VOLUME 17 ISSUE 2

THE Actuary

SPRING 2020

Cut through the noise

Actuarial systems evolution

Focused on solutions

PERSEVERE WITH PASSION

Meet SOA member and advocate for diversity, Rolande S. Mbatchou, ASA, MAAA

SOCIETY OF ACTUARIES®

THE Actuary

VOLUME 17 ISSUE 2

SPRING 2020

FEATURES

CONTENTS

PHOTO: JOEL MAISONET

POSTMASTER: Send address changes to the SOA, c/o Communications Department, 475 N. Martingale Rd., Schaumburg, IL 60173-2226.

CONTENTS

DEPARTMENTS

SOA PRESIDENT
Andrew D. Rallis
FSA, MAAA
arallis@metlife.com

SOA STAFF CONTACTS
Anna Clarke
Senior Director, Marketing & Communications
aclarke@soa.org

Cheré LaRose
Director of Member & Candidate Communications
clarose@soa.org

Julia Anderson Bauer
Publications Manager
jandersonbauer@soa.org

Jacque Kirkwood
Magazine Staff Editor
jkirkwood@soa.org

Erin Pierce
Magazine Staff Designer
epierce@soa.org

The Actuary is published bimonthly (February/March, April/May, June/July, August/September, October/November, December/January) by the Society of Actuaries, 475 N. Martingale Rd., Suite 600, Schaumburg, IL 60173-2226. Periodicals postage paid at Schaumburg, IL, and additional mailing offices. USPS #022-627.

The Actuary is free to members of the Society of Actuaries. Send subscription requests to *customerservice@soa.org*.

CREATIVE SERVICES
Kathleen Hagan
Associate Director of Content

Enrique "Rick" Cruz
Senior Art Director

GLC
glcdelivers.com

CONTRIBUTING EDITORS
Matthew Cuculic, FSA, CERA
mcuculic@humana.com

Sara Goldberg, FSA, MAAA
sara.goldberg@genre.com

Kelly Hennigan, FSA, CFA
Kelly.Hennigan@venerableannuity.com

Jason Hiquet, FSA, CERA
jhiquet@deloitte.com

Olga Jacobs, FSA, MAAA
olga_jacobs@uhc.com

Tim Koenig, ASA, MAAA
timothy.koenig@pwc.com

Dave Snell, ASA, MAAA
Dave@ActuariesAndTechnology.com

Martin Snow, FSA, MAAA
martin@atidot.com

Ricardo Trachtman, FSA, MAAA
ricardo.trachtman@milliman.com

Nimmie Veerappen, FSA, FCIA
nveerappen@rgare.com

EDUCATION CONSULTANT
Kory Olsen, FSA, CERA, MAAA
Kory.Olsen@PacificLife.com

INTERNATIONAL CONSULTANT
Bill Sayre, FSA, MAAA
bill.sayre@milliman.com

ADVERTISING INFORMATION
Inquiries about advertising should be directed to:
M.J. Mrvica Associates Inc.
2 West Taunton Avenue
Berlin, NJ 08009
Phone: 856.768.9360
Fax: 856.753.0064
Email: dmather@mrvica.com

ISTOCK.COM/ABU

Fitting In

BY MARTIN SNOW

> ***As actuaries, we need to do a better job at maintaining the relevance and sustainability of our profession and work products.***

BILL GATES, CO-FOUNDER OF MICROSOFT, SAID, "The advance of technology is based on making it fit in so that you don't really even notice it—so it's part of everyday life." By this measure, the technology that actuaries use has not succeeded. We notice it all the time. For example, how often is a project late because the model took longer than anticipated? In this issue of *The Actuary*, we explore what can be done to make technology fit in so that it is not noticed. But first, let us start with some historical context.

The evolution of actuarial practice over the past 40 years has been revolutionary. Teams of actuaries now build, implement and maintain extremely sophisticated models using technology that does things we only could have dreamed of in the 1980s. We work on requirements that did not exist—and were totally unanticipated—as well as in newly created disciplines. Sometimes we try to do this on our own, but increasingly we are working with software engineers and other professionals. What problems do we face in addressing the resultant challenges, how can we summarize our accomplishments, and where do we go from here?

As models started to play a larger and larger role in our professional lives, there were those who predicted the shrinkage of our profession. After all, the computer could do everything we used to do. How wrong they were, as our profession continued to grow at a rapid clip. The ability to understand what needs to be done and analyze the results, together with the increasing needs and requirements, has required more and more actuaries. Today, some say artificial intelligence (AI) and machine learning (ML) will replace actuaries. I believe they will likewise be proven wrong as more actuaries than ever will be required to understand what can be done, assure this work is done effectively, and interpret the results. The machine will never do all of this on its own.

What then is the challenge? Well, as models have become more sophisticated and requirements more detailed, many actuaries have become caught up in the minutiae of doing their existing jobs correctly. Often this is necessary to produce an accurate, reliable and replicable result for the audience. But this has come at the expense of communicating our findings effectively, making our processes repeatable and producible, and being ahead of the curve on the next big thing that will truly add value that we should be working on. In other words, we need to do a better job at maintaining the relevance and sustainability of our profession and work products.

Oftentimes when we hear of these types of challenges, we are told about all of the soft skills we need to improve on—and there's no doubt that we, as a profession, need to improve our communication and collaboration skills. But in this issue of *The Actuary*, we show concrete steps—hard skills—we can take to make us more relevant and more sustainable in the face of the emerging challenges of the 2020s.

The article by Tom Peplow, MSc, stresses the need for actuaries to work closely with other professionals (e.g., software engineers). Different groups of professionals can all learn from each other and develop a much broader knowledge base. Actuaries have realized that combining our skills with those of software engineers results in models that are more dependable, more reliable, better controlled and more flexible—models that are more likely to stand the test of time. Rich Lauria, FSA, CFA, MAAA, shares the history of ERM and how it has affected and been impacted by model development—for example, how regulations have evolved and how models have kept pace. Andy Smith, FSA, MAAA, writes about the numerous challenges with building models today and some potential solutions. These three articles position us well to respond to the challenges of model development, model maintenance, version control and model updates during this time of rapidly changing technology. Perhaps most important, an actuary gets to see how other actuaries have successfully responded to many of the challenges being faced today.

Then we come to today's new frontier: ML and AI. The importance of and need for a multidisciplinary team to tackle the challenges in this space cannot be emphasized enough. With a properly assembled and managed team, the whole is much greater than the sum of its parts. If you want the diamond team, think about how diamonds are formed: You put a lump of graphite under enormous pressure for a long time, and then you have diamonds. Well, in ML and AL, you need domain expertise, technical proficiency, data science expertise, strong interpretive skills and strong communications skills. Taken together, you have the diamond team.

But actuaries may wonder how they can put all of these skills to practical use. Dave Czernicki, FSA, MAAA, and his coauthors share many of the use cases in which actuaries have an interest today. These use cases can all be practically implemented with the right team and resources. Yet the advantages of AI and ML go even further. In the article that Adam Haber and I wrote, we show an innovative way to get a much better handle on post-level term mortality and how to interactively model the shock lapse, mortality and post-level term premium jump.

The techniques described in this issue of *The Actuary* can bring big strategic advantages to those who successfully implement them. We wish you the best as you read these articles and assure you that the authors would be pleased to answer any questions you may have. We hope you agree that successful implementation of the ideas presented will make the technology fit in, so you don't notice it. ■

ABOUT THE WRITER

MARTIN SNOW, FSA, MAAA, is vice president, chief delivery officer and chief actuary at Atidot, a provider of predictive analytics solutions to the life insurance industry. He is currently a contributing editor for *The Actuary* magazine. He can be reached at ***martin@atidot.com***.

COVID-19—SOA Updates Page

The Society of Actuaries (SOA) continues to work to keep you—our members, candidates, staff and volunteers—informed about the latest developments and changes caused by the COVID-19 pandemic that swept through the world early this year. When we received the first wave of changes to the exam schedules, followed by the policies that changed the way we could conduct our in-person meetings, we knew we had to create a way to mitigate the impact of these changes on our stakeholders to the best of our ability. So, we created the COVID-19 SOA Updates Page as a place where you can easily find all the information you need when you need it, in one place.

This page contains updates on exams and new schedules; professional development event cancellations and new virtual sessions; and a variety of podcasts and research reports from the SOA Research department and the Health Section. These reports include the SOA Research Brief: Impact of COVID-19 that is updated regularly and accompanied by a podcast that discusses the brief.

Most recently, the Health Section released the COVID-19: Updates on Impact podcast in which Jackie Lee, FSA, MAAA, and Rebecca Owen, FSA, FCA, MAAA, discussed the impact of COVID-19—from quarantine and shortages to homeschooling, small businesses, group therapy and isolation.

Find the latest updates at *SOA.org/covid-19*.

ISTOCK.COM/LAYRITTEN

THE ACTUARY—WEB EXCLUSIVE

Actuaries are reacting personally and professionally to the spread of the novel coronavirus and the respiratory disease COVID-19 across the globe. How should actuaries in particular think and spend their time in regard to the pandemic? Read "Actuaries in the Time of Coronavirus" by Robert Eaton, FSA, MAAA, at ***bit.ly/Actuary-COVID-19***.

Fifty States, Fifty Stories

Should we "repair" or "replace" the Affordable Care Act (ACA)? To help us answer this question, the Health Section Council of the SOA launched the ACA@10 Strategic Initiative last year. This initiative consisted of a data-driven research project, entitled "Fifty States, Fifty Stories: A Decade of Health Care Reform Under the ACA." This research, which was authored by Paul Houchens, FSA, MAAA; Lindsy Kotecki, FSA, MAAA; and Hans Leida, Ph.D., FSA, MAAA, looks at measures of success for the ACA from a number of different perspectives. Read the report for data-driven observations. It can be accessed at *bit.ly/50-States-Stories*.

The ACA@10 Web-Exclusive Series

March 23, 2010, the day the ACA was signed into law, was a day of great promise for everyone without health insurance—it promised access to affordable health care. Ten years later, the question is: Was that promise kept? Certainly, it is an achievement that there are now 20 million more people insured than there were in 2010. Yet, there are still 30 million Americans who are uninsured, and many more who are struggling with paying premiums and the cost-share on their existing coverage. As the 2020 U.S. election draws near and the electorate decides where we want to go from here, we need to be able to understand the ACA's real-world application more fully. Read the articles in this web-exclusive series for details surrounding the ACA: *bit.ly/ACA-10*.

Artificial intelligence (AI) and machine learning (ML)—collectively referred to as AIML—have been hot topics lately. And for good reason: These technologies are bringing significant advances that are reshaping the world as we know it. From driverless cars to detecting cancer, AIML already has presented tremendous breakthrough opportunities to myriad industries, and the adoption and expansion of business applications are only expected to accelerate. In the financial services sector, FinTech firms have developed concepts leveraging AIML such as chatbots, automated document processing and deep hedging (hedging strategy informed by ML algorithms). Banking and insurance companies increasingly are injecting funds into these efforts.

Cut Through

BY DAVE CZERNICKI, PETER CARLSON, JEAN-PHILIPPE LAROCHELLE AND JONATHAN DeGANGE

the Noise

Machine learning and artificial intelligence are a leap forward for life and annuity actuarial modeling

Actuaries are taking notice. The Society of Actuaries (SOA) made predictive analytics a component of its strategy,[1] extended the candidate's curriculum and launched a related Actuarial Innovation & Technology Strategic Research[2] initiative. Andrew D. Rallis, FSA, MAAA, president of the SOA, established AIML as an area of focus in his presidential luncheon speech at the 2019 SOA Annual Meeting & Exhibit.[3]

AIML likely will disrupt our work like other technologies have, with an ever-increasing pace of change. Just as we can't fathom what our work would have looked like without technologies we now take for granted—such as distributed processing, first principles models or even basic data storage capabilities—AIML soon will be a common component of the actuarial toolkit.

With all of that said, it can be hard to cut through the noise surrounding AIML and get a clear understanding of the concrete applications this emerging technology can have for the life and annuity insurance industry. While keeping up with the latest disruptions brought by AIML is important, given its transformative impact on the insurance industry, our objective with this article is to share some concrete and present-day applications in actuarial modeling from which actuaries can harness the power of AIML.

Overview of AIML Techniques

AI can be thought of as the combination of various branches—ML, natural language processing, speech and robotics—which together can simulate intelligent behavior. Using a set of techniques that often overlaps with predictive analytics, ML algorithms take inputs provided by AI and return outputs simulating the AI "thinking" behavior. ML techniques have risen in prominence and effectiveness in today's world because of the depth and breadth of readily available data—a key ingredient in training and informing ML algorithms.

ISTOCK.COM/CREATIVAIMAGES

Various types of algorithms commonly employed in AIML are relevant for actuarial modeling. This article explores practical examples that leverage specific classes, such as artificial neural networks (ANNs), generalized linear models (GLMs) and k-means clustering, to provide a sense of how those algorithms work and can be put into practice. Readers should note that many other types of AIML algorithms exist (and still others have yet to be invented). Some examples include decision trees, random forests, gradient boosting machines (GBMs) and support vector machines. Actuaries interested in exploring these techniques can benefit from an extensive AIML community where documentation, research and open-source libraries are available.

While most AIML algorithms can be processed without extensive computing power, the training aspect can be demanding. This is especially true when the application requires the use of deep learning, which is a subset of AIML that uses neural networks with multiple layers that can reach higher levels of performance given enough data. Often, actuarial applications of AIML require the use of deep learning to attain the level of precision sought by actuaries.

Companies that want to successfully use this technology need to develop expertise and capabilities to effectively identify and design potential algorithms for new use cases, establish a conducive environment, and identify risks and controls. Further, while it is common for AIML algorithms to lack transparency, various methodologies are available to explain AIML.[4]

Augmenting Actuarial Models With ML

A key opportunity is embedding AIML in our models. While this can appear counterintuitive, as some may think of AIML as an *alternative* methodology, AIML can significantly boost the speed and accuracy of our models by replacing specific components. The premise is AIML

Figure 1 Simplified Illustration of the Calculation Components of a Variable Annuity Projection Model

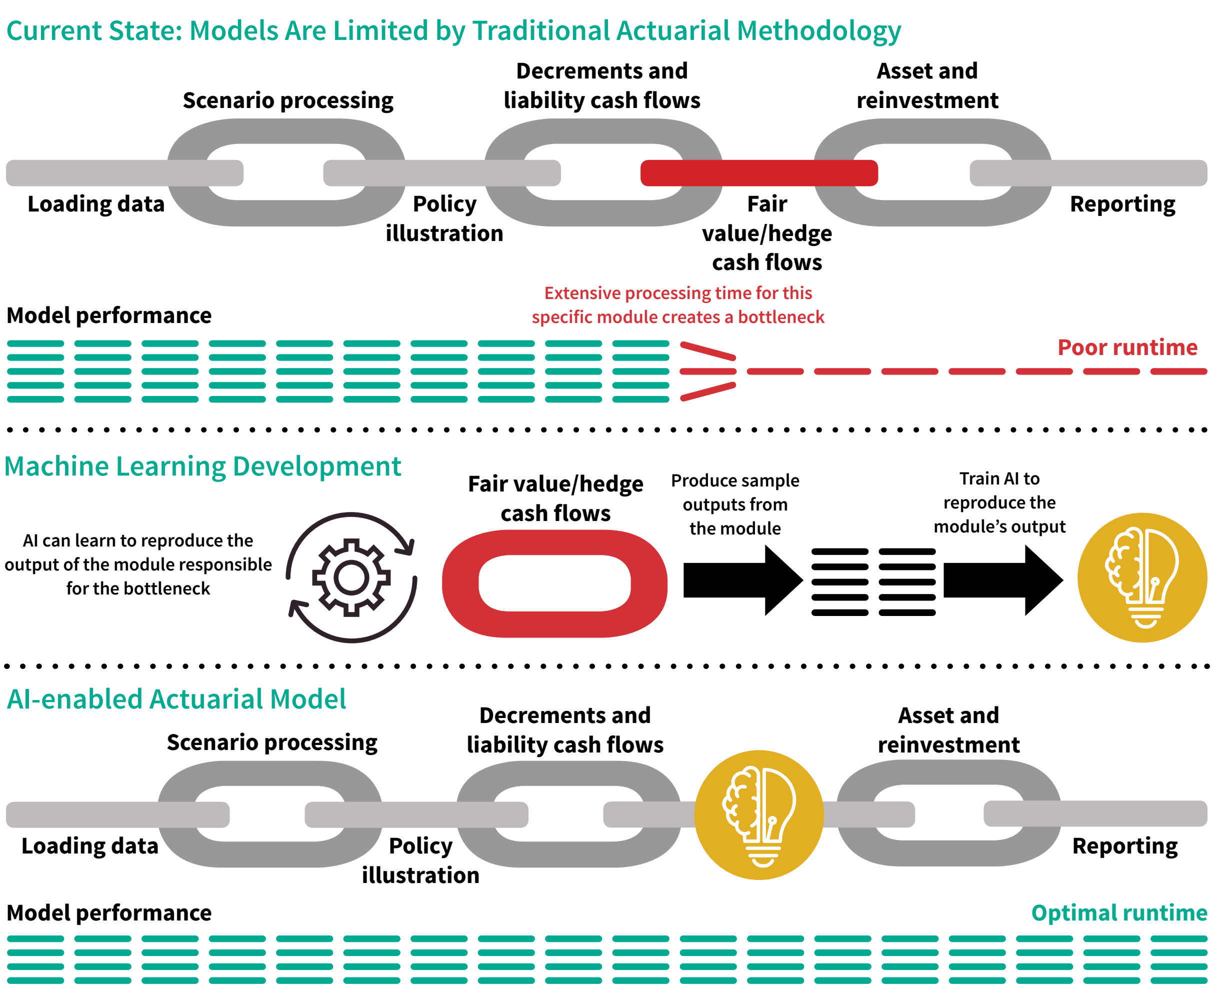

can reproduce a specific actuarial calculation faster than an actuarial model would when processing it under first principles. See Figure 1 for an example.

