The Modeling Platform

April 2016
Issue 3

Model Efficiency in the U.S. Life Insurance Industry

By Tony Dardis

In a life insurance industry where there seems to be an insatiable desire for analytics, practitioners are increasingly under pressure to produce more numbers, and to produce them quicker.

This in turn has led to a cottage industry in itself—the world of "model efficiency."

Model efficiency refers to the development of financial models that yield results with a minimum of time and effort. Model efficiency might be achieved through well-written code, creative application of mathematical and actuarial techniques, or state-of-the-art technology.

Its emergence as a field of practice in the life insurance industry is rooted in the fact that models of life insurance companies can be complicated:

- Life insurance products, and the assets backing them, can be complex.
- Millions of contracts may need to be modeled, each contract with its own special characteristics.
- Long-term projections are involved and may be performed over hundreds of thousands of different projection paths.

In time, model efficiency has emerged as both a science and an art, involving elements of creative thinking and technical know-how.

In this article I walk through the history of model efficiency in the U.S. life insurance industry, examine the current application of model efficiency, and look ahead to what we might expect to see in the future. I would like to think that, for those who stay with me through this article, you will get some useful insights that may in turn prove to be helpful to you in developing your own efficient models.

A HISTORY OF MODEL EFFICIENCY IN THE U.S. LIFE INSURANCE INDUSTRY

Figure 1 gives a summary timeline of key milestones for model efficiency in the U.S. life insurance industry over the years.

Before the widespread availability of computing power, insurance calculations were performed manually. Commutation

