Curve fitting for calculating SCR under Solvency II

Practical insights and best practices from leading European Insurers

Leading up to the go live date for Solvency II, insurers in Europe are in search of practical solutions for calculating their Solvency Capital Requirement (SCR) under Solvency II. There are a range of options available to insurers to calculate SCR from Standard Formula to partial internal models to full internal models. It is generally accepted that full internal models promise a greater range of advantages to insurers compared to the other alternatives. One barrier for firms considering a full internal model approach may be the perceived technological requirements. An internal model is essentially a huge Monte Carlo simulation that requires precise data management and the processing of tens of thousands of scenarios. This approach has historically been time consuming, expensive and computationally demanding, but new options, such as curve fitting offer dramatic performance improvements at a more acceptable cost.

As a new technique, there is a shortage of practical knowledge generally available to insurers facing similar challenges about the use of applying a curve fitting methodology to help calculate SCR. The purpose of this paper is to provide an insight into how some of Europe's leading insurers are applying Curve Fitting in practice, the benefits that are being derived and some of the practical challenges that need to be addressed.

We hope that readers will find this a useful contribution as they prepare for Solvency II and we would like to thank L&G and Aviva for their involvement to make this possible.

Kind regards,
Dr. Andrew Aziz
Head of IBM Risk Analytics - Insurance and Buy Side Institutions

In an ideal world, insurance companies would run full stochastic scenario simulations across all their assets and liabilities in order to measure their risks and to calculate their economic and regulatory capital and other metrics. Even the biggest computer installations in the industry, however, cannot currently perform a full stochastic analysis in a time frame that would be meaningful, especially for large groups. So companies need a way of either compressing their portfolios – particularly their liabilities – or finding proxies that can make the task more manageable.

One approach that is finding increasing favor, especially among companies preparing internal models for Solvency II, is curve fitting – also known as formula fitting. The idea is to develop a formula that mimics the behavior of best estimate liabilities as calculated by an actuarial system under a range of scenarios. To put it another way: to find the curve that best fits the data points of the best estimate liabilities. In addition to liability estimation, some organizations are choosing to apply curve fitting to the asset side of the balance sheet while others are modelling assets directly.

The resulting curve or formula is then applied to the full range of scenarios – typically 100,000 or more, incorporating virtually all the relevant risk factors. Because applying the formula is a straightforward calculation, it is far quicker than simulating all the scenarios, helping reduce a task that could potentially take years to a matter of hours.

Producing a curve or formula involves four main steps. First, a number of training or calibration scenarios – typically 20-25 – are run in the actuarial system to obtain a set of target best estimate liabilities. These will be the data points for the curve fitting process. Next, a prototype formula or curve that might fit the data points is created. Note that the polynomial form will need to be appropriate for the number of risk factors under consideration, and should include cross terms if there are interactions of risks which have additional impact on the liabilities. The third step is to try to fit the curve to the data points. Rarely will it be right the first time, so the formula needs to be adjusted until it reproduces as accurately as possible the best estimate liabilities for each of the calibration scenarios. Finally, a number of out-of-sample scenarios are run in the actuarial system, and the formula is then tested on these out-of-sample runs to see how accurately it recreates the best estimate liabilities.

Naturally, this simplified description can hide a number of challenges and limitations. Nevertheless, curve fitting can have a number of advantages over other proxy modeling methods, which has led some of the largest UK insurers to adopt the approach for their current regulatory reporting and to plan to make it central to their Solvency II internal models.

London-based Aviva aims to have a single integrated capital model for the whole Aviva group that follows best practice and reflects as accurately as possible the risks of the business. “This model will be used for a multitude of purposes including risk management, calculating the solvency capital requirement for regulatory reporting under the Solvency II regime, strategic capital allocation and planning, merger and acquisition activity, etc.,” says Tim Thornham, technical development director in the economic capital team at Aviva. The model must not only represent Aviva’s exposure to its key risks as accurately as possible, but must do this rapidly enough for it be fully integrated into the company’s analytical and decision-making processes.

