RESEARCH
INSURANCE ERS

Author

Steven Morrison, PhD

Moody's Analytics Research

Contact Us

Americas	+1.212.553.1653 clientservices@moodys.com
Europe	+44.20.7772.5454 clientservices.emea@moodys.com
Asia-Pacific	+85.2.2916.1121 clientservices.asia@moodys.com
Japan	+81.3.5408.4100 clientservices.japan@moodys.com

The Unintended Consequences of Scenario Post-processing in the Valuation of Insurance Liabilities

Overview

In this paper we explore the use of scenario re-weighting as a method for post-processing scenario sets to reflect calibration targets without having to recalibrate the model.

While post-processing techniques can be quite flexible in their ability to match targets, they may result in unintended changes to distributional assumptions that are not included in the set of calibration targets. Using simple examples, we demonstrate how a scenario set's ability to match a set of vanilla asset prices does not uniquely define the resulting prices of more exotic liabilities (or assets).

Ultimately, a liability valuation model provides a way of translating assumptions about the distribution of underlying risk factors into liability values. Market asset prices are an important input into such models, but in practice many other subjective modelling assumptions also affect liability values. By recalibrating the model (rather than reweighting), the user can ensure that all such assumptions are understood and controlled.

CONTENTS

1. Introduction 3
2. The Challenge of Multiple Risk Factors 3
3. The Challenge of Path Dependency 7
4. Summary 9
References 10

1. Introduction

The standard approach to calibration of Economic Scenario Generator (ESG) models involves changing the model parameters so that the resulting scenarios are consistent with a set of calibration targets (including market asset prices in the case of market-consistent valuation). This can present several challenges for the modeller: Simpler models, with few parameters and highly constrained distributional shapes, are often fast to calibrate but limited in their flexibility to match target prices. On the other hand, models that are more flexible in their ability to match target prices in principle may be slow to calibrate in practice. Fast and reliable recalibration is particularly important when this needs to be carried out many times, for example in projection of market-consistent value for the purpose of capital assessment.

As an alternative to this calibration process, a number of post-processing techniques exist that modify a base set of scenarios so that the modified scenarios meet the calibration targets. Post-processing usually involves a combination of changing the scenario output itself and/or re-weighting i.e. assigning different weights to different scenarios (rather than assuming they are equally weighted). Such techniques can be quite flexible in their ability to match targets regardless of the base scenario set and model used to generate it. Furthermore, the processing can often be automated in a highly efficient way, so that the modeller can be confident of meeting their calibration targets quickly.

Armed with the ability to match any calibration targets quickly via post-processing, does this replace the need to reclibrate via the more traditional route of changing model parameters? In this paper, we consider this question by exploring the differences in properties of scenario sets produced using these two different methods. Using some simple examples, we will illustrate how different methods can agree on certain properties and disagree on others, leading to different prices for instruments not included in the set of calibration targets.

2. The Challenge of Multiple Risk Factors

One of the challenges of market-consistent valuation of insurance liabilities is that, in general, their values depend on the assumptions not just of the marginal distributions of individual risk-factors but also on the dependency between them. However, the prices of assets with reliable, liquid, market prices (for example options) often only provide information about marginal distributions of individual risk factors. A calibrated Economic Scenario Generator also embeds assumptions about dependency that are not necessarily implied by market prices but can be chosen according to the user's views (usually using analysis of historical data). A natural question then arises as to how these dependency assumptions change as scenarios are post-processed to match features of the marginal distributions implied by market option prices.

In this section, we consider the valuation of an option that depends on the dependency between two risk factors: an exchange option on two indices. We calibrate the model to match prices of options on each individual index, via full recalibration of the model and via a standard post-processing technique of Weighed Monte Carlo (Avellaneda, Buff, Friedman, Grandchamp, Kruk, & Newman, 2001). How do the resulting exchange option prices compare between the two calibration methods?

To carry out this valuation, we generated 10,000 'base' risk-neutral scenarios for the two indices. Each index was assumed to start at 1 and follow a simple Geometric Brownian Motion (GBM) model, each with a volatility of 20%, and a correlation of 0.5 between them. We also assume zero interest rates. In this simple model the Black-Scholes implied volatility surface is flat at 20% for both indices.

