MOODY'S ANALYTICS

RESEARCH
INSURANCE

Fitting Proxy Functions for Conditional Tail Expectation: Comparison of Methods

Authors

Aubrey Clayton PhD
Moody's Analytics Research

Steven Morrison PhD
Moody's Analytics Research

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.com

Japan +81.3.5408.4100
clientservices.japan@moodys.com

Overview

Proxy function methods have been used in the insurance industry for nearly a decade, primarily in nested-stochastic applications requiring a calculation of a distribution mean, such as the projection of an option value defined as a risk-neutral average of discounted cash flows.

Previous research notes (Clayton et al., 2013; Clayton and Morrison, 2016) have shown the extension of these methods to projecting dynamic hedges. Other research (Morrison, Tadrowski, and Turnbull, 2013; Clayton et al., 2016) has described a method of constructing proxy functions for Conditional Tail Expectation (CTE), used in particular to project run-off reserve and capital requirements.

However, all of the above applications have used some variant of Ordinary Least Squares (OLS) regression for proxy function fitting, which requires unbiased (or nearly unbiased) estimates of the distributional statistic of interest as an input. For a fit to a mean, the process of constructing unbiased estimates is straightforward, but for “tail statistics” such as a CTE, finding an unbiased estimator in general can be difficult or impossible.

This note details alternative methods for fitting proxy functions to CTE, employing quantile regression in combination with OLS among other techniques. We compare methods according to quality of fit for an example portfolio of variable annuities.

Table of Contents

1. Introduction

We consider the general problem of fitting a proxy model to describe the Conditional Tail Expectation (CTE) of a probability distribution – that is, expected value of a random variable given that it exceeds a given threshold – conditional on the values of some other given variables. A typical application might be projecting forward a prospective capital requirement supporting a book of insurance business under various economic scenarios over the coming year; the required capital in many cases is defined as a CTE(70) or CTE(90) of the present value of future deficits under a stochastic "run-off" calculation of assets and liabilities. Thus, projecting this capital requirement is a nested-stochastic problem, with "outer" scenarios specifying the possible evolution of markets over one year, and "inner" scenarios branching off to give the conditional run-off capital in each outer scenario.

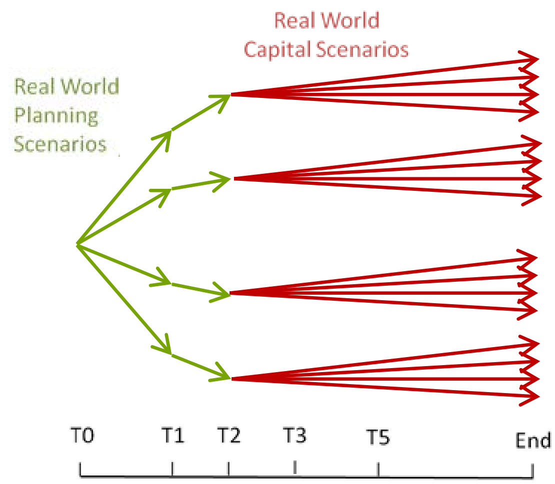

As described in Clayton *et al.* (2016), this is mechanically similar to the nested-stochastic problem of projecting market value under different economic scenarios, with the key difference being that the latter application calls for an *average* of discounted cash flows over *risk-neutral* inner scenarios whereas we are interested in the *CTE* over (typically) *real-world* inner scenarios. For the risk-neutral/average version of the problem, Least Squares Monte Carlo (LSMC) has proven to be an extremely effective technique, greatly reducing calculation time relative to a full stochastic-on-stochastic calculation with only a minor loss of accuracy.[1] The core insight of LSMC is that the inner statistic of interest (e.g., average over conditional risk-neutral distribution) should have a *continuous functional relationship* to the outer risk-variables (e.g., one-year economic scenarios); by crudely estimating the value using a reduced number of inner scenarios and then applying function-fitting techniques such as polynomial regression, we can approximate this functional relationship with a "proxy function."

In general we would expect a similar argument to hold here as well. Our aim in the present work is to examine how much of the LSMC methodology can be translated from the risk-neutral/average regime to the real-world/CTE one, and in particular what changes to the function-fitting process might be necessary or desirable. The main challenge we are confronted with is that the standard fitting technique, Ordinary Least Squares (OLS) regression, is only suitable for extracting the conditional *mean* behavior of a response variable given some explanatory variables. Thus, in order for OLS to be applicable, we must somehow transform our CTE problem into an average problem after all, for example by using inner samples to produce an *unbiased* (or nearly unbiased) estimator of the CTE. This is the first method we consider and also the method previously used in Morrison, Tadrowski, and Turnbull (2013) and, with some improvements, in Clayton *et al.* (2016).

Alternatively, we can use other function-fitting techniques. We explore two alternative methods here: one using quantile regression in combination with OLS to extract the conditional quantile/CTE, and the other using quantile regression alone to fit a parametric distribution in the tail of the conditional distribution.

After describing these techniques in generality, we show the results of applying them to an example portfolio of variable annuities with a view towards projecting reserve/capital requirements at a given future time. Here we assume a fixed *scenario budget*, measured by the total number of scenarios passed to an actuarial cash flow model, and attempt to answer the question of which function fitting technique makes best use of that budget to produce the highest quality proxy function fit.

[1] See Elliot (2016) or Morrison and Tadrowski (2015) for recent surveys and descriptions of the LSMC approach.

2. Methodology

For the remainder we will assume we have a scenario generator capable of producing samples of "outer" risk-variables, which we denote $\boldsymbol{x}$, and an "inner" variable, denoted y, from the conditional distribution $p(y|\boldsymbol{x})$. A practical example might involve $\boldsymbol{x}$ representing the yield curve and equity markets at time 1, and y being the present value of future run-off deficits conditional on the given realized market.

We assume that the ultimate metric of interest is the CTE of this distribution at some given confidence level τ:

$$CTE(\tau) \coloneqq E_x[y|y > q_\tau]$$

where q_τ is the τ-quantile; $P[y < q_\tau|x] = \tau$.

This value then depends only on $\boldsymbol{x}$. In keeping with the usual LSMC approach, to estimate the functional relationship between the CTE and $\boldsymbol{x}$, we generate a set of *fitting stresses* $\boldsymbol{x}_i, i = 1, \dots, N$, and assume we can produce samples from the appropriate corresponding conditional distributions.

2.1 Method 1 – Ordinary Least Squares

This approach mimics the usual methodology for fitting a proxy function to a mean. For each outer fitting stress, we generate some number M of *inner samples* $y_1, \dots, y_M$ from the distribution $p(y|x)$ and then use these to construct an estimator for the distribution CTE. If the estimate is unbiased – meaning the expected value of the estimator over its distribution is the true CTE – then OLS regression can extract the functional relationship between the true CTE and the underlying risk factors. In effect, this transforms the problem back to that of fitting a proxy function to a mean of a distribution (now the distribution of the CTE estimator itself) given a single sample in each fitting stress. The hard work then becomes finding the right estimator.

A reasonable guess would be to use the in-sample CTE as an estimator:

$$\widehat{CTE}(\tau) = \frac{1}{M(1-\tau)} \sum_{j=M\tau+1}^{M} y_{(j)}$$

where $y_{(j)}$ denotes the jth order statistic of the sample; however, as shown in Manistre and Hancock (2005), this estimator is in general negatively biased, by an amount that depends on the particular distribution and size of the sample. As a result, for small inner sample sizes, the fitted proxy function will exhibit large systematic error in the final validations.