Figure 1
A Brief History of Model Efficiency in the U.S. Life Insurance Industry

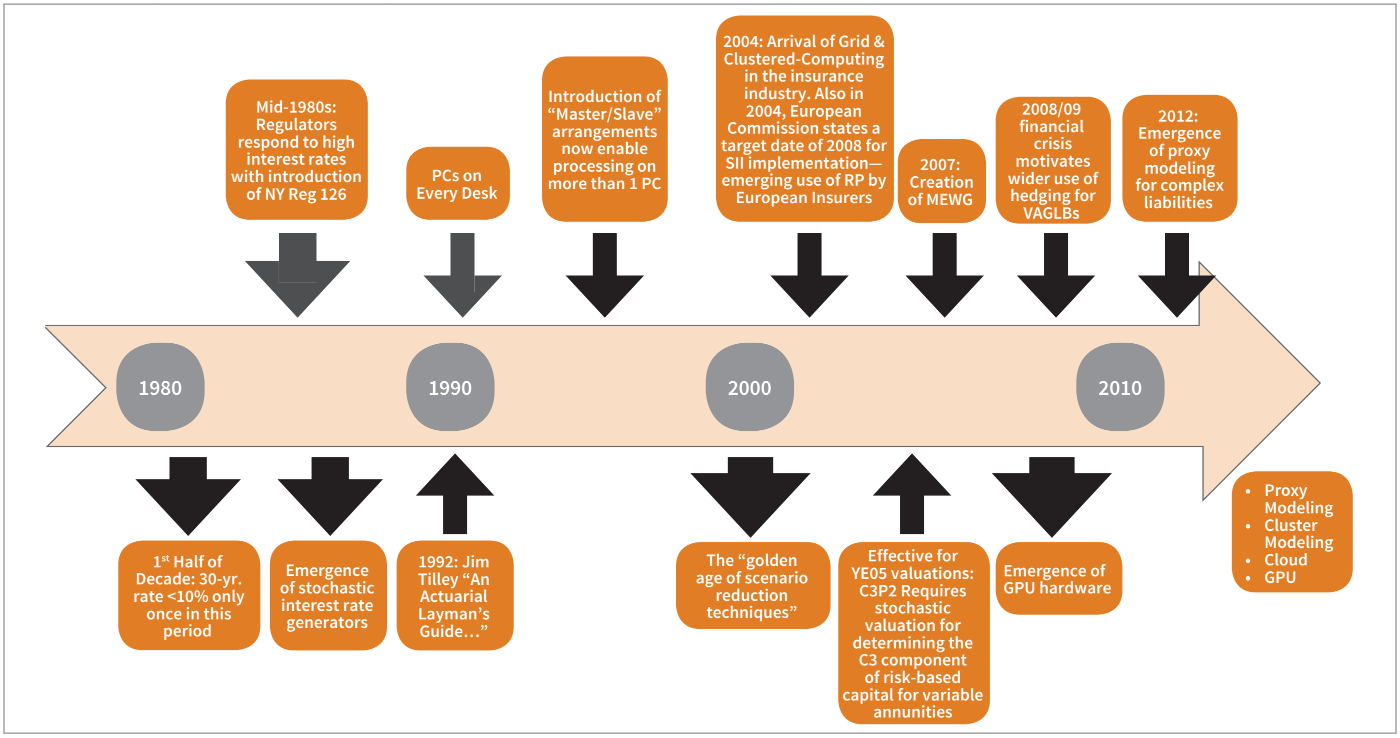

functions were a primary tool for actuarial work until well into the 1980s. These could summarize in a single number the present value of discounted life contingent cash flow projections for a given rate of interest. The precise form and value of a particular function would depend on the nature of the projection and on age and duration, as well as the rate of interest; large books of commutation tables were a fixture in life offices. Commutation functions were undoubtedly ingenious, a testament to the actuarial profession of yesteryear, and indeed might be viewed as a model efficiency technique, saving the actuary from tedious calculation. (See Figure 2.) They also fostered a wonderful common language within the actuarial profession, and continue to serve as a convenient shorthand even in today's high-tech world.

Figure 2
Commutation Functions as a Modeling Efficiency Technique

$$Whole\ Life\ Policy = A_x = \frac{M_x}{D_x} = \frac{\sum_{t=x}^{\infty} C_t}{D_x} = \frac{\sum_{t=x}^{\infty} v^{t+1} d_t}{v^x l_x}$$

The emergence of the desktop micro-computer or "PC" in the early 1980s effectively revolutionized actuarial work. PCs could be accessed at will, and enabled multiple iterations of code so that a production process could be put in place at a fraction of the time it would take to get a mainframe process fully up and running. Although the early PCs certainly had significant limitations around data storage and processing power, they paved the way for actuaries to build models of insurance business and readily perform cash flow projections. This allowed for the development of "profit models" in product pricing.

From this point on, the world of actuarial modeling moved forward quickly. Changes in the economic environment necessitated better models, the development of increasingly complex insurance products and asset/liability profiles stressed the models and drove the need for efficiency, and advances in technology enabled increased usage of actuarial models.

Economic Environment and Increasing Complexity

Only once during the **first half of the 1980s** did the 30-year rate in the U.S. dip below 10 percent. High interest rates drove disintermediation, in which policyholders drew down the value of their policies, forcing insurers to sell assets at a loss. At the same time, products were being issued with high interest rate guarantees, putting insurers at risk of decline in interest rates. Regulators responded to this emerging risk by requiring cash flow testing for annuities using deterministic scenarios. This was the first widespread use of interest rate scenarios in an actuarial setting. By the mid-1980s, New York regulators required testing using the now familiar "New York 7" scenarios, first for annuities and eventually for most life insurance business.

During the **second half of the 1980s**, we started to see a handful of insurers using stochastic interest rate scenarios in addition to the deterministic scenarios required for cash flow testing under New York Regulation 126. Actuarial literature started to deal with the development of stochastic interest rate generators. Jim Tilley's classic 1992 paper, "An Actuarial Layman's Guide to Building Stochastic Interest Rate Generators," heralded the start of widespread stochastic modeling in the U.S. insurance industry.

2005 saw the introduction of the C3 Phase II requirement for variable annuity capital, adopted by the National Association of Insurance Commissioners (NAIC) on Oct. 14, 2005, and effective for year-end 2005 valuations. This required stochastic valuation for determining the C3 component of risk-based capital for variable annuities. Suddenly model efficiency became of critical importance to all VA writers.

