OCTOBER 2016

MOODY'S ANALYTICS

RESEARCH
INSURANCE

Proxy Methods for Run-off CTE Capital Projection: A Life Insurance Case Study

Authors

Aubrey Clayton PhD
Steven Morrison PhD
Moody's Analytics Research

Ronald Harasym FSA FCIA CFA
Andrew Ng FSA CFA
New York Life Insurance Company

Contact Us

Americas
+1.212.553.1658
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia (Excluding Japan)
+85 2 2916 1121
clientservices.asia@moodys.co

Japan
+81 3 5408 4100
clientservices.japan@moodys.com

Overview

Approximation techniques such as Least Squares Monte Carlo have proven to greatly reduce calculation time for stochastic-on-stochastic problems requiring an estimate of a risk-neutral distribution mean (e.g., option value). However, required reserves and capital in the North American insurance industry are often defined by a Conditional Tail Expectation (CTE) of run-off deficits under a real-world stochastic projection.

Previous research demonstrated a theoretical extension of proxy methods to projecting CTE measures for simple example products. In this paper, we show a practical application to forecasting capital requirements for real portfolios of participating whole life and annuity business, carried out in a joint research project between Moody's Analytics and New York Life Insurance Company.

Table of Contents

1. Introduction

Reserve and capital requirements, either for regulatory or internal purposes, for North American insurers are commonly calculated through a stochastic run-off projection of assets and in-force liabilities, with the measure of interest being defined by a "tail statistic" of a distribution of accumulated deficits over the lifetime of the liability. A prevalent measure of risk in this context is the Conditional Tail Expectation (CTE), defined as the average over the portion of the deficit distribution beyond a given threshold. So, for example, the CTE(70) measure corresponds to the average of the worst 30% of deficits, the CTE(90) corresponds to the average of the worst 10%, and so on. Depending on the extremity of the tail threshold, a single calculation of this sort may require tens of thousands of Monte Carlo scenarios to achieve a desired level of accuracy.

An emerging business requirement, however, is the ability to calculate such capital measures not only at a single point in time but at *future* times under various economic stress projections or planning scenarios. Thus, what was already an onerous stochastic problem can become an intractable nested-stochastic calculation, with scenario requirements in the millions and beyond. Though the ultimate business application is quite different, the nested-stochastic structure of the problem bears a strong conceptual similarity to the problem of projecting market-consistent balance sheets (for example, for the purpose of computing 1-year value-at-risk), for which proxy methods have shown widespread success in recent years.

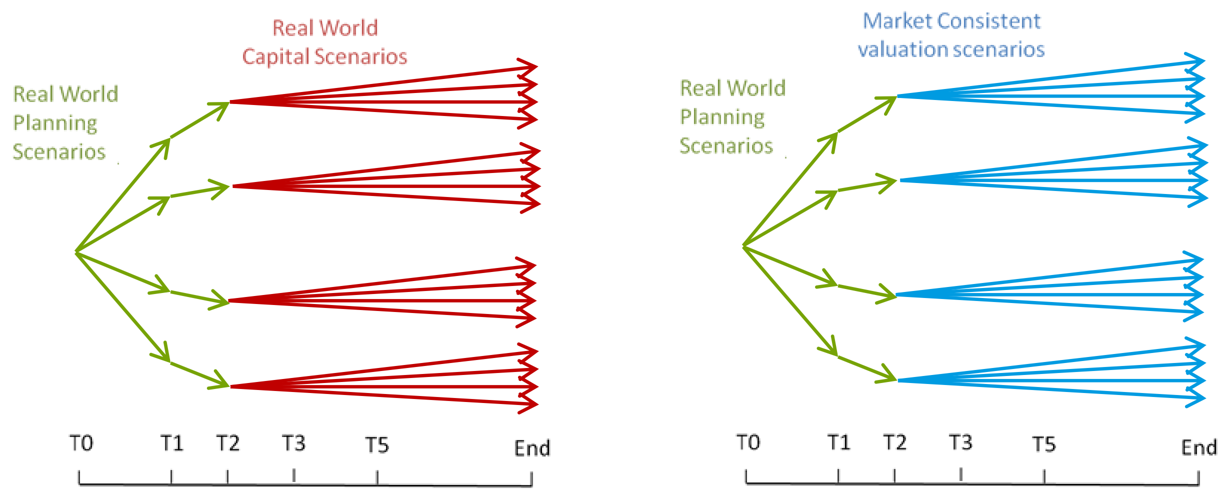

The essential idea of the Least Squares Monte Carlo (LSMC) proxy approach is to exploit the fact that similar projection/stress scenarios measured by some set of explanatory economic risk variables should produce similar valuation results, and thus one can posit a continuous functional relationship of the form

$$MarketValue_t = f(\{Risk\ Variables\}_t)$$

Assuming a simple function form (such as a polynomial) and applying regression techniques allows the functional relationship to be extracted from a set of "training scenarios," where the innermost loop of stochastic scenarios has been dramatically reduced in number to give deliberately *inaccurate* estimates of the metric of interest. Provided that these "crude estimates" are unbiased, the errors are averaged out by the function-fitting process and the estimated function converges to the true theoretical function. Once the functional relationship is known, re-computing the market value under any stress involves a quick function call rather than a stochastic simulation.