On the other hand, larger sample sizes M will naturally come at the cost of reducing the number of available fitting stresses N that can be evaluated, assuming a constant overall scenario budget, $N \cdot M$. This can also result in a lower quality proxy fit. So for this estimator, a balance must be found between running enough inner samples to reduce the size of the estimator bias while still running enough fitting stresses to fill out the risk factor space.

For example, Morrison, Tadrowski, and Turnbull (2013) compared fitting results for CTE(70) using 10, 100, and 1,000 inner scenarios in each of 10,000, 1,000, and 100 fitting points, respectively, and concluded the 1,000 outer x 100 inner scenario allocation gave the best balance between bias and variance of the estimator.

An improvement is to use a "bias-corrected" estimator derived from the exact bootstrap method as described in Kim and Hardy (2007):

$$\widehat{CTE}(\tau) = c^T(2I - w^T)y_{:M}$$

where $y_{:M} = \left(y_{(1)}, \dots, y_{(M)}\right)^T$ is the vector of rank-order statistics of the samples, w is the weight matrix:

$$w_{i,j} = j\binom{M}{j}\left[B\left(\frac{i}{M}; j, M-j+1\right) - B\left(\frac{i-1}{M}; j, M-j+1\right)\right]$$

with $B(x; a, b) \coloneqq \int_0^x t^{a-1}(1-t)^{b-1}dt$ being the incomplete beta function, and

$$c = \frac{1}{M(1-\tau)}(0, \dots, 0, 1, \dots, 1)^T$$

has zeroes in the first $M\tau$ elements.

This can allow nearly unbiased estimators for the CTE to be constructed even from small samples, as in Clayton *et al.* (2016), which estimated the CTE(99) from sample sizes on the order of $M = 50$.

2.2 Method 2 – Quantile Regression Plus OLS

Quantile regression is an alternative method of estimating the relationship between a response variable y and a set of predictors $\boldsymbol{x}$, when what is ultimately desired is a particular quantile of the conditional distribution.[2] In contrast to OLS, which estimates the conditional mean by minimizing the sum of squared residuals, quantile regression works by minimizing the objective

$$\sum_{i=1}^{N}(y_i - \boldsymbol{x}_i\boldsymbol{\beta})\left(\tau - I_{(y_i - \boldsymbol{x}_i\boldsymbol{\beta}<0)}\right)$$

as a function of the coefficients $\boldsymbol{\beta}$, where τ is the quantile level of interest. This objective function is motivated by the observation that the τth quantile q_τ of the distribution of any random variable Y minimizes the value

$$E\left[(Y-u)\left(\tau - I_{(Y-u<0)}\right)\right]$$
$$= (\tau-1)\int_{-\infty}^{u}(y-u)dF_Y(y) \; + \tau\int_{u}^{\infty}(y-u)dF_Y(y)$$

(Differentiating with respect to u and setting equal to 0 shows the minimum occurs at u_0 if and only if $F_Y(u_0) = \tau$, i.e., $u_0 = q_\tau$.)

Given a set of fitting stresses and samples from the inner conditional distributions, we can then use quantile regression to extract the conditional quantiles, such as the 70th or 90th percentile. Note that the above minimization problem is not as simple as that of OLS (which reduces to a system of linear equations); instead, the problem is typically reformulated as a linear program in terms of the positive and negative parts of the β coefficients:

$$\beta^+ = \max(\beta, 0)\,, \beta^- = -\min(\beta, 0)$$

for which simplex and interior point solution methods are available (Koenker, 2005, pp. 181-190).

These conditional quantiles can be useful statistics in themselves, as in Morrison (2017), which considered the problem of estimating the conditional 99.5th percentile for the purpose of projecting 1-year VaR. For our purpose of projecting CTE, however, the quantiles can be an intermediate step, since by definition the CTE(τ) is the mean of the distribution conditional on exceeding the τth quantile. Therefore, considering *only* those samples that exceed the estimated quantile, we can apply OLS regression to extract the conditional CTE.

To summarize, the steps to fitting a proxy for the CTE using this method are:

1. For $i = 1, \dots, N$, generate fitting stresses $\boldsymbol{x}_i$ and individual samples y_i from the inner conditional distributions $p(y|\boldsymbol{x})$
2. Use quantile regression to fit coefficients of the quantile model $q_\tau(y|\boldsymbol{x}) = \boldsymbol{x}'\boldsymbol{\beta_Q}$
3. Select only those samples for which the inner sample exceeds the predicted quantile, that is, where $y_i > \boldsymbol{x}_i'\boldsymbol{\beta_Q}$. Suppose these samples have indices $i_1, \dots, i_n$. Assuming step 2 has successfully fit the quantile model, we should have n approximately equal to $N(1-\tau)$.
4. Perform Ordinary Least Squares regression to fit the model $y \sim \boldsymbol{x}'\boldsymbol{\beta_C}$, only on the selected points $(\boldsymbol{x}_{i_1}, y_{i_1}), \dots, (\boldsymbol{x}_{i_n}, y_{i_n})$. This gives an estimate of the mean of the response distribution conditional on exceeding the given quantile, that is, the CTE.

In contrast to the "pure" OLS method described in the previous section, this approach can use as little as *one* inner sample per fitting stress, no matter what CTE level is desired. However, it comes at the cost of having to do two regressions – one for the quantile and one for the CTE – introducing multiple possible sources of error in the process. In addition, the quantile regression step can be computationally expensive, depending on the complexity of the model (number of coefficients) and number of fitting stresses.

[2] See Koenker (2005) for a complete introduction.

2.3 Method 3 – Distribution Matching

One final method we consider uses quantile regression alone to understand the shape of the desired distribution in the tails. In principle, if we had access to the quantiles q_τ of a distribution for *all* τ, we could calculate the CTE at any level using Acerbi's Integral Formula:[3]

$$CTE(\tau) = \frac{1}{1-\tau}\int_\tau^1 q_\beta d\beta$$

In practice, however, we will never have this much information (which amounts to the full cumulative distribution function of the response variable). Instead, we consider repeating the quantile regression process a small number of times at different quantile levels and matching these to a given parametric distribution, for which the CTE can be calculated either analytically or numerically.

For example, to estimate the CTE(70) of the distribution, we might first estimate the 70th and 90th percentiles $q_{0.7}$ and $q_{0.9}$ by means of two quantile regression fits with $\tau = 0.7$ and $\tau = 0.9$. Then, assuming the distribution is approximately normal in this tail region, we could find parameters μ, σ of a normal distribution $N(\mu, \sigma^2)$ with quantiles matching our given estimates. Writing the normal distribution in terms of a standard normal variable Z as

$$\mu + \sigma Z$$

makes this straightforward, since we then have two equations in the two unknown parameters:

$$\mu + \sigma\Phi^{-1}(0.9) = q_{0.9}$$
$$\mu + \sigma\Phi^{-1}(0.7) = q_{0.7}$$

The $CTE(\tau)$ of the normal distribution for any level τ can then be calculated directly using the formula:[4]

$$CTE(\tau) = \mu + \frac{\sigma}{1-\tau} * \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\Phi^{-1}(\tau)^2\right)$$

Similar methods would apply to other distributional assumptions, such as lognormal or generalized Pareto, as we explore in the next section.