In **2008-09**, the financial crisis alerted insurers to the very real risks associated with variable annuities with guarantees and motivated much wider use of hedging as a means to better manage these risks. This is itself further exacerbated the processing requirements of variable annuity business, already brought to a new level by the introduction of C3 Phase II a few years earlier, e.g., the need to do full- scale hedge effectiveness testing to get credit for dynamic hedging in the C3 Phase II valuation.

Innovations in Efficiency

Around **the mid-1990s**, actuaries started to look at how to run fewer scenarios that were representative of a larger set for stochastic modeling, as a means of managing the run-time issues associated with running many scenarios. In the decade to follow, some very imaginative ways of tackling the issue were proposed and in some cases used successfully in practice—the "golden age of scenario reduction techniques":

- Alistair Longley-Cook proposed a rather novel approach using least squares to fit 1,000 stochastic scenarios to the New York 7 scenarios in his paper "Probabilities of 'Required 7' Scenarios (and a Few More)," *The Financial Reporter* (July 1996).
- In her paper "Representative Interest Rate Scenarios," *North American Actuarial Journal*, vol 2, no. 3 (July 1998), Sarah Christiansen developed a practical approach to picking representative scenarios from a stochastically generated set by testing multiple subsets from the full set and choosing the subset that best meets various criteria (e.g., best matches the mean of each term rate in the scenario, extremes, standard deviations, etc.).
- Yvonne Chueh used distancing techniques to establish a reduced scenario set in her paper "Efficient Stochastic Mod-

eling for Large and Consolidated Insurance Business: Interest Rate Sampling Algorithms," *North American Actuarial Journal*, vol. 6, no. 3 (July 2002), a methodology that has been used quite widely throughout the industry over the past decade in one form or another.

- In their 2005 paper "Variance of the CTE Estimator," *North American Actuarial Journal*, vol. 9, no. 2 (April 2005), John Manistre and Geoff Hancock propose an approach for choosing a representative scenario set that works well in the tails of distributions.

In **2004**, the European Commission stated a target date of 2008 for new Solvency II regulations to be effective, upping the bar and demanding that companies squeeze yet more out of their models. In response, ING pioneered the approach of using a replicating portfolio as means of reducing liability model runtime associated with Solvency II-type capital calculations; other European insurers followed suit.

In **2007**, the American Academy of Actuaries established the Model Efficiency Work Group, or MEWG, as it was affectionately known as by its members, a subgroup of AAA's Life Financial Soundness/Risk Management (FS/RM) Committee, which in turn was responsible at the time for making proposals for the implementation of a principles-based approach (PBA) to reserves and capital for life insurance in the United States. The FS/RM Committee recognized that for some companies, the requirements of PBA could lead to onerous calculation requirements, and wanted to have expert input from a separate group focusing on ways to mitigate this burden and hence make calculations more manageable without compromising on accuracy. This group had some successes in promoting model efficiency as a practice area around the industry, notably:

- Two surveys of model efficiency practices in the U.S. life insurance industry, the first being published in November 2007, summarizing responses from 30 companies, and the most recent being published in April 2013, based on responses from 51 companies.

- The publication of "Modeling Efficiency Bibliography for Practicing Actuaries," last updated December 2011, which lists the publicly available documents in the area, categorized according to "actuarial modeling techniques" and "technology," and a handful of further subdivisions within each of these broad headings (more on this below).

- At the Society of Actuaries' Life 2008 Spring Meeting in Quebec City, a "conference within a conference" was presented, with a series of four panels related to model efficiency under the banner of "Introduction to Modeling Efficiency and Scenario Reduction." This was the first large-scale "event" for the U.S. actuarial profession dedicated to the topic of modeling efficiency.

In **2012 and beyond**, increasing attention is paid to the use of liability proxy models, whereby the value of a liability is expressed as a polynomial function, as a way to speed up run model time by not having to run a full "heavy" actuarial model.