Before we can explain how this works, first think of an actuarial model as being a collection of various calculation components, or modules. Here a module could be an account value roll-forward or a nested calculation. A significant amount of calculations typically takes place when processing an actuarial model. Actuaries are used to facing long model runtimes and needing to resort to approximations.

Actuaries interested in implementing AIML in actuarial models should first identify what module(s) is/are responsible for the undesirable runtime. Once the module is identified, the AIML algorithm can be trained to accurately predict the module's output. A key advantage of this application is the control of methodology and data volume. Further, developing the AIML provides an opportunity for greater understanding of the specific calculations targeted.

As a practical example, we can train a neural network to proxy variable annuity (VA) fair value (FV) and Greeks for hedge cash flow projections. This neural network can take inputs such as the VA FV module within the actuarial model (e.g., policyholder fund allocation, market variables) and other relevant measures (e.g., moneyness). Using common training techniques, the neural network will draw relationships among those variables as it learns from the data provided.

A conceptual illustration of such a neural network can be seen in Figure 2. This illustrative network first takes the inputs provided and infers information about the equity risk, interest rate risk, time value of money and contract-specific contributions in the first layer of artificial neurons—mathematical functions conceived as a model of biological neurons. Each neuron first weighs the information from the inputs and applies a predefined activation function. Then, this network processes the information inferred from the first layer and transforms it again to provide an estimate of the FV and Greeks.

With VA statutory reform and the changes to generally accepted accounting principles (GAAP) in the United States, many companies are looking for capabilities to project FV and their hedge strategy to gain capital credits and a better understanding of their GAAP balance sheet. This has been a challenging task in our experience, and

Figure 2 Illustrative Neural Network

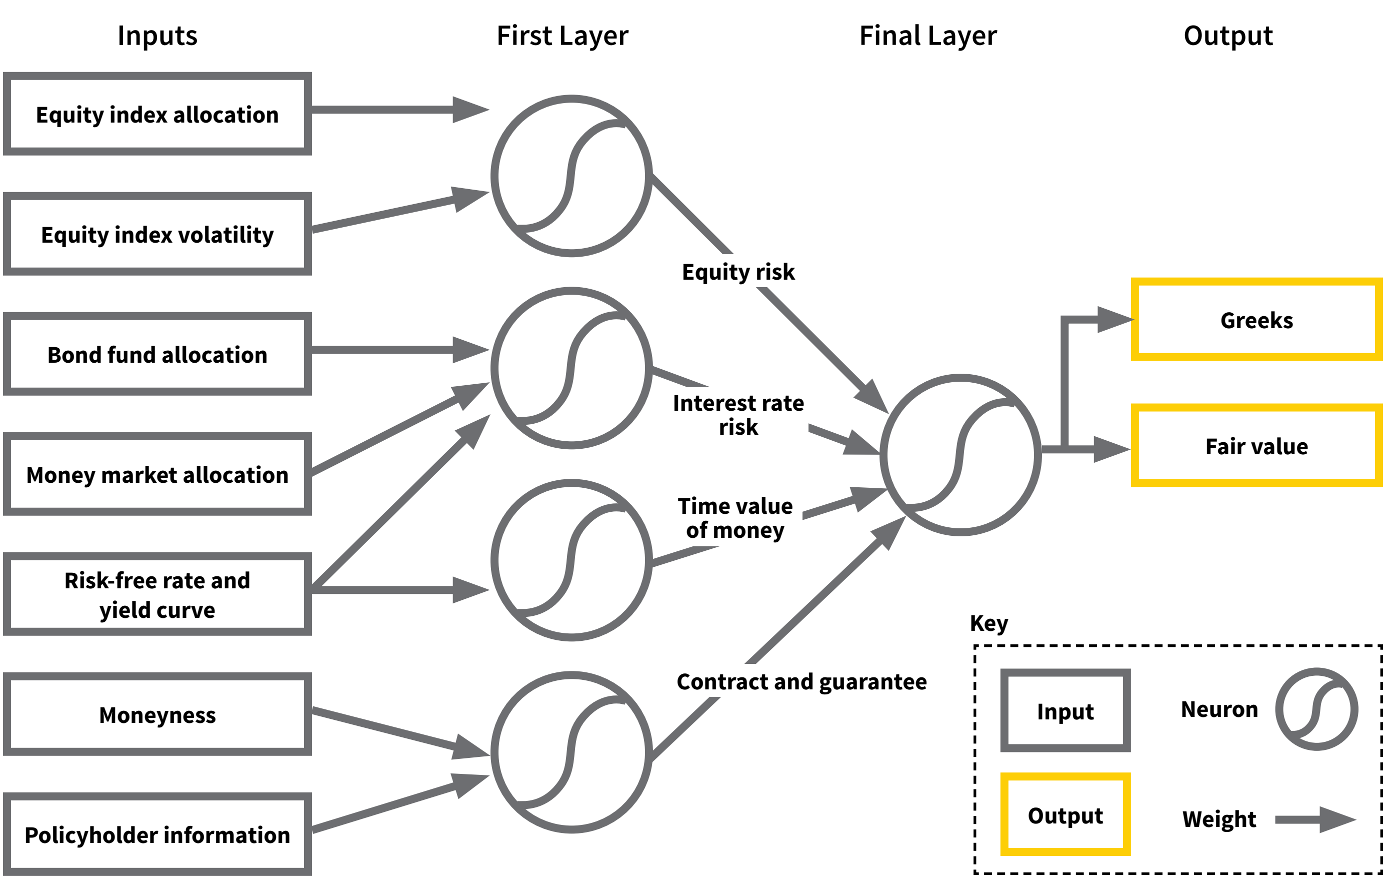

we have found that neural networks and other algorithms such as GBMs are able to provide reliable proxies.

Potential use cases for this application include many areas where our actuarial models must perform nested calculations. This includes not only FV and Greeks, but also projecting reserves, setting the parameters of exotic crediting strategies and many others.

Policyholder Behavior and Other Actuarial Assumptions

Actuarial assumptions are a critical component of our actuarial results and forecasts: They establish our view about how key insurance risks such as mortality, morbidity and policyholder behavior will unfold in the future. Generally, actuaries start by establishing a hypothesis as to the drivers of the risk, develop assumptions by leveraging data and other sources, and monitor actuals to expected. While traditional actuarial methods used in developing those assumptions certainly remain valid, actuaries responsible for developing assumptions may gain further insights by applying techniques rooted in predictive analytics and ML.

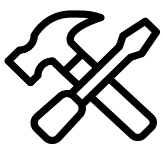

AIML soon will be a common component of the actuarial toolkit.

For example, if we take the behavior of variable and fixed indexed annuity policyholders, actuaries typically start by looking at qualitatively intuitive dimensions such as policy year duration and rider moneyness, among other items. By using predictive analytics methodologies, actuaries can quickly evaluate the relevance of a great number of variables and identify patterns and relationships that might have otherwise gone unnoticed. This is in addition to the natural objective of reducing the actuals to expected across the identified policyholder attribute dimensions.

Actuaries can use ML algorithms and other methodologies such as principal component analysis to learn how policyholders behaved in the past (based on historical data) and infer which variables best predict how policyholders will utilize their contract options. While actuaries would have a good sense for the key drivers, this process can provide additional insights. In some cases, actuaries could even use the algorithm as an actuarial assumption.

Actuaries using these techniques should apply judgment when formalizing relationships and differentiate between correlation and true cause-and-effect. Actuaries who decide to use AIML algorithms should be careful to select an approach that is fully interpretable and avoid the temptation to strictly minimize actuals to expected. In particular, we have found GLMs to be a popular option, given they provide greater transparency relative to other methods.

Compressing Model Points

It is common for actuaries to seek to reduce the number of model points (e.g., in-force policy records or new business cells) to be processed by an actuarial model for certain use cases because of runtime limitations. The actuary may be seeking to process a large scenario set or project complex nested stochastic reserves. Attempting to process all model points in those situations could very well be a time-consuming, or even impossible, task.

Currently, a conventional approach consists of selecting representative policies by creating predefined segments across the data, based on the actuary's judgment. For instance, the actuary evaluates which features of the model points are most important and then defines subsections across the range of potential values that these representative characteristics provide. This could cover five-year increments in attained age, gender, moneyness and so on.

Alternatively, actuaries may explore the use of clustering techniques such as k-means clustering to drive the selection of representative policies. Actuaries who use these techniques can base their selection of representative policies off of the sets of characteristics that effectively group the in force in policies that behave in a similar fashion. Using these techniques, actuaries likely will find they can significantly reduce the number of clusters with minimum loss of precision.

Other Applications

There are various other applications where the use of predictive analytics or ML can be considered for actuaries in the life and annuity insurance industry.

Process and Controls

ML can be used to implement new checks and controls. This can be particularly useful for production processes such as quarterly actuarial valuation, allowing actuaries to identify errors earlier or that would otherwise have been missed. For instance, ML could be trained to identify suspicious policies within the in force before starting actuarial calculations. Similarly, actuaries could train ML programs to identify incorrect results using previous valuation quarters.

Actuarial System Conversions

As most modeling actuaries know, migrating actuarial applications from a legacy system to a new platform can be a daunting task. Often, most time and efforts are spent identifying and reconciling records between the two platforms. Techniques based on data science, such as principal component analysis, can help actuaries identify differences faster and accelerate the conversion.

Replacing Actuarial Model Runs

In a similar fashion to the actuarial model augmentation previously described, actuaries also could train ML algorithms to predict specific results from actuarial model runs. This could include specific metrics, reserves or setting rates for insurance products, allowing actuaries to get insights faster for any application where frequent and/or manual model runs are needed.

Financial Planning & Analysis (FP&A) and Asset/Liability Management (ALM)

Lastly, ML can be used to accelerate and streamline the process to generate forecasts for management as well as enhance the ALM and/or hedging functions.

> *Companies that want to successfully use this technology need to develop expertise and capabilities to effectively identify and design potential algorithms for new use cases, establish a conducive environment, and identify risks and controls.*

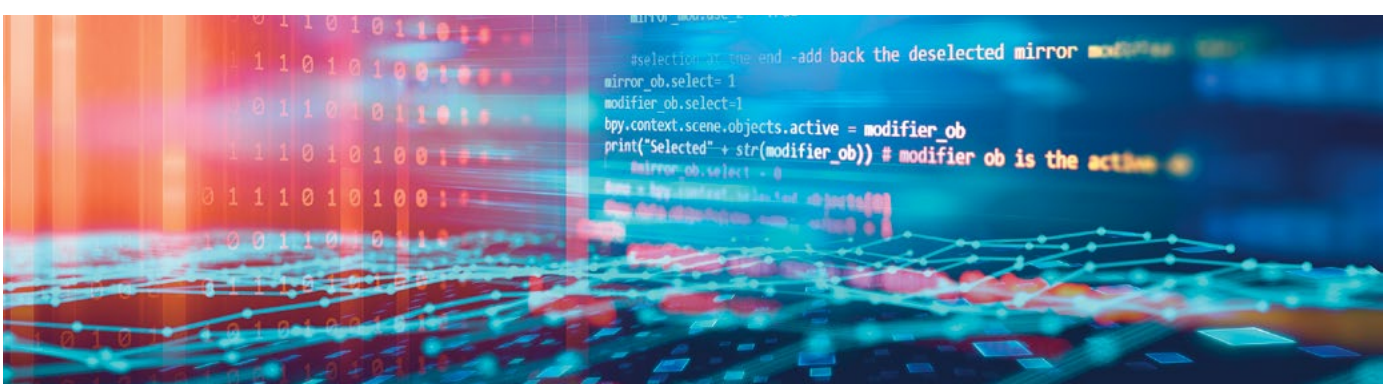

ISTOCK.COM/MONSITJ

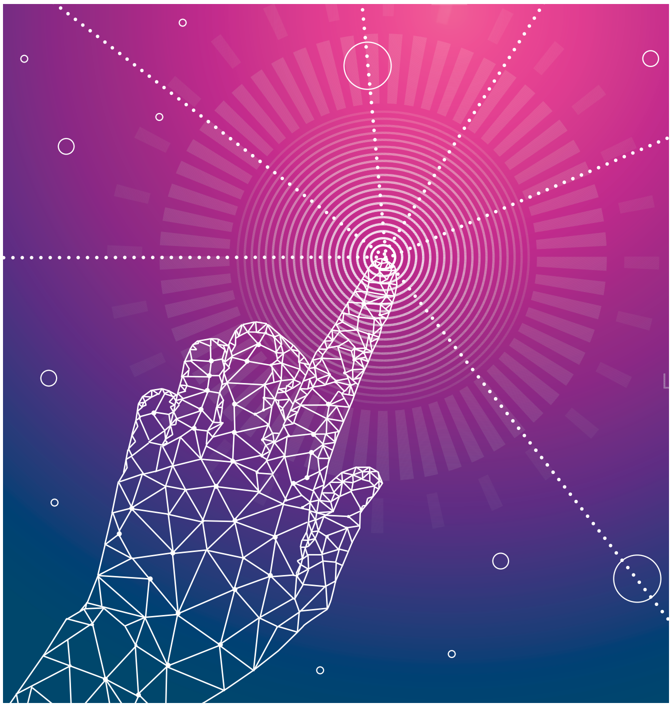

SHUTTERSTOCK/VLADIMIR VIHREV

Limitations and Key Risks Involved

While we believe AIML and predictive analytics can bring significant value to the actuarial function, there are important implications that adopters should keep in mind when using this technology. Some examples are:

- **Interpretability.** A critical aspect adopters need to keep in mind is the lack of interpretability associated with most ML models. While neural networks, GBMs and random forests—when trained properly—can reproduce a data set with a high degree of fidelity, users cannot easily trace the calculations leading to the prediction. While various tools and methodologies can be used to interpret these models, they do not provide as much transparency as actuaries are used to.
- **Overfitting.** Another critical aspect is the risk of overfitting. If not monitored, models can learn irrelevant details or reproduce noise in the data. This is a key risk for most applications of ML outside of actuarial work and for the actuarial assumption use case highlighted in this article. Various methodologies exist to avoid or minimize overfitting.
- **Labeled data.** Lastly, ML requires large volumes of data, and actuaries may not have enough data on hand to adequately train the applicable algorithms. For applications leveraging policyholder or other "real" data, actuaries may seek to supplement their company database with external sources. For the applications where we aim to reproduce specific actuarial calculations, users may run into challenges to produce a sufficient data set from the actuarial models, or they may need to develop new automation procedures.

Putting It All Together

There is a wide range of potential applications for the modeling actuary looking to leverage AIML and predictive analytics. These techniques can provide efficient solutions to complex, nonlinear functions and classification problems common in modeling-related applications.