Morrison, Tadrowski, and Turnbull (2013) first described how a similar LSMC technique could be used to extract a proxy function for a 1-year projection of a CTE reserve measure, and Morrison, Turnbull, and Vysniauskas (2013) extended the technique to multi-year projections:

$$Reserve_t = g(\{Risk\ Variables\}_t)$$

Their analyses were limited to a fairly simple example product with a guaranteed minimum return on a corporate bond investment. In this paper we demonstrate the practical application to a real-world problem of projecting capital requirements for a block of insurance business.

To illustrate the method for extreme tail measures, for this project we considered the example problem of quickly calculating a Stochastic Required Capital ("SRC") measure at a horizon one year beyond the in-force date under a wide range of market stresses (interest rates, equities, and so on.), with capital defined as the CTE(99) of the distribution of present-value of maximum balance sheet deficiencies over a 40-year run-off projection.[1] To accomplish this, we attempted to produce a proxy function for the SRC as a function of the variables comprising the market stress that would validate well against full-blown stochastic calculations, at a significant savings in overall scenario budget.

For the purposes of the project, we limited our analysis to two individual product groups as well as the collective pool of both:[2]

- Participating Whole Life ("OL").
- Fixed Deferred Annuities ("FDA").
- Aggregate of OL and FDA ("Aggregate").

Section 2 outlines the project methodology, with emphasis on how this differs practically from the usual LSMC approach for market-value projection, and section 3 shows the proxy validation results and discusses example applications of the capital proxy functions.

2. Methodology

The process of designing and calibrating a proxy function generally involves four major steps:

- Identifying relevant risk factors that are allowed to vary and generating "fitting points" to fill out the risk factor space.
- Producing "crude" estimated values at each fitting point using a small number of inner scenarios.
- Fitting a proxy function through the crude estimates using regression or other function fitting techniques.
- Validating the resultant functions at a relatively small number of "validation points," using a large number of inner scenarios to construct an accurate value at those points.

These steps are common to both the mean-value proxy problem and the CTE proxy problem. However, each step required some modifications for projecting CTE as we outline below.

2.1 Stresses and Real-world Recalibration

The economic risk variables defining the one-year market projection for this case study were:

- Nominal yield curves, described by two principal components of yield curve movements.
- US equity returns (S&P 500 index).
- Corporate credit spreads, described by the one-factor credit risk model in the Moody's Analytics Economic Scenario Generator (ESG).

Thus, in total there were four dimensions in the risk factor space, the ranges of which were chosen to cover a wide variety of extreme 1-year stresses. A spanning set of 2,048 fitting points was selected out of this space using Sobol sampling. In addition, 81 validation points were hand-chosen to test the proxy function performance. By design, many of the validation points were near the extremes of the fitting space. Figure 1 below shows the locations of the fitting and validation stresses within the risk factor space.

[1] The definition of required capital given here and used throughout this paper is for illustrative purposes only and does not necessarily reflect the capital measures used by New York Life for any purpose.

[2] The results shown in this paper are for illustrative purposes only and do not necessarily reflect any particular line of business operated by New York Life.