Technological Innovation

Many of the early actuarial models had been built by individual actuaries using the programming language APL either on a mainframe or PC environment. During the late '80s and early '90s, commercial software running on the PC became increasingly common. In the **mid-1990s**, commercial actuarial systems originally designed to run on a single PC evolved to take advantage of more than one computer. This was typically realized in a "master/slave" arrangement whereby the software running on one PC, the master, was programmed to off-load some of its work to the same program running in slave mode on other PCs. While this paradigm lacked the sophisticated resource management of cluster, grid and cloud technologies to follow—typically leveraging only a handful of computers—this early form of distributed computing offered a means of significantly reducing elapsed runtime.

Computer processing power has been increasing steadily since the introduction of the PC with the escalation in chip speed. Other technological breakthroughs in computing capacity began to take root in the **early 2000s**, with the introduction of Intel's Hyper-Threading Technology in **2002** and the subsequent emergence of dual-core and, later, quad-core central processing units (CPUs) beginning in **2005**. These advances opened up the ability to run sequences of instructions concurrently on a single computer, although it would take some vendors several years to natively leverage these capabilities within their software applications.

2004 saw the arrival of grid- and clustered-computing technology in the insurance industry, including DataSynapse GridServer and Milliman C-Squared, with Windows Compute Cluster from Microsoft (now HPC Server) and Symphony from Platform Computing following shortly after. These solutions enabled developers of actuarial systems to distribute workload over hundreds of CPUs. In addition to raising the ceiling on distributed computing capacity by one or two orders of magnitude compared to master/slave arrangements, grid technology opened the door for information technology specialists to play a role in the adoption, configuration and maintenance of actuarial systems that had previously been localized to end-user workstations.

The mid-to-late **2000s** heralded the introduction of graphics processor unit (GPU) hardware as a practical means of off-loading highly parallel computations from conventional CPUs. This powerful and innovative technology was made accessible to software developers through technologies like

OpenCL from Khronos Group, CUDA from NVidia and DirectCompute from Microsoft, and to quantitative and model developers through integration with systems like MatLab from MathWorks.

The launch of Amazon Web Services in **2006** and Microsoft Azure in **2010** ushered in the era of cloud computing as a means of accessing CPU resources on a scale capable of far exceeding the capacity available in most on-premises solutions. With technologies like Elastic MapReduce from Amazon, HPC Pack with Azure "burst" capabilities from Microsoft and GridStep Cloud Edition from Milliman, grid-enabled models could access the CPU cycles (and storage) needed for seriatim valuation, nested stochastics, and forward and backward projections—simultaneously.

THE CURRENT STATE OF MODEL EFFICIENCY

Having stepped back and viewed model efficiency from a historical perspective, where are we today and what does the future hold? In this section of the article, I'll consider today's perspective under three headings:

- Model efficiency taxonomy
- Model data-building techniques: replicating portfolio, proxy modeling and cluster modeling
- Technology

Model Efficiency Taxonomy

One of the outcomes of the work done by MEWG was the establishment of a general taxonomy, designed to categorize the various model efficiency techniques and thus to provide a framework for a common dialogue among practitioners. This framework is summarized in Figure 3.

Figure 3
Model Efficiency Taxonomy

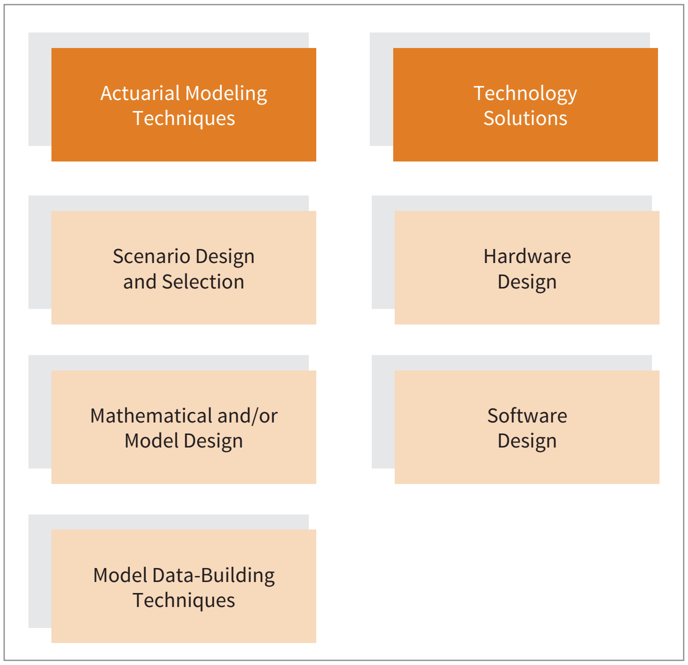

Under the heading of **Actuarial Modeling Techniques**:

- **Scenario design and selection** covers how we choose or design our scenarios and includes the wide array of scenario reduction techniques.
- **Mathematical and/or model design** covers how the choice of a mathematical approach to a model can simplify calculations and/or reduce the time required to perform these calculations. For example, for runs that require an estimate of future market prices, instead of having to generate a set of market consistent scenarios at each future point in time, we may be able to use a closed form mathematical solution, such as Black-Scholes.
- **Model data-building techniques** include traditional approaches to building actuarial models, involving the development of model points designed to bucket seriatim data by homogenous groupings, such as issue age, contract duration, contract features, etc., plus the use of emerging state-of-the-art techniques to make models even more efficient, such as cluster modeling (more on this later).

Under the heading of **Technology Solutions**:

- **Hardware design** covers the broad spectrum of using today's technology to its fullest extent to best meet actuarial processing needs, including the use of grid and cloud computing.
- **Software design** covers using efficient programming code design to best meet the requirements of a particular appli-

cation, and optimized for the underlying technology. For example, for multicore solutions with multithreading, and for GPU applications, the software must be written with considerable technical knowledge of the hardware in mind. Another issue here is that the choice of language can significantly impact performance, e.g., Excel workbooks implemented using spreadsheet formulas in combination with Visual Basic for Applications (VBA) logic are commonplace on actuarial desktops, but not readily scalable to a production environment and VBA code runs much slower than compiled C or C++ or .Net languages like C#.

It is also worth noting there are some approaches that are not readily categorized under any of the headings highlighted above, such as a hybrid modeling approach developed by Steve Craighead, which uses a mix of representative scenario techniques to create training data for a predictive model.[1] Emerging techniques such as proxy modeling may also be viewed as falling under the "hybrid" category.

Indeed, implementing efficient models often calls for combination approaches, and real success requires both actuarial and technology expertise. The former is necessary to leverage techniques such as replicating portfolio, proxy models and cluster models. The latter is necessary to ensure the selected algorithms execute as quickly as possible and model performance is not degraded by inefficient access to data.

Another related point to make here is that while model efficiency may require a blend of approaches, there is a balance to be struck between increasing complexity and ensuring a smooth process that can be maintained on an ongoing basis. Generally speaking, the more aggressive the approach to maximize performance, the more complex the process with associated potential increases in bug rate, as well as a shrinkage in the pool of sufficiently knowledgeable resources to debug, maintain and enhance the model.

The remainder of this article focuses on the areas where we have seen the most rapid developments in model efficiency most recently, looking separately at developments from the actuarial modeling and technology perspectives, and looks ahead to what we may expect to see in the future.

Recent Developments in Actuarial Modeling Techniques: Replicating Portfolio, Proxy Modeling and Cluster Modeling

So far as actuarial modeling techniques are concerned, the most rapid developments in recent years have been around the use of the replicating portfolio (RP) and, most recently, proxy modeling and cluster modeling.

The essence of the RP technique is to find a basket of assets that matches the value of a liability inventory over a wide range of shocks and then use this portfolio as a surrogate for the value of the liabilities in further analysis. The advantage of this approach is the analysis of this RP will be more manageable than working with the liability models, especially if the assets in the basket have closed-form solutions for market valuation.

As an example, the liabilities model could be run through 100 different shocked scenario sets to come up with the sensitivity of the liability market value to 100 different shocks, and the RP would be calibrated to those results. Then the RP could be run through a much larger number of shocks than would be practical for the liabilities, to come up with conditional tail expectations.

