SOCIETY OF ACTUARIES

Article from:

Risks & Rewards

August 2013 – Issue 62

LAYERING YOUR OWN VIEWS INTO A STOCHASTIC SIMULATION—WITHOUT A RECALIBRATION

By Tony Dardis, Loic Grandchamp and David Antonio

Editor's Note: This article summarizes the authors' presentation at the May 2013 Life and Annuity Symposium.

There is a useful technique that is gaining some popularity amongst practitioners that enables "own views" to be superimposed on a stochastically generated scenario set without having to recalibrate the underlying model(s). The technique won't be appropriate for all applications, but it may be effective for some, such as superimposing alternative estimates or creating "what-if" scenarios and stress tests.

The approach involves the use of a statistic known as entropy. The value of entropy is maximized when equal weighting is given to each individual scenario in a given scenario set. Thus the objective of the "own views" exercise is to re-weight scenarios so that they hit your specific target, while maximizing the value of entropy.

Mathematically, for a set of N scenarios each with weight W_i the entropy S of a scenario set is defined as follows:

$$S = -\sum_{i=1}^{N} w_i \ln w_i$$

There is potentially an infinite combination of weightings that would hit any specified target that we may have. The objective of the entropy technique is to find the optimal weights W_i so that we maximize S while hitting our specific target.

A simple example will help reinforce the concept.

Let us assume that we have calibrated a one-year interest rate model so that on average it hits 3 percent. We then use the model to generate five scenarios as follows:

SCENARIO RATE	ONE-YEAR PROJECTED
1	2.0%
2	2.0%
3	3.0%
4	4.0%
5	4.0%
Average	3.0%

The maximum entropy for these scenarios will always be achieved by equally weighting these scenarios, and is calculated as follows:

FIVE SCENARIOS EQUALLY WEIGHTED

Scenario	One-year projected rate	Weight	Entropy
1	2.0%	0.20	0.3219
2	2.0%	0.20	0.3219
3	3.0%	0.20	0.3219
4	4.0%	0.20	0.3219
5	4.0%	0.20	0.3219
Avg/total	3.0%	1.00	1.6094

Let us now say that we have an alternative view as to what will transpire in the future and instead would like to re-weight the scenarios so that on average we hit a lower rate—say, 2.5 percent. Our first inclination might be to give weighting only to the lowest rates from our original set of five scenarios, with entropy calculated as follows:

TARGET = 2.5 PERCENT; CHOOSE ONLY THE VERY LOW WEIGHTS

Scenario	One-year projected rate	Weight	Entropy
1	2.0%	0.25	0.3466
2	2.0%	0.25	0.3466
3	3.0%	0.50	0.3466
4	4.0%	0.00	0.0000
5	4.0%	0.00	0.0000
Avg/total	2.5%	1.00	1.0397

However, weighting only the very low rates is missing a lot of very important information about the overall distribution of the rates and this becomes apparent from the entropy value—a much lower number than what we started with for the original scenario set that as equally weighted. So let's now consider what happens if we give some weighting to all the scenarios, while still hitting the "own views" target of 2.5 percent:

TARGET = 2.5 PERCENT;
SOME WEIGHTING TO ALL SCENARIOS

Scenario	One-year projected rate	Weight	Entropy
1	2.0%	0.35	0.3674
2	2.0%	0.35	0.3674
3	3.0%	0.10	0.2303
4	4.0%	0.10	0.2303
5	4.0%	0.10	0.2303
Avg/total	2.5%	1.00	1.4256

This simple example demonstrates the key features of using the maximum entropy technique:

- Maximum entropy is where equal weighting is given to all scenarios.
- Minimum entropy is where one scenario is given a weight of one, and all other scenarios a weight of zero.
- The optimization algorithm favors solutions where the weight is as evenly distributed across scenarios as possible. This ensures we don't overweight any particular batch of scenarios and thus ensures we retain as much as possible the features of the original probability distribution.