Suppose we reweight these base scenarios so that each index has an increased implied volatility of 30% (at strikes 0.8, 1 and 1.2 and maturity of 5 years), and additionally each asset exactly passes the martingale test at 5 years. Note that since the number of calibration targets (here, 9[1]) is much less than the total number of scenarios (10,000), there are numerous choices of weights that will exactly match all targets. Weights are usually chosen so that they are, in some sense, 'as close as possible' to the original (equal) weights. Here, we reweight using the popular 'maximum entropy' objective (discussed further below).

As you might expect, reweighting to exactly match the higher implied volatility targets (relative to the implied volatilities of the original scenario set) tends to put more weight on more extreme returns, as illustrated in Figure 1 which shows the weights plotted against the (log) index values at year 5.

Figure 1: Weights to match higher volatility target (produced using maximum entropy method) as a function of log index values

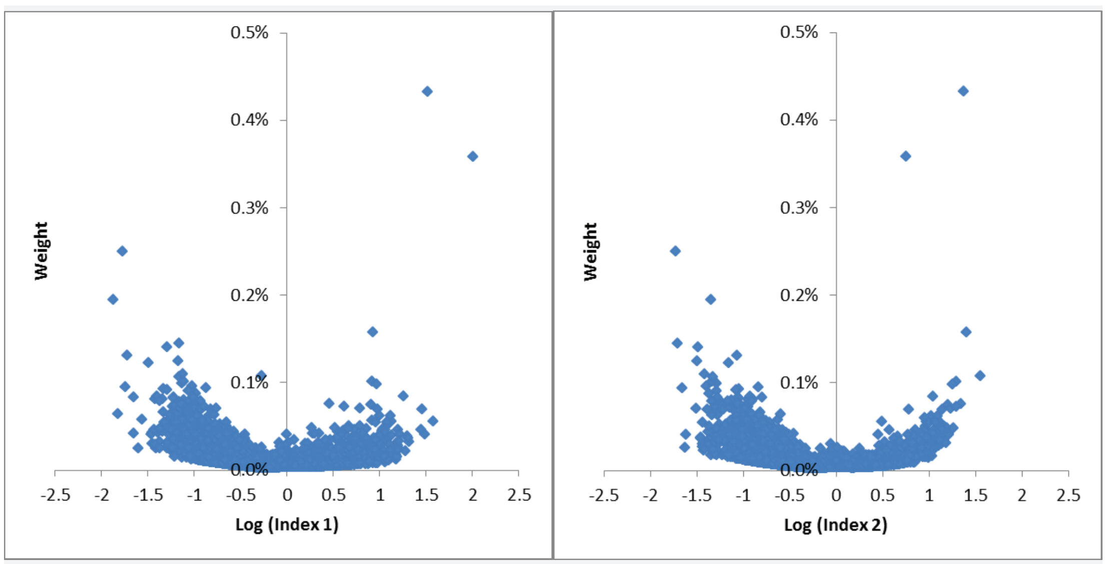

Note that the equal weighting corresponds to all scenarios having a weight of 0.01%, and reweighting produces new weights quite far from this (for example, the most heavily weighted scenario has a weight of 0.43%). We can quantify how far we have had to change the weights relative to their initial uniform weighting via a single number: the 'entropy', or equivalently the 'effective number of scenarios'[2]. In this case the effective number of scenarios is 7,161 (while the initial scenario set had 10,000 scenarios). The 'maximum entropy' technique adopted here seeks to maximise the effective number of scenarios while matching all constraints.

The scenarios have been re-weighted to exactly match the prices of a number of vanilla options on the two indices, at 5 years. Now we consider the impact on pricing the corresponding exchange option i.e. an option to exchange one of the indices for the other, at 5 years. While vanilla option prices only depend on the marginal distributions of the two indices, the payoff of the exchange option depends on the difference between the two indices and hence its price also depends on the assumed dependency between them. Intuitively, the higher the correlation between the two indices, the lower the variation in the difference between them, and hence the lower the price of the option to exchange one for the other (and similarly, the lower the correlation, the higher the exchange option price).