From accelerating model calculations, developing actuarial assumptions and compressing model points, to providing enhanced controls, accelerating actuarial system conversions and substituting repetitive model runs, AIML has the potential to revolutionize our actuarial capabilities. ■

References

1 Society of Actuaries. Strategic Plan. *Society of Actuaries, https://www.soa.org/programs/strategic-planning* (accessed January 26, 2020).

2 Society of Actuaries. Actuarial Innovation & Technology Strategic Research. *Society of Actuaries*, *http://www.soa.org/programs/act-innov-tech* (accessed January 26, 2020).

3 Society of Actuaries. Presidential Luncheon Speech, SOA 2019 SOA Annual Meeting & Exhibit. *Society of Actuaries*, *https://www.soa.org/about/governance/board-and-leadership/2019-rallis-annual-meeting-speech* (accessed January 26, 2020).

4 Murdoch, W. James, Chandan Singh, Karl Kumbier, Reza Abbasi-Asl, and Bin Yu. 2019. Interpretable Machine Learning: Definitions, Methods and Applications. *Proceedings of the National Academy of Sciences of the United States of America* 116, no. 44:22,071–22,080.

ABOUT THE WRITERS

DAVE CZERNICKI, FSA, MAAA, is a principal at Ernst & Young LLP, where he leads the actuarial modeling consulting services. He can be reached at ***Dave.Czernicki@ey.com***.

PETER CARLSON, FSA, MAAA, is an executive director at Ernst & Young LLP, where he leads the firm's insurance value management practice. He can be reached at ***Peter.Carlson1@ey.com***.

JEAN-PHILIPPE LAROCHELLE, FSA, CERA, MAAA, is a senior manager at Ernst & Young LLP. He leads the firm's life and annuity actuarial analytics practice. He can be reached at ***JeanPhilippe.Larochelle@ey.com***.

JONATHAN DeGANGE is a senior manager at Ernst & Young LLP. He specializes in machine learning, natural language processing and computer vision. He can be reached at ***Jonathan.DeGange@ey.com***.

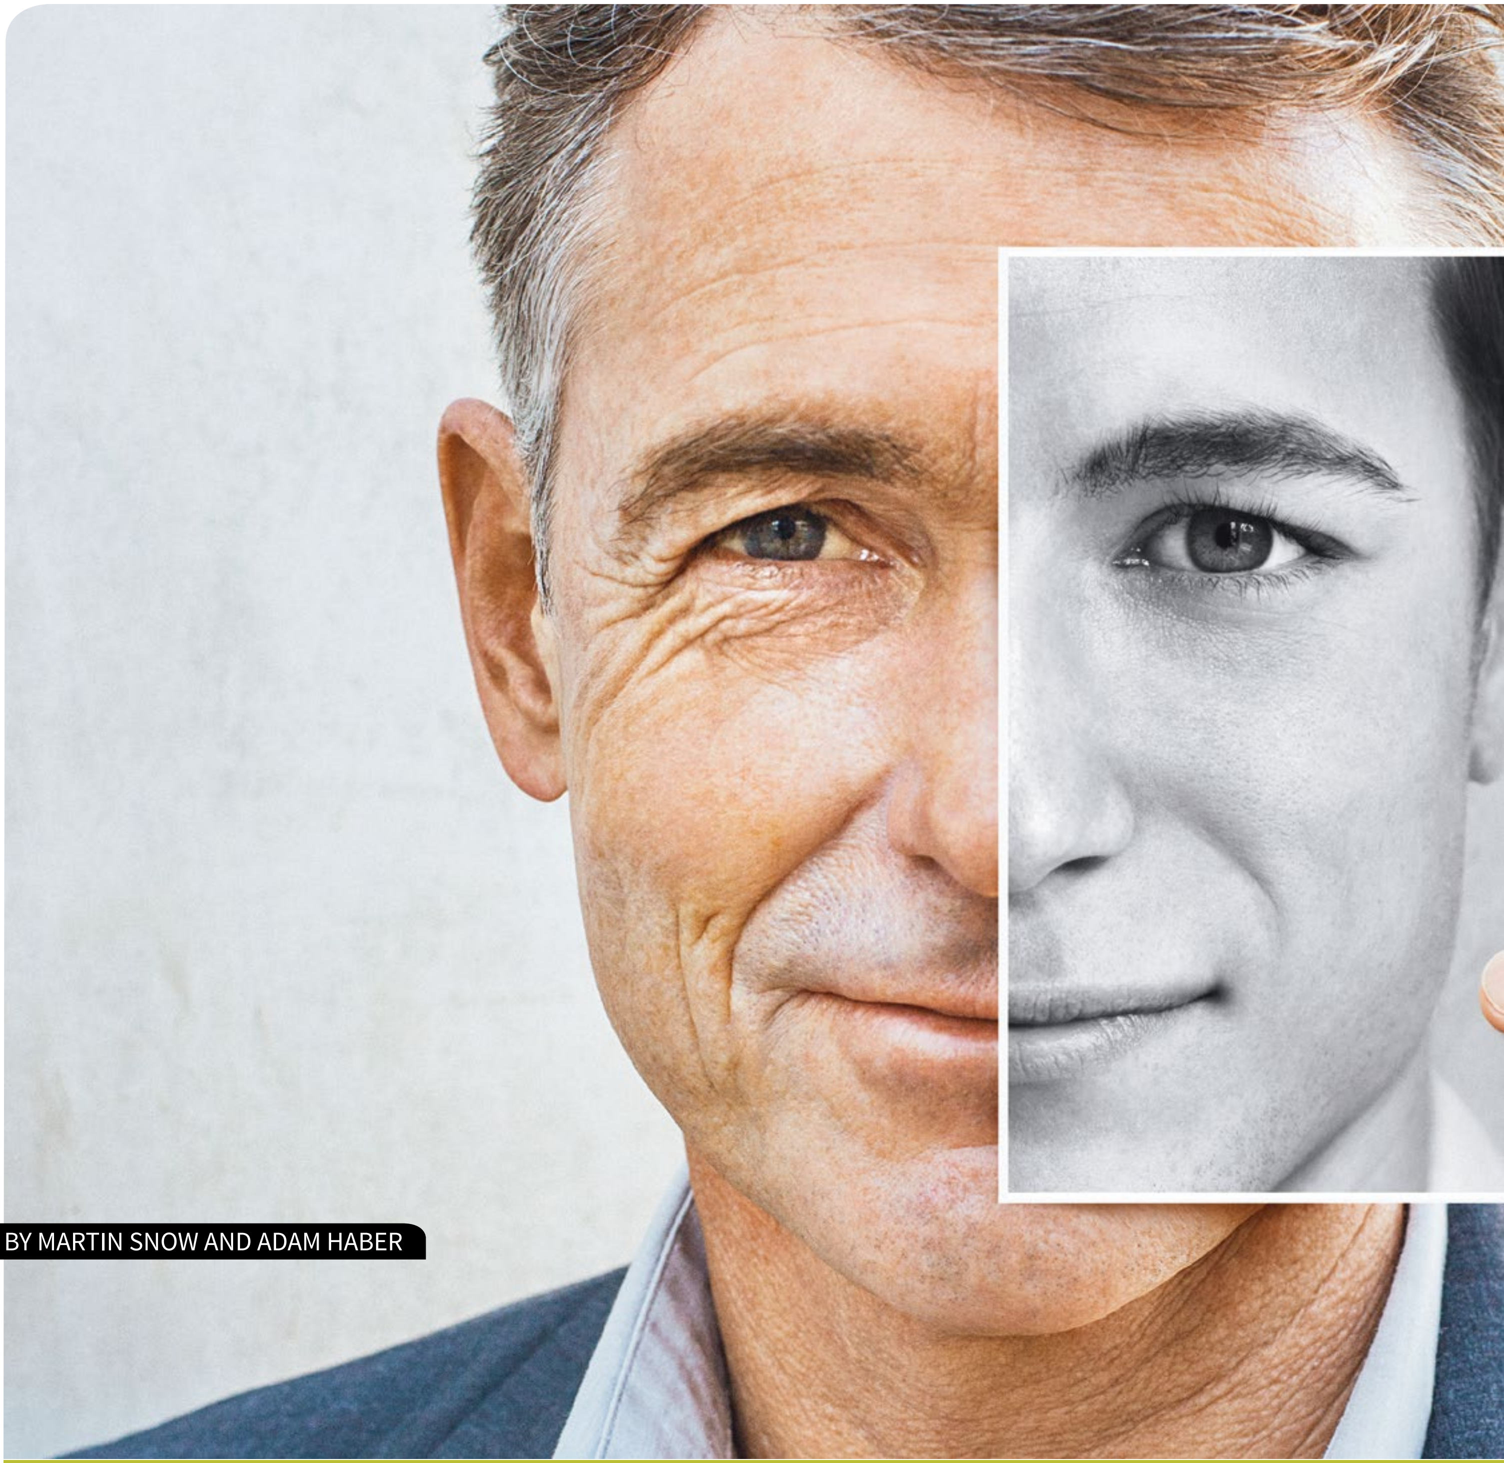

BY MARTIN SNOW AND ADAM HABER

Better With Age

GETTY IMAGES/DIMITRI OTIS

Predicting mortality for post-level term insurance

Actuaries have a long and storied history of providing the joint mathematical and business foundation for the insurance industry. Yet, advanced predictive analytics techniques with machine learning (ML) and artificial intelligence (AI) have not made it into the standard toolkit of the typical actuary. Insurers and actuaries could reap major strategic benefits if they were to significantly increase their use of these advanced predictive techniques. In this article, we focus on mortality and lapse studies as one example.

Post-level term (PLT) insurance presents a unique set of challenges when it comes to predicting mortality and lapse experience. After a set period of, say, 10 or 20 years when the policyowner paid level premiums, the premium will rise annually. Customers will be highly motivated to identify all of their other options. Healthier individuals will have good alternatives and lapse their policies; the less healthy ones will remain. The greater the premium increase, the greater this effect will be—resulting in the classic mortality spiral.

How can we get a good quantification of the interrelationship between premium increases and lapse and mortality experience? By building a predictive analytics model—more advanced than those previously developed[1,2]—to set lapse and mortality assumptions, and price and value PLT insurance. Our model will statistically integrate heterogeneous customer cohorts,[3] improve credibility in cohorts with sparse claims data, and provide a more complete understanding of the impact of premium changes on mortality rates. We can only imagine the additional improvements to insurer pricing and financial reporting that could be achieved with broader applicability of these techniques beyond PLT.

Our PLT Model

Our PLT model comprises three advanced predictive methods:

1 | **An innovative application of a statistical multivariate framework to model PLT lapse and mortality.** This multivariate model reflects the causal structure (and almost immediate impact) of PLT lapsation and premium changes on mortality (PLT causal structure[4]) and provides better guidance for setting PLT premiums. Taking the causal structure into consideration is especially important when answering predictive “what if” questions (e.g., what happens to mortality if we change premiums by X percent).[5]

Consistent with our plan to model the lapse rate as a major driver of the dependence of mortality rates on premium level, we make assumptions in our model about the underlying data-generating processes:

- Whether a policyholder lapses at the end of the level term period is a stochastic function of various characteristics such as age, gender, risk class, face amount and the change in premium.
- This function may include complex dependencies among variables. For example, the effect of different face amounts on lapsation may vary by age, gender and so on.
- The differences in both base and shock lapse among cohorts cause perceptible differences in mortality levels.

❷ | **The statistical technique of "partial pooling" to increase the credibility of sparsely populated cohorts.** This is especially important when the volume of available data (especially mortality data) differs substantially by cohort, leading to differences in credibility—including cohorts with very limited credibility.

Partial pooling is a principled middle ground between complete pooling, which fits a single model for the entire population and ignores variations, and no pooling, which fits a single model for each cohort and ignores similarities shared among cohorts. Partial pooling is also known as hierarchical partial pooling.

Partial pooling enables us to share information (borrowing strength) among cohorts, regularize[6] our model and account for different cohort sizes without incorporating ad hoc solutions. The data for each observed cohort informs and adds credibility to the probability estimates for all of the other cohorts. The extreme estimates are driven toward the population mean ("shrinkage" in Bayesian statistics) with significant lessening of variability that may have been created by noise in the data. This phenomenon is closely related to the concept of bias-variance trade-off,[7] in which the tightness of fit to the observed data is reduced, so the derived estimates serve as better predictors. Partial pooling leaves us with better estimates, reduced variability and improved credibility.

Partial pooling smooths mortality estimates, which by itself is not new in actuarial science—different graduation techniques have been developed and implemented over the years. The distinct advantage of partial pooling is that it achieves the same goal by explicitly sharing information among cohorts in a principled way (guided by domain knowledge and analysis of the data), and it can improve credibility in sparsely populated cohorts.

❸ | **The integrative statistical approach of Bayesian inference[8,9] to quantify differences in experience among cohorts with different exposure levels.** The generative nature[10] of Bayesian modeling enables the incorporation of expert knowledge into the models in the form of model structure and informed priors.[11,12] Bayesian models produce crucial uncertainty estimates (unlike the point estimates supplied by more traditional maximum likelihood approaches) needed for informed decision-making—especially with sparse mortality data. We use Bayesian multivariate modeling of lapse and mortality, but we do not include a numerical comparison of the Bayesian and non-Bayesian approaches in this article due to space considerations.

ISTOCK.COM/CECILIE ARCURS

There are two key elements of our **mortality-lapse** model. The first is a nonlinear regression lapse model inspired by previous Society of Actuaries (SOA) studies.[13,14] We added partial pooling of parameters across cohorts to increase accuracy, credibility and predictability. We changed the link function of the model from log to logit to ensure per-cohort lapsation is bounded by the exposure (previously it was possible for the model to predict more lapses than exposures, i.e., an actual-to-expected ratio > 1).

The second key element of our model is that it is a Bayesian version of the Dukes MacDonald (DM) **mortality model**.[15,16] In this version, we model the effectiveness parameter as a nonlinear function of the cohort characteristics (e.g., age, risk class, gender, etc.), use priors that reflect actuarial knowledge regarding plausible parameter values of G (e.g., a reasonable prior might put more weight on values of G closer to 1 than 0),[17] and infer the posterior distribution of G from the data (the distributions over model parameters after conditioning on the data). We use the nonlinear regression lapse model previously described to estimate a distribution of lapse rates by cohort. Mortality is estimated by integrating over two variables: the joint distribution of base/shock lapse rates and the effectiveness parameter, thereby completing the mortality-lapse model.

Our Model in Action

To implement the model, parameters for both the lapse and mortality models were estimated using Stan, a state-of-the-art platform for statistical modeling and high-performance statistical computation.[18] We validated the estimates Stan provided with both Bayesian model comparison methods, such as leave-one-out (LOO) and Watanabe–Akaike information criterion (WAIC),[19] and actual-to-expected (A/E) ratios.

The SOA data[20] we used for our modeling, consisting of 8,066 different customer cohorts, is summarized in Figure 1.

To quantify and validate the impact of the new Bayesian tools presented, we conducted an analysis. First, for the multivariate modeling of lapse and mortality, we examined three variants of DM mortality estimates:

1. Assume **fixed** base lapse rates before the PLT period, fixed total lapse rates at the end of the level term period, and fixed effectiveness parameters. Optimal values for base and total lapse rates and the effectiveness parameter were found by using a standard gradient descent optimization algorithm. The lapse and effectiveness parameters do not vary by cohort though the select and point-in-scale mortality do vary by cohort.
2. **Empirically** assess from the data both the base and total lapse rates by cohort. The effectiveness parameter was fixed. It was optimized using grid search.[21]
3. Use a partially pooled **model** to estimate both base and total lapse rates that vary by cohort. The distribution of the effectiveness parameter was inferred from the data itself using NUTS,[22] an adaptive extension of the Hamiltonian Monte Carlo Markov Chain algorithm.[23]

In each of these variants, expected mortality is computed based on the five input parameters to DM: effectiveness, base lapsation, shock lapsation, select mortality and point-in-scale mortality. The select and point-in-scale mortality used in the computation of expected mortality were selected from standard tables. We compared the actual deaths for each method in each cohort to the expected, and we then computed a weighted error as the mean absolute deviation of the predicted A/E ratio from an A/E ratio of 1, weighted by exposure. Figure 2 on page 22 shows the results.[24]

A model such as this can be continually improved. For example, we know mortality is often a bit higher for lower socioeconomic classes. Building in this knowledge may result in an A/E ratio closer to 1. Similarly, upper-income policyholders may have the ability to anti-select, which also could be built into the next model iteration. The Bayesian framework used is especially well-suited to the incorporation of this type of expert knowledge.

