Article from

The Modeling Platform

April 2016
Issue 3

A Framework for ESG Validation

By Bruce Rosner

With all of the modeling changes that have been taking place in the last few years, many companies have undertaken a full validation of their economic scenario generators (ESGs). In this article, I discuss the facets of such a review. As a standard cookie-cutter approach is not possible, I focus on key conceptual topics and questions that should always be a part of the review process, and then cover some of the tools and techniques available.

For the purposes of this discussion, let's first lay out a scope of the review. The scenario generation process might follow this sequence of steps:

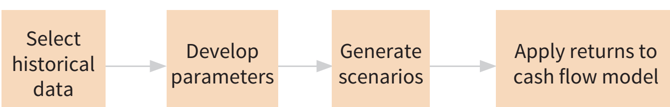

Note that the flow will often vary. For example, some risk-neutral models will calibrate the scenarios directly to market prices, there may be model efficiency techniques that introduce additional complexity, and other wrinkles may emerge in the process.

A review would typically include the whole chain and not just check whether the scenario generator produces reasonable scenarios. The review should also help determine whether the generator is using the input parameters in a manner consistent with how they were developed, and that the cash flow model is interpreting the scenarios correctly. It is actually far more common for the process to break at these points of interpretation rather than within one of the subprocesses.

Having established the start and end points of the review, several questions should be asked.

Underlying distribution

- Is the distribution of the underlying process fit for purpose and permitted under applicable guidelines?
- Empirically, are the final scenarios produced by the model consistent with the underlying distribution?

Parameterization

- Are the assumptions fit for purpose (e.g., calibrated to market data for risk-neutral vs. historical data for real world)?
- Where applicable, are the initial parameters consistent with market/historical data?
- Was good judgment exercised in the development of parameters that do not have clear data (e.g., mean reversion patterns)?

Scenario generation

- Are the scenarios produced by the model consistent with the assumptions that were entered?
- Where applicable, did the scenarios effectively reproduce market prices?
- Are the scenarios consistent with applicable calibration criteria?

Other downstream items

- Were scenario reduction techniques fit for their purpose?
- Were scenario reduction techniques implemented correctly?
- Did the company use enough scenarios for the downstream results to converge?
- Is the downstream model interpreting the scenarios correctly?

The remainder of this article will highlight some practices and pitfalls that have emerged in real-life cases.

DISTRIBUTIONS FIT FOR PURPOSE

Equity Index Distribution

The most common equity index distribution used by insurance companies is the lognormal distribution, which is convenient for many reasons. This type of distribution is simple, and it allows for parameters to easily be developed and for analytical solutions for some problems (such as the Black-Scholes option pricing formula).

However, it does not accurately reproduce historical distributions; in particular, it is known to underestimate the likelihood of significant losses. The lognormal distribution also assumes constant volatility, which is contrary to observed market prices for options. For some applications, based on the accuracy requirements and the product being modeled, this approach may be appropriate. However, many insurance industry applications focus on the extreme events, and, as a result, the lognormal distribution is sometimes thought to be overused.

Alternative distributions for equity models include the Heston model, jump diffusion models or regime-switching models.

Interest Rate Distribution

Practices around interest rates vary widely. Common models include:

1. Short-rate models (for example, the Cox-Ingersoll-Ross model, which allows for a drift, mean reversion and volatility proportion to the square root of the short rate)
2. Principal component analyses (in which the level, slope and curvature, and other shape changes, are directly simulated based on historical movements in the yield curve)

Consider the following: Does it incorporate mean reversion? Does it allow for changes in the shape of the yield curve? Ultimately, is it effective at modeling the specific risks in the product?

For further reading on types of interest rate models, take a look at December 2013 guidelines by the Canadian Institute of Actuaries' Committee on Life Insurance Financial Reporting that classifies different types of "acceptable" models: "Calibration of Stochastic Risk-Free Interest Rate Models for Use in CALM Valuation."[1] Focus on Section 7, which is about medium-term rate guidance.

PARAMETERS' CONSISTENCY WITH MARKET/HISTORICAL DATA

There are several ways to derive parameters, and each of these methods has different implications for appropriate validation techniques.