The practical application of RP first began at ING, spearheaded by Tom Wilson (now chief risk officer at Allianz) in the mid-2000s as a way to support their internal economic capital calculations based on looking at the performance of the market consistent value of the balance sheet in the tails under many real-world stochastic scenarios, similar to the Solvency II view on capital. The ING RP approach soon became de rigueur with a number of the other large multinational insurers.

RPs have certainly proved to be very useful for certain applications, but not for all. One of the limitations of the RP approach is that if based purely on liquid and analytically tractable instruments, there can be accuracy limits for some products. The need for ever-increasingly complex tail-risk orientated calculations have become more prominent—such as the calculation and projection of economic capital—and has led practitioners to look at alternative approaches to approximating liabilities. One such alternative approach is proxy modeling, which is already being used extensively in Europe to help manage the calculations required for Solvency II, and is beginning to get some traction in the United States.

The essence of the proxy modeling approach is that a function (proxy) is fitted to the liabilities, with that function expressing the liabilities in terms of the underlying risks to which the liability is exposed. Thus, for a variable annuity portfolio, we might say our liabilities are a function of equity returns, movements in the yield curve, and equity and interest rate volatility. The exercise then becomes a question of fitting that function to the liabilities to give a result that is accurate, even in the tails, and does not require frequent re-fitting.

There are a number of ways of fitting such functions but all involve essentially four steps:

- **Step 1.** Determine what risks to consider and generate "fitting points." This the key part of the exercise; getting this wrong may mean you end up with a meaningless function. After determining the risks to consider, you need to establish the points to which you are going to fit the proxy model. Figure 4 summarizes the issue under consideration here.

Ideally, you would like to be able to cast a very wide net, and calculate many accurate values (from the underlying heavy model) against which a function can be fitted. In practice, it may not always be possible to generate as many fitting points as we would like as the underlying heavy model may take too long to run. For some applications, the optimal solution might be to create a series of carefully selected, accurately calculated values and use interpolation (as in the case of radial basis functions and curve fitting); for other applications, we might be able to get good results by calculating approximate values for every point across the risk space and fitting a proxy function to these approximate values, as in the case of least squares Monte Carlo (more on all of this in a moment).