For partial pooling when measuring mortality rates, we fit a nonlinear regression model to publicly available mortality data[25] with and without partial pooling of model parameters and held all else (e.g., the data and the characteristics being analyzed) constant. We compared the partially pooled model to both regularized and nonregularized nonlinear regression models using R's glmnet package.

We ran the models with different characteristic subsets to validate that our results are not characteristic-dependent. Almost always, the models without partial pooling of parameters yielded implausible estimates for cohorts with

Figure 1 Experience Used in the Model

	Base Lapse	Shock Lapse	Mortality
Incidence of event	150,000	250,000	610
Exposure	2,100,000	370,000	80,000

Source: Society of Actuaries. Lapse and Mortality Experience of Post-Level Premium Period Term Plans. *Society of Actuaries*, 2009, *https://www.soa.org/resources/research-reports/2010/research-shock-lapse-report/* (accessed January 27, 2020).

Figure 2 Mean Absolute Deviation of Actual/Expected Ratios

		Fixed	Empirical	Model	Exposure
Gender	Female	1.40	1.46	1.06	32,125
	Male	1.93	1.42	1.01	47,354
	Weighted error	0.71	0.44	0.03	
Premium Mode	Annual	2.44	1.15	0.86	10,490
	Semiannual	1.64	1.32	0.86	7,323
	Quarterly	1.98	1.40	0.98	18,000
	Monthly	1.50	1.66	1.18	42,792
	Weighted error	0.75	0.50	0.13	
Risk Class	Super preferred nonsmoker	2.99	1.55	1.27	6,674
	Preferred best nonsmoker	0.89	0.94	0.67	13,511
	Preferred residual nonsmoker	1.74	1.18	0.94	11,385
	Nonpreferred nonsmoker	1.71	1.34	0.89	18,972
	Preferred smoker	1.65	1.29	0.89	2,950
	Nonpreferred smoker	2.84	2.80	2.00	1,688
	Undifferentiated nonsmoker	1.83	1.51	1.11	18,457
	Undifferentiated smoker	1.89	2.12	1.35	5,839
	Weighted error	0.78	0.41	0.19	
Face Amount Band	< 100k	1.80	1.80	1.17	7,801
	100k–249k	1.91	1.46	1.10	37,049
	250k–999k	1.49	1.20	0.83	30,416
	1m+	2.22	1.58	1.13	4,211
	Weighted error	0.75	0.40	0.13	

especially low exposures or claims, sometimes deviating from the population mean by more than four orders of magnitude. On the other hand, the mortality rates in the partially pooled model were much closer to the population mean on an exposure-controlled basis. Outlier behavior of the magnitude seen when partial pooling was not used was not observed.

Mortality rates in the partially pooled model were much closer to the population mean on an exposure-controlled basis.

Figure 3 Model Validation Comparison

	Expected Log Predictive Density (ELPD)	Standard Error of Difference (SED)
With partial pooling	0.0	0.0
Without partial pooling*	–9.0	3.5

*For this row, we show values for the regularized (nonpartial pooling) model that gives the best results.

When comparing models using Bayesian selection methods,[26] the partially pooled model had significantly better LOO cross validation and WAIC scores, as shown in Figure 3.[27]

When predicting mortality rates for cohorts with relatively small exposures (~5 percent of the mean per-cohort exposure, 153 cohorts out of 8,000), the nonpooled models yielded mortality estimates that are less than 0.01 percent of the mean mortality rate (interestingly enough, over-estimation was not observed). This under-estimation resulted from improper handling of small sample sizes. These results held even with the regularized models, which are very similar to models with graduation.[28]

On the other hand, models with partial pooling did not produce such extreme estimates because of the beneficial impacts of shrinkage. Proper handling of mortality estimates in cohorts with small exposures is critical, as such cohorts will almost certainly exist when modeling data at high granularity.

Conclusion

This article explored innovative approaches to modeling PLT lapse and mortality. A multivariate PLT lapse and mortality model improves mortality estimates and sheds new light on the interactions among changes in premium, persistency and mortality. Because management would have the information it needs in real time, this transforms pricing, reserving and "what if" analysis.

Partial pooling shares information among cohorts, accounts for different cohort sizes, regularizes estimates and improves credibility. When there are multidimensional cohorts with sparse data, partial pooling can provide unique insights into policyholder behavior, which is very valuable for insurers looking to manage risks and finances and optimize top-line growth.

The Bayesian model allows us to capture our prior knowledge of the data-generating process, such as the reasonable values of the effectiveness parameter. Such a model will be practical and implementable—and not just a nice theoretical toy.

The methods discussed in this article are valuable for answering a wide range of sophisticated actuarial questions. Actuaries and insurers will want to consider how advanced methodologies such as the innovative lapse-mortality model, causal inference and Bayesian decision theory could be used to address crucial challenges. Now that the availability of computational resources facilitates the implementation of these advanced methodologies, insurers face a new imperative. These techniques can be extended to general lapse and mortality studies along with other aspects of the insurer experience.

ISTOCK.COM/DEAN MITCHELL

We look forward to seeing the improvements in pricing and reserving (such as for principles-based reserving) and the increases in credibility that will emerge from greater use of these techniques. ■

References

[1] Society of Actuaries. 2014 Post Level Term Lapse & Mortality Report. *Society of Actuaries*, 2014, *https://www.soa.org/resources/experience-studies/2014/research-2014-post-level-shock/* (accessed January 27, 2020).
[2] Society of Actuaries. Lapse Modeling for the Post-Level Period—A Practical Application of Predictive Modeling. *Society of Actuaries*, 2015, *https://www.soa.org/resources/research-reports/2015/lapse-2015-modeling-post-level/* (accessed January 27, 2020).
[3] In this article, a cohort means a group of policyholders who are the same age, gender, risk class, premium mode, face amount band and premium jump band.
[4] PLT causal structure means an in-depth understanding of the causal relationships between PLT pricing and experience.
[5] Due to space considerations, we do not include an in-depth causal analysis in this article.
[6] In this context, regularization means constraining model parameters, therefore reducing the risk of overfitting and improving the ability of the model to generalize.
[7] James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning With Applications in R*. New York: Springer.
[8] A field of statistics in which uncertainty is quantified with probability, and quantities of interest are expressed as posterior expectations, taking into account prior/expert knowledge and how likely different models are given the observed data.
[9] Makov, U., A. Smith, and Y.-H. Liu. 1996. Bayesian Methods in Actuarial Science. *Journal of the Royal Statistical Society. Series D (The Statistician)* 45, no. 4:503–515.
[10] A generative model describes how a data set is generated in terms of a probabilistic model. New data can be generated by sampling from this model.
[11] An informed prior is a distribution that adds information to the statistical inference by incorporating external knowledge.
[12] Vose Software. Informed Prior. *Vose Software*, 2017, *https://www.vosesoftware.com/riskwiki/Priorfrompreviousanalysis.php* (accessed January 27, 2020).
[13] Supra note 1.
[14] Supra note 2.
[15] The traditional non-Bayesian DM model uses these inputs: base lapse rate, PLT lapse rate, select and point-in-scale mortality, and an "effectiveness" parameter G, which represents the distribution of healthy lives and antiselectors between excess lapsers and persisters. The mortality of the persisters is then estimated using the concept of conservation of total deaths.
[16] Dukes, Jeffery T. Term Mortality and Lapses. *Society of Actuaries*, August 2005, *https://www.soa.org/globalassets/assets/library/newsletters/product-development-news/2005/august/pdn-2005-iss62-dukes.pdf* (accessed January 27, 2020).
[17] Doll, Douglas C. Mortality Anti-selection—Different Versions of Dukes/MacDonald. *Society of Actuaries*, July 2003, *https://www.soa.org/globalassets/assets/library/newsletters/product-development-news/2003/july/pdn-2003-iss56-doll-a.pdf* (accessed January 27, 2020).
[18] Carpenter, Bob, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A Probabilistic Programming Language. *Journal of Statistical Software* 76, no. 1:1–32.
[19] LOO and WAIC are standard tools used by data scientists and statisticians to measure out of sample accuracy (i.e., predictive capability) of the model. Colloquially, LOO estimates how well a statistical model can predict each member of the original data set, assuming the model was trained on all of the data except for the member being predicted. WAIC is a generalization of the Akaike Information Criterion (AIC) that analyzes model accuracy relative to how many parameters the model uses.
[20] Society of Actuaries. Lapse and Mortality Experience of Post-Level Premium Period Term Plans. *Society of Actuaries*, October 2009, *https://www.soa.org/resources/research-reports/2010/research-shock-lapse-report/* (accessed January 27, 2020).
[21] Grid search allows for the search of all possible values, and G is selected to be equal to the point in the grid that most closely reproduces the observed number of deaths.
[22] Hoffman, Matthew D., and Andrew Gelman. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *Cornell University*, November 18, 2011, *https://arxiv.org/abs/1111.4246* (accessed January 27, 2020).
[23] Betancourt, Michael. A Conceptual Introduction to Hamiltonian Monte Carlo. *https://arxiv.org/pdf/1701.02434* (accessed January 27, 2020).
[24] Due to the sparse nature of the data, a multidimensional table (e.g., showing mortality rates for combinations of premium mode and risk class) could be misleading for the empirical analysis. Therefore, multidimensional results are not presented. However, the partially pooled model could give good estimates.
[25] Supra note 20.
[26] Vehtari, A., A. Gelman, and J. Gabry. 2017. Practical Bayesian Model Evaluation Using Leave-One-Out Cross-Validation and WAIC. *Statistics and Computing* 27, no. 5:1,413–1,432.
[27] ELPD and SED are standard tools used to measure the relative accuracy of the models using LOO and WAIC. The best model gets a score of 0 as it is being compared to itself. For the other models, the larger the magnitude of ELPD, the greater the divergence from the results of the best model. The SED gives a measure of confidence in the ELPD.
[28] Girosi, Federico, and Gary King. 2008. *Demographic Forecasting,* 66. Princeton: Princeton University Press.

ABOUT THE WRITERS

MARTIN SNOW, FSA, MAAA, is vice president, chief delivery officer and chief actuary at Atidot. He can be reached at ***Martin@Atidot.com***.

ADAM HABER is a data scientist at Atidot in Tel Aviv. He can be reached at ***adamhaber@gmail.com***.

The Actuarial FOUNDATION

Only 20 states require a personal finance course to graduate high school (Council of Economic Education)

Average math literacy scores among American 15-year-olds have dropped since 2009 (Programme for International Assessment)

Only five states are proficient at teaching financial literacy (Champlain College)

BRIDGING THE GAP

15 Million

Number of students who have accessed our programs

873

Students registered in 2019 for the Modeling the Future Challenge—a six-fold increase from the previous year

263

Number of tutors actively tutoring in Chicago, Hartford, Lincoln, New York City, St. Paul, Seattle and more

Visit
www.actuarialfoundation.org/educate-connect-change/
for more information.

FEATURE

Riding the ERM Wave

A brief history of ERM modeling and what's coming next

BY RICH LAURIA

It may be hard to believe, but enterprise risk management (ERM) has become mainstream. ERM processes exist today at essentially every major insurer operating in well-established insurance markets. Regulators across the globe have some type of own risk and solvency assessment (ORSA) requirement, facilitated by the International Association of Insurance Supervisors' adoption of Insurance Core Principle 16. Evaluating insurer ERM is now embedded into the rating criteria of each major credit rating agency.

The actuarial profession has made significant contributions to the advancement of many key components of the ERM process, including risk identification, assessment and decision-making. However, its main achievements, arguably, have been in the area of risk quantification. As I reviewed all of the ERM modeling practices I have utilized over the course of my career, I was struck by their ambition and comprehensiveness, as well as the challenge in making them transparent and intuitive to key stakeholders.

Logical and Not-so-humble Beginnings

Before getting into where ERM modeling stands today, it is instructive to review how it all started. As with most evolutionary processes, there is no single find in the archives that definitively establishes a genesis of the first ERM model. Rather, different parts of the profession moved toward its conceptualization from their own unique vantage points. In short, the roots of today's ERM models can be found in previous actuarial models constructed for different purposes. Those models were leveraged and modified to paint the larger mosaic of risk being sought after.

Life and annuity actuaries have long performed total company projections, built from the ground up using policy-level information. Output includes future premiums, claims, withdrawals and surrenders, investment income, expenses and other

ISTOCK.COM/PETESPHOTOGRAPHY

key financial statement items. Aggregation across blocks of business provides a total legal entity viewpoint under specific assumptions. New York Regulation 126 and the National Association of Insurance Commissioners (NAIC) Actuarial Opinion and Memorandum Model Regulation were established in the 1980s, leading to the modern-day Appointed Actuary and the use of cash-flow testing in asset adequacy analysis. This provided life insurers a view into interest rate risk exposures. Life valuation actuaries went beyond prescribed scenarios, utilizing economic scenario generators (ESGs) to capture exposure to other financial risks, such as equity market volatility. Beyond reserve testing, these models were readily adapted through adding capital requirements and profits released adjustments to calculate the embedded value of the business, providing useful information on the drivers and risks of long-term value.

Concurrently, actuaries at property and casualty (P&C) insurers began creating ground-up models of business projected over multiple years. These dynamic financial analysis (DFA) models also project key financial items that capture exposure to uncertainty in both claim and investment experience. They also utilize ESGs and tie claim costs to parameters such as inflation and interest rates. DFA models differed from their life and annuity counterparts in that they incorporated new business into projections. Insurers found several nonregulatory applications for them, including capital management, reinsurance, pricing, and mergers and acquisitions. Indeed, DFA models contained many of the seeds of modern-day ERM models.

The ERM Modeling Opportunity

The ERM wave began gaining momentum at the turn of the millennium, inspired by many risk events that occurred during that time. Examples include the dot-com boom and bust; several major accounting scandals; the Sept. 11, 2001, terrorist attacks; and the advent of cyberattacks. The failure of the hedge fund Long-Term Capital Management illustrated that an individual firm of sufficient size, scope and interconnectivity could threaten the viability of financial markets. In the insurance industry, low interest rates combined with other adverse experiences exposed the risks of variable products and long-term care. It was recognized that risks beyond underwriting and investment needed to be considered, and scenarios involving multiple risks manifesting concurrently needed to be assessed.

Furthermore, company management was looking for risk-adjusted performance measures to aid strategic decisions and ensure they were compensated for the risks they took. Everything from product enhancements to business dispositions would be subject to a risk lens. Risks would be prioritized for better resource allocation. Budgets would be sized appropriately for audit, compliance and cybersecurity. Protection levels would be analytically set for dynamic hedging, corporate insurance and reinsurance programs.

Regulators and credit rating agencies would be assuaged that company risks were known and well-managed, with the knowledge impacting business decisions. They thought that with more efficient financial examinations and better credit ratings, perhaps lower capital requirements would result. Focus on these stakeholders helped shape initial ERM modeling efforts around capital.

The Rise of Economic Capital

Most ERM modeling discussed in actuarial literature focuses on economic capital (EC). In this context, EC initially was defined as the amount of capital an insurer needs to cover its risks based on a specified security standard, usually without consideration of any external constraints, such as desired credit ratings or regulatory minimums. Ignoring these constraints does not seem economical at all, and as discussed later, more EC models today are incorporating them.

Early forms of EC were developed based on a modified-factor approach, similar to NAIC risk-based capital (RBC). Most risks are quantified by applying a factor to a measure of risk exposure. For example, the credit risk of a fixed-income investment is calculated by applying a factor to the book or market value of that investment, where the factor captures the relative risk of the investment based on its individual characteristics. Similarly, mortality risk typically is captured by applying a factor to the net amount of each policy at risk, with the factor varying by the characteristics of the insured and the policy itself.

Other risks are quantified based on the modeling results of blocks of business with exposure to those risks. The nature and complexity of the risks do not lend themselves well to a factor approach. These typically include market risks such as interest rate, equity and foreign exchange that can be captured using cash flow testing models in conjunction with an ESG. Another category is catastrophe risk, arising from either natural or manmade perils, which is often captured using third-party vendor models. The C3 phase 1 and phase 2 portions of the life RBC formula and the catastrophe risk charge in the P&C RBC formula are the regulatory capital parallels.