- Real-world generators using historical data: This is the most straightforward category, but there still is a fair amount of judgment to be applied, including how far back to collect data, and how much to rely on the data. One approach is to recalculate the parameters for at least a selection of indices.
- Risk-neutral generators using market data to directly fit parameters (for example, deriving implied volatility from market put option prices): One needs to exercise judgment over how credible the prices are at longer durations, using put option prices that are at a similar level of "moneyness" to the liability in question, and how to develop a long-term volatility assumption.
- Risk-neutral models in which parameters are derived through a calibration process: This may occur in addition to the step above (using market data to directly derive parameters). At this point, there is no further derivation to review; however, one can check how well the model reproduces market prices, and that would function as a single validation process that reviews the final scenario set itself, implicitly checking both the underlying distribution and the parameters at the same time.

In other cases, the results can be visually analyzed for reasonableness, and the analysis can also check whether the model is capable of producing the types of environments that have historically been observed.

The validation should check whether there is a documented rationale for any judgment in the process, and that the calculations are consistent with that documented rationale.

EMPIRICAL TESTING OF SCENARIOS

The final scenario outputs (typically, periodic returns, by scenario) can be summarized into meaningful analytics for validation. For example, in the case of a simple lognormal distribution, it is possible to derive the mean, standard deviation and correlation directly from the scenario data, and check whether those are consistent with the desired distribution.

In other cases, the results can be visually analyzed for reasonableness, and the analysis can also check whether the model is capable of producing the types of environments that have historically been observed. For example, how often does the interest rate generator produce upward-sloping and downward-sloping yield curves?

It is also possible in some cases to apply statistical tests such as the Kolmogorov-Smirnov to check whether the resulting distribution follows the expected distribution. In the simplest case, the lognormal distribution can be tested directly through a one-sided test[2] against a theoretical distribution using specified parameters. Also, any distribution can be tested by creating an independent tool and using the two-sided test to check whether the two ESGs produce results that are statistically the same.

One instance in which these tests catch issues is when a company interprets normal parameters as lognormal parameters and vice versa.

As noted earlier, for a class of risk-neutral models, the ultimate test of the scenarios, including the underlying process and parameters, is how effectively they reproduce market option prices.

SCENARIO REDUCTION TECHNIQUES

It can be very difficult to test scenario reduction. An ideal validation approach would involve a company producing results using the full scenario set and a reduced scenario set, under various market conditions. The results can then be directly compared to determine if the reduced run is close enough for the stated purpose of the analysis. However, real life sometimes calls for more simplicity. For example, if a company is simply unable to produce the full run, the only alternative may be to have a reduced run passing statistical tests and rerun the reduced version a small number of times to demonstrate that the final results converge.

As for pitfalls, the techniques must be considered carefully in light of the situation. If a company is only interested in the mean, a technique where a different set of scenarios is used for each policy may be appropriate. However, if the company is interested in any other point on the distribution, this approach will create invalid results.

A good scenario reduction technique might introduce a bias into the result but the company may consider it acceptable anyway. An example of this is where the technique is known to produce a more conservative result, and the application permits additional conservatism.

CLOSING THOUGHTS

An economic scenario generator can have many facets, and any validation process should consider the reasonableness of the methodology and outputs, as well as how the results are used, and compliance with applicable regulations. It is a critical point in a company's modeling infrastructure, and as a practices evolve, the organization should take the time to review and lock it down as securely as it does with the valuation system and other highly controlled systems. ■

The views expressed herein are those of the author and do not necessarily reflect the views of Ernst & Young LLP or the global EY organization.

Bruce Rosner, FSA, MAAA is a senior manager with Ernst & Young in New York. He can be reached at *bruce.rosner@ey.com*.

ENDNOTES

1 Document 213107 at the Canadian Institute of Actuaries, *http://www.cia-ica.ca/docs/default-source/2013/213107e.pdf.*

2 A one-sided test checks whether the empirical distribution is consistent with a theoretically correct distribution. A two-sided test checks whether two empirical distributions are consistent with each other.