Figure 1: Fitting and validation stresses

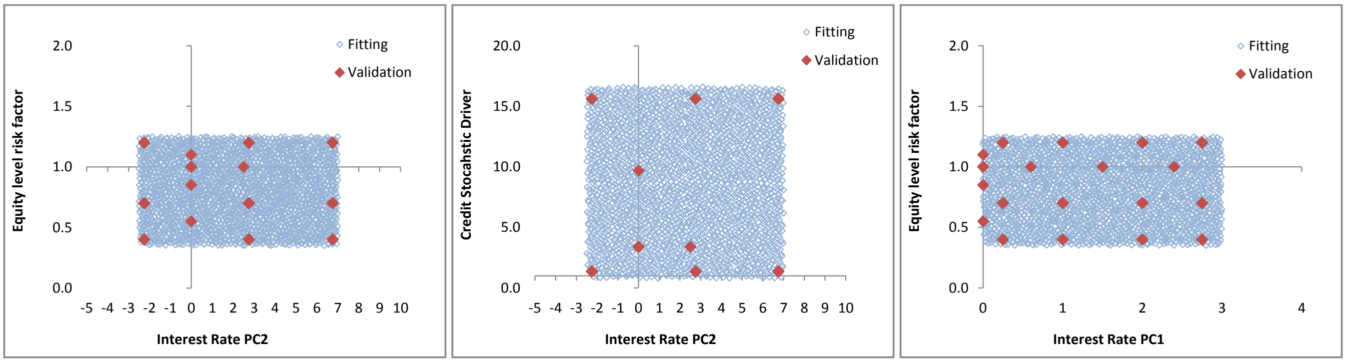

Estimates for the required capital under each stress were to be calculated using 40-year scenarios (at quarterly outputs) of the ESG, initialized at each point in the risk factor fitting space. As usual, this required recalibrating the ESG models to be consistent with the chosen risk factor values. However, unlike market-value LSMC applications, the "inner" ESG models in this case were *real-world* models instead of risk-neutral. Therefore, the necessary recalibration involved updating real-world *embedded views* for quantities such as interest rate term premia or projected inflation.

Since this recalibration process had to be repeated a large number of times, an algorithm was defined to translate the initial yield curve (in forward-rate terms) into a target for mean projected interest rates, and the Moody's Analytics Calibration Tools were used to calibrate the ESG model parameters to those targets.

Figure 2 illustrates one such example yield curve stress and the corresponding forward-rate curve and target mean interest rate path.

Figure 2: Example real-world recalibration (target mean interest rate path)

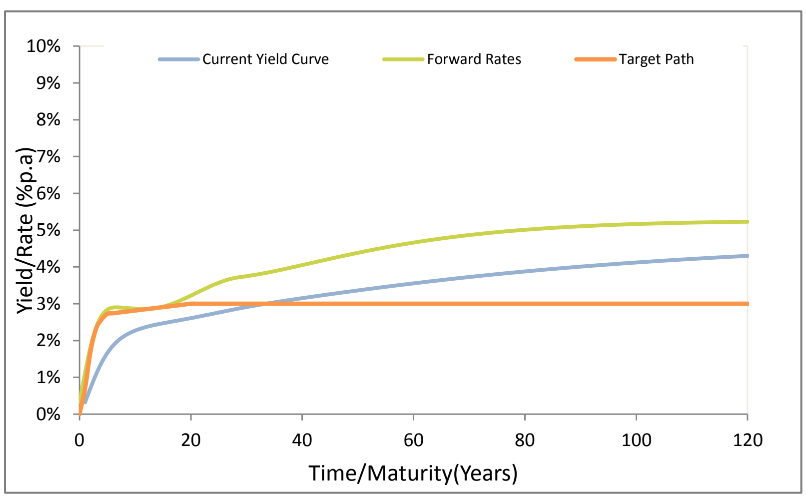

Similar processes were defined for all of the other model parameters required for real-world projection.

2.2 First Year Paths

The desired projection horizon was defined to be one year beyond the in-force date. Consistent with every 1-year stress, the liability portfolio also needed to be updated accordingly, including the extra year of policyholder aging along with changes to account values and guarantees, and so on. As with any complex projection problem, we should expect the projection of capital requirements to be *path-dependent*, meaning different paths over the year could result in different capital numbers despite sharing the same end state. Previous proxy methods for market-value projection have addressed the path-dependency problem either by treating the stress as instantaneous or, as in the hedging application considered in Clayton and Morrison (2016), addressing the proxy problem at the *policy level*, for which the paths of economic risk factors may be summarized by the policy account value, moneyness, and so on. Portfolio-level results are then built up by summing the policy-level values in effect at a given time.

The policy-level approach would not work for CTE capital projection, however, since the CTE statistic is non-additive. For the purposes of this project, therefore, we defined a path *interpolation rule* to describe the first year paths deterministically as a function of the 1-year stresses to the market variables. This was done via linear interpolation of the economic risk variables (yield curve principal components, and so on) between time = 0 and time = 1. The interpolated risk variables were then passed to the ESG to generate a consistent set of all required outputs over the first year. Thus, each group of fitting scenarios corresponding to a given risk factor stress shared a common first year path, and the cash flows determining the required capital estimates were discounted to year 1 of the simulation.

2.3 CTE Estimates

As described above, the required capital measure for this exercise was defined as the CTE(99) of the distribution of present-value of maximum projected balance sheet deficiencies, in contrast to a market-value measure given by a mean. Whereas a proxy function for the mean might require as little as one scenario for estimation, the extreme nature of this tail statistic posed a particular technical challenge, as a naïve approach to estimating the CTE(99) – even for proxy fitting purposes – might require many hundreds or thousands of inner scenarios to be passed to the actuarial cash flow models, causing the overall run-time to escalate quickly.

