculation I considered 365.25 days to be an exposure. I do not expect choices on this matter to have a material impact.

Because I had split the data into policy years, I could join the corresponding mortality rate for the record’s smoker status, gender, attained age and duration from the tidy VBT2015. Again, we see the utility of having mortality tables in a tidy format. A column of random numbers is generated and used to simulate deaths.

```
#mortality_data is modified as described in Experience Study Calculations SOA publication.
mortality_with_table <- mortality_data %>%
  inner_join(VBT2015,
            by = c('tobacco', 'gender', 'attained_age', 'duration'))
mortality_with_table$rand = runif(nrow(mortality_with_table))
mortality_with_table <- mortality_with_table %>%
  mutate(sim_death = if_else(rand <= q, 1, 0))
```

I performed this simulation 1,000 times and generated a distribution of results. It may be desirable to aggregate the results at varying levels of granularity. I would recommend aggregating the data at the most granular level and summing these results to

produce the less granular results. This avoids performing multiple aggregations on the data set. Since I used random numbers in the simulation, I set a seed to ensure that others can run the code and get the same results.

```
set.seed(1)
female_ns_col <- rep(0, 1000)
female_s_col <- rep(0, 1000)
male_ns_col <- rep(0, 1000)
male_s_col <- rep(0, 1000)
female_col <- rep(0, 1000)
male_col <- rep(0, 1000)
ns_col <- rep(0, 1000)
s_col <- rep(0, 1000)
```

```

total_col <- rep(0, 1000)

for(i in 1:10000) {

  mortality_with_table$rand = runif(nrow(mortality_with_table))
  mortality_with_table <- mortality_with_table %>% mutate(sim_death = if_else(rand <= q*-
exposure, 1, 0))

  #aggregate at most granular basis of analysis
  single_simulation <- mortality_with_table %>%
    group_by(gender, tobacco_mapped) %>%
    summarise(sim_death_amt = sum(sim_death*face))

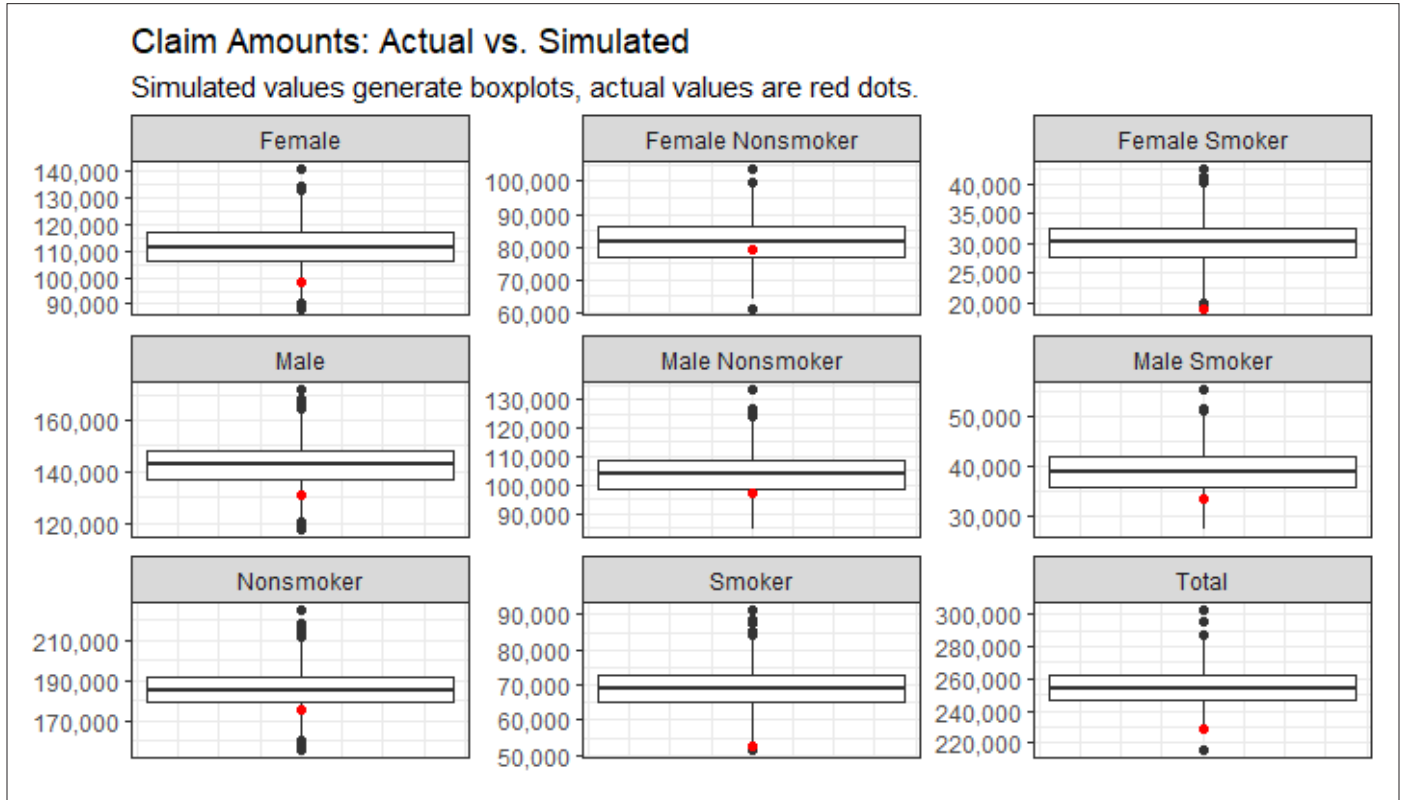
  female_ns <- single_simulation[1,3] %>% unlist()
  female_s <- single_simulation[2,3] %>% unlist()
  male_ns <- single_simulation[3,3] %>% unlist()
  male_s <- single_simulation[4,3] %>% unlist()

  female_ns_col[i] <- female_ns
  female_s_col[i] <- female_s
  male_ns_col[i] <- male_ns
  male_s_col[i] <- male_s
  female_col[i] <- female_ns + female_s
  male_col[i] <- male_ns + male_s
  ns_col[i] <- female_ns + male_ns
  s_col[i] <- female_s + male_s
  total_col[i] <- female_ns + male_ns + female_s + male_s
}

results <- tibble(female_ns = female_ns_col,
  female_s = female_s_col,
  male_ns = male_ns_col,
  male_s = male_s_col,
  female = female_col,
  male = male_col,
  ns = ns_col,
  s= s_col,
  total = total_col)

