Article from

The Modeling Platform

August 2020

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

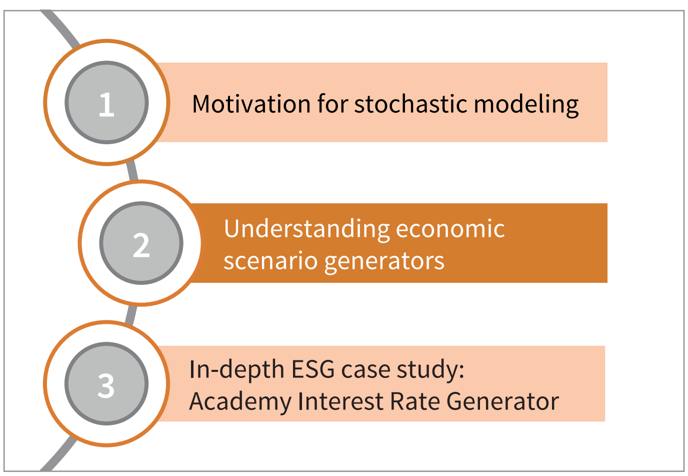

ESGs, like other models, come in all shapes and sizes, with a wide array of applications. Decisions made in ESG selection and parameterization fundamentally impact the nature of results. It is therefore crucial to understand and challenge these decisions as an end-user. Only by asking the right questions, 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