To address this, the process of estimating the required capital at each fitting point was accomplished in stages, with the goal of significantly reducing the number of scenarios ultimately passed to the cash flow models:

First, the ESG was used to generate a batch of 10,000 scenarios at each risk factor point, with model parameters updated per the recalibration step described in the previous section. These scenarios were summarized with a set of descriptive statistics (percentiles, moments) of key output variables and other derived quantities at each projection time. A *stratification rule* was then applied to find a subset of 1,200 scenarios out of 10,000 with corresponding statistics that matched the master set as closely as possible.

Next, these 1,200 scenarios were further summarized by *risk statistics* (cumulative equity return, yield curve movement, and so on) in order to exclude them from consideration in the tail calculation. That is, we appealed to the intuitive concept that if a given scenario, based on summary information, is clearly not one in which the balance sheet deficiency is likely to be "large," for the purposes of calculating the CTE(99), this scenario's *actual* deficiency is irrelevant. We need only concern ourselves with the "bad" scenarios, which can at least be roughly discerned without computing how bad they are exactly, and thus a *scenario filtering* rule can be defined to reduce the 1,200 scenarios to a smaller number without affecting the CTE measure much. Obviously, this was product-dependent and relied heavily on an expert understanding of the business in question, but previous experience had shown the filtering algorithm to produce accurate results with approximately a 4:1 reduction in scenario set size. Thus, for each fitting point, instead of estimating the CTE(99) by the average of the worst 12 out of 1,200, we could consider averaging the worst 12 out of approximately 300, corresponding to the CTE(96) of the reduced subset.

Finally, to further reduce the number of inner scenarios per fitting point, and introduce the deliberate inaccuracy to be smoothed out by the proxy fitting process, we randomly sampled a subset of the filtered scenarios at a ratio of 6:1. Thus, instead of 300 scenarios, we used a randomly selected subset of 50. These were passed to the actuarial cash flow model and balance sheet deficiencies returned.

A naïve estimate for the distribution CTE(96) using this set of 50 scenario results would be the corresponding sample CTE(96), i.e., the average of the worst 2 scenarios. However, for samples of this small size, the sample CTE estimator has a well-known downward bias (Manistre and Hancock, 2005), which, after we passed the estimates through the proxy fitting engine, would show up as systematic error. Instead, we made use of the *bias-corrected CTE estimator*, derived in general in Kim and Hardy (2007). This estimator uses the "exact bootstrap method" to associate a set of non-uniform weights to the ranked sample values for the estimate, instead of a simple restricted average.

Figure 3 below summarizes the steps of producing a CTE estimate from the 10,000 scenarios corresponding to a single risk factor stress.

Figure 3: CTE estimate workflow (single fitting stress)

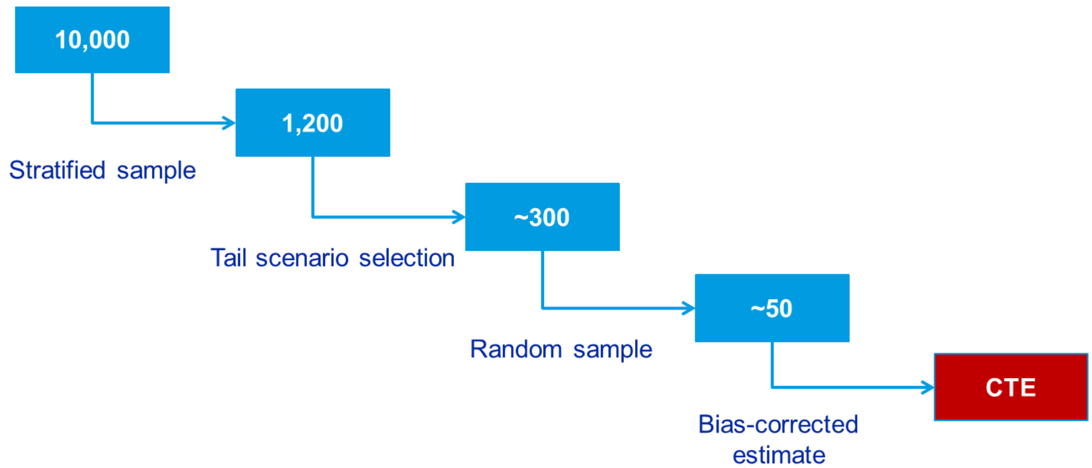

This process was repeated for each of the 2,048 fitting stresses. The estimates were then used to fit proxy functions for the CTE as a polynomial function of the explanatory risk variables (yield curves, equities, and credit spreads) using the stepwise regression algorithm in the Moody's Analytics Proxy Generator software. The proxy fitting was carried out for each of the two product groups individually as well as the aggregate of the two. In contrast to market value, the required capital is a sub-additive measure (owing to diversification benefit), so the aggregate capital function needed to be estimated separately from its constituent parts.