London-based Legal & General Group’s has a similar ambition. When it looked at its specific objective – calculating the value-at-risk (VaR) of its group balance sheet for a 1-in-200 year event over a one year time horizon for virtually all its risks and correlations across the whole group – and contemplated using traditional modeling methods, the scale of the challenge became apparent. “To run just our with-profit model takes an hour, even if we squeeze every bit of efficiency out of it. To run it 100,000 times would take 10 years,” says Stuart Carroll, director of savings actuarial at L&G.

Performance of the liability models is a concern for all firms building internal models for Economic Capital and Solvency II” says Curt Burmeister, Head of Buy Side Products, Risk Analytics, IBM. “One of our clients estimated that to calculate the 1-in-200 year VaR on its life insurance balance sheet using its traditional actuarial models would take more than 45 years.”

Given these facts, it was clear to the insurers that they needed a proxy or 'lite' modeling method they could use at the company or group level to model their overall balance sheet more frequently and that would complement the traditional 'heavy' modeling they do at the business unit level quarterly or even less frequently. "What we wanted to do was to separate our heavy models from our group representation, and then run the lite group models under all the relevant risk factors, which in our case is 200 or so," says Carroll.

In their search for effective lite models for company or group level modeling, especially with Solvency II internal models in mind, insurers in the UK and Europe have so far mostly opted for either curve fitting or replicating portfolios. As the name suggests, replicating portfolios are created by selecting a portfolio of assets from a candidate asset set, designed to replicate the behavior of complex liabilities. However, while this approach can model market risk, it is difficult to apply to insurance risks.

According to Thornham of Aviva, "Replicating portfolios can work well as a replication approach for market risks, but it is not clear how to construct asset portfolios to replicate non-market risks." This means that insurers need to calculate non-market risks using a different approach. By contrast, curve fitting is a way to represent the value of the life insurance company's balance sheets as a function of virtually all the risk factors affecting it, not just market risk. "So curve fitting allows a consistent approach to modeling life insurance risks and hence, a more integrated approach to overall aggregation and modeling of the interdependencies between market and non-market risks," he says. "In practice, we then go on to aggregate the life risks, modeled using curve fitting, with the non-life risks, modeled within an aggregate dynamic financial analysis environment and sampled empirically. This provides our total balance sheet exposures across life and non-life in a consistent set of calculations."

"Replicating portfolios can work well as a replication approach for market risks, but it is not clear how to construct asset portfolios to replicate non-market risks."

Carroll at L&G sees curve fitting as a generalized form of replicating portfolios. "Any proxy model, whether it be a replicating portfolio or other form, can be reduced to a polynomial of the risk factors. Curve fitting cuts to the chase and goes straight to the polynomial without having to create a candidate universe of assets and select an appropriate portfolio, which can be quite an art," he says. Creating a replicating portfolio requires a deep understanding of liability behaviors and a vast library of possible assets to replicate the liabilities whereas curve fitting is more intuitive as well as more mechanical than replicating portfolios.

Although curve fitting has several advantages as a proxy modeling technique for liabilities, it nevertheless presents its own challenges.

"Finding curves which accurately represent the underlying loss behavior for different risk types of complex products is a challenge," says Thornham. The challenge is particularly notable for products with complex guarantees, or where management actions taken in different scenarios create highly non-linear loss behavior. Thornham states that the solution is to do more analysis on the losses in order to understand and refine the calibrations, which is an ongoing process of constant improvement.

Another challenge is how to update the curves to make sure they remain an appropriate representation of the underlying losses between full recalibrations. One solution is to build into the curve fitting models, a set of update parameters that allow the curves to be scaled or otherwise adjusted in response to changes in exposure that result from external market changes or internal business developments, which take place between full recalibrations.

Also, although curve fitting does not require the deep knowledge of assets and their behavior as is the case with replicating portfolios, it nevertheless helps to have a good understanding of the insurance liabilities that are being fitted. (Practitioners point out that the technique can also be applied to the valuation of assets.) Knowing the business and how it should respond to a risk or combination of risk drivers helps in making a good choice of calibration scenarios. The calibration process in turn enhances the understanding of how a business responds to risk drivers, and it becomes a virtuous circle.