In each case, the performance of this method would rest on the quality of fit in each of the quantile regressions and the reasonableness of the distributional assumption in the tail. Each of these could depend sensitively on which particular quantiles were used in combination to estimate a given CTE level and how extreme that level was. This method also carries the additional expense of needing to perform multiple quantile regressions.

3. Case Study Results

To compare performance among the different methods described in the previous section, we consider the problem of projecting reserve and capital requirements for a representative block of variable annuities. The block in question consists of approximately 75,000 policies with a mixture of accumulation, withdrawal, and death benefit guarantees[5] at various levels.

At any given point in time, a distribution of present value of future deficits is defined by a 40-year run-off projection using real-world scenarios, from which the required reserve is defined as the CTE(70) and the required capital as the CTE(90). We are interested in projecting these requirements forward 5 years in the future under various economic scenarios. For this example, we consider interest rate risk only, including the sensitivity to the initial yield curve. Specifically, we have outer stresses consisting of the following variables:

- Initial yield curve, described by two principal component shocks ("PC1" and "PC2").
- The change in the yield curve during the first 5 years, defined as a parallel movement over time ("YC_Change").

[3] Acerbi and Tasche (2002)

[4] See Hardy (2006) for a derivation.

[5] Specifically, GMAB, GMMB, and GMDB with optional GLWB and GMWB

- The speed of the above yield curve movement, described as the time over which the parallel movement occurs ("YC_Period"); if this completes before the 5 year period, the yield curve is assumed to be constant for the remaining time.

These stresses, although constrained to fairly simple yield curve movements, allow a wide range of possible yield curve paths of the kind required for stress testing or capital planning purposes. The chart below illustrates one such possible path:

Figure 1: Possible yield curve evolution

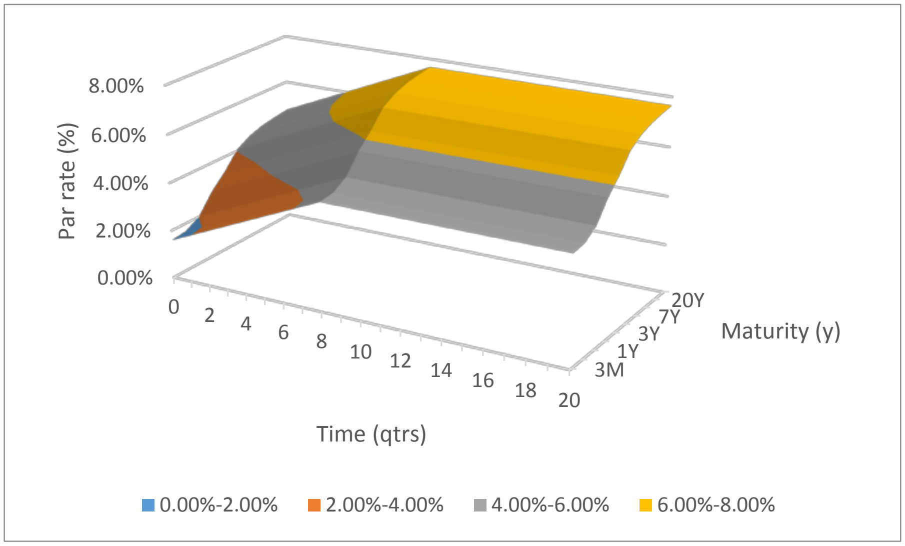

Ordinarily, to compute reserve and capital requirements at the end of such a path, we would recalibrate an inner real-world scenario generator and run a large number – on the order of 10,000 – of conditional scenarios. We would then calculate future cash flows and the CTE of the aggregated results. If we then required ~100 planning scenarios, we would need approximately 1 million total scenarios to be passed to the actuarial cash flow model.

Instead, for this example, we assume a fixed total scenario budget of **100,000 scenarios**, to be allocated differently for each proxy fitting method.

To validate the performance of the proxy functions, we compute full nested stochastic results using 3,000 inner scenarios for a set of 100 validation points. These points are divided in two categories: multivariate validation points, chosen by Sobol sampling from the 4-dimensional risk factor space described above, and univariate validation points, constructed by fixing all but one variable in turn at a chosen value (e.g., median) and allowing the remaining risk factor vary. A selection of the validation points is shown in the charts below:

Figure 2: Multivariate Validation Points

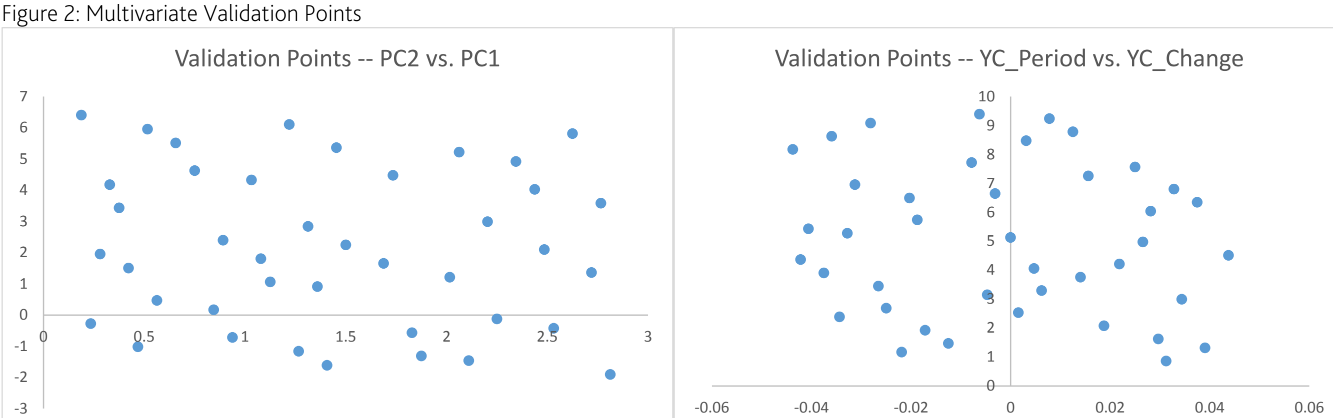

Figure 3: Univariate Validation Points

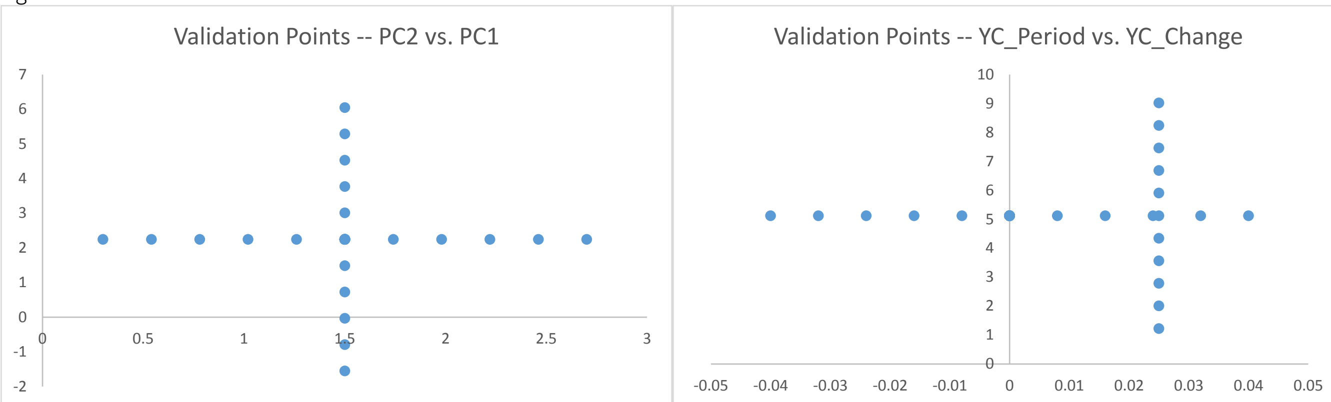

3.1 Method 1 – Ordinary Least Squares

We consider two different allocations of the total 100,000 fitting scenarios. First, we use the scenarios to construct 1,000 outer fitting points with 100 inner samples each and fit proxy functions to the bias-corrected estimates of the CTE(70) and CTE(90) using OLS. The charts below show proxy function value vs. validation value; perfect agreement is indicated by the line $y = x$.

Figure 4: Proxy vs. Validation for 1,000x100 scenario allocation OLS

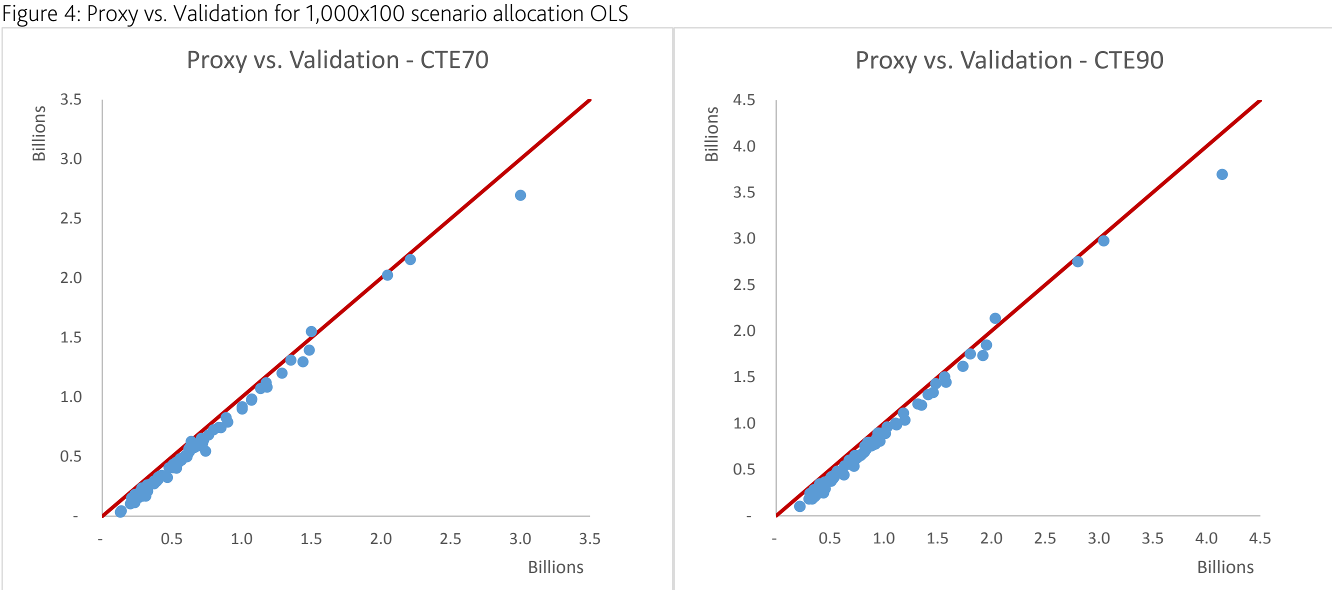

The proxy functions show some deviation from the validation values, particularly for the more extreme values. This is likely the result of the reduced number of fitting points (1000) and the subsequent inability to fully cover the 4-dimensional risk factor space. Also, the fits show some systematic negative bias, indicating that the proxy function is consistently underestimating the true CTE by a small amount.

Second, we consider a scenario allocation with 10,000 outer fitting points with 10 inner scenarios each.

Figure 5: Proxy vs. Validation for 10,000x10 scenario allocation OLS

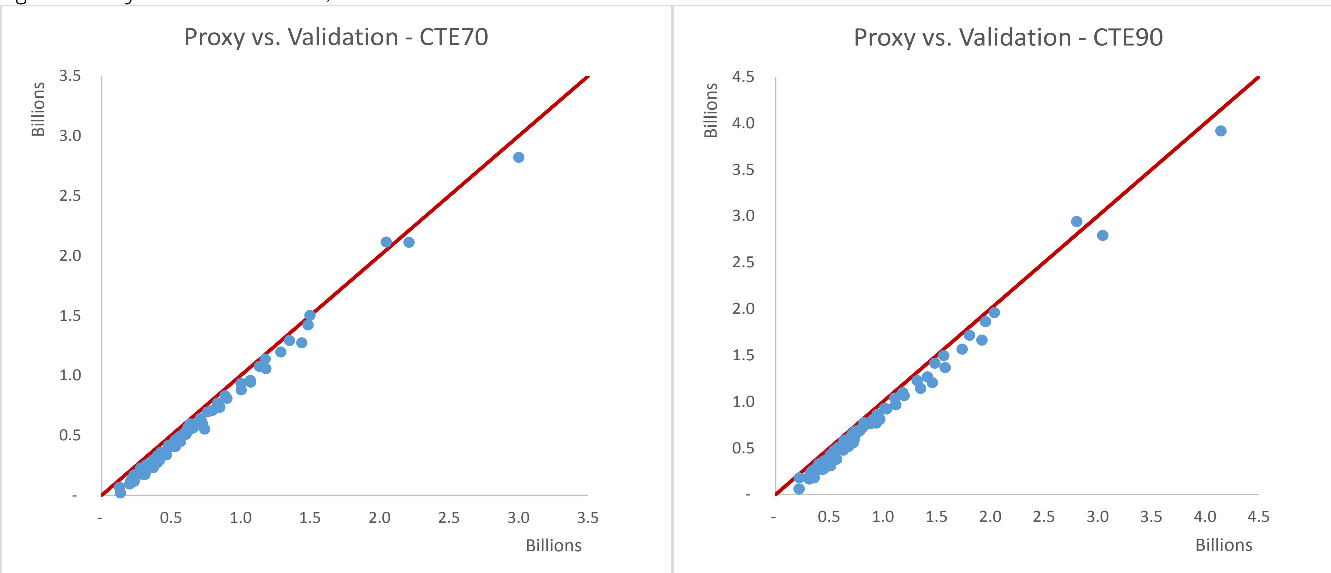

The deviation at the extreme values has improved relative to the previous fits, with a slight increase in the amount of systematic bias in the proxy function. With the in-sample CTE estimator, ordinarily a sample size this small (10 scenarios) would result in a very biased function, particularly for the CTE(90) proxy; however, with the bias correction described in the previous section, we are able to keep the bias relatively small.