A key modeling decision is the model time horizon. Since life insurers write long-term liabilities, one approach is to capture the total risk over the run-off period of policies in

ISTOCK.COM/RAMIHALIM

Risk events that coincided with the ERM wave include the dot-com boom and bust; several major accounting scandals; the Sept. 11, 2001, terrorist attacks; and the advent of cyberattacks.

force, plus one to three years of new business. This requires the insurer to hold the present value of capital needed over the run-off period—this is the approach taken in C3 RBC phases 1 and 2. The primary alternative is a one-year horizon. This amounts to “just in time” capital and is the horizon for the standard formula in Solvency II. Usage of the shorter time horizon usually is accompanied by a higher security standard—the solvency capital requirement in Solvency II targets a 99.5 percent value at risk (VaR). The model horizon and security standard are harmonized with how factors are developed.[1]

Individual risk amounts are aggregated using a variance-covariance matrix reflecting the expected co-movement of risks. Note the square root of sum of squares approach used in the RBC formula is a special form. In early iterations of EC, operational risks were often the last step. They were an add-on calculated as either a percentage of premium (e.g., RBC) or as a gross-up of aggregate capital (e.g., 15 percent). This simplified approach reflects the lack of operational risk data and the heterogeneity of such risks.

One early example in ERM modeling took place shortly after the turn of the century and involved the U.S. subsidiary of a major European financial services company with large banking operations. Because of the overwhelming amount of enterprise risk capital due to the market and credit risks of the European parent, the marginal EC requirements of the U.S. insurance operations were a fraction of its gross capital charges. The EC model assumed no correlation between underwriting and asset risks, and the bulk of the diversification benefits generated were allocated to the U.S. insurance operations. The EC model result indicated that the U.S. insurance operations needed capital well below what the rating agencies expected.

Getting the agencies to agree with the company view hinged on alignment of model correlation assumptions and the fungibility of capital flows across legal entities. Obtaining such buy-in is a multiyear process that includes clear demonstration that the model is driving business decisions (the “use test” criteria).

Recent EC Model Developments

Since that time, EC models have evolved in a number of ways.

1. **Models employ fewer factors and model more risks directly.** Some models have abandoned the factor approach altogether. A subset of these calculate EC holistically by modeling all risks concurrently and then allocating capital to each risk source and business unit based on the results and security standard selected (there are well-established mathematical procedures for performing this, such as the Ruhm-Mango-Kreps capital allocation algorithm). The increase in direct modeling has been accompanied by increased model testing, governance and validation. In particular, correlations and extreme events are scrutinized for their plausibility.
2. **Modeling advancements have occurred for certain risks common to many organizations.** More data has been collected on credit risk and operational risks, such as cyber, financial statement errors and litigation. Both trends allow for more precise risk measurement.
3. **EC is calculated for each legal entity on a stand-alone basis** unless it is demonstrated that capital shortfalls will be reimbursed by another entity in the holding company family on a timely basis.
4. **The modeling time horizon coincides with the business plan.** This allows for the synchronization of capital planning with projected growth. New business projections are included.
5. **Capital requirements also are projected over the plan horizon,** under normal conditions and stress scenarios.

ORSA requirements have driven much of these trends. Nos. 1 and 2 from this list align closer to the principles embedded in Section 2 of the ORSA manual. The other three enhancements address expectations for Group Capital Assessment spelled out in Section 3.[2]

Treasury-based Approach to EC

There has been convergence toward a methodology that brings ERM into the treasurer's office and largely mirrors the Federal Reserve's stress tests of large banking groups. This approach supplies critical information on the riskiness of future cash flows within a given holding company system. The steps in this approach are:

ISTOCK.COM/SANFEL

Studies of past corporate failures have shown that **60–65 percent** were caused by strategic risks, and another **20–30 percent** were caused by operational risks. Yet these risks can be overlooked when focusing exclusively on capital.

1 | Project required capital for each entity under the business plan over the horizon based on external constraints from rating agencies and regulators. This necessitates setting ratings and capital ratio targets.

2 | Project actual capital for each entity under the business plan over the horizon based on future profitability projections from the businesses and capital management actions from treasury operations.

3 | Assess any capital redundancies and deficiencies under the plan.

4 | Repeat the three prior steps for each risk scenario that management has specified for inclusion as part of its risk appetite statement on capital adequacy.

5 | Based on the modeling results, determine in consultation with the businesses and treasury operations whether the business plans and/or capital management actions need to be adjusted to satisfy risk appetite.

I led the initial implementation of such an approach in 2012, and with some procedural refinements, it remains in place today. It has been the cornerstone of ORSA quantification for the company, and it has provided valuable input into share repurchase, reinsurance purchase, and mergers and acquisition decisions. The model output resonates with the C-suite and the board of directors. The approach has been reviewed favorably by two leading ERM consulting firms.

Going Beyond Capital

The past decade has witnessed the emergence of a comprehensive approach to quantify all risks, including strategic and operational risks. Studies of past corporate failures have shown that 60–65 percent were caused by strategic risks, and another 20–30 percent were caused by operational risks. Yet these risks can be overlooked when focusing exclusively on capital.[3]

The value-based ERM approach, authored by ERM thought leader Sim Segal, marries traditional ERM techniques with value-based management. Risk is defined as any deviation—upside or downside—in company value from its baseline amount as calculated under the company's strategic plan. Company value is the present value of free cash flows to the company's owners discounted at its cost of equity. Scenarios are created for all key risks based on "Failure Modes and Effects Analysis." They include estimated impacts to key financial items from their baseline values to recalculate company values under each scenario.[4]

The value-based ERM model is simultaneously a dynamic strategic planning tool and an economic capital model, both of which capture key volatility and can be run rapidly to inform decision-making at the highest levels. The model projects statutory financials; required capital ratios; and key metrics such as company value, capital ratios and so on at the business subsegment level, rolled up to segment, legal entity and total company levels.[5]

The value metric highlights risks less emphasized under capital-centric measurement. In my experience, the value metric helped identify and quantify risks that were not pure loss events (e.g., asset write-downs during a financial crisis or natural catastrophe claims). Risks of changes in regulation (e.g., health care reform), systems obsolescence driving loss of distribution, adverse litigation outcomes damaging company brand and disruptive competitors are all modeled as reductions in future profitable business as well as increased claims and expenses. This adds a critical dimension to ERM.

What's Next?

ERM modeling will continue to evolve, driven internally by the need for more accurate and timely risk information and by external stakeholder requirements for reporting on risks and capital. Practitioners will continue refining existing methodologies built on pro forma projections. Data analysis of risk events will become more formalized, providing ERM modelers regularly updated assumptions. Advancements in machine learning (ML) and artificial intelligence (AI) may be the impetus of future models.

ERM model leaders will continue to evaluate the trade-offs between the transparency and intuition of simpler approaches, and the robustness and detailed insights offered by complex models. Both deterministic and stochastic scenarios will play a role, with the former leaned on for risk messaging and the latter used for quantifying nonlinear risks (e.g., options). Both will be used for setting risk appetite and limits.

Capital will remain the primary focus for many insurer ERM programs. However, models will adapt to capture other metrics such as earnings volatility and value impacts. The result will be a robust dashboard of output addressing the multiple perspectives to be considered in effectively managing enterprise risk. ■

References

[1] Swain, Robin, and David Swallow. 2015. The Prudential Regulation of Insurers Under Solvency II. *Quarterly Bulletin*. Bank of England.

[2] ORSA Subgroup of the Financial Condition Committee. 2014. *Own Risk and Solvency Assessment (ORSA) Guidance Manual*. Kansas City, Missouri: National Association of Insurance Commissioners.

[3] Segal, Sim. 2011. *The Corporate Value of Enterprise Risk Management*. Hoboken, New Jersey: John Wiley and Sons.

[4] Ibid.

[5] Ibid.

ABOUT THE WRITER

RICH LAURIA, FSA, CFA, MAAA, is associate academic director of the enterprise risk management program at Columbia University. He was previously senior vice president in enterprise risk and actuarial roles at Assurant Inc. He can be reached at ***rl2764@columbia.edu***.

Actuarial Systems EVOLUTION

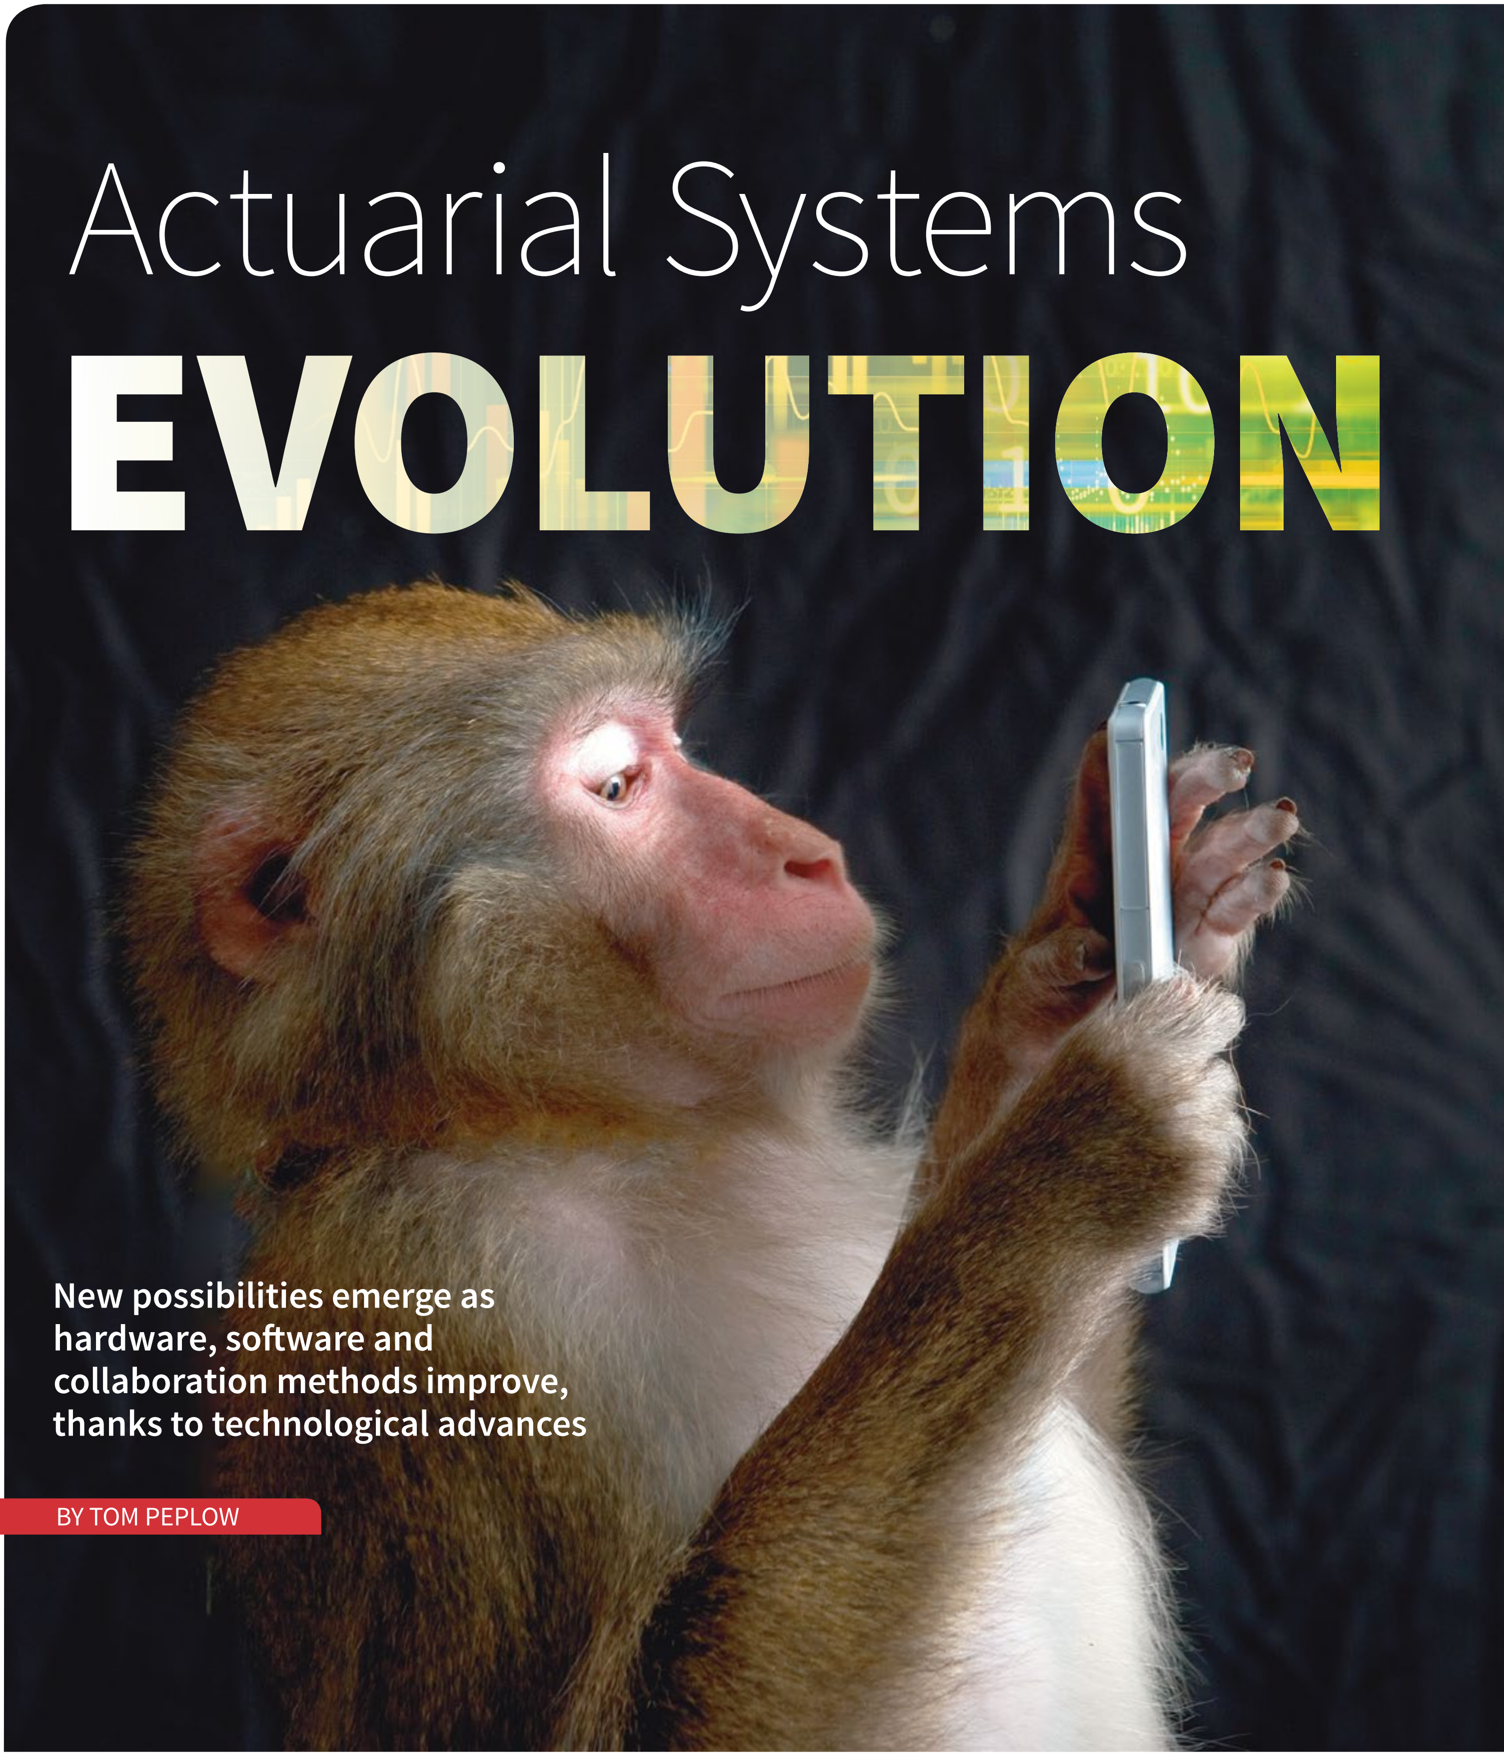

New possibilities emerge as hardware, software and collaboration methods improve, thanks to technological advances

BY TOM PEPLOW

In this article, I provide a software engineer's perspective on keeping pace with both the actuarial and software industries. Software engineering is a young industry compared to actuarial science, although the two lines of work share a common mathematical and scientific foundation.