Figure 4
The Essence of Best-in-Class Proxy Modeling—Covering the Risk Space

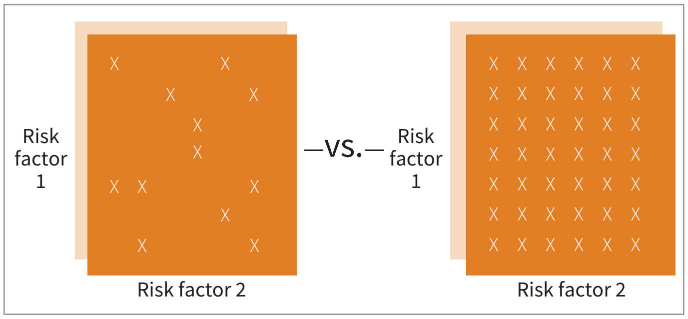

- **Step 2.** Calculate the target metric for each fitting point by using the underlying heavy model. The choice of target metric is dependent on the application under consideration, e.g., market value of liability, or net liability value (assets less liabilities).
- **Step 3.** We will now have a series of fitting values for each fitting point (which in turn represents a combination of the various risks that we are exposed to). We can now readily fit a function through those points, via an optimization routine using agreed fitting criteria. Simple curve fitting techniques, where a relatively simple curve is fitted to a series of accurately calculated values and then linear interpolation used to fit values outside the fitted value, can be a useful "starter pack." However, for better results that permit interpolation in a high dimensional space, the radial basis function approach has been demonstrated to give extremely good results. Under the radial basis approach, we establish a series of accurately calculated points, these points having been carefully selected to cover the entire spectrum, and then all sample points get considered when interpolating to a given point, enabling us to better capture the shape of complex underlying functions. Another approach is least squares Monte Carlo, where we generate fitting values that cover the entire risk space but each individual fitting point on its own is inaccurate—the success of the methodology hinges on being able to generate so many points that when you fit a curve through the points, the errors on average cancel each other. Least squares Monte Carlo works very well where the points are unbiased and independent, but is more challenging when applied to estimating CTE-based measures such as those required for Actuarial Guideline 43/C3 Phase II, for which one may need to address bias in estimates.[2]
- **Step 4.** The final part of the process is to validate the proxy function. This will include looking at out-of-sample validations. The process here is similar to what was done for the fitting points, i.e., we first establish the validation points we want to test, and we then have to go back to the heavy model to calculate the targeted metric values for each point. The validation is then a straightforward task of comparing the value generated from the heavy model versus the value generated by the proxy function. Figure 5 illustrates the point nicely—if proxy values are plotted on the y-axis, and actual values on the x-axis, then a straight line diagonal at 45 degrees is reflective of a perfect fit. Plotting the values actually generated at each validation point then gives us a very simple at-a-glance view of how good the proxy model is—and is also something senior management can quickly understand. Another key test is to perform dynamic validations of the proxy function—how well does the function behave over time? This type of dynamic test isn't something that is necessarily always done rigorously when developing proxy functions, but it needs to be.

Figure 5
The Proof-in-the-Pudding: Validating Proxy Models

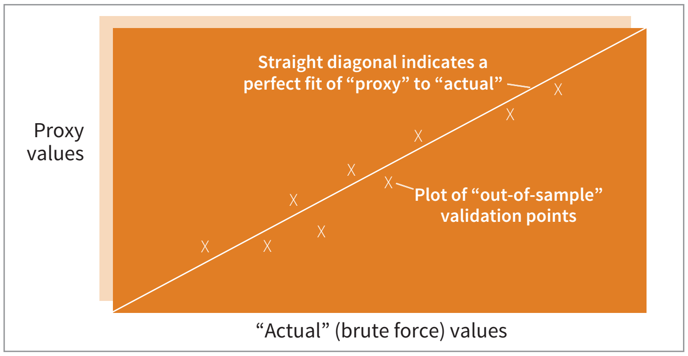

Another important development in recent years in the area of model data building has been the emergence of the cluster modeling technique. This has been applied by a number of U.S. insurers very successfully, and may be viewed as a straightforward extension of more traditional actuarial model point development.

Cluster modeling involves establishing the importance of individual data points and mapping less important points into the more important ones and continuing that reduction process until the desired number of model cells has been reached. The process thus begins at the individual policy record level, mapping policies of lesser importance into those policies of greater importance, and continuing that mapping process until the desired level of compression has been achieved. "Importance" in this instance is defined as size times distance, where size would be typically face amount for life insurance and account value for deferred annuities, and distance is determined relative to a policy's or cell's nearest neighbor with reference to whatever we deem to be the key metrics that characterize the policy, e.g., present value of future profits per unit, or reserves at projection date per unit.

There have been some papers written on cluster modeling and some presentations given at industry events—for an excellent introduction to the topic, a Milliman report by Avi Freedman and Craig Reynolds, "Cluster Analysis: A Spatial Approach to Actuarial Modeling" (2008),[3] is well worth reading.

Cluster modeling continues to hold much promise for the industry. Moreover, to the extent runtime issues will still exist around proxy modeling in having to calculate "actual" values for fitting and validation purposes, cluster modeling can be a very useful supplement to proxy modeling, and again we are aware of some practitioners in the industry considering application of the techniques in tandem—a kind of "hybrid" approach to model efficiency.

Recent Developments in Technology

Technology advances rapidly, and insurers who fail to keep pace with those developments face the combined risks of increased inefficiency and decreased competitiveness. There are many aspects of emerging technology that are exciting, but perhaps two that hold the most promise for insurance companies to make difficult and time-consuming calculations more manageable are cloud computing and GPU technology.