1 Each asset has 4 targets (3 implied volatility targets plus the martingale property) and additionally we constrain the weights to sum to one.

2 The entropy is defined as $entropy = -\sum_{i=1}^{n} w_i ln(w_i)$ where w_i is the weight on scenario i (out of n scenarios in total) and the effective number of scenarios defined as $exp(entropy)$. Note that for equally weighted scenarios $w_i = \frac{1}{n}$, the effective number of scenarios is n, while at the other extreme as all weight becomes concentrated on a single scenario the effective number of scenario tends to one.

Using the original 10,000 scenarios (assuming volatilities of 20% for both indices, and a correlation of +0.5), and assuming equal weights, we estimate that the exchange option has a price of 0.180[3]. After reweighting to match increased implied volatilities of 30% the estimated exchange option price increases to 0.200. However, rather than reweighting, if we recalibrate by increasing the model's volatility parameters to 30% (and retain a correlation parameter of +0.5) and generate a new set of (equally weighted) scenarios, we get a price of 0.269[4]. So the price obtained via reweighting is 26% lower than that obtained via recalibration. Equivalently, the 'implied correlation'[5] required to match the price using the reweighted scenarios is higher at +0.72 (rather than +0.5).

Figure 2 plots the joint scenarios for the two log index values, with the symbols sized by the new weights. We see that the effect of reweighting is to put more weight on extreme joint scenarios, which explains the higher implied correlation (and hence lower exchange option price) in this case[6].

Figure 2: Joint scenarios for log index values at year 5 (area of symbol proportional to new weights); Original scenarios generated using 20% volatility and reweighted to match 30% volatility target

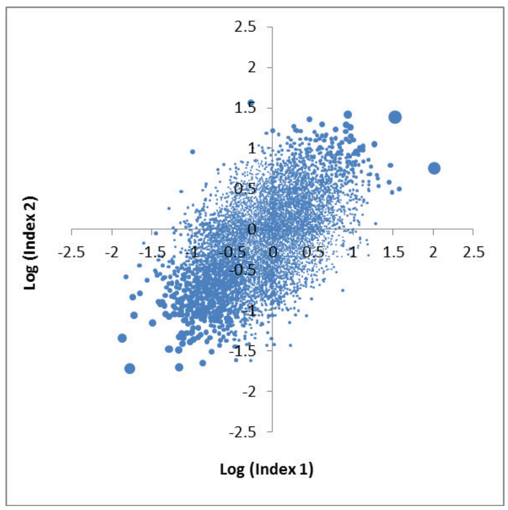

The problem is that the reweighting adopted here doesn't explicitly control for correlation - it just controls for volatilities, and means, of each asset. In reweighting the correlation does change, but does so in an uncontrolled way. Furthermore, the change in correlation depends on the base scenarios that we started with. To illustrate, suppose we repeat the reweighting exercise starting with a different set of 10,000 risk-neutral scenarios, produced using an increased base volatility assumption of 40% for both indices, and a correlation of 0.5 (as before) between them. In this case, reweighting to match the volatility target of 30% results in higher weights in the body of the distributions (rather than at the extremes), since the target volatility is now lower than in the base scenario set. Figure 3 shows the scenarios for the log index values, with the symbols sized by the new weights. In this case, the effect of reweighting is not only to lower the volatility of each index but also to decrease the correlation between them.

3 This is the estimated price using one particular set of 10,000 scenarios. The theoretical price, calculated using Margrabe's formula (Margrabe, 1978) is 0.177. The small difference in price here is due to sampling error given the finite number of scenarios used.

4 The theoretical price (calculated using Margrabe's formula) is now 0.263.

5 The implied correlation here is the correlation parameter in a joint GBM model that would match the exchange option price given an assumption that both volatilities are 30%.