One of the first users of the term "software engineer" was Margaret Hamilton, who worked at NASA on the Apollo space program. There is a beautiful story of how she re-engineered the software for guidance computers. Her daughter caused a system crash while playing in the simulator, which made Margaret rethink asynchronous messaging processing. This change resulted in the astronauts landing safely on the moon, as the computer was not overwhelmed by erroneous input and could focus on the task at hand.

The reason I share this story is because software engineers and actuarial scientists have learned a great deal from each other. Margaret's daughter prompted her to ask, "What if that did happen?"—a question that actuaries strive to answer. Software engineering needed to grow up fast. It did, and as a result, it solved problems that benefit the way actuaries work today.

The actuarial platform I work on continues to evolve from its inception 25 years ago. Complexity has grown to support new types of insurance, riders and so on. Yet it must remain resilient enough to support policies issued many years prior. Over the years, regulations and the economic environment have changed in ways that couldn't have been predicted. Black swans, ideally a 1-in-200-year event, have been more like a 1-in-10-year event. In that

Over the years, regulations and the economic environment have changed in ways that couldn't have been predicted.

GETTY IMAGES/RYOUCHIN

time, there also have been significant paradigm shifts in technology. The technology evolution has enabled much of the regulatory change, but it also caused some of the economic shifts. Technology has changed the way we live and the shape of many industries. This pace of disruption is increasing—a startup can become an incumbent and displace the status quo quickly; Uber and Airbnb are the most obvious examples.

There are two distinct aspects of the actuarial system evolution on which I would like to reflect. The first is looking at what has become possible thanks to the evolution of hardware and software; the second is the evolution of how actuaries work together. In both cases, I see a fusion of ideas that gives me hope that the insurance industry will continue to evolve and adapt to serve the needs of a world that is very different than the one it originally set out to serve.

Computation

Models in the Desktop Era

The problem with the mainframe computer was the fact that it was a shared resource. You had to be efficient and precise. There was nothing worse than coming to work in the morning to find an error message where your results should have been displayed. The personal computer (PC) was liberating—the actuary was free. As insurers started rolling out desktop computers to their staffs, actuaries were given a capable computing environment in a box under their desk. It didn't take them long to start making that central processing unit (CPU) hot calculating reserves and pricing new products. But Moore's Law,[1] the observation that processing power doubles every two years, meant the humble PC was keeping up with the needs of actuaries.

The new challenge was that actuaries needed to learn to program. Vendors introduced domain-specific languages (DSLs), so actuaries could be more expressive and productive than they would be in general-purpose languages, such as C. Vendors implemented standard domain concepts, such as the double-entry accounting structure, and they delivered standard libraries to meet regulatory needs, further increasing the leverage of the individual actuary. This all reduced the barrier to entry and helped the actuarial department concentrate on innovation.

Moore Processing Power

It didn't take long for the PC to start to running out of the horsepower needed to run models in a timely fashion, due in large part to increased product complexity, regulatory needs and more data. Thankfully, servers with multiple CPUs and more memory had become ubiquitous, so actuaries could move their models to servers.

The DSLs enabled vendors to leverage multi-CPU architectures without actuaries needing to change all of their business logic. It's also relatively straightforward to parallelize an actuarial model. There were two embarrassingly parallel distribution mechanisms available, assuming you were happy with some limitations/simplifications. Distributing calculations by economic scenario meant the stochastic requirements could be met by running scenarios in parallel. It also was possible to parallelize a single scenario by liability/asset cell.

The single server soon became the limiting factor. Fortunately, the vendors' engineers could take advantage of the same distribution mechanics used to leverage multiple CPUs in a single server to scale out across multiple servers.

The interesting observation is that the actuarial modeling platform—at least the one for which I am responsible—started with the tagline: "Built by actuaries, for actuaries." At this point, the needs of the actuarial department created a diverse product development group, one of actuaries working with software engineers. The addition of engineers accelerated the rate of innovation: It modernized the platform and led to a product built for actuaries with actuarial input.

Cloud

Despite the computational wins of the multiple-server solution, in a way, the actuarial department found itself back in the days of the mainframe. Models needed a shared compute infrastructure to execute, and teams were back to waiting in line to get an answer. It was not as bad as it once was—at least actuaries could test and debug models locally before running them—but it was still expensive and frustrating to wait for the answer. What if there were unlimited compute capacity and no more waiting in line?

The cloud is delivering this promise. As you lease servers by the second, it is possible to scale a grid to hundreds of thousands of servers in minutes. Owning this many servers simply would not be economical, and even if that capital investment were justified, utilization rates would be extremely low. The cloud has the economy of scale to make this investment worthwhile.

This is similar to the evolution of software in other industries. RenderMan, the technology behind the *Toy Story* movies, has scaled out to meet the needs of modern photorealistic computer-animated movies like the *Lion King*. This technology is cloud-compatible, allowing anyone access to huge compute capacity to render their ideas

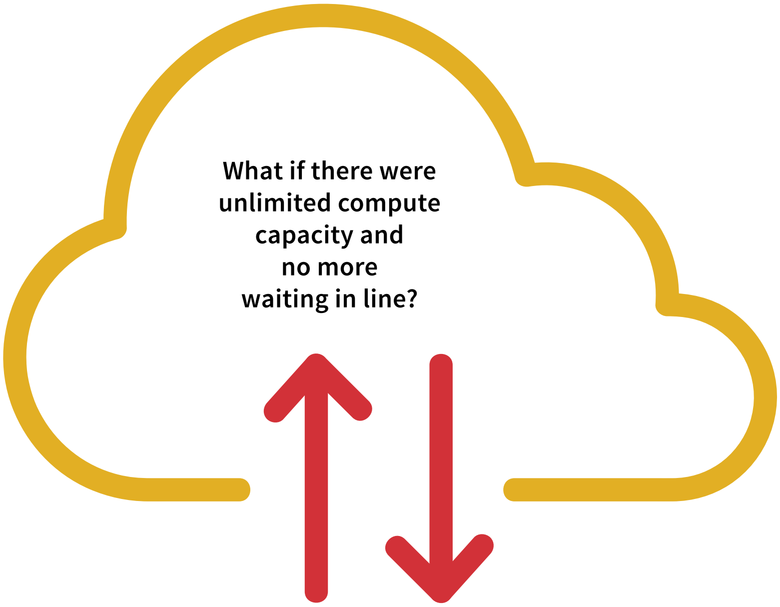

a reality with very little upfront investment. Access to super computers is no longer a barrier to innovation.

Modeling Complexity and Controls

The Challenges of Model Development

In parallel to the technical evolution, the process around modeling—especially for financial reporting—has become more sophisticated. The need for a team of actuaries to collaborate during model development has increased due to the need to draw upon specialized expertise and the sheer volume of model updates to keep pace with emerging regulations and product innovation. There are various operating models, but the requirement for more collaboration on a single model created the significant challenge of organizing and merging changes from many contributors across the business.

Chief actuaries further raised the bar when they became frustrated by different sources of truth being used to answer different business questions. For example, answers from the forecasting models differed from valuation model answers. Delivering on the diverse modeling requirements required collaboration. But at the same time, modeling teams needed to consolidate models to provide consistency across the business, which was forced, in part, by principle-based regulations such as Valuation Manual 20 (VM-20).

Regulators have become increasingly interested in the rigor with which models are built and maintained, so the actuarial development team members are required to demonstrate that controlled change processes are embedded and followed. Fortunately, leaning on the engineering teams for diversity of thought (as well as leveraging centralized and shareable resources) provides an alternative way to work. There are easy wins to the workflows of

COMMODITY VS. SPECIALIZED COMPUTE

Actuarial modeling vendors have not yet been able to translate their existing domain-specific languages (DSLs) to execute efficiently on graphics processing unit (GPU) cards. GPUs are specialized microprocessors designed to render high-resolution images and video. This specialism can be leveraged to perform floating-point arithmetic at higher degrees of parallelism than a central processing unit (CPU).

The expression of business rules does not easily translate to the kinds of operations suited to GPUs, so there is little penetration of these options. DSLs for GPUs leak too much of the implementation detail, which creates a chasm between mass-market opportunity and the promise of the increased performance. Crossing the chasm is expensive and requires specialized talent, so the application of this hardware is limited to narrow use cases, such as economic scenario generation and narrow modeling use cases (e.g., hedge calculations with variable annuity embedded guarantees).

The pace of innovation in the cloud has provided an alternative, using commodity CPU instructions. The density of CPUs in a server has increased, as has available memory and memory bandwidth. The restrictions of input/output (I/O) communication between the CPU, its cache, its memory and hard disk, and across networks, are less impactful thanks to solid-state drives (hard disks without a mechanical head that provide much better random-access behavior and more I/O operations per second) and remote direct memory access (RDMA) networks, which allow direct memory access across servers at high throughput and low latency.

Combined, this delivers significant performance at a lower cost than GPUs due to more optimal thermal efficiency and lower power consumption. Even though transistor density is no longer doubling each year (as predicted by Moore), cloud innovations are finding ways to stay ahead of demands without a paradigm shift to alternative hardware. As the cloud evolves without any capital investments, vendors quickly can bring new options to customers without complex changes to the models.

updating models and running production models, which actuaries can gain by leveraging what engineering teams do.

All this means that modeling software needed to:

1 | Evolve to be more open to concurrent editors.
2 | Simplify bringing that work together.
3 | Reflect transparency to show what has changed over time.

Collaborative and Auditable Model Development

The nice thing about version-control systems, like GitHub (see sidebar), is they enable additional use cases: collaboration, auditability and traceability. Collaboration becomes easier as these systems make it easy to combine changes and handle conflicts where several people change the same model. This reduces the overhead of multiple contributors. The same mechanisms for tracking changes from one person to another creates an immutable audit trail that allows users to see what changes happened between two points in time. That audit trail links requirements to changes and then results, creating end-to-end traceability. This is an auditor's dream, and it frees actuaries to innovate within a controlled framework.

Collaboration becomes easier as these systems make it easy to combine changes and handle conflicts where several people change the same model.

GITHUB

Git is a distributed version-control system to help manage concurrent changes made by teams of developers working collaboratively on source code. It provides a single source of truth of all code changes over time.

GitHub revolutionized how the open-source software community collaborated on open-source projects. This, in turn, transformed software teams within established organizations to be more collaborative and open internally. GitHub made it easy for developers to host Git and layered on new concepts, such as pull requests (for code reviews), that engaged individuals in collaborative problem-solving and created a community where people learn, share and work together to build software.

Making change control a first-class citizen in model development has simplified the GitHub approach, much like SharePoint has done for business users collaborating on Microsoft Word documents. Rather than asking actuaries to learn Git—and externalize models in a source-control system—vendors have been able to make these best practices simple and less obtrusive by integrating them into their tools. They also can provide semantic merge capabilities thanks to the simplicity of their DSLs, so merge conflict resolution is less onerous than merging unfamiliar files.

Producing Results

The model's business logic and configuration is just one part of the end-to-end production process for creating actuarial results. Models typically reference externalized assumptions, which are set during a separate experience analysis process. There are also inputs that change from period to period—for example, assets and liabilities, which often undergo transformation from one format to another. And there are economic scenarios, either fed in or generated by integrated economic scenario generators (ESGs). There often is a mechanism, after model runs are completed, for including nonmodeled results and manual adjustments. New regulations, such as Long Duration Targeted Improvements (LDTI) and International Financial Reporting Standard (IFRS) 17, require historical modeled results and transactional data as inputs, which adds a new dimension of change management not previously in scope for actuarial reporting. To add

ISTOCK.COM/4X6

even more pressure, the regulator defines tight reporting timelines, which requires insurers to be faster and more agile. This has driven vendors to enable automated processes that can easily evolve over time.

Providing traceability across all areas of change, including data lineage, is extremely powerful for change attribution analysis and essential for building confidence in the modeled result. Demonstrating Sarbanes Oxley (SOX) controls during an audit is onerous if this process is not managed holistically, and now with historical input, that causality chain spans several years. Managing this in a piecemeal way requires that all participants in the value chain provide a mechanism for end-to-end traceability.

Combining a controlled development process for models, data transformations and iterations with data (asset, liabilities, assumptions and economic scenarios) extends the traceability of the change control system to production results. Results are trustworthy thanks to an automated and locked-down execution environment that prevents tampering and human error. Distributed computing provides the means to deliver these results quickly, and the cloud reduces capital expense by scaling to meet the demands of increasing complexity and regulatory needs.

Conclusion

When I reflect on how the software product I helped build has evolved, and the use cases we now support, I feel very proud. The pride comes from the fact that we created a diverse team of people who empowered actuaries and helped them handle the increasing demands of their jobs. Our clients no longer must choose between compliance and innovation—we provide tools that support freedom to innovate within a controlled framework. These tools allow actuaries to collaborate, provide access to unlimited cloud computing and more data, derive more insight, and reimagine the way they work for the better.

That partnership has meant the clients we serve are more equipped to build better insurance products. Without these products, people would not be able to retire or provide for their loved ones when they die. So I also feel a great sense of pride that the software I help deliver fulfills a huge need in society.

These efforts are critical, because insurance provides financial security against so many risks: premature death, disability, poor health, living longer and so on. In the United Kingdom, where I am from originally, 55 percent of Generation X are at a high risk of not achieving a moderate level of income in retirement. A report by The Phoenix Group[2] contrasts the situation from the prior generation: Members of Generation X lose £13,000 in state pension over their lifetimes, and they occupy rented accommodations at a rate 8 percent higher than baby boomers, leaving them with less disposable income and fewer assets.

As software has evolved, so have the needs of retirees. My goal is to continue to innovate in partnership with my actuarial friends and accelerate the transformation of the tools they use, so they are empowered to solve the vast array of issues related to financial security—including the transformation of insurance—so everyone has access to the retirement they deserve. ■

New regulations require historical modeled results and transactional data as inputs, which adds a new dimension of change management not previously in scope for actuarial reports.

References

[1] Moore, Gordon E. 1965. Cramming More Components. *Electronics* 38.

[2] Phoenix Group. "Generation VeXed" Faces a Grey Retirement. *Phoenix Group*, November 7, 2019, *https://www.thephoenixgroup.com/~/media/Files/P/Phoenix-Group-v3/Press%20releases/191107%20Phoenix-PPI%20Report%20-%20Generation%20VeXed%20Solving%20The%20Retirement%20Puzzle%20FINAL.pdf* (accessed January 28, 2020).

ABOUT THE WRITER

TOM PEPLOW, MSc, is a principal in Life Technology Solutions at Milliman, where he is the director of Product Development. He can be reached at ***Thomas.Peplow@milliman.com***.