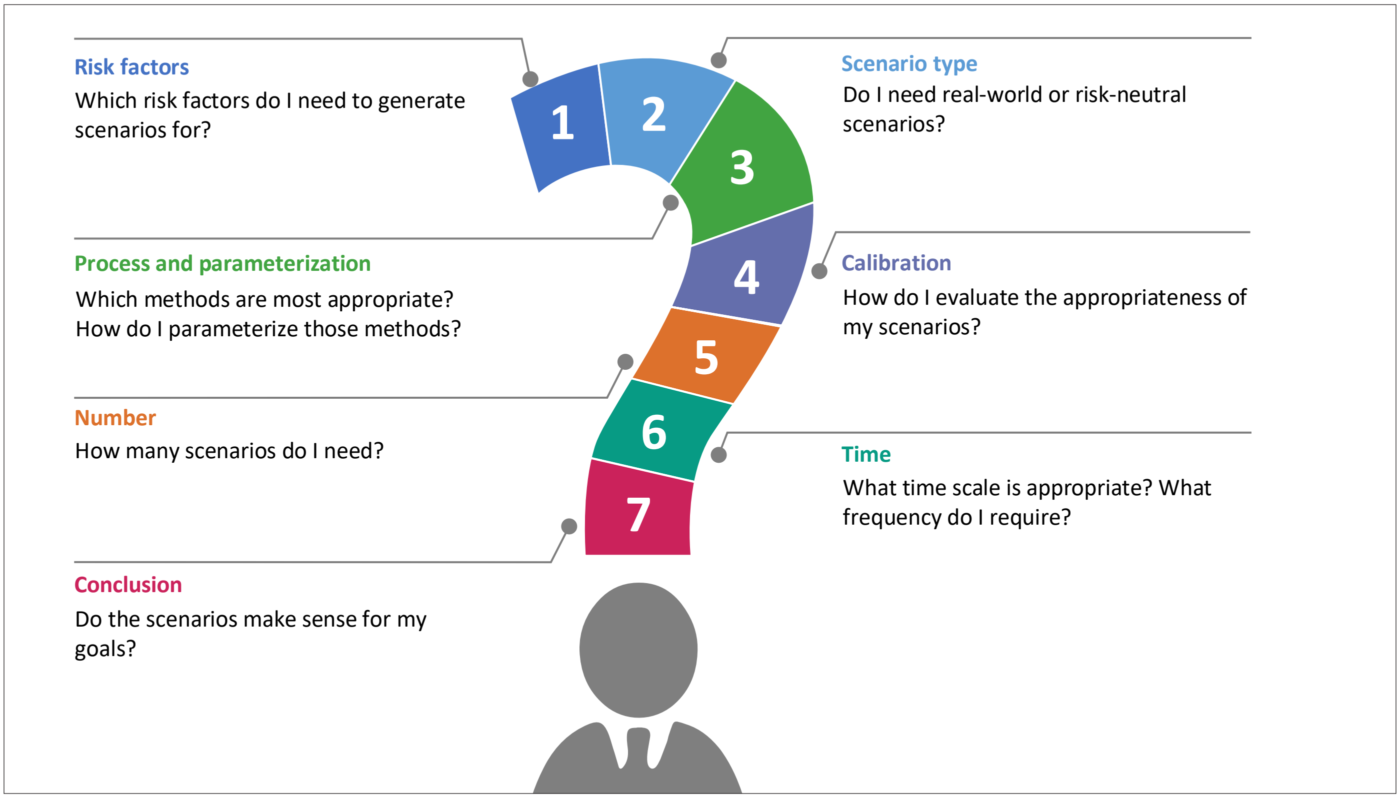

Figure 3
Example ESG Cascade Structure

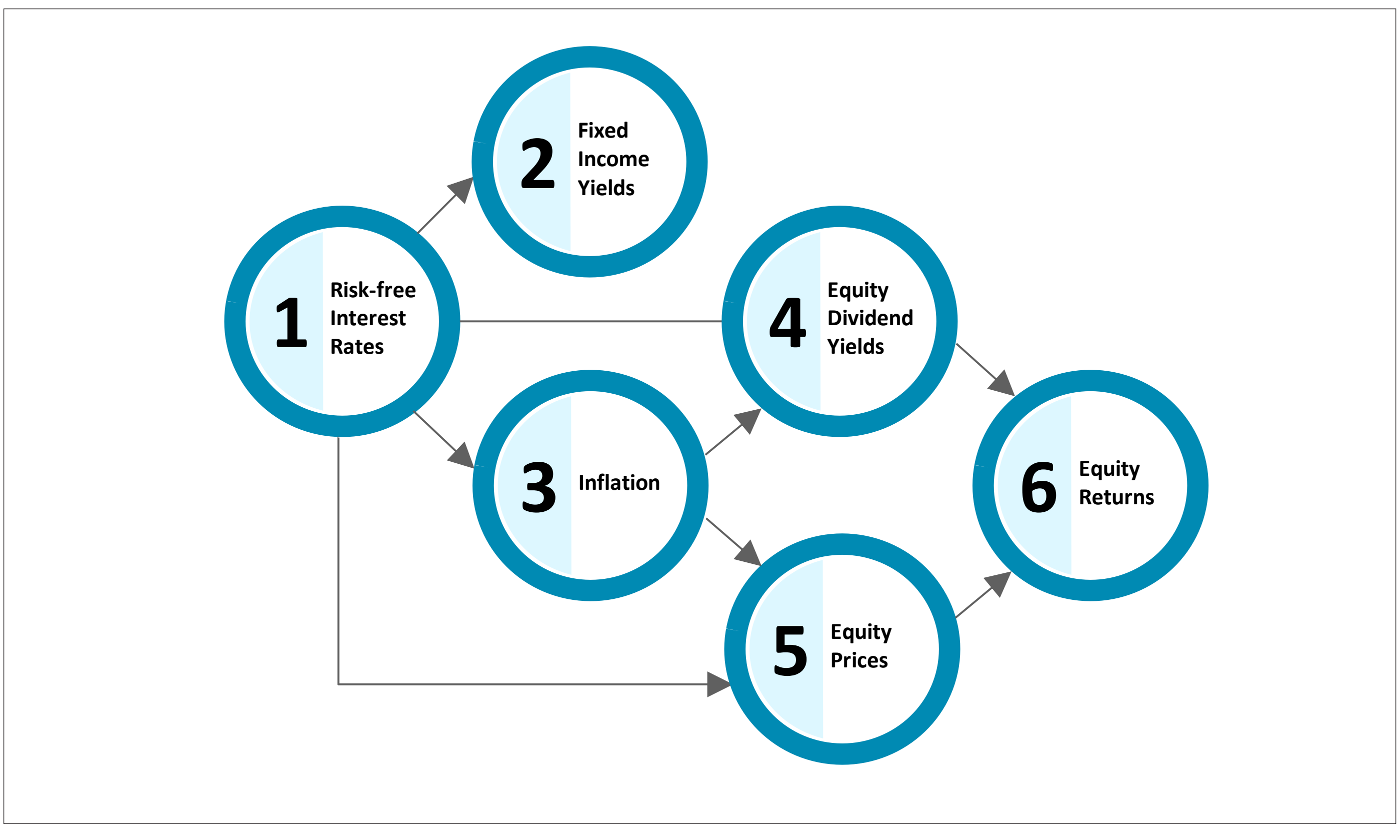

2. SCENARIO TYPE

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

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

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

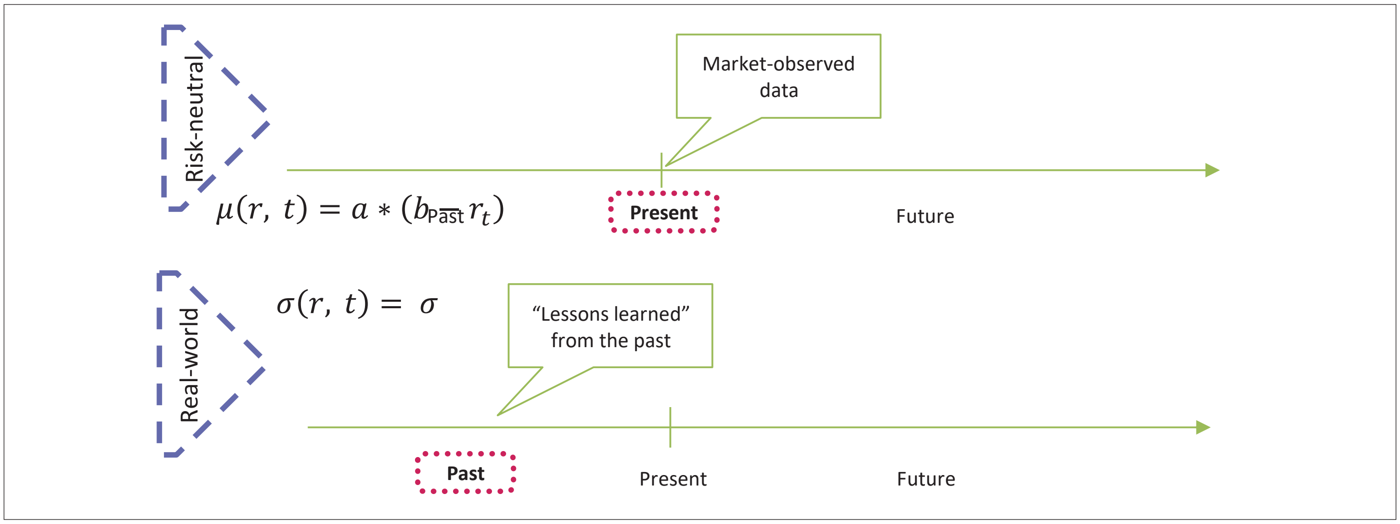

$$\mu(r,\ t) = a * (b - r_t)$$

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

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 Assessment of fit can be subjective, as one cannot match the entire market Goodness of fit should be aligned with level of risk exposure (e.g., if pricing optionlike assets, matching option prices is a priority)	Do not allow risk premium for investing in risky assets 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