We are already seeing widespread use of cloud to help manage very large data and processing requirements in many aspects of the financial services industry, e.g., to conduct day-to-day banking. The insurance industry is far from exhausting the cloud potential, but some insurers are realizing immediate benefits from cloud with regard to computational throughput and large-scale data management. At the 2015 ERM Symposium, Jim Brackett of Milliman presented a very useful talk on some of the developments around cloud (and other) technology.[4]

Also at the 2015 symposium, on the same panel as Jim, we heard Iouri Karpov, of Prudential Financial, give a fascinating presentation on what is emerging around GPU and how the technology could potentially be used more widely.[5]

> In this age where the answer to almost any question can be found on the Internet with the click of a button, it seems inconceivable that insurance companies will not soon be doing things much quicker.

Virtually every modern video game console, computer and smart phone has a GPU, and developments in the area have been largely driven by the ever-increasing demand for improved high-definition standards in gaming. General-purpose GPU computing refers to other scientific and business applications. GPU computing works extremely well where there are numerically intensive and parallelizable calculations that need to be done. Clearly, this holds much promise for many insurance-based calculations, such as the massive parallelization that's done in an insurance valuation involving running a significant number of policies with similar payoff definition across multiple scenarios.

Iouri's presentation at the ERM Symposium, based on his own use of GPU technology to help with some of his work at Prudential, certainly created some buzz among those who attended the session. While application of GPUs in the industry currently remains in its infancy, and there are practical issues to address around productionizing a GPU process, it certainly seems to hold huge promise.

WHAT WE MIGHT EXPECT TO SEE IN THE FUTURE

Model efficiency is a field of practice that should continue to develop, as the growing appetite for usable and up-to-date analytics continues unabated.

In this article, we have discussed the emergence of proxy modeling and cluster modeling, and cloud and GPU technology, as powerful developments we can expect to see more widespread use of. But as ever in life, it is probably going to be something we are not even aware of today that ends up taking model efficiency up another level.

In this age where the answer to almost any question can be found on the Internet with the click of a button, it seems inconceivable that insurance companies will not soon be doing things much quicker. Providers of risk and actuarial platforms who make the investment now in the emerging methodologies and technologies stand to take a dominant position and will shape the way the industry does things for many years to come.

As for MEWG, there remains a core group of practitioners that have continued to share information and do some research, but MEWG under the AAA is today more of a “sleeping dog”; it remains “officially” in existence by name but has not performed any work for some time, and there are currently no members of the group, other than its chair, myself, continuing to be listed as a contact point. That said, there is a new subteam of the Society of Actuaries’ Modeling Section assisting the Section Council with model efficiency matters that plans to be proactive in the area of research and seminars. This new subteam is again being led by myself, and so far we have recruited Paula Hodges from Ameritas and Mike Beeson from Pacific Life to be on the team. The mandate for the team is somewhat informal, but in essence the objective is to assist the Modeling Section Council in its mission to support the basic and continuing educational, research, networking and other specialized needs of its members within the specific area of model efficiency. If you are interested in getting involved, we would be delighted to hear from you; please contact me at *anthony.dardis@milliman.com*.

Acknowledgments

The author wishes to thank the following for their help in the preparation of this article:

- Anna Berezovskaya, consulting actuary at Milliman
- Jim Brackett, technology lead at Milliman
- Craig Reynolds, consulting actuary at Milliman
- Phil Schechter, actuary at AXA

Tony Dardis is a consulting actuary with Milliman, located in Chicago. He can be reached at *anthony.dardis@milliman.com.*

ENDNOTES

[1] See Steven Craighead, “PBA Reserves and Capital Modeling Efficiency: Representative Scenarios and Predictive Modeling,” *The Financial Reporter*, 73 (June 2008).

[2] For a discussion of bias in estimates of a CTE value, see B. John Manistre and Geoffrey H. Hancock, “Variance of the CTE Estimator,” *North American Actuarial Journal*, vol. 9, no. 2 (2005): 129–156.

[3] *http://www.milliman.com/uploadedFiles/insight/research/life-rr/cluster-analysis-a-spatial-rr08-01-08.pdf.*

[4] *http://www.ermsymposium.org/2015/presentations/C-12-Brackett.pdf.*

[5] *http://www.ermsymposium.org/2015/presentations/C-12-Karpov.pdf.*