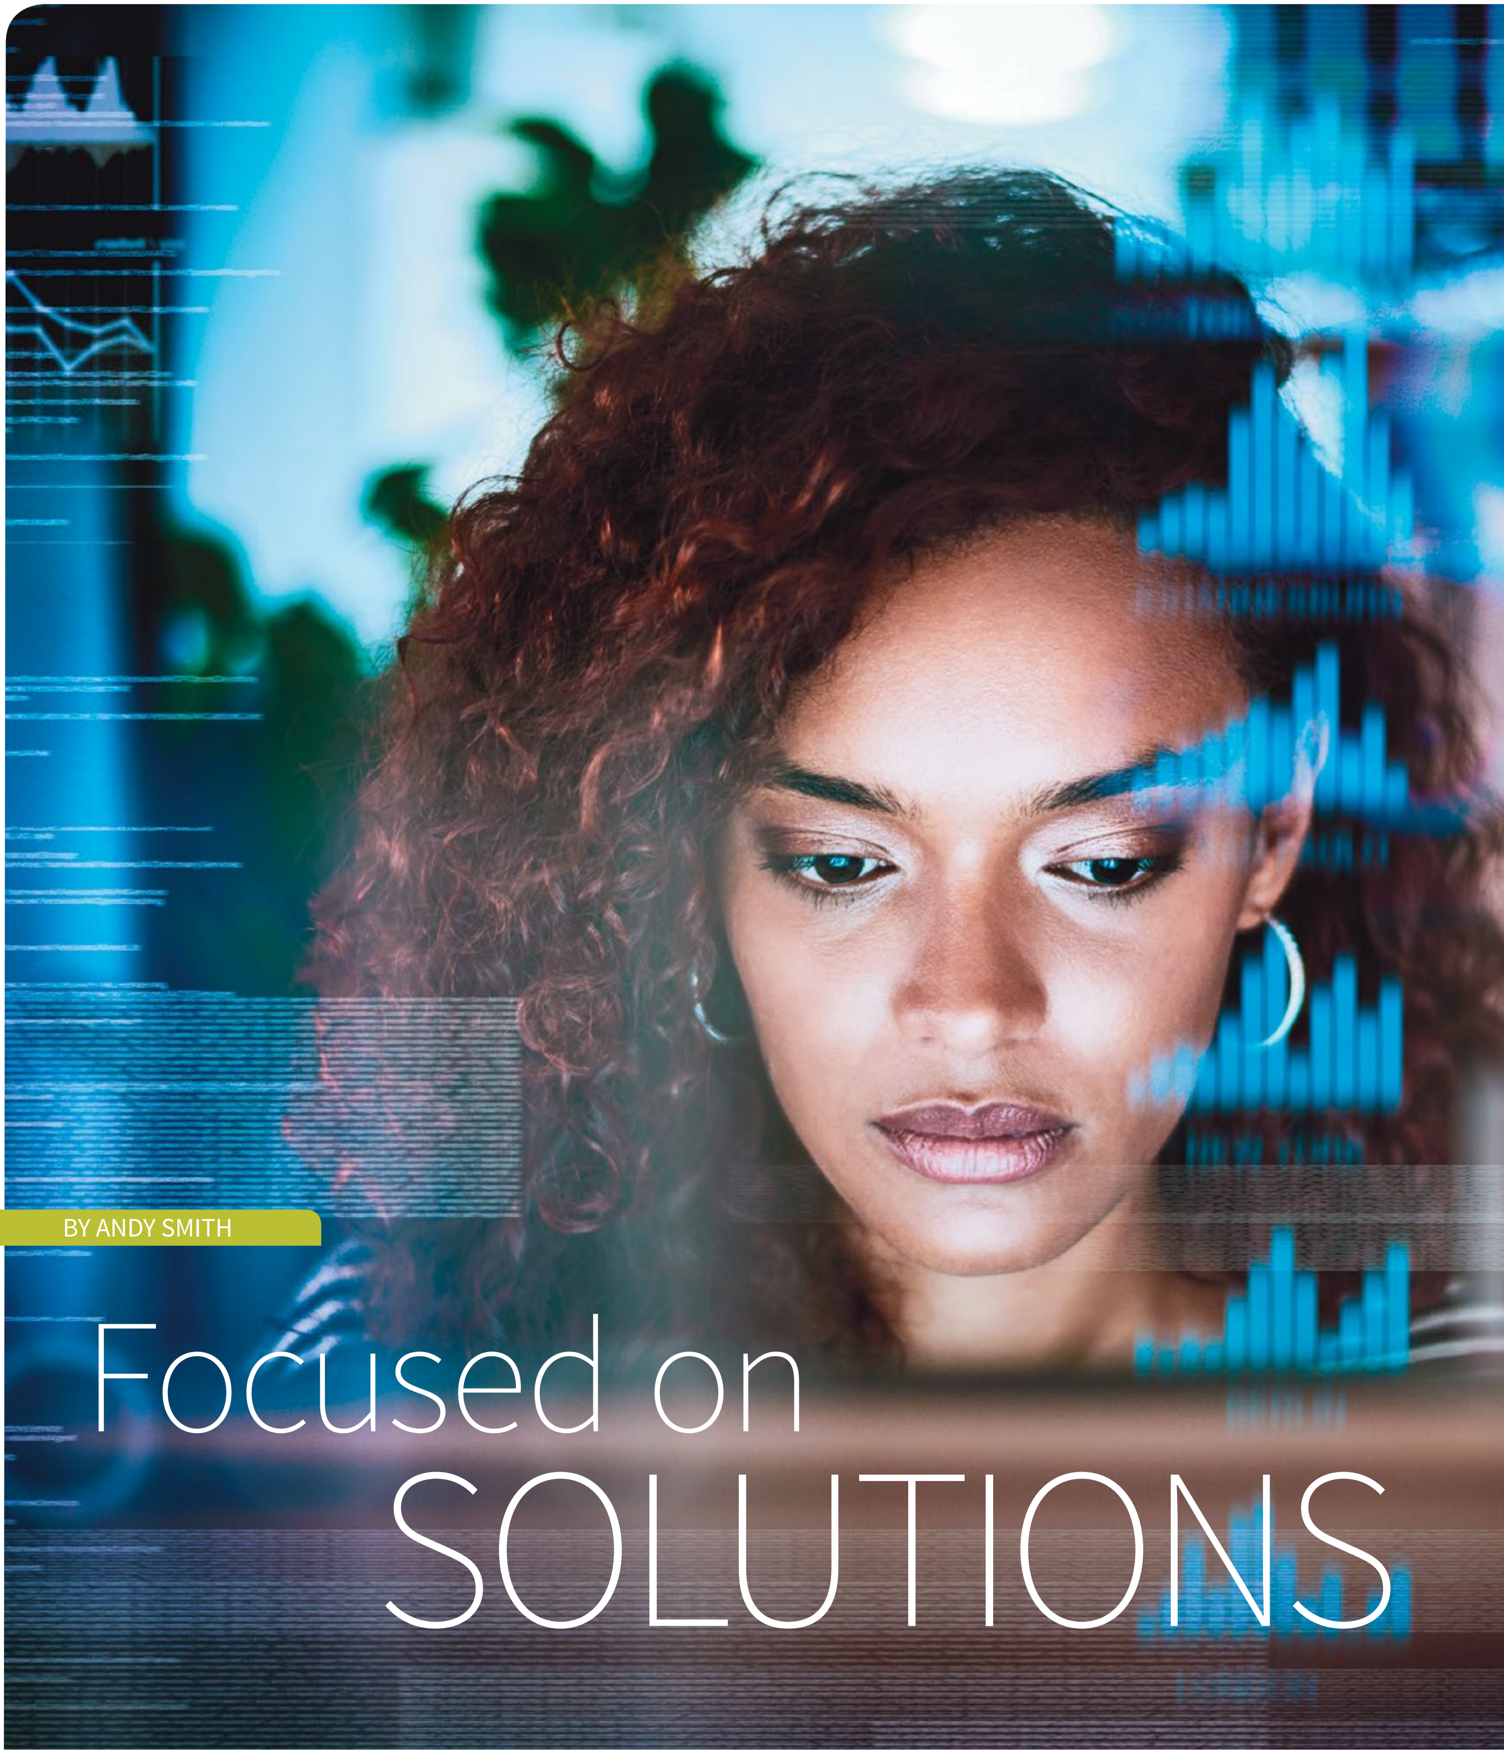

BY ANDY SMITH

Focused on SOLUTIONS

GETTY IMAGES/DELMAINE DONSON

How modern technology can help manage and run complex models

The biggest problem with maintaining a *good* model is that it is an awful lot of work. Within an actuarial model life cycle, there are several common challenges that seem to come up again and again. These challenges can be grouped into four buckets:

1. Calculation management
2. Data management
3. Execution
4. Results

But, as the saying goes, anything that is worth doing at all is worth doing well. So, this article will provide an introduction and short overview of a few modern technical solutions to help with common problems that arise in each of these four areas.

Calculation Management

Calculation management is the process of managing *what* a model does—as in, what mathematical calculations are being performed, in what order and what level of detail is associated with those calculations. The management of these calculations is often a large focus of model governance practices, and this is to ensure calculations are doing what they are supposed to and are being used appropriately. This is a key part of the new Actuarial Standard of Practice (ASOP) 56 on modeling that goes into effect in October 2020.

The management of a model's calculations can take on very different practical forms depending on the underlying modeling platform being used. A closed-box vendor platform may focus more on the specific configuration as well as the version of the vendor platform being utilized. Thus, there is often an intrinsic link between vendor development cycles and model management cycles. On the opposite end of the spectrum are home-grown models, where the calculations are developed and coded in-house (this includes self-built models contained in Excel workbooks). In these situations, the calculation management cycle will more likely be driven by business needs.

In both of these situations (closed-box and home-grown), the management of changes is critical for ensuring model updates are made appropriately.

Version-control Systems

Version-control systems, also known as revision control, track changes to files and documents. These systems provide workflows for the management of changes being made to files and documents, and they allow for the tracking of historical versions, as well.

Some commonly used examples include Amazon Web Services' (AWS') CodeCommit, Microsoft's Team Foundation Server, Git and Subversion—just to name a few. When properly used, they provide a complete trail of all changes made to a software system or set of files through the use of a common shared repository. These systems frequently are used in the software development world. However, they are not applied as often to modeling applications.

GETTY IMAGES/DELMAINE DONSON

Managing model data may be handled through version-control systems that track changes and versions of various assumptions. An added layer of complexity comes in when there are multiple sets of assumptions that feed into the same inputs within a model.

Integrating a version-control system with existing models can be difficult, depending on how models are organized and stored. Since these systems rely on files and file structures to manage changes, some types of files—most notably binary data files—are harder to maintain because the version-control system doesn't necessarily understand how to translate the binary data into meaningful business implications. The systems can still easily track different versions, but they may not provide meaningful insight into the nature of the changes.

This is also true for Excel workbooks, which are, by default, stored in either a binary or compressed format. Several companies now provide Excel-specific version control and tracking with similar capabilities to record different user revisions to formulas, data and workbook designs.

Data Management

Data management is the process of managing the data that is used to pass through the model calculations. It includes inputs to the model, assumptions used by the model, tables and other fixed external data sources, as well as the model results.

In many types of models, managing data also can be handled through version-control systems to track changes and versions of various assumptions. However, an additional layer of complexity comes in when there are multiple sets of assumptions that feed into the same inputs within a model. For example, you may have a best-estimate mortality assumption that is used for planning or pricing purposes, but you may have a different mortality assumption for valuation that includes conservatism. Some modeling systems might require these assumptions to be stored in the same *relative* file location of the base model, and this can cause problems when multiple versions of assumptions are needed. Good management of the model assumptions is key.

There also can be interdependencies among different assumption inputs that need to be managed as sets of data. For example, a complex dynamic lapse formula on an account value-based product may contain different sets of inputs for base-lapse vs. dynamic-lapse functions. But the full-lapse assumptions from a business perspective consist of a combination of all of these various inputs working in coordination with one another. This can make it more difficult to manage specific assumptions using just a version-control system.

Execution

Generally speaking, most people want models that:

1 | Produce accurate results.
2 | Use precise methods for all features.
3 | Run quickly.

But often we find ourselves forced to pick just two out of these three characteristics. As a model starts to become complex enough, runtime will almost inevitably become an issue. As actuaries, we have a tendency to seek greater amounts of information with ever-increasing layers of complexity and detail. This leads to longer-running models.

There are many different methods that can help with runtime, and they largely fall into three categories:

1 | Reduce model complexity.
2 | Reduce model data.
3 | Increase computational capacity.

Reducing model complexity means to reduce the number of calculations that need to be performed for each model segment, thus reducing the strain on hardware. Reducing model data includes various model-reduction techniques, including clustering, model segmentation and reduced projection periods. Increasing computation capacity includes solutions such as purchasing newer and more robust hardware, distributive processing and cloud computing. I will only address the last of these.

Over the past few years, it seems that computer processors are not getting much faster. Technology improvements that allowed the electronic circuits on microchips to be made increasingly smaller drove advancements in computing power for many years. The rate at which this occurred allowed processors to double in speed roughly every 18 months—this is known as Moore's Law. But today's microchip design is approaching the physical limitations of what is possible. Instead, newer chips include more cores, which are available for processing data in parallel to continue increasing computation power. This is why methods for distributive processing have increased.

Distributive Processing

Distributive processing provides mechanisms for spreading workloads across multiple sets of hardware for parallel processing of calculations. This comes in several forms of increasing complexity. The simplest—handled by most modern-day laptop and desktop computers—is multithreading. This takes several executable processes and splits the workload across the multiple cores, or threads, on a single processor microchip.

Grid computing takes this to the next level by splitting workloads across multiple computers connected via a closed network. These grids can be scalable, but because they are physically connected within a single company, the maintenance can be costly. Some grids also require special connections to manage the data movement among them, which can lead to complexity.

Cloud Computing

Cloud computing provides some of the answers to the problems with distributive processing. Cloud providers have a vast pool of resources available for computation. These providers have reached economies of scale with the hardware they provide, which is hard to compete with for even the largest companies.

Cloud technologies also have evolved over the years. At first, cloud servers were nothing more than an outsourced server space for hosting applications. They were not much different from owning your own computer, other than the fact that they were stored in someone else's server room.

Virtualization has allowed for a more flexible solution. Virtual machines (VMs) are an abstraction of hardware in a given computer. Multiple VMs can be set up on a single computer and share the same underlying set of hardware, allowing for separation of multiple independent servers without the need for each to have dedicated hardware. This allows for large, high-capacity servers to provide resources for multiple hosts simultaneously, allowing for higher utilization of the total available hardware.

The downside to a VM is there is a lot of redundancy in the setup. Each VM contains its own operating system, system files and drivers to interact with the hardware, which isn't necessary on a single shared machine. Containers have solved this problem. They act similarly to VMs, but they don't use that redundant overhead. A container can be set up so that it contains only the base set of code and specific required dependencies, which makes it much smaller. This allows for more operations to fit into a single set of hardware. One of the most common container systems used today is called Docker.

There are many container management systems that allow for management and deployment of containers across a large pool or grid of servers for dynamic management of resources. This includes systems like AWS Elastic Container Service (ECS) and, more recently, Kubernetes ("koo-burr-NET-eez"). These systems orchestrate a large number of smaller processes or resources and can effectively manage the distribution of this workload across the servers, regardless of the type of workload each individual container needs to perform.

Now, cloud providers also provide what they call "serverless offerings." This is, of course, a misnomer, as servers are still processing the underlying applications. The difference is in the additional abstraction from the user of *which* servers are being used. Some common examples include Azure Functions, Google Cloud Functions and AWS Lambda. In these systems, code is executed against large pools of available hardware, and there is no explicit setup to define a server, VM, container or other aspects of the underlying infrastructure. Instead, the cloud provider handles all of the details of managing it.

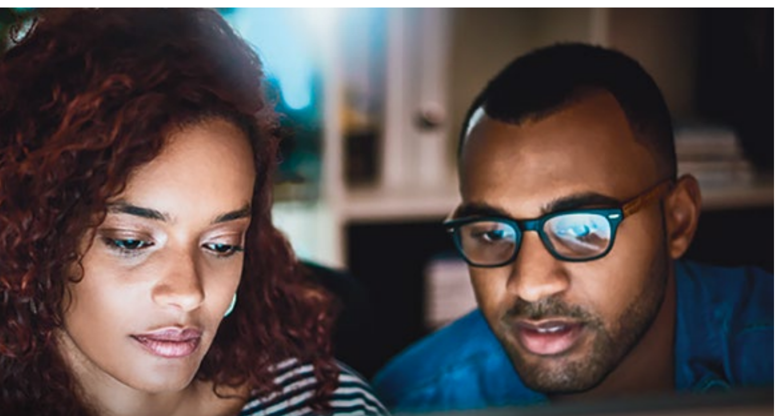

Business intelligence tools can help give insights into data by providing query and visualization tools that are easy to use and simple for end users to manipulate.

GETTY IMAGES/DELMAINE DONSON

Automation

Models are used for many purposes within organizations. They can have many upstream and downstream effects depending on how data flows into and through them, and is used downstream from them. This can result in a larger number of steps required to facilitate the modeling process.

Robotic process automation (RPA) is pretty much exactly what the name says. Robots (of the software-only kind) perform automated processes for you. If you have ever recorded and run a macro in Excel, you've built a very basic form of RPA: a clearly defined, repeatable process where the computer can perform a series of steps you've defined, often much faster than you could do on your own.

Modern RPA takes this a step further. The biggest problem with an Excel macro is that it lives inside one particular Excel file and has limited reach for what it can do, especially if you require steps outside of Excel. Modern RPA systems are installed on centralized systems, can be shared across multiple user processes and are software-agnostic, meaning they can run automated processes in your email system, customer relationship management (CRM) platform, Excel or your home-grown custom administration system equally well. Integrating an Excel macro that could do all of that would be quite a feat.

Currently, RPA systems generally are delegated to handling clearly defined, manual, repeatable processes. Most of these systems have not yet reached the maturity where they can take on complex decision-making processes. But increases in machine learning (ML) and artificial intelligence (AI) are moving them in that direction. Perhaps, sometime in the future, more complex tasks also will be delegated to the robots!