```


Figure 5
Boxplots Showing Actual vs. Simulated Claim Amounts




Displaying Results

In Figure 5 I have created boxplots from the distribution of simulated claim amounts and show a red dot that represents the actual outcome of the simulated data. This is like the Monte Carlo testing done in the National Association of Insurance Commissioners Model Regulation XXX and I used it to see if the simulated data has a mortality rate consistent with the mortality rate used to generate the data.

CONCLUSION AND THOUGHTS

I have shown how to simulate data and use it to communicate actuarial methods. There are many natural language documents that ultimately serve as specifications for software. I often find myself confused about how exactly I should implement something, which is why I advocate for implementations when giving technical specifications.

Open-source implementations can serve as guidelines for what to do but would not seek to replace the law. I can imagine the specific details of a company would require the freedom to exercise judgment. Freedom to exercise judgment comes at a price, though. The more freedom granted in structuring assumptions, the more potential variance there is in results due to differences in methods. ■



Matthew Caseres, ASA, is a master's student in computer science at Georgia Institute of Technology. He can be reached at matthewcaseres@outlook.com and on GitHub at github.com/ActuarialAnalyst.



STAY UPDATED

Your source for all things modeling!



SECTION COMMUNITY

Connect with like-minded individuals by joining our members-only [LinkedIn SOA Modeling Section group](#). Join the conversation or share your opinion of present and future modeling practices. Discussion topics include model development, validation, controls and governance, operation and efficiency, scenario generation, stochastic modeling, assumptions in modeling such as mortality and longevity, and more. Join the discussion today!

Every vote counts! Be sure to vote in the [2020 Section Council Elections](#). Read the candidate bios and submit your vote Aug. 17–Sept. 4.

PROFESSIONAL DEVELOPMENT

Given today's environment, the trajectory of modeling technologies and applications, and the mission-critical nature of modeling within insurance organizations, the time has come for the chief modeling officer. [Listen to our latest podcast](#) to hear April Shen, FSA, MAAA, interview Van Beach, FSA, MAAA, on the role of chief modeling officer in insurance companies. Learn about the responsibilities of a chief modeling officer and the organizational and institutional benefits of this role.

Get access to more info at SOA.org/sections/modeling