Figure 8
Martingale Tests for Risk Neutrality

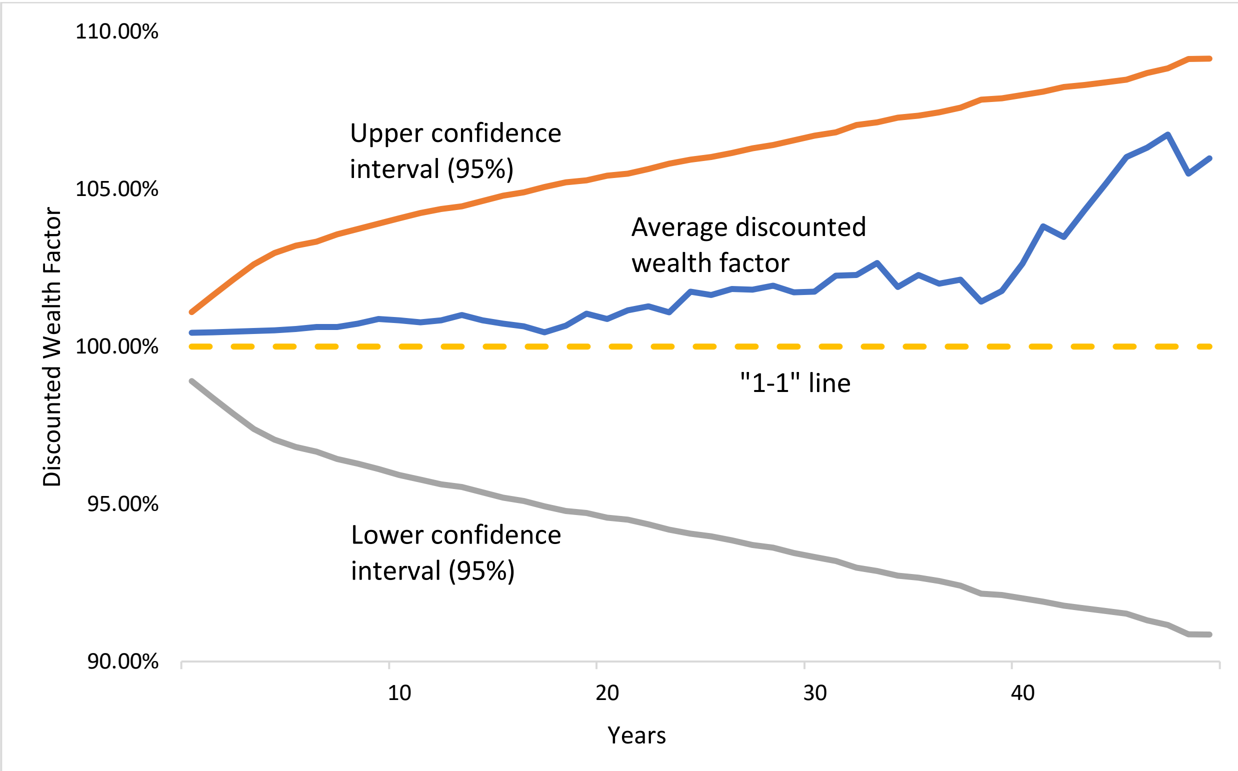

counted wealth factor for the risk-neutral set is reasonably close to 1 and therefore passes the test for risk neutrality.

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

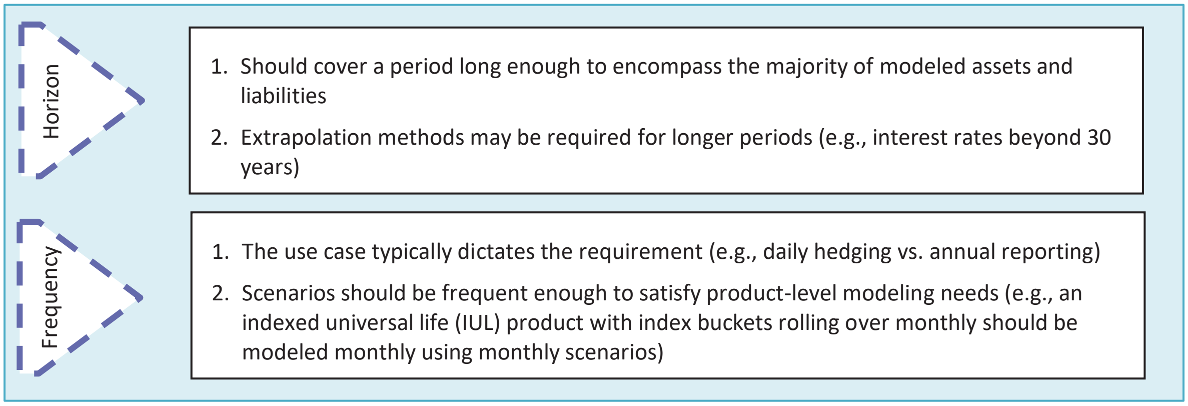

key is to strike a balance between the number of scenarios that achieve robust results without becoming unmanageable.

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.*