3.2 Method 2 – Quantile Regression Plus OLS

Here, we use quantile regression to fit the 70th and 90th conditional quantile of the distribution and then OLS on the residuals to fit the CTE. Since the extra regression step introduces an additional potential source of error, we validate both the quantile and CTE proxy functions to show the quality of fit at each stage.

Figure 6: Proxy vs. Validation for Quantile Regression fits

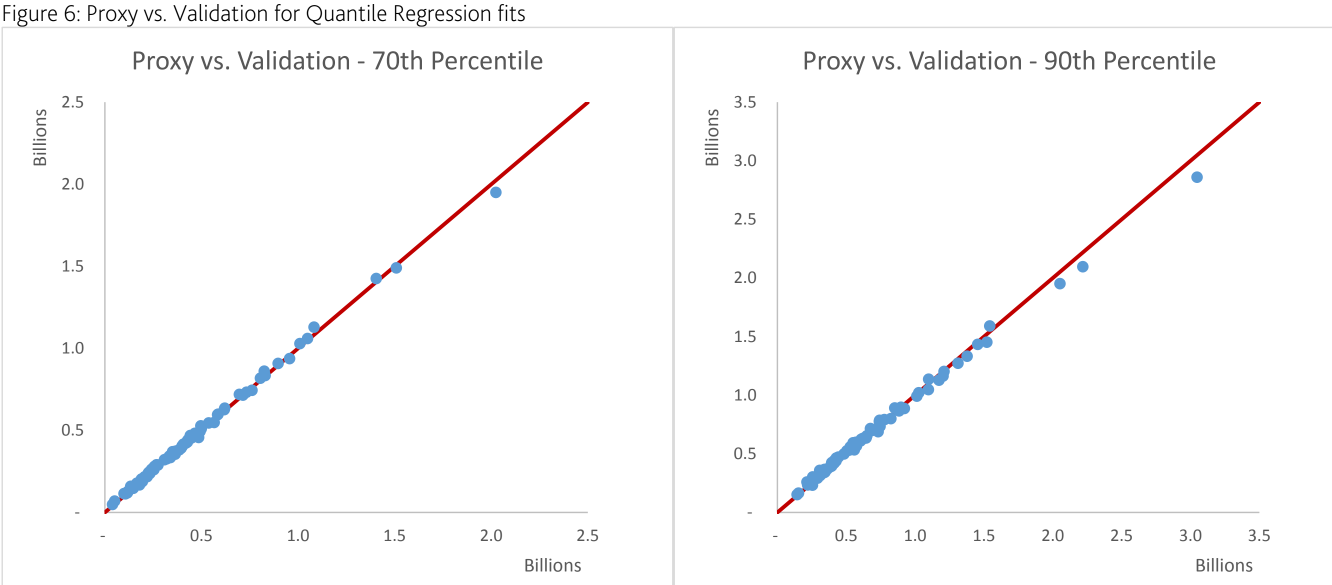

Figure 7: Proxy vs. Validation for CTE using Quantile Regression plus OLS

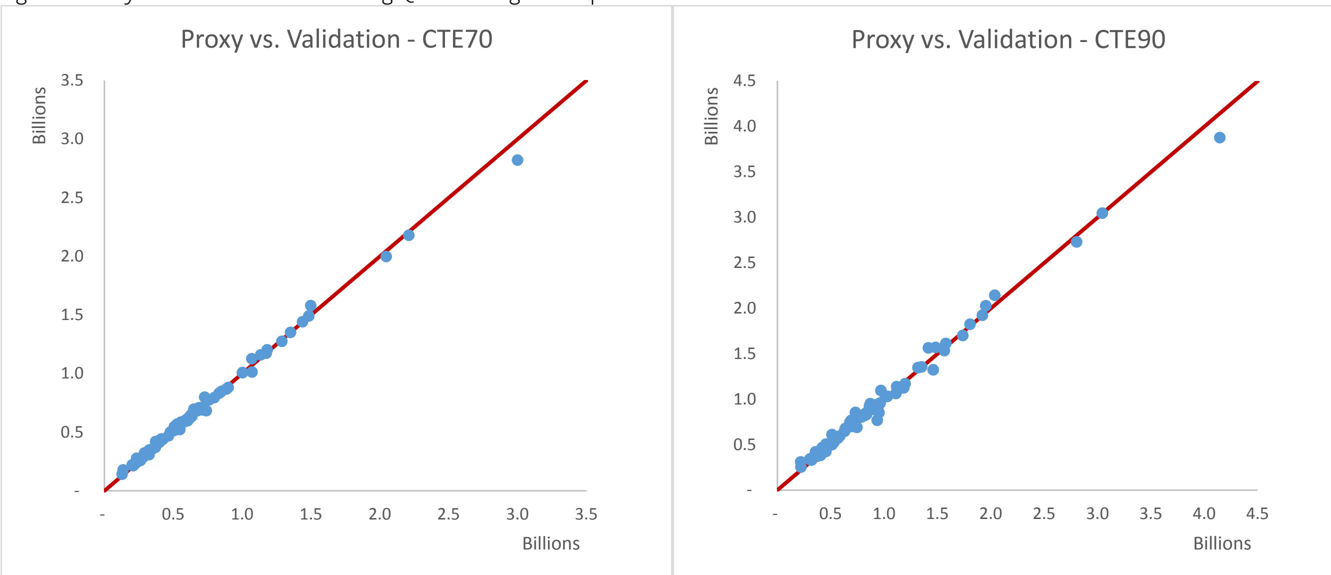

To further investigate the behavior of the proxy functions, we consider the univariate validation points in isolation. This illustrates the overall behavior of the quantile and CTE proxies as functions of individual risk factors; shown here are dependencies with respect to the yield curve change and interest rate PC1 shock.