6 As an alternative to the 'implied correlation', we can measure Pearson's correlation coefficient between the log index values (modified to account for the non-equal weights). In this case, Pearson's correlation coefficient is 0.69 (compared to an implied correlation of 0.72). In comparing the different methods here we prefer to quote the implied correlation as this directly relates to an option price.

Figure 3: Joint scenarios for log index values at year 5 (area of symbol proportional to new weights);
Original scenarios generated using 40% volatility and reweighted to match 30% volatility target

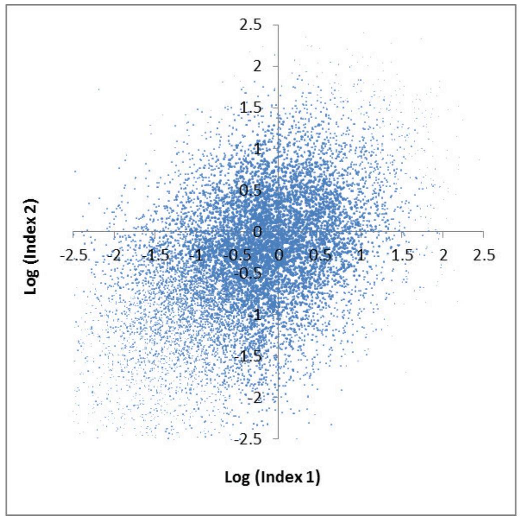

In this case, the effective number of scenarios is 8851, and the exchange option price is 0.301 (corresponding to an implied correlation of 0.34), which is 12% higher than the price obtained via recalibration.

To summarise, Figure 4 compares exchange option prices produced by three different methods. Recall that all methods have been calibrated to match the same implied volatilities of 30% on each individual index, and yet all methods produce different prices for the exchange option. These price differences can be attributed to different dependency assumptions embedded within the scenario sets produced by the different methods.

Figure 4: Exchange option prices using re-weighting vs full recalibration

Method	Effective number of scenarios	Exchange option price	Price difference relative to full recalibration	Implied correlation
Full recalibration (assuming correlation unchanged at +0.5)	10,000	0.269		0.48
Reweighted scenarios (starting with base of 20% vol)	7,162	0.200	-26%	0.72
Reweighted scenarios (starting with base of 40% vol)	8,851	0.301	12%	0.34

Note: Implied correlation using full recalibration is slightly different from +0.5 due to sampling error.

Of course, we could explicitly control for correlation (in addition to volatilities and means), for example by including the exchange option price in the set of constraints, though the effective number of scenarios can be expected to reduce every time we add a new constraint. In this particular example, including the 5-year implied correlation in the set of constraints results in the effective number of scenarios reduces from 7,162 to 6,764 (starting with base of 20% vol) and from 8,851 to 8,719 (starting with base of 40% vol), so the single correlation target can be achieved with only a small change to the effective number of scenarios. However, in more realistic examples, the number of such targets required is likely to make this impractical (for example, requiring a very small effective number of scenarios).

3. The Challenge of Path Dependency

The exchange option example illustrated how the prices of assets used in the calibration process do not uniquely determine all of the distributional assumptions in the model. In the exchange option example, the calibration assets only determine the marginal distributions of individual indices, and not the correlation between them.

Another type of dependency that often impacts on insurance liability values is the dependency between risk factors at different points in time. Insurance liabilities are often path dependent. For example, guarantee levels may be periodically updated according to the recent performance of backing assets, or backing assets may be periodically rebalanced in an effort to manage their ability to meet future liabilities. In such cases, the assumed dependency between risk factors at different points in time will impact on the liability values. However, the prices of vanilla options only provide information about the marginal distribution of the underlying risk factors at the maturity of the option, and not on the joint distribution at different times. Given our experience with the exchange option example, we might expect that post-processing and recalibration might produce different dependencies between risk factors at different times, and hence different prices for path dependent options.

In this section, we consider the pricing of a particular path dependent instrument: a variance swap, paying a fixed amount (at year 5) and receiving the realised variance (calculated using the previous 5 annual returns on some index). We assume that the index follows a Geometric Brownian Motion model with a volatility of 20% (and interest-rate of zero) and generate 1,000 'base' scenarios at annual time-steps. The Black-Scholes implied volatility surface is flat at 20% (at all maturities).