In terms of the calibration range, it is important to cover a sufficiently wide proportion of the distribution of risk driver values. So if it is determined that a 1-in-200 year shock for UK equity would be a drop of 40 percent in a related index, then the calibration range for the risk driver 'UK equity' should cover at least a 40 percent price fall.
The number of calibration points should be carefully chosen so that it gives palpable insights without the danger of over-fitting.

In L&G's view, curve fitting is not simply a form of linear regression where it is a matter of trying to fit a line to a set of random variables. "We think curve fitting is a problem of estimation. We believe that the liability curves behave in a predictable way given particular inputs – in other words, a specified set of inputs will consistently generate a unique set of outputs," says Carroll. And since it is a problem of estimation, it is possible to apply approximation theory, which is helpful in finding both the best function and the best calibration scenarios. "There is a mathematical way of identifying those points that not only helps us define the curve we are fitting, but it also helps to define the maximum error – the difference between the point on our formula-fitted curve and our true curve where our error is at the maximum and provides us with a simple measure of the goodness of fit of the curve," he says.

Like all modeling approaches, curve fitting has its limitations.

One of the weaknesses of curve fitting is that it is only as good as the heavy models you use, so if there are vulnerabilities in the heavy models, particularly under a stressed environment, it may produce nonsensical results. This can be overcome by carrying out response function analyses on the fitted curves, where any suspicious kinks could point to deficiencies in the heavy models. Others point out that the heavy models may give only limited information about the impact of management actions on with-profit business, and may not include reinsurance if it is only modeled at the group, not business level.

No proxy modeling method is without its limitations and challenges, and the significant benefits of curve fitting means that interest in the approach is growing as UK and European insurers look at developing an internal model for Solvency II. A survey in February 2011 of 60 UK insurers by the Economist Intelligence Unit on behalf of Deloitte showed that 29 percent planned to use curve fitting as the primary proxy modeling method for their internal model compared with 5 percent who plan to focus on replicating portfolios, with 35 percent planning to use both (and 31 percent undecided).

The Curve Fitting methodology has been used for a number of years by many of the leading firms in the UK. For example, Aviva has been using curve fitting as its primary technique for economic capital and ICA modeling for over four years. As more firms use curve fitting, the methodology continues to improve. "Although the basic methodology is well established, insurance companies are evolving and enhancing the curve fitting techniques as they apply it to new problems such as Solvency II" says Burmeister.

In terms of implementing curve fitting, one of the most time and resource consuming aspects is preparing the training scenarios. These need to be run in the business line actuarial systems and can require substantial computing resources. Once the training scenarios are produced, fitting the curve is a relatively quick process.

"One of the weaknesses of curve fitting is that it is only as good as the heavy models you use."

It can be done on a spreadsheet or in a mathematical software package. Solvency II however, requires that internal models are auditable, robust, scalable and embedded in the insurer's businesses processes.

"Curve fitting is an effective way of creating a liability proxy, but productionizing the methodology is a challenge," says Burmeister. A large company is likely to want actuaries in various parts of its business to help with the work of creating the curve fitting calibration, and this can present a challenge in terms of managing the workflow of the process. There are then the downstream challenges. "Once you have produced the 100,000 or more scenarios using the formulas, how do you generate the Solvency II numbers from this – the solvency capital requirement, minimal capital requirement, risk margin, etc? Fitting the curves is actually only the first step in a long chain of events that must be controlled and managed," says Burmeister.

It is because of these challenges that insurers such as Aviva and L&G have looked to IBM to provide a controlled industrialized environment for their curve fitting-based modeling and aggregation.

In summary, curve fitting is a closed form solution that replicates the results of full liability models under a large number of random scenarios, producing results in a fraction of the time. "That means we can run many more sensitivities – what-if scenarios – to improve our understanding of the risk profile of the business," says Thornham of Aviva. Carroll of L&G adds: "As we use the technique more, hopefully it will give us deeper insights into the risks in our business and support our decision-making in a more active and timely way than is possible with traditional actuarial models."

