A smart approach to the validation of stochastic dynamic ALM models

Grzegorz Darkiewicz-Moniuszko, Ph.D., FRM
Ed Morgan, FIA

Milliman

For companies using deterministic, liability-only embedded value models, a full understanding of the results (and how they depend on the assumptions made) is usually within reach, and validation of these models may be relatively straightforward. However, model requirements have been changing at an accelerating pace, with the evolution of reporting measures such as Market Consistent Embedded Value, Solvency II, and now IFRS 17, and in some markets these requirements necessitate the introduction of stochastic dynamic Asset Liability Management (ALM) models. These involve complex modelling and a jungle of rules and assumptions which may not be easily understandable and can impact important disclosures in a highly complex and unpredictable way. In these cases, validation of the model is no longer straightforward and new, innovative approaches to validation are needed.

Introduction

Market consistent valuation has become a standard in reporting for the insurance industry. In many markets this has meant that the use of complex stochastic dynamic ALM models can now be unavoidable, given the existence of significant blocks of business with embedded options and guarantees, such as participating products. Working with this type of model creates important challenges, but the complexity of the models can also be seen as a potential opportunity to manage the business with an enhanced insight into the risks.

One of the key aspects in this process is to gain trust in the calculations and have confidence that the reported figures are correct. Given the significant complexity of stochastic dynamic ALM models this is not straightforward, as there is a heightened likelihood of errors in the code and/or the model may not be used correctly.

In our consulting work we have come across many client models which produced fundamentally wrong figures. But we have also seen models that produce results which, at first sight, are counterintuitive, but on further investigation prove to be justified. Similarly, there can be cases where the results are strange, not due to errors as such, but simply because the logic of the model is driving unintended behaviour (e.g., the build up of large negative cash balances). Performing an independent validation of a model is therefore often a critical step in developing confidence in the results of that model and flushing out any issues.

Performing a robust validation of stochastic dynamic ALM models is rarely straightforward however. Historically, validations often took the form of replicating selected functionality from the model within Excel. However, this task can be a lot of work, and often it may not be possible to replicate all features of an ALM model within a reasonable timescale. As a result, validation exercises are often not comprehensive enough to arrive at a high level of confidence that the model is working well. Reliance then has to be placed on top-down reasonableness checks on the results. Given the number of 'moving pieces' in the overall model, setting up such checks can often require leaps of actuarial faith.

Another issue is that ALM models often include approximations to facilitate the modelling or to increase performance, such as simplifications in the modelling of asset-liability interactions, simplified definitions of management actions, and less frequent passing of information between assets and liabilities. It is typically difficult to check the materiality of such approximations, particularly across stochastic results.

We have seen quite a lot of frustration with this type of situation, with companies having to rely on models with which they are not fully confident. This is why we have developed an innovative validation approach for dynamic ALM that is based on running two models in parallel: the client's own ALM model, and an auxiliary model developed by Milliman using an innovative tool that we have designed for this purpose: Agile ALM.

While comparing the results of independent models has always been a valuable validation technique, it has hitherto been impractical for stochastic dynamic ALM models. The 'trick' which enables our approach is that we have developed a way of separating the asset model from the liability model while maintaining practically at the same level of accuracy as that of the full dynamic ALM run. This separation of the asset and liability models is highly efficient as it allows us to take advantage of the client's existing liability modelling and focus the validation directly on the ALM modelling and interactions.

Current approaches to validation of stochastic dynamic ALM models

Even though stochastic dynamic ALM models can be very complex, they are often also a very useful tool in everyday management of the business, for example:

1. Valuing options and guarantees for regulatory purposes and other disclosures, like Solvency II reporting, IFRS 17 reporting, and Market Consistent Embedded Value
2. Pricing based on the economic value of newly written products
3. Finding a strategic asset mix and steering the investment process (ALM)
4. Determining the cost of hedging the financial risks of life insurance products
5. Developing dynamic management actions aimed, for example, at stabilising the value of the business and managing unwanted risks (e.g., the risk of excessive lapses)

Before stochastic ALM models can be fully relied upon they should undergo rigid testing and validation, as they generally have a very complex design and a high diversity of parametrisations and settings (including different investment options, management actions, dynamic policyholder behaviour, etc.).

In some cases, this complexity can be due to complicated asset classes or liability features. However, in our experience most issues with ALM models do not arise within the standalone modelling of the assets or the liabilities, but instead within the complex interactions of those assets and liabilities, and the points at which investment decisions are triggered (for example, at the moment of asset sales).

Validation of such models is therefore rarely based on checking the programming code as, on the one hand, this could be extremely time consuming, and, on the other hand, one might miss important findings showing up more readily in an analysis of the results of the model rather than in the code (which in many cases may be quite dense and difficult to follow).

Validation of dynamic ALM models will therefore typically centre around reasonableness checks on the results (including sensitivity analyses) and replication of some of the features of the model within Excel. However, this usually only allows for part of the picture to be revealed. As a result, while it is usually possible to verify how individual assets are modelled in terms of calculating market values, book values or book returns, and to check the modelling of liability features for the most material insurance products, the comprehensive validation of dynamic interactions between assets and liabilities, including asset sale strategies, is often too difficult or too time consuming to carry out in any detail.

A new approach to validation using independent recalculation

Imagine you are required to perform a validation of an ALM model and are in the fortunate position of having a second, equivalent ALM model available to you which can be run on the same portfolio of business and produce results for side-by-side comparison. In this case the validation can be much easier as, from the start, one can clearly identify potential issues by spotting figures which disagree and investigate the reason for these differences.