Let's now consider a practical, and very topical, example. The American Academy of Actuaries provides a basic interest rate and equity scenario generation capability on its website *actuary.org*. This is made available to practicing actuaries as a means of meeting the reserving and capital requirements of variable annuity business under Actuarial Guideline 43 and C3 Phase II which require a stochastic valuation. Related to this, the Academy has also posted a set of 10,000 interest rate and equity scenarios, which practitioners can download without needing to use the generator itself. Some actuaries have gone on to use the generator and/or scenarios for applications beyond meeting the statutory requirement, and while there are very significant limitations to this (the Academy generator was originally developed as a "starter pack" for purposes of enabling companies with relatively simple asset-liability profiles to meet

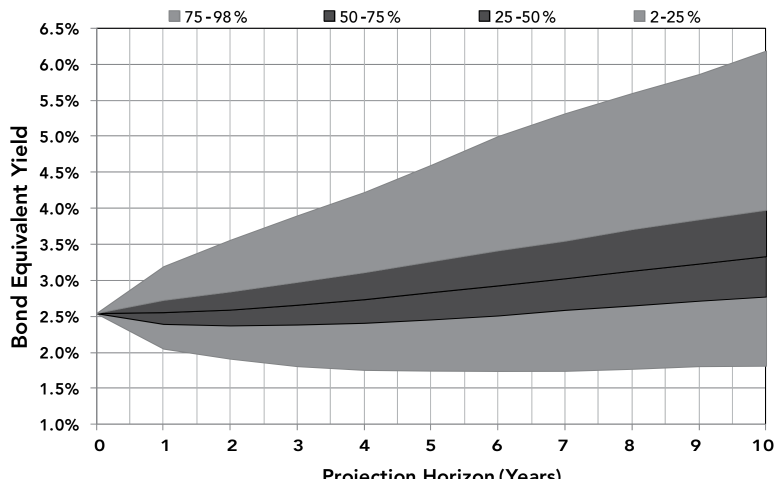

CHART 1:
10-YEAR PROJECTION OF THE 20-YEAR TREASURY BOND EQUIVALENT YIELD USING THE AMERICAN ACADEMY OF ACTUARIES' SCENARIO GENERATOR INITIALIZED TO THE TREASURY YIELD CURVE AT 12/31/2012, 10,000 TRIALS.

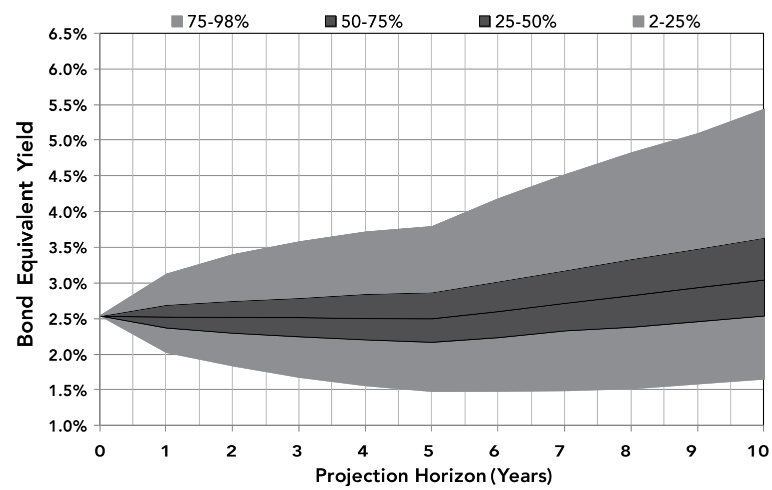

CHART 2:
10-YEAR PROJECTION OF THE 20-YEAR TREASURY BOND EQUIVALENT YIELD USING THE AMERICAN ACADEMY OF ACTUARIES' SCENARIO GENERATOR INITIALIZED TO THE TREASURY YIELD CURVE AT 12/31/2012, 10,000 TRIALS, BUT REWEIGHTED USING MAXIMUM ENTROPY TO KEEP RATES CONSTANT ON AVERAGE FOR THE NEXT FIVE YEARS.

CONTINUED ON **PAGE 26**

the newly emerging statutory stochastic requirement), there may be some applications, such as testing of a new product campaign where interest rate risk is the only market driver of the business, where this is appropriate. This in turn leads to the natural question: if I download the Academy's 10,000 scenarios, is there a way I can rebalance these so that they produce an average interest rate path that is different to what is assumed in the Academy calibration?

The entropy technique can be used extremely effectively in such an example, and has great flexibility. In Chart 1 (pg. 25, top) we show the distribution of the 20-year Treasury bond equivalent yield projected over a 10-year horizon under 10,000 Academy scenarios. These scenarios were generated from the Academy generator, initialized to the Treasury yield curve at 12/31/2012.

As will be immediately apparent, the average path of the 20-year rate under the Academy calibration immediately sets off on an upward trend which persists throughout the projection period. What if our "own view," however, was that given the current economic climate, and the very high expectation that the government will persist in a monetary policy that continues to keep interest rates at extremely low levels, a much more realistic expectation is that on average rates will remain at or very close to today's levels for at least the next five years?