Figure 8: Proxy and Validation vs. yield curve change risk factor

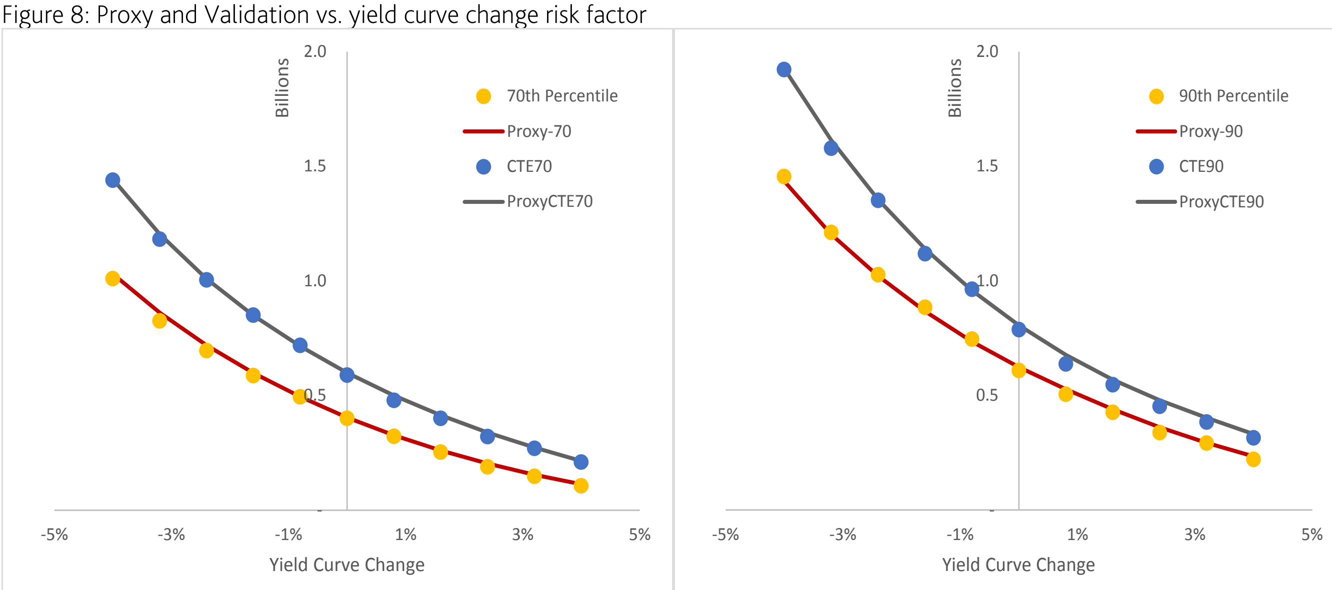

Figure 9: Proxy and Validation vs. yield curve PC1 risk factor

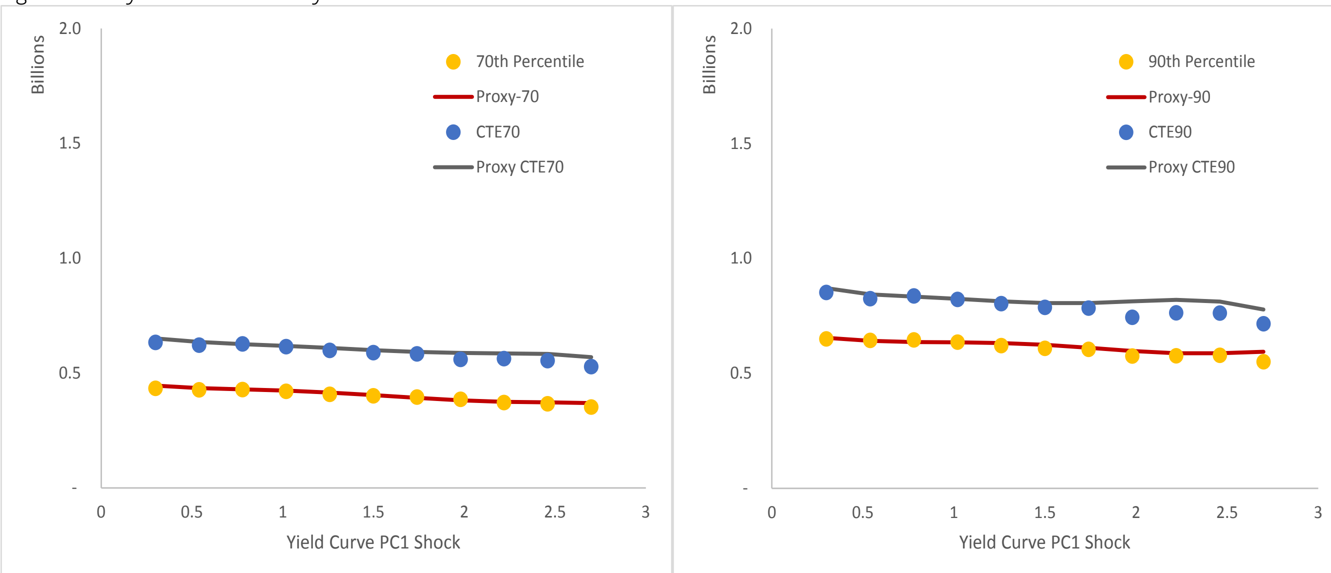

Overall we see very close agreement between the proxy function and validation values, and a consistent overall functional relationship between proxies and risk factors.

3.3 Method 3 – Distribution Matching

First, we fit a Gaussian distribution to the estimated 70th and 90th percentiles as described in the previous section and use the analytical formula for the Gaussian CTE to estimate the CTE(70) and CTE(90).

Figure 10: Proxy vs. Validation for Gaussian distribution fit to 70th and 90th percentiles

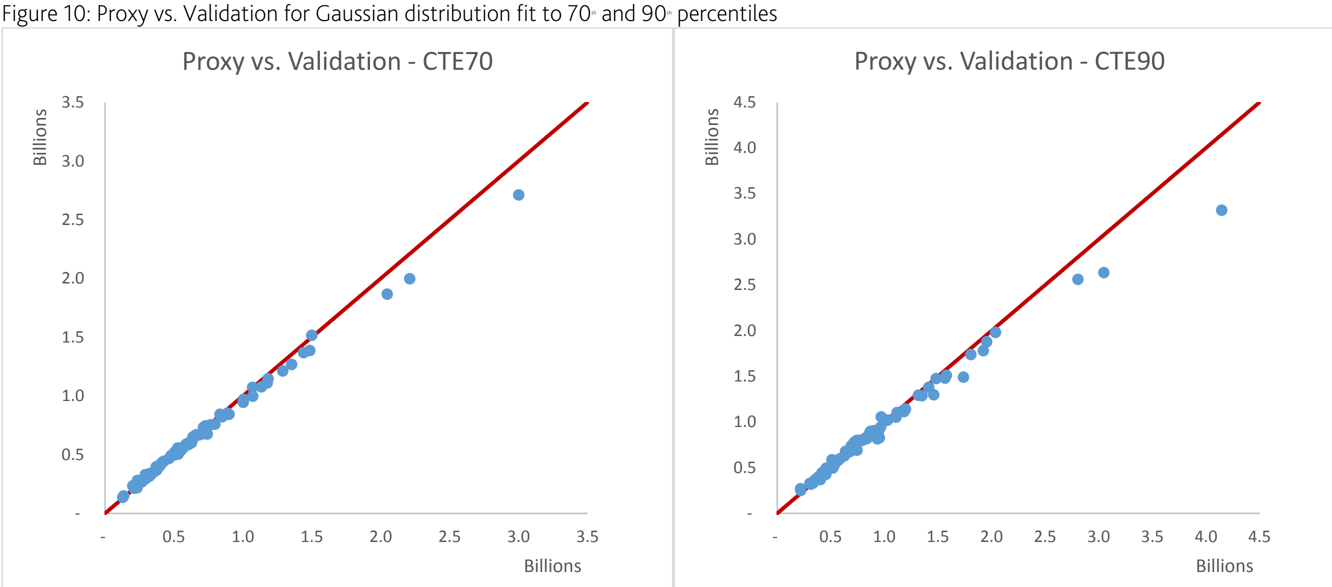

Both charts show the proxy function underestimating the relevant statistic, particularly for more extreme values. This indicates the Gaussian approximation may break down due to the distribution becoming more fat-tailed in the extreme stresses.

As an alternative, we consider fitting a lognormal distribution at the same estimated quantiles as above, now using the analytical formula for the lognormal CTE.