For example, if over the course of a few runs you see that total investment returns on assets are the same in both the base and validation models then there is practically no need to trace book returns for any individual asset. Moreover, you can also be reasonably sure that the realising of capital gains and losses when assets are sold is being performed correctly. On the other hand, if you spot a difference in investment return at the portfolio level, you can look at more granular results to detect for each asset class, or at each moment of the ALM projection, where the problem is arising. This would make validation so much easier and more comprehensive!

The 'only' problem with this approach is that creation of the second ALM model through re-modelling the portfolio would make this validation very resource-intensive—in effect becoming as big a job as building the original model in the first place!

In our experience, companies will generally have achieved a good level of confidence in their deterministic liability models. This suggests a more practical approach to validation: If the liability component of your model has already been well validated, and if the complexity of the results is in reality generated by its interactions with the asset component, then you can turn your primary focus to validation of this latter aspect.

Our new approach provides a way to do this: generating a parallel set of check results by taking advantage of the (well validated) liability model and re-running it against an independent model of the assets that reproduces the ALM interactions.

With our approach, asset models and management rules are standardised in a separate tool—Milliman Agile ALM. Agile ALM differs from typical ALM models as liabilities are not modelled but instead import liability cash flows into the tool for each economic scenario and projection period.

However, the tool is more than simply an asset model, as it also captures the interactions between assets and liabilities in the same way as a full dynamic ALM model. It does this via a process of iteration to ensure a correct fitting of the assets and liabilities. As shown in our example below, this simple process is in fact very powerful and provides a robust independent validation of the results calculated by the base model. In particular, Agile ALM involves:

- Economic scenarios which are imported from an external economic scenario generator
- Assets held at the start of the projection which are input on an asset-by-asset basis
- Investment strategies which are set up (for example) in terms of strategic asset mix and the duration of reinvestments
- Parametrizations which are similar to those used by typical ALM models, for example, management rules for the order of sales of assets, and/or assumptions that the realisation of gains and losses are made to achieve a certain target yield.

The only assumption is that the deterministic liability model has already been subject to its own validation.

Case study of validation using Milliman Agile ALM

We were asked by a major multi-national insurer to validate the company's calculation of the time value of financial options and guarantees (TVFOG), for which the company uses a popular third-party financial projection tool. Milliman's analysis was based on calculations made running the company's dynamic ALM model and also its equivalent liability-only model. The results were based on the company's own stochastic economic scenarios. We also received information on the company's asset holdings and management rules regarding the realization of gains, profit sharing, and other related items.

FIGURE 1: RATIO OF CONVERGED SCENARIOS

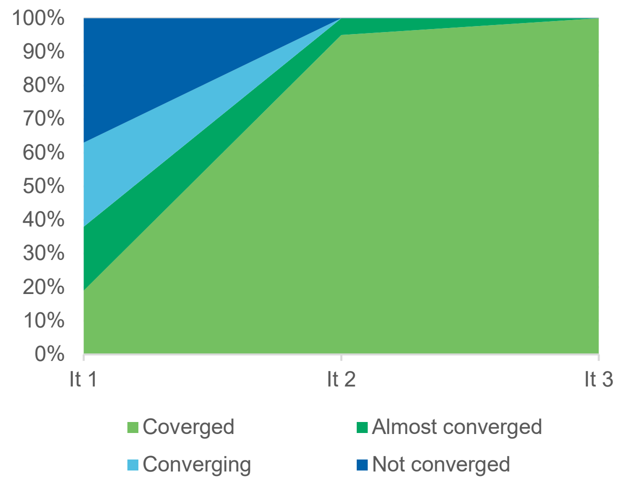

We set up Milliman Agile ALM using information provided by the company, together with cash flows produced by running the company's liability-only model. A process of 'fitting' the Agile ALM model was made via an iterative process whereby Agile ALM took the liability cash flows as inputs and calculated asset book yields. These book yields were then used in a new run of the liability-only model to calculate new liability cash flows. The process was repeated until convergence was reached and the results no longer varied. We find that this process typically takes two to three iterations, but many scenarios in fact converge more quickly. As shown on the graphic above, the green area represents the ratio of scenarios which converged to the target value; yellow — almost converged; orange — converging; and red not converged (to illustrate the speed of convergence in this case study).

Initially the analysis indicated some differences between the results produced by the company's dynamic ALM model and the results produced using the company's liability-only model with Agile ALM. These differences were investigated and some key issues with the company's ALM model were identified. Once these issues were fixed, we re-ran the analysis and then achieved extremely close results between Agile ALM and the company's ALM model, with a difference in calculated TVFOG of less than 1% and a difference in stochastic best estimate liability (BEL) of 0.05%.

The analysis provided a practical and powerful way to:

- Identify and correct issues with the company's dynamic ALM model.
- Provide a strong and independent validation of the company's TVFOG calculation.

Conclusion

Given the importance of results calculated using stochastic dynamic ALM models for regulatory and other financial reporting, as well as for risk and capital management, a practical, reliable, and efficient way to perform model validation can be extremely valuable. Independent calculation of TVFOG or BEL can also be very useful in other circumstances (for example, M&A due diligence).

Furthermore, we note that using Agile ALM to validate existing models is only the first step. With the Agile ALM analysis already set up, it is then also possible, for example, to test alternative strategic asset allocations or different management rules with a goal to optimising the ALM process or other risk management processes. Another application can be to estimate the TVFOG for alternative product designs.

Overall, we believe this new approach is an extremely valuable addition to the toolkit for financial modelling under a wide variety of circumstances.

CONTACT

Grzegorz Darkiewicz-Moniuszko
grzegorz.darkiewicz@milliman.com

Ed Morgan
ed.morgan@milliman.com