3. Case Study Results

In this section we present the results of applying the proxy function methodology described above to the two product groups as well as the aggregate business. At each chosen validation stress point, we compare proxy function value with the results of a "brute force" calculation using 50,000 real-world scenarios. [3]

Figure 4: Proxy vs. validation results (required capital)

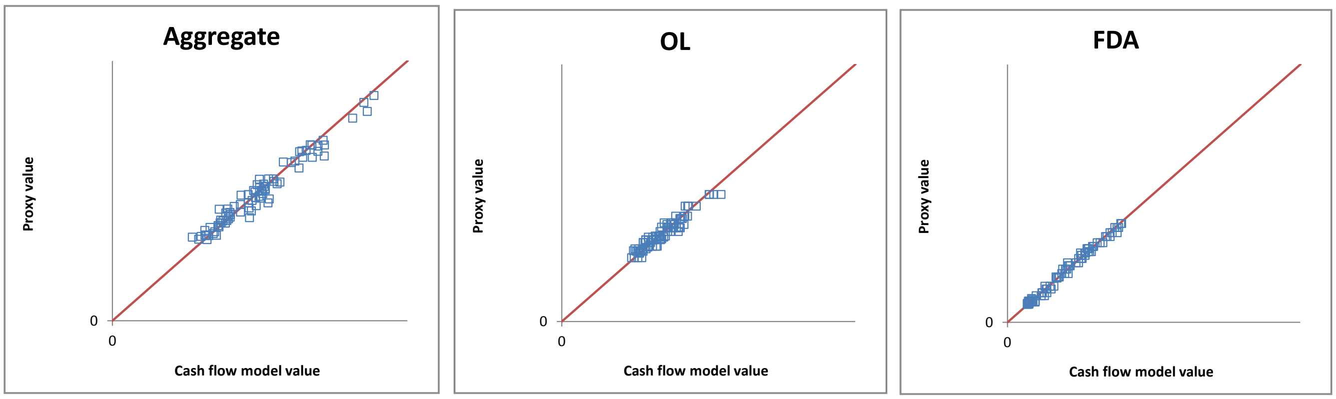

We observe generally good agreement between proxy function and cash flow model values, with some divergence in the aggregate capital numbers likely attributable to divergence in the OL component. Agreement for the FDA capital requirements is comparatively excellent.

Examining the validation errors more closely indicates that the greatest errors occur at the more extreme values of the risk factors, consistent with the reduced density of fitting points at the edges and corners of the fitting space. Figures 5, 6 and 7 show validation errors for each function plotted in the 4-dimensional risk factor space, with the size of the dot proportional to the size of error.

Figure 5: Validation error (Aggregate)

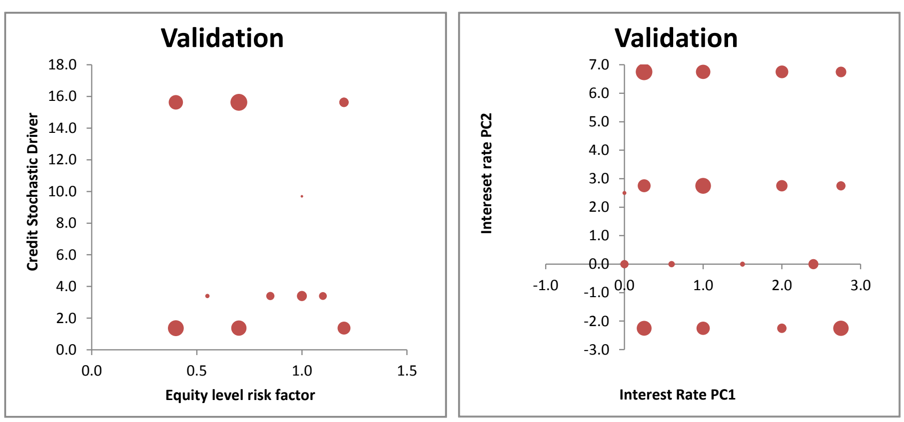

[3] The proxy function results shown here and in the following have been numerically modified but are illustrative of the relationships between variables.