As is usually the case with swaps, we set the fixed payment such that the swap has zero value initially, which for this model and calibration implies a fixed payment of 0.04 (i.e. $20\%^2$). Using the 1000 base scenarios, the average realized variance is 0.041, resulting in an estimated variance swap value of 0.001. The slight difference from zero can be attributed to sampling error.

Now suppose that we reweight the base scenarios so that each index has an increased implied volatility of 30% (at strikes 0.8, 1 and 1.2 and maturities 1,2,3,4,5 years) and additionally each asset exactly passes the martingale test at 1,2,3,4,5 years. We reweight using the maximum entropy objective. The resulting effective number of scenarios is 609. If we calculate the average realized variance under these new weights, it increases to 0.049 ($22\%^2$), resulting in a variance swap price of 0.009. This is an increase in price relative to the base scenarios, but not nearly as large an increase as you might expect given the increased implied volatility of 30%. Indeed, if we recalibrate the model to reflect the 30% implied volatility assumption, the theoretical value of the variance swap increases to 0.05 ($=30\%^2 - 20\%^2$)[7]. In this case, the reweighting technique produces a variance swap price that is 82% lower than the price obtained via full recalibration of the GBM model.

As before, we also consider the effect of starting with a different set of base scenarios, produced using a GBM model with a 40% volatility assumption. Reweighting these higher volatility scenarios to reflect the 30% implied assumption results in an average released variance of 0.145 ($38\%^2$) and a variance swap price of 0.105. The effective number of scenarios in this case is 870. In this case, the reweighting method produces a price for the variance swap that is around twice that obtained via full recalibration of the GBM model.

7 The price using the 1000 scenarios here is 0.053, and the difference can be attributed to sampling error.

To summarise, Figure 5 compares variance swap prices produced by three different methods. Recall that all methods have been calibrated to match the same implied volatilities of 30% (at all maturities), and yet all methods produce different prices for the variance swap. Put another way, all methods target the same marginal distribution of the index (at 1,2,3,4,5 years) but they differ in the distribution of realized variance of annual returns (which depends on the index path over years 1,2,3,4,5).

Figure 5: Variance swap prices using re-weighting vs full recalibration

Method	Effective number of scenarios	Average realised variance	Square Root (Average realised variance)	Variance swap price	Price difference relative to full recalibration
Full recalibration	1,000	0.093	30%	0.053	
Reweighted scenarios (starting with base of 20% vol)	609	0.049	22%	0.009	-82%
Reweighted scenarios (starting with base of 40% vol)	870	0.145	38%	0.105	98%

As before, we can attribute these price differences to differences in the dependency between risk factors (in this case, values of the index at different times). For example, Figure 6 shows joint scenarios for the (log) index values at 1 and 2 years, with symbols sized by the new weights. Starting with a 20% volatility assumption and reweighting to achieve a higher volatility of 30% is achieved by putting more weight on the extreme scenarios at both 1 and 2 years, which also increases the correlation (left hand chart), while starting with a 40% volatility assumption and reweighting to achieve a lower volatility of 30% is achieved by putting more weight on the scenarios in the body of the distributions at both 1 and 2 years, which has the effect of decreasing the correlation (right hand chart)[8].

Figure 6: Joint scenarios for log index values at years 1 and 2 (area of symbol proportional to new weights); Original scenarios generated using 20% volatility (left hand chart) and 40% volatility (right hand chart)