Figure 11: Proxy vs. Validation for lognormal distribution fit to 70th and 90th percentiles

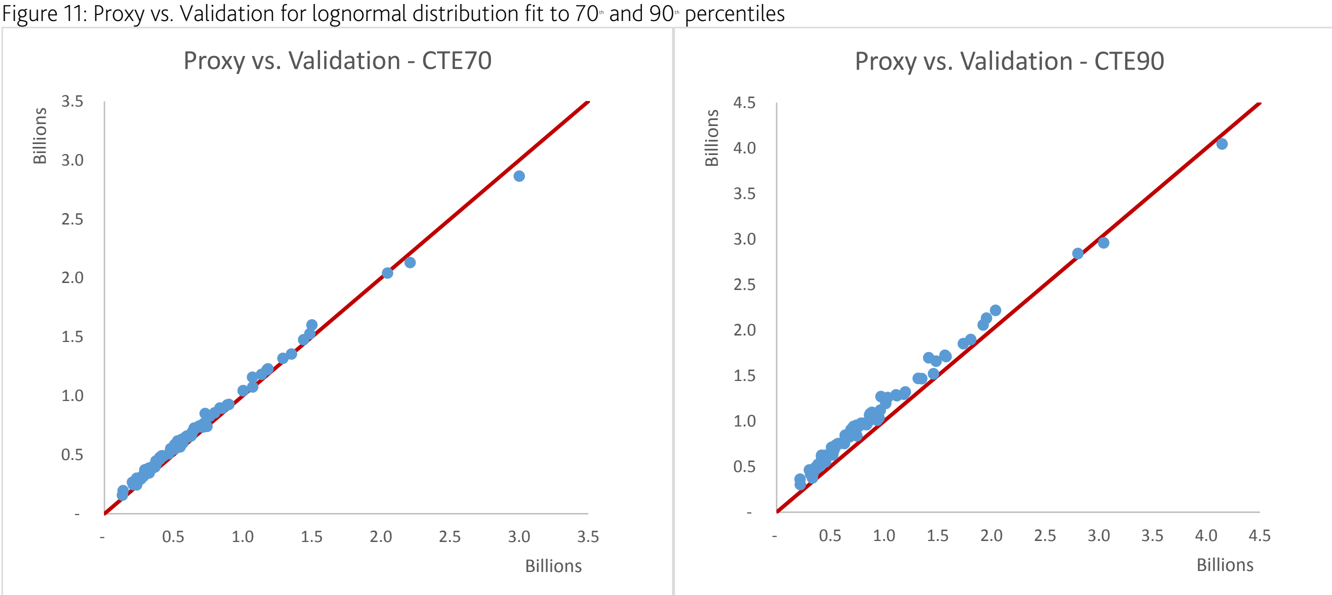

The quality of fit is improved at the higher regions, indicating we have captured the shape of the tail more accurately for those stresses. However, it has come at the cost of an *overestimate* for many of the other stresses, suggesting the shape of the tail is less fat than lognormal for those stresses.

Finally, we consider fitting a 3-parameter family using the generalized Pareto distribution. In order to solve for the three parameters, we require another estimated quantile, for which we use a quantile regression fit to the 80th percentile, constructed in the same manner as the 70th and 90th.

Figure 12: Proxy vs. Validation for generalized Pareto distribution fit to 70th, 80th, and 90th percentiles

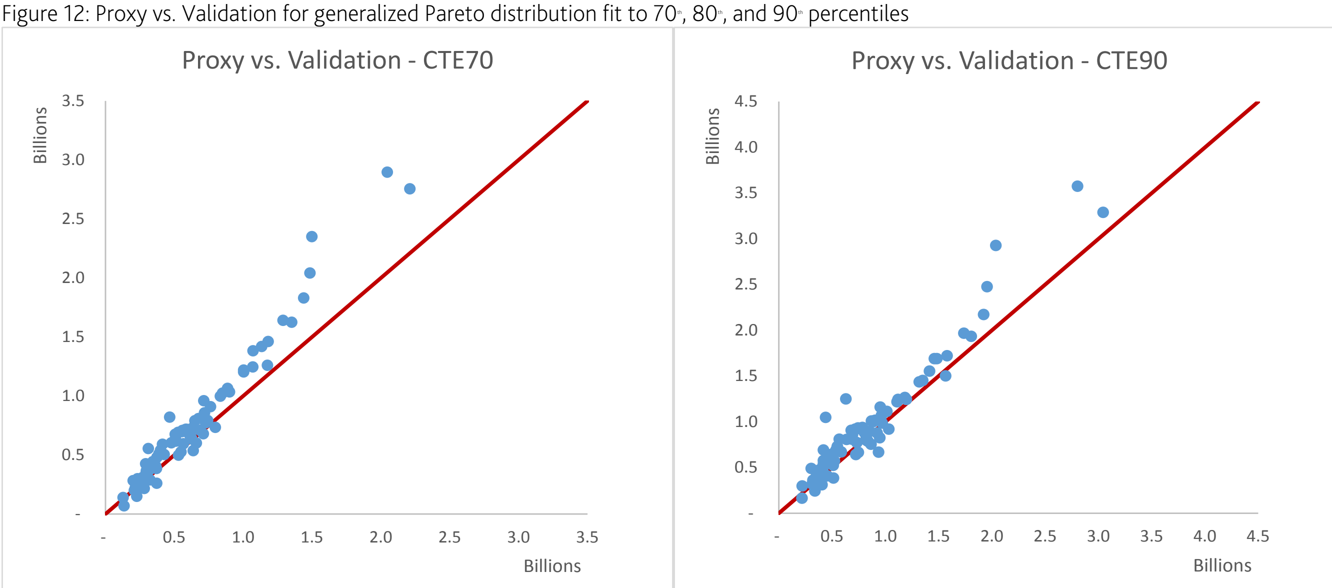

Here, the quality of fit is generally poor, indicating either that the Pareto distribution assumption is not reliable on this region of the tail or that the fitting process is unstable, perhaps due to the many sources of error.

4. Discussion

In principle we can compare results of the various fitting methods in many ways. The charts of the previous section already demonstrate two possible types of error that can be present in a proxy function fit: *bias* owing to some systematic error present in the fitting process (a biased estimator, omitted explanatory variable, or incorrect distributional assumption), and *variance* indicating a failure of the proxy function to converge to the true functional relationship (too few fitting points, too many coefficients, error accumulated over multiple regressions).

To include both types of error simultaneously, for each of the fits to the CTE(70) and CTE(90) above, we calculate proxy function error on a *root mean square (RMS)* basis over the set of multivariate validation points:

$$RMS = \left[\frac{1}{n}\sum_{i=1}^{n}(Proxy_i - Validation_i)^2\right]^{1/2}$$

which has the appealing property

$$RMS^2 = Bias^2 + Variance^2$$

We also compare methods according to *average absolute error*:

$$AvgAbsoluteError = \frac{1}{n}\sum_{i=1}^{n}|Proxy_i - Validation_i|$$

Finally, to put the above error measures on a meaningful scale, we express both relative to a base value, taken to be the value of the relevant statistic (either CTE(70) or CTE(90)) at the median validation point, that is, center of the risk factor space in all dimensions.