Figure 6: Validation error (OL)

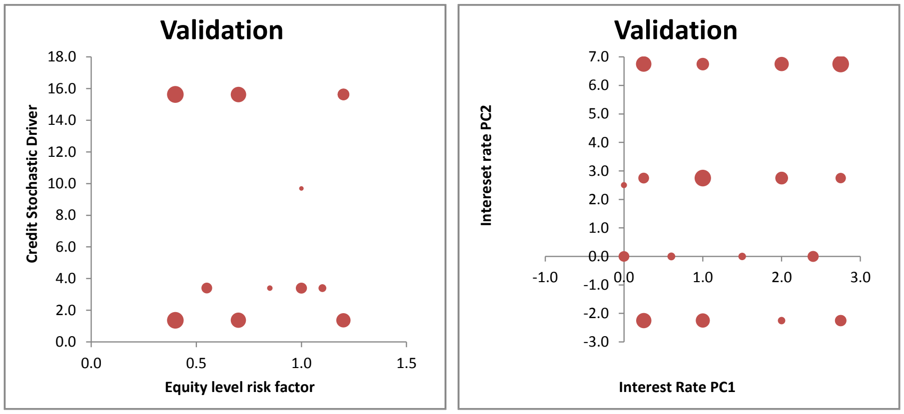

Figure 7: Validation error (FDA)

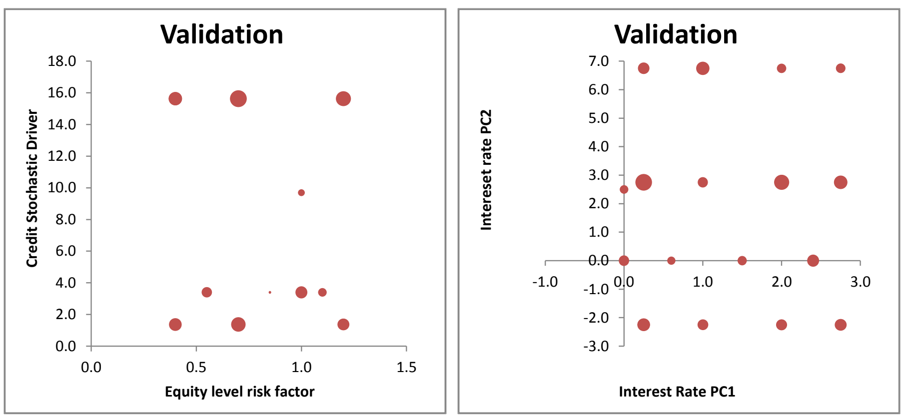

Using the proxy functions, we are able to quickly assess the effect that each risk factor has on the required capital. For example, holding all risks constant apart from the second yield curve factor shows a complex non-monotonic relationship between capital and interest rates.

Figure 8: Univariate analysis (required capital vs. yield curve factor 2)

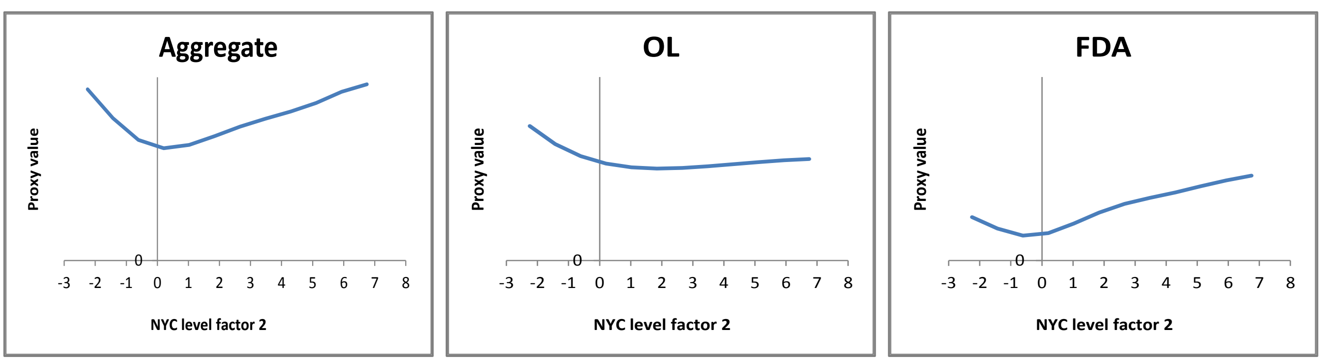

Furthermore, the proxy functions allow for rapid analysis of the behavior of required capital with respect to joint stresses of *multiple* risk factors. Regions where risks compound each other in perhaps non-intuitive ways are readily identifiable.