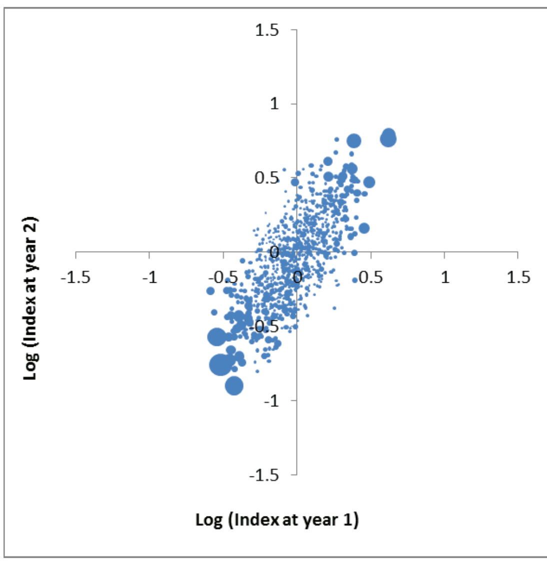

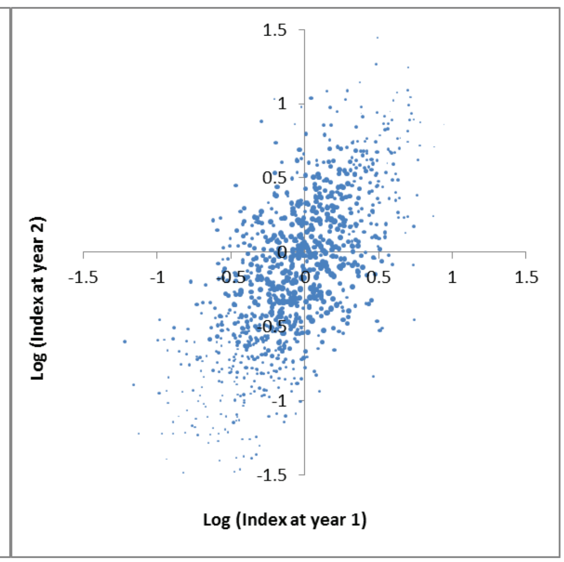

In this example including the realized variance in the set of targets reduces the effective number of scenarios from 609 to 306 (starting with base of 20% vol) and from 870 to 698 (starting with base of 40% vol).

8 The theoretical correlation here is $\frac{1}{\sqrt{2}}$ =0.71, while the reweighting technique produces estimated correlations of 0.84 (20% base volatility assumption) and 0.55 (40% base volatility assumption).

4. Summary

In this paper we have explored the use of scenario re-weighting as a method for post-processing scenario sets to reflect calibration targets without having to recalibrate the model. Using simple examples, we demonstrated how a scenario set's ability to match a set of vanilla asset prices does not uniquely define the resulting prices of more exotic liabilities (or assets). Vanilla option prices usually provide information about marginal distributions of individual risk factors at specific maturities. However, insurance liabilities often depend on assumptions about dependency both between different risk factor types, and across time. In general, such dependency assumptions are not determined by market prices and judgment needs to be applied in the setting them.

While post-processing techniques can be quite flexible in their ability to match targets, and processing can often be automated in a highly efficient way, they may result in unintended changes to distributional assumptions that are not included in the set of calibration targets. Furthermore, the examples discussed in this paper show how these distributional assumptions change significantly depending on the 'base' scenario set used.

Of course, one can include additional calibration targets in the re-weighting process, adding targets that reflect certain additional distributional assumptions where the user has strong views (such as dependencies). However, given the typical number of risk factors of interest to insurers, and the requirement for their paths, the number of such targets required is likely to make this impractical (for example, requiring a very small effective number of scenarios). And in practice, no matter how many targets are included, there will always be aspects of the distribution that are uncontrolled.

Ultimately, a liability valuation model provides a way of translating assumptions about the distribution of underlying risk factors into liability values. Market asset prices are an important input into such models, but in practice many other subjective modelling assumptions also affect liability values. By recalibrating the model (rather than reweighting), the user can ensure that all such assumptions are understood and controlled.

References

1. Avellaneda, M., Buff, R., Friedman, C., Grandchamp, N., Kruk, L., & Newman, J. (2001). Weighted Monte Carlo: a new technique for calibrating asset-pricing models. *International Journal of Theoretical and Applied Finance*, 4: 1-29.

2. Margrabe, W. (1978). The value of an option to exchange one asset for another. *Journal of Finance*, 33: 177-186.

 SP36898/101215/IND-104