The results for the CTE(70) proxy fits (ranked by RMS error) are summarized in the chart and table below:

Figure 13: Proxy method errors – CTE(70)

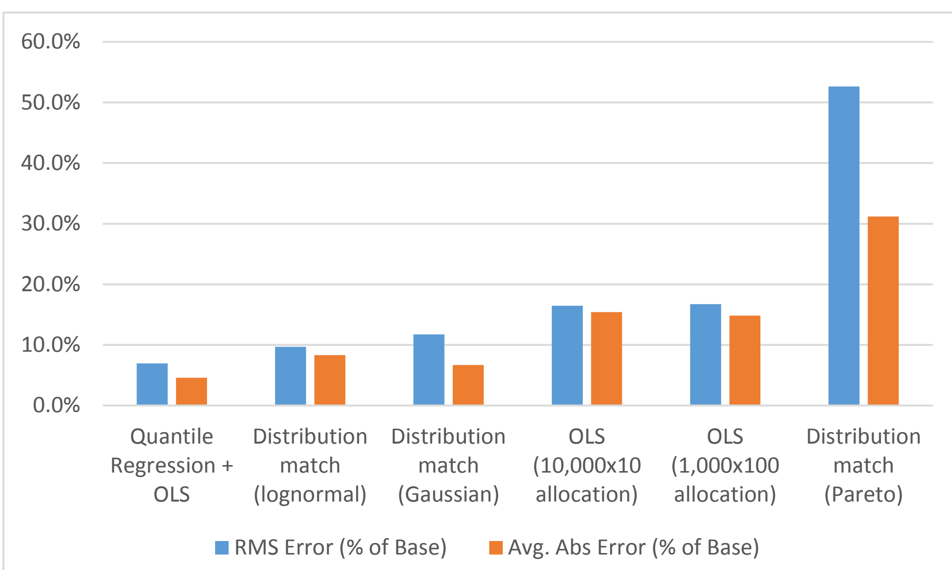

Method	RMS Error (% of Base)	Avg. Abs. Error (% of Base)
Quantile Regression + OLS	7.0%	4.6%
Distribution match (lognormal)	9.7%	8.3%
Distribution match (Gaussian)	11.7%	6.7%
OLS (10,000x10 allocation)	16.5%	15.4%
OLS (1,000x100 allocation)	16.7%	14.9%
Distribution match (Pareto)	52.7%	31.2%

The results for the CTE(90) proxy fits are summarized below:

Figure 14: Proxy method errors – CTE(90)

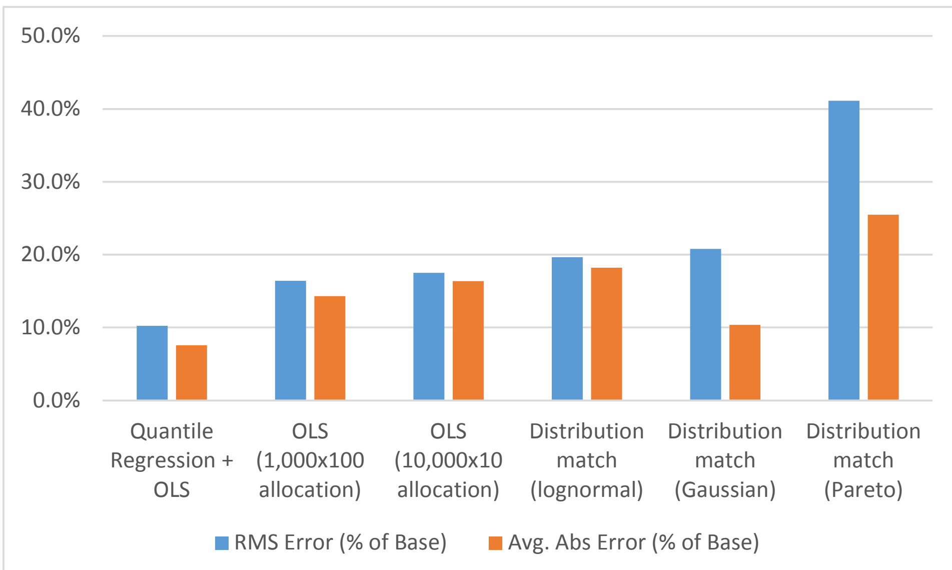

Method	RMS Error (% of Base)	Avg. Abs. Error (% of Base)
Quantile Regression + OLS	10.2%	7.5%
OLS (1,000x100 allocation)	16.4%	14.3%
OLS (10,000x10 allocation)	17.5%	16.4%
Distribution match (lognormal)	19.7%	18.2%
Distribution match (Gaussian)	20.8%	10.4%
Distribution match (Pareto)	41.1%	25.5%

With respect to either the RMS or average absolute error, the method of quantile regression combined with OLS gave the most accurate results overall for both tail measures. For the CTE(70) fits, the method of distribution matching – either Gaussian or lognormal – performed nearly as well, but the results were significantly worse for the CTE(90) fits, indicating a failure of these distributional assumptions further out in the tail of the distributions. The pure OLS fits performed moderately well under either scenario allocation considered, suggesting the increase in bias from using fewer inner scenarios was approximately made up for by the improvement from using more outer scenarios. The Pareto distribution fits were uniformly worse than all of the above.

Based on these results and the high quality of fit to the functional relationships shown in the univariate validations, combined with the general stability of the approach, we would recommend the quantile regression + OLS method as the primary technique for fitting to a CTE on the order of those considered here, between CTE(70) and CTE(90).

For statistics further out in the tails (e.g., CTE(99) or CTE(99.5)), some care must be taken, since this method effectively discards all fitting data points whose estimated values are less than the relevant estimated quantile before passing them to the OLS regression. So, for example, for a fit to the CTE(99), after fitting to the 99th percentile we would exclude all but approximately 1% of the fitting points, reducing the number of points from 100,000 to something like 1,000 for the OLS fit. It seems likely that this would result in lower quality proxy fits and that the other methods (OLS and distribution matching) may perform relatively better.

In any case, the addition of quantile regression to the suite of available tools provides much-needed flexibility for the task of fitting proxy functions for CTE. Either in combination with Ordinary Least Squares, as in our preferred approach, or on its own, quantile regression can give important information about the conditional behavior of tail risk measures with respect to underlying risk factors. These come at the cost of a different and more computationally difficult regression method, formulated as a linear program instead of a system of linear equations, which must be run separately for each desired quantile. Nevertheless, we expect this approach to be a key component of any proxy-fitting exercise involving projecting CTE or other tail risk measures in the future.

References

Acerbi, Carlo, and Dirk Tasche. "On the coherence of expected shortfall." *Journal of Banking & Finance* 26, no. 7 (2002): 1487-1503.

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

Clayton, Aubrey, Steven Morrison, Ronald Harasym, and Andrew Ng. "Proxy Methods for Run-off CTE Capital Projection: A Life Insurance Case Study." Moody's Analytics (2016).

Clayton, Aubrey, Steven Morrison, Craig Turnbull, and Naglis Vysniauskas. "Proxy functions for the projection of Variable Annuity Greeks." Moody's Analytics (2013).

Elliot, Martin. "Making Proxy Functions Work in Practice." Moody's Analytics (2016).

Hardy, Mary R. "An Introduction to Risk Measures for Actuarial Applications." *[Construction and Evaluation of Actuarial Models Examination Study Note]*. Society of Actuaries (2006).

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

Koenker, Roger. Quantile Regression. Cambridge University Press (2005).

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

Morrison, Steven. "Solvency In Sight – New tools for understanding the impact of investment decisions on capital." Moody's Analytics (2017).

Morrison, Steven, and Laura Tadrowski. "Proxy Model Validation." Moody's Analytics (2015).

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).