Figure 9: Multivariate sensitivity analysis

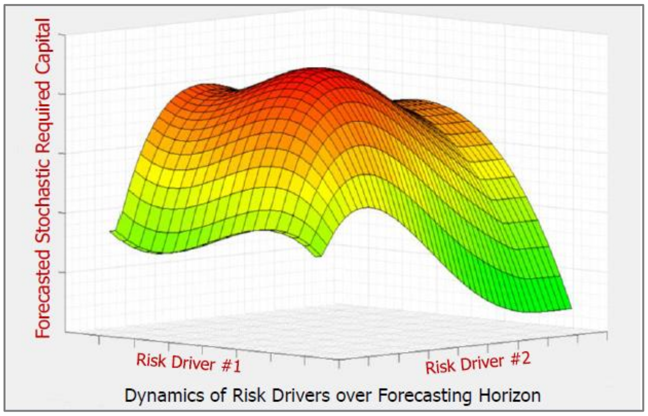

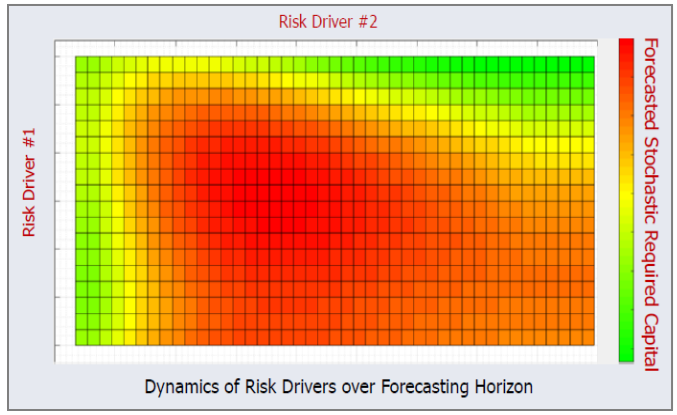

Combining the functions for individual product groups and the aggregate, we can quickly gauge the diversification benefit on capital as a function of any of the chosen risk factors or under various stressed conditions. Figure 10 illustrates one such analysis.

Figure 10: Diversification benefit (aggregate capital compared to sum of individual capital requirements)

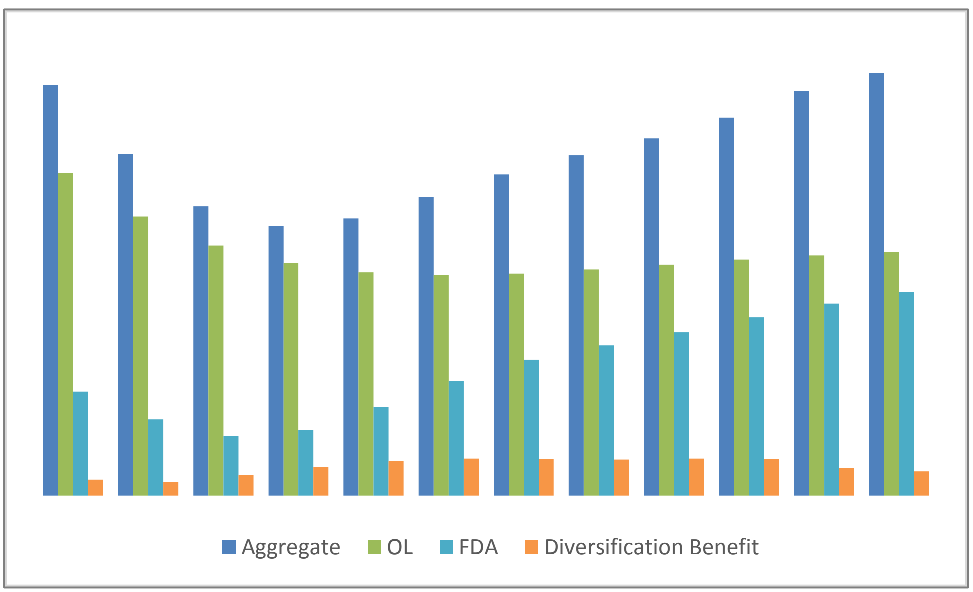

Finally, we consider what effective reduction in total scenario budget we have achieved using the proxy function approach. Using a baseline assumption of 50,000 scenarios for an accurate validation, we see that the 81 validation stresses by themselves would have required a total of 50,000 * 81 = **4.05 million** scenarios. By contrast, using the combination of scenario filtering and proxy estimation described above, we have attained the proxy function values at these and all other desired points with a total of 2,048 * 50 = **102,400** fitting scenarios, or approximately **1/40th** of the total actuarial model run-time. Naturally, the savings are made even larger the more projected stresses are required, for example in the multivariate analyses above, which requires hundreds of combination stresses to fill out.

As an alternative, we may wonder what the effect might have been of doing a "brute force" calculation using merely 10,000 instead of 50,000 scenarios, for an obvious 5:1 reduction in run-time, and forgoing the proxy fit altogether. Somewhat surprisingly, on a relative basis, we find that this reduction in scenarios alone introduces *consistently more error* than the proxy function approach does. Figure11 illustrates this point for a subset of the validation stresses.