Our first port of call might be to download the Academy's interest rate generator from actuary.org and recalibrate that so that it hits our target. The path of the 20-year rate in the generator can be controlled using two parameters: long-term mean reversion level and a speed of mean reversion. There isn't, therefore, sufficient flexibility to target a more general path for interest rates and the user is also constrained because he can only directly control the evolution of the 20-year rate. It may also not be obvious to the user what parameters should be input to achieve the desired path. Perhaps entropy can help?

In Chart 2 (pg. 25, bottom) we show a revised distribution of the 20-year Treasury bond equivalent yield projected over a 10-year horizon that starts with the 12/31/2012 10,000 Academy scenarios, but reweights using maximum entropy in order to maintain the current level of the 20-year rate over the next five years.

There are some very interesting features of the new distribution that should be highlighted:

- The entropy technique has worked beautifully in hitting our "own views" path on average.
- The overall characteristics of the probability distribution in terms of dispersion and tails are similar under the original and the reweighted scenario sets.
- It will be noted that the lower band of the reweighted set at the second percentile level is well below the original set. This doesn't mean to say that we are weighting scenarios that were outside the original scenario set, but rather that we are now giving a lot more weighting to scenarios that were originally outside the second percentile level. This highlights another important characteristic of the entropy method—it will not work where our target falls outside any of the scenarios that were originally generated. In this example, this would mean that we may not be able to use entropy to target ultra-low 10-year rates, e.g., at or close to 1 percent for a prolonged period.
- The entropy of our reweighted scenario set corresponds to a set of 8,353 equally weighted scenarios. This number, called the effective number of scenarios is a useful statistic of the entropy method which allows practitioners to gauge how far apart the original and reweighted sets are. However, the technique should not be viewed as a scenario reduction technique.

Note that while this article has focused on looking at interest rate scenarios and how we can reweight according to "own views" around a target path, other variables and target metrics could be used equally effectively. For example, we might be more interested in setting a target for returns rather than yield, and perhaps it is equity rather than interest rate scenarios that are of most interest. Indeed, theoretically it would even be possible to take a pre-generated real-world scenario set and optimize the weights to pass martingale tests and hence create a set of risk neutral scenarios. While we wouldn't necessarily recommend such an approach—this would be a considerably more complex exercise than having a relatively simple target such as a different long-term interest rate path, and also creating a set of risk neutral scenarios directly from a genuine calibration is a much easier task than creating a set of real-world scenarios directly from a full calibration process—it still highlights how flexible the entropy approach can be.

Another point to make about the attraction of the entropy technique is that it avoids the scope for negative scenario weights, which other methodologies might not handle so satisfactorily. That is to say, because of how the entropy value is calculated, looking at the log of weights, it won't accept negative weights. So trying a "trick" such as weighting a scenario you really like by, say, 1.2, and one you don't like by -0.2, just wouldn't work.

While the entropy technique holds much promise for certain uses, it comes with a number of words of warning:

- It is not a model recalibration, and is not a substitute for recalibration.
- Although the integrity of each individual scenario from the original set is maintained, validation work is needed when more than one risk variable is being modeled, e.g., equities as well as interest rates. What does a new target for one variable mean in terms of targets for other variables, and what is the impact on correlations?
- As we get further out into the tails, we need to be increasingly careful. Generally, the entropy technique will be very effective for metrics that are close to the central estimate, e.g., CTE(70), but less effective for metrics that are focused on the tails, e.g., CTE(95).
- A large original scenario set will be needed for effective re-sampling.
- Not all "own views" targets can be achieved, i.e., they may fall outside the range of the original set.

With that said, if these limitations are recognized and understood, there may be a number of applications for which the technique can be effective:

- Weighing up the relative merits of strategic decisions where risk and return and focused on the inner tails of the distributions, e.g., decision to launch product X versus product Y, or testing of a variety of different asset mixes.
- Testing the relative impact of different "own views.
- Stress testing and sensitivity analysis.
- Ensuring ownership and consistency of economic assumptions used throughout various business units.

REFERENCES

Avellaneda et al. (2001) – Weighted Monte Carlo: A New Technique for Calibrating Asset-Pricing Models. International Journal of Theoretical and Applied Finance, 4, 1-29.

Tony Dardis is director, Product Specialist with Moody's Analytics Inc, responsible for sales and client support for the company's North American insurance client base.

Loic Grandchamp, PhD, FCAS, MAAA, is P&C product manager in Enterprise Risk Solutions, Insurance, at Moody's Analytics. He can be contacted at loic.grandchamp@moodys.com.

David Antonio, PhD, is a quantitative specialist specializing in Economic Scenario Generation and the Barrie & Hibbert product suite at Moody's Analytics.