Burmeister concludes: "Curve fitting is a clever yet simple solution. It is what you build on it that counts and different companies are developing it in different ways."

"Curve fitting is an effective way of creating a liability proxy, but productionizing the methodology is a challenge."

IBM® Algorithmics Economic Capital and Solvency II solution

Algorithmics Economic Capital and Solvency II solution is specifically designed to help meet the needs of insurers irrespective of their size, complexity and level of sophistication. The solution comes in three editions which are robust solutions that cover the three pillars of Solvency II by addressing the calculation of solvency capital, governance and reporting. Built around industry "best practices" by working with leading insurers, the solution is pre-configured to help to lower project risk and give companies more confidence to meet their Solvency II deadlines. All three editions provide the client with the flexibility to incorporate liability cash flows generated either by a third party actuarial system or by IBM® Algo Financial Modeler®. Looking beyond the Solvency II deadlines, as insurers' needs change, the migration path between the editions offers the opportunity to scale up to more powerful and complex analytics on IBM's award-winning, tried and tested platform.

About Business Analytics

IBM Business Analytics software delivers data-driven insights that help organizations work smarter and outperform their peers. This comprehensive portfolio includes solutions for business intelligence, predictive analytics and decision management, performance management, and risk management.

Business Analytics solutions enable companies to identify and visualize trends and patterns in areas, such as customer analytics, that can have a profound effect on business performance. They can compare scenarios, anticipate potential threats and opportunities, better plan, budget and forecast resources, balance risks against expected returns and work to meet regulatory requirements. By making analytics widely available, organizations can align tactical and strategic decision-making to achieve business goals.

For more information

For further information please visit www.ibm.com/business-analytics.

Request a call

To request a call or to ask a question, go to www.ibm.com/business-analytics/contactus. An IBM representative will respond to your inquiry within two business days.

Notice

The information contained in this documentation is provided for informational purposes only. Although efforts were made to verify the completeness and accuracy of the information contained in this document, it is provided "as-is" without warranty of any kind, Express or Implied. In addition, this information is based on Algorithmics' current product plans and strategy, which are subject to change by Algorithmics without notice.

Algorithmics will not be responsible for any damages arising out of the use of, or otherwise related to, this document or any other materials. Nothing contained in this document is intended to, or shall have the effect of creating any warranty or representation from Algorithmics (or its affiliates or their suppliers and/or licensors); or altering the terms and conditions of the applicable license agreement governing the use of Algorithmics software. References in this publication to Algorithmics products or services do not imply that Algorithmics intends to make them available in all countries in which Algorithmics operates.

For any reference to an Algorithmics software program, the software program can be used to help the customer meet compliance obligations, which may be based on laws, regulations, standards or practices. Any directions, suggested usage, or guidance provided by the software program, or any related materials, does not constitute legal, accounting, or other professional advice, and the customer is cautioned to obtain its own legal or other expert counsel. The customer is solely responsible for ensuring that the customer and the customer's activities, applications and systems comply with all applicable laws, regulations, standards and practices. Use of the software program, or any related materials, does not guarantee compliance with any law, regulation, standard or practice.

Any information regarding potential future products and/or services is intended to outline Algorithmics' general product and service direction and it should not be relied on in making a purchasing decision. Any information mentioned regarding potential future products and services is not a commitment, promise, or legal obligation to deliver any material, code, functionality or service. Any information about potential future products and services may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for Algorithmics' products or services remains at Algorithmics' sole discretion.

© Copyright IBM Corporation 2012

IBM Corporation
Software Group
Route 100
Somers, NY 10589
USA

Produced in the United States of America
November 2012

IBM, the IBM logo, Algorithmics and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

The content in this document (including currency OR pricing references which exclude applicable taxes) is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs. THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation.

Please Recycle

YTW03228-USEN-00