Figure 11: Comparative error analysis (10K scenarios vs. proxy function)

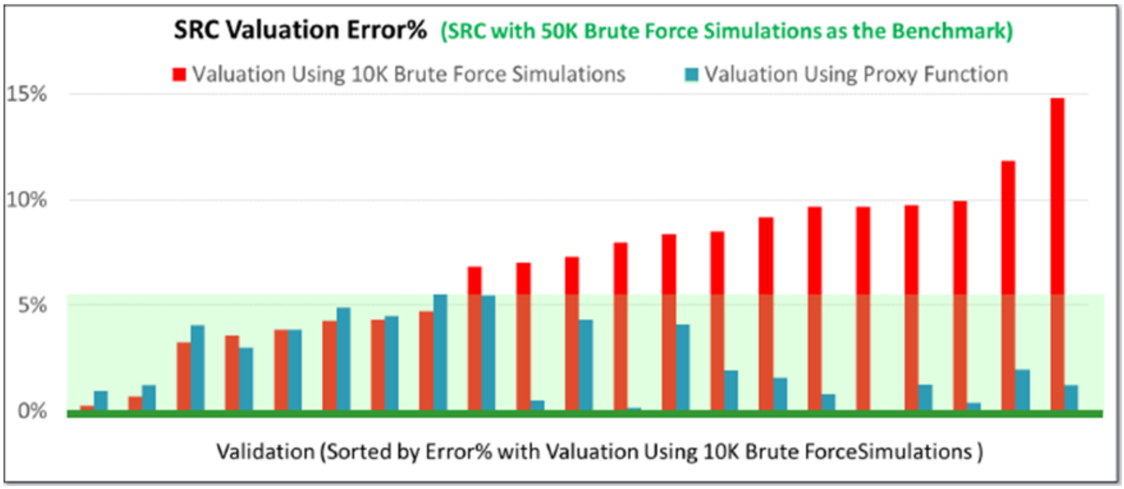

The proxy function errors are generally within 5%, whereas the relative errors from the 10,000 scenario estimates can be as high as 15%. In other words, despite the proxy function having been initially calibrated using only 50 scenarios per risk factor stress, the *effective number* of scenarios represented by the proxy function is something more than 10,000, measured according to its relative accuracy. The LSMC approach has made such efficient use of the available information that it is capable of extracting an impressively large amount of information from a comparatively much smaller data set.

4. Conclusions

We have demonstrated the successful application of proxy methods to the problem of projecting a CTE capital measure over a one year horizon for a large and complex block of insurance business. The required capital metric in this case is an extreme tail statistic – the CTE(99) of the distribution of balance sheet deficiencies – but with the LSMC proxy approach we are able to replicate this value to within a small tolerance over a wide range of economic stresses, at a substantial reduction in run-time compared to full nested-stochastic calculation.

The proxy fitting procedure follows the same general template as other proxy applications, but with several adaptations to suit this particular problem, most notably:

- The use of real-world models and recalibration processes for the “inner” scenarios
- The interpolation of the first year paths to account for the path-dependency of the required capital metric
- A much larger number of inner scenarios for each risk factor fitting point, together with a stratification, filtering, and random sampling process to reduce the number of these that are actually passed to the actuarial cash flow model
- Bias-corrected estimators to produce CTE estimates from the cash flow results that are suitable for proxy fitting

Incorporating these changes and building the necessary tools to make the proxy method work relied on expert knowledge of the products and stochastic models under consideration, as well as close collaboration between the Moody's Analytics and New York Life teams.

Once in hand, the proxy functions facilitate a number of rapid analyses of the capital position of the business, including forecasting capital under various “what if” scenarios, decomposing the effect of the various economic risk factors on the capital requirements either individually or in combination, and showing the diversification benefit between different product groups under economic stresses. Any of these analyses is likely to be prohibitively costly with brute-force Monte Carlo simulation. The slight compromise in accuracy due to the use of proxy methods is seen to compare favorably to other possible remedies, such as reducing the number of Monte Carlo scenarios from 50,000 to 10,000. Meanwhile, the proxy approach requires only a single batch of scenario runs for calibration and offers a dramatic savings in total model run-time.

References

Clayton, Aubrey, and Steven Morrison. "Proxy Methods for Hedge Projection: Two Variable Annuity Case Studies." *Moody's Analytics* (2016).

Kim, Joseph Hyun Tae, and Mary R. Hardy. "Quantifying and Correcting the Bias in Estimated Risk Measures." *ASTIN Bulletin* 37 (2007): 365-386.

Manistre, John B., and Geoffrey H. Hancock. "Variance of the CTE estimator." *North American Actuarial Journal* 9.2 (2005): 129-156.

Morrison, Steven, Laura Tadrowski, and Craig Turnbull. "One-year projection of run-off conditional tail expectation (CTE) reserves." *Moody's Analytics* (2013).

Morrison, Steven, Craig Turnbull, and Naglis Vysniauskas. "Multi-year Projection of Run-off Conditional Tail Expectation (CTE) Reserves." *Moody's Analytics* (2013).