Results and Analysis

Understanding what comes out of your models is important—really important. In fact, it's most likely the reason why the model exists in the first place. However, getting to the truly useful information can sometimes be difficult.

Business intelligence tools can give insights into data by providing query and visualization tools that are easy to use and simple for end users to manipulate. Behind the scenes, these tools require some setup to connect to all of the data sources. But once done, and the data are loaded and available, they can be powerful tools for diving into the data, discovering important results, and conveying those results to others quickly and easily.

Multiple tools are available for providing business intelligence and data visualization. Some commonly used ones include Tableau and Microsoft's BI Query. Various visualization libraries are available for R and Python programmers. Using data visualization tools for analytics can provide more powerful insight into results.

Conclusion

These are just a few of the technologies that companies already use to manage some of the common problems with managing and running models. Incorporating some of these technologies into your modeling processes can provide more efficient model management, execution and analysis. ■

ABOUT THE WRITER

ANDY SMITH, FSA, MAAA, is co-founder and CEO of Slope Software. He can be reached at ***Andy@slopesoftware.com***.

Persevere With Passion

Q&A with Rolande S. Mbatchou, ASA, MAAA, managing actuary, Provider & Network Data Science, at Health Care Service Corporation

How did your upbringing influence your career choice?

My dad has a Ph.D. in economics, and my mom has a master's degree in science education. They encouraged my siblings and me to excel in mathematics, because they believe mathematics is the foundation of all sciences. So, growing up, my dad used to make us solve math problems all the time. My dad even installed a giant chalkboard in one of our living rooms. I still remember spending hours solving math problems for fun on that chalkboard.

Math became a passion of mine at a very young age. Another amusing memory from when I was in elementary school was that I used to play with my younger sister and pretend I was the math teacher. She would sit down and listen to me teach her math. I am not sure if she enjoyed it then, but she eventually got a Ph.D. in statistics, so I think my lessons were somehow helpful!

The strong math foundation I had as a child greatly influenced my career choice. But that was not the only thing—I also loved business. I started my first popsicle business at age 10. I created the mix of water with syrup and sugar, and I had my siblings and cousins work for me for free to help me sell the popsicles in our neighborhood. My business skills were not sharp then, but I knew early on that I wanted to work with numbers in a business setting.

Why and how did you become an advocate for diversity in STEM fields?

I was about 12 years old when I realized there was a diversity gap in STEM fields. I grew up in a predominantly black suburb of Paris. The schools in my neighborhood were known to always be on strike and had a poor level of STEM education. My parents, who were firm believers in STEM education, paid for my siblings and me to attend a private school in downtown Paris with a strong science track. We were some of the few black students in the school. By the time I got to high school, I was admitted into the science track, and I was the only black student on that track for three full years. The [demographic] statistics were not any different when I moved to Chicago for college to study actuarial science, or when I started working as an actuary. It seemed this was the norm—but I knew I would one day do something about it. It was clear to me that if my neighborhood friends had the same STEM career exposure, mentorship and opportunities I had, they would have succeeded in STEM fields.

PHOTO: JOEL MAISONET

Rolande S. Mbatchou can be reached at ***Rolande_S_Mbatchou@bcbsil.com***.

ONLINE VIDEO EXCLUSIVE!
Visit *TheActuaryMagazine.org/Persevere-With-Passion* to watch a video with Rolande S. Mbatchou.

> ***If you are discouraged because you failed an exam, do not give up! Get up and try again! Learn from those who were able to pass successfully and restrategize.***

In college, I was a champion for diversity as I tutored middle school students in math in Austin, a Chicago neighborhood traditionally underrepresented in STEM. Fast-forward to 2018, and I was leading multiple initiatives as an actuary at my company to recruit more black actuarial talent. I introduced the International Association of Black Actuaries (IABA) to actuarial leaders and recruiters, and I made sure we attended the IABA Annual Meeting both in 2018 and 2019. I also reached out to historically black colleges and universities (HBCUs) with actuarial science programs to assess if we could partner with them. We decided to focus our efforts on Morgan State University due to the larger actuarial student pool. I also connected actuarial leaders with key partners in my company's Diversity and Inclusion team, and we partnered with them at multiple STEM events targeted to public high school students in Chicago. Currently, I mentor high school and college students to encourage and inspire them to pursue careers in STEM, and I am very engaged as an advocate for diversity.

Who has inspired you in your career?

The people who have inspired me the most are those who have believed in me. That includes my parents, siblings, family, close friends, high school principal, university professors and mentors (both professional and personal).

I recall a time as a teenager when I wanted to drop out of high school. I think I was suffering from an identity issue (part of it was that I did not like being the only black woman in my class), so I wanted to attend one of my neighborhood schools instead. The principal called me to her office and, instead of scolding me, she encouraged me. She told me that I was very smart and that if I persevered and worked hard enough, I would be able to inspire others through my success. I never forgot the words she spoke to me. Even in my actuarial career, there were many times when I wanted to give up—especially due to the exam process—but the words of encouragement from those who believed in me echoed in my head, telling me to keep pushing and to press on. That is why I believe so strongly in mentorship, and that is why I am trying to be a mentor to other students.

What excites you about your job?

My current actuarial rotation at Health Care Service Corporation (HCSC) is within the Provider Data Science team. What excites me the most is that we are solving real problems, and I can clearly see the impact of my work. Currently, I am working with doctors to create an emergency room (ER) predictive model for accountable care organizations (ACOs) that will help them better manage the health of patients who are most likely to end up in the ER. Managing ER patients has been a real cost challenge for many provider groups, and my team is building a solution to help them. I use my knowledge of health insurance (learned through past actuarial rotations and exams), data analytics and now machine learning (ML) to build valuable tools for our customers. That is very exciting!

What innovative techniques make you look forward to changes in the actuarial profession?

I think it is important that actuaries learn the benefits of ML and familiarize themselves with the different types of ML models (supervised, unsupervised and reinforcement learning). ML is not for every use case, but when applied correctly, it has great predictive power due to its ability to more accurately learn from data compared to traditional actuarial models.

The Society of Actuaries (SOA) added a Predictive Analytics exam, which is a plus for actuarial students. An actuary colleague and I started an "Automation & Machine Learning" group at our company. We periodically invite employees to present on how they automate their work using tools such as Python, and share how they successfully built ML algorithms to solve health care problems. I look forward to seeing more actuaries use ML for their predictive work and tools such as Python to more effectively automate their work processes.

What would you say to encourage anyone interested in becoming an actuary?

Perseverance is an important virtue! Becoming an actuary requires a lot of perseverance, mostly due to the rigorous exam process. Nevertheless, looking back, with one exam left to get my fellowship, I can truly say it was all worth it. I learned so much through these exams, and they helped me build confidence as an analytical professional. I am not afraid to tackle some of the most challenging work problems because I keep telling myself that they cannot be harder than actuarial exams. So, if you are discouraged because you failed an exam, do not give up! Get up and try again! Learn from those who were able to pass successfully and restrategize.

I also think it is important to educate yourself about the different tracks an actuary can pursue. As a freshman in college, I read the entire *beanactuary.org* website. I took two actuarial exams my sophomore year and obtained an actuarial internship my junior year. During my time at DePaul University in Chicago, I founded the Actuarial Club and invited speakers from different companies, and I held actuarial study sessions. I was creating opportunities around me to ensure that I would secure a job and pass the exams.

I also attended actuarial conferences and events to network with professionals in the field, which provided me with opportunities to ask them a lot of questions. During my internship at Allstate Financial, I scheduled lunches with all of the actuarial directors in the company—no one rejected my invitation; they were all willing to give me professional advice.

Preparation was key for me, and I was determined to learn and get inspired. I encourage every actuarial student to reach out to actuaries. Do not be afraid to attend actuarial events and conferences, and network with other professionals. It will help you build your confidence and grow your communication skills, which helps tremendously when interviewing for jobs.

PHOTO: JOEL MAISONET

What attributes make a good leader?

There are many attributes that make a good leader, such as communication, confidence, commitment, accountability and creativity. But my top characteristics are integrity and the ability to inspire others.

Former U.S. President Dwight Eisenhower once said, "The supreme quality of leadership is unquestionably integrity." I think it is important—whether in the workplace or wherever one leads—that people can trust their leader's judgment, actions and guidance. To me, integrity means that we act—even when no one is watching—the same way we tell others we are acting when they are not watching. This applies to an organization, as well. Do the employees feel the leaders are truly abiding by the corporate values, as marketed to external stakeholders? This is a critical question that addresses corporate integrity. It is important for good leaders to question the impact of their actions before making any critical decision, and to assess whether their decisions are aligned with the values of the organization.

Another important attribute of a good leader is the ability to inspire others. We all have been inspired by someone who took the time to mentor and invest in us. A good leader should inspire others to believe in themselves and in their ability to reach their own goals and aspirations and achieve excellence in all that they do. Great leaders make themselves available for others to learn from their experience and open doors of opportunity for others whenever they can.

What is your definition of success?

Success is the journey of the discovery of one's purpose—as it aligns with one's skills and interests—and the ability to take critical steps toward achieving that purpose. ■

The New Guide to SOA Exams

BY STUART KLUGMAN

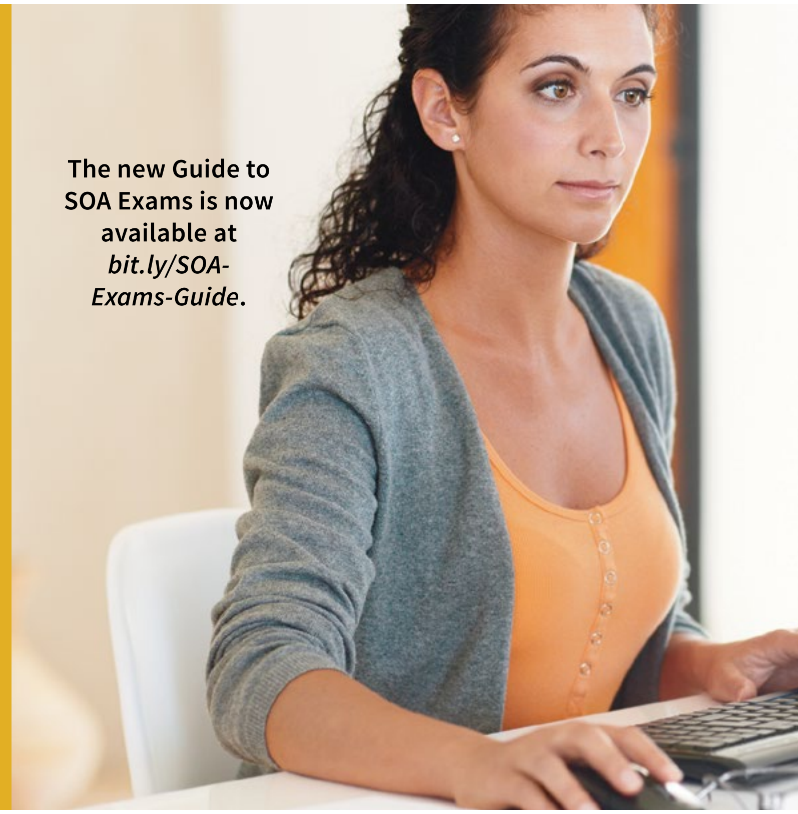

The new Guide to SOA Exams is now available at *bit.ly/SOA-Exams-Guide*.

ISTOCK.COM/GLOBALSTOCK

Have you ever wondered how Society of Actuaries (SOA) exams are constructed? Or the difference between a question that asks for a recommendation vs. one that asks for an analysis? Or how pilot questions are used on multiple-choice exams? Or how pass marks are set? Or what goes into a model solution?

The answers to these questions and more are in the new Guide to SOA Exams, now available to access from many of the education webpages at *SOA.org*. While you will need to read the full guide to get all of the answers, this article will explain why it was produced and reveal one change we have made.

For years, the SOA education webpages contained numerous articles covering various aspects of the candidate experience—from exam development and grading to tips for success and disciplinary procedures. An initial attempt to consolidate information resulted in the previously published Guide to Written-Answer Examinations. That document had an exclusive focus on fellowship written-answer exams, and was later expanded to cover the written-answer portion of what is now the Long-Term Actuarial Mathematics (LTAM) exam.

In 2019, it was decided to add other relevant information to the guide, in particular, coverage of multiple-choice exams and e-Learning assessments. For candidates, those interested in becoming candidates and those who work with candidates, the new Guide provides a single point to access information. While the Guide is designed to be a comprehensive document, it is important to note that some documents must remain separate (though they are described in the Guide). Examples are the Code of Conduct for Candidates and the Confidentiality and Discipline Procedures for Computer-Based Testing for Candidates.

With the publication of the Guide, we are announcing a new way of reporting results for candidates who are unsuccessful on written-answer examinations. Prior to 2020, failing candidates received question-by-question feedback on a 0–10 scale that was meant to be interpreted in the same way as the 0–10 scores given for the entire exam. Such a scale requires a passing score for each question. Because the exam committees do not set pass marks for each question, we had to infer

the pass mark for each question. This led to candidate misinterpretations and confusion.

Beginning with the 2020 written-answer exams (both fellowship and the associateship LTAM and Predictive Analytics exams), a new form of reporting will be used. For each question, candidates will be provided the percentile rank of their score. This will make it immediately obvious how a candidate performed relative to others on each question. It is important to note that a weighted average (by exam points) of the percentiles will not yield the overall percentile rating.

We want this document to be as informative and useful as possible. Please feel free to email me with any questions or comments about the Guide. ■

ABOUT THE WRITER

STUART KLUGMAN, FSA, CERA, is senior staff fellow, Education, at the Society of Actuaries. He can be reached at ***sklugman@soa.org***.

Notice of Disciplinary Determination

On July 18, 2019, the Society of Actuaries (SOA) convened a Discipline Committee to review a matter referred by the Actuarial Board for Counseling and Discipline (ABCD) related to the conduct of Michael W. Frank, FSA, MAAA, EA. The Discipline Committee concurred with certain findings of the ABCD that Mr. Frank materially violated the *Code of Professional Conduct* (*Code*). The Committee felt that Mr. Frank's work product did not meet reasonable standards that would be expected from a member of the SOA. The Discipline Committee further determined that discipline is warranted, and that Mr. Frank should be suspended from SOA membership for two and one-half years, after which Mr. Frank may pursue readmission.

Mr. Frank is a sole practitioner who provides actuarial services to private defined benefit plans through third-party administrators. This matter arose from litigation involving improper Internal Revenue Service (IRS) filings made by one of Mr. Frank's clients. Although Mr. Frank was dismissed from the lawsuit, it uncovered numerous deficiencies in actuarial services provided by Mr. Frank, resulting in this disciplinary action.

Findings Regarding Precept 1 of the *Code of Professional Conduct*

The Discipline Committee concluded that Mr. Frank materially violated Precept 1[1] of the *Code of Professional Conduct* because he did not exercise proper skill and care when preparing governmental filings. He knowingly prepared, used and reported false information on matters relating to employee benefit plans and actuarial services, including backdating governmental filings, so as to reduce the likelihood of others detecting problematic figures presented on the filings and avoid potentially significant IRS penalties. When presented with inconsistent plan information from his principal, Mr. Frank failed to reconcile and resolve those inconsistencies, and he neglected to make necessary modifications to his valuations upon obtaining updated information. Mr. Frank did not provide certain actuarial services necessitated by federal regulation, such as certifying the plan's Adjusted Funding Target Attainment Percentage (AFTAP). Mr. Frank was negligent in his professional obligations by allowing his principal to direct and limit his work to exclude responsibilities required under the law. When he allowed these limitations, Mr. Frank failed to inform his principal of the consequences of not obtaining an AFTAP certification and, for at least one year, completed a valuation that failed to consider those consequences. The Discipline Committee concluded that these material errors collectively demonstrated an underlying lack of skill and care in Mr. Frank's work as an actuary.

Findings Regarding Precepts 3 and 4 of the *Code of Professional Conduct*

The Discipline Committee agreed with the ABCD that Mr. Frank materially violated Precepts 3[2] and 4,[3] which require an actuary to observe applicable

Continued on page 50

DISCIPLINE NOTICE

Continued from page 49

Actuarial Standards of Practice (ASOPs). Mr. Frank did not follow ASOP #23 when he failed to question inconsistencies in the data provided by the principal or disclose those limitations and their implications. Proper resolution of the inconsistencies may have had a material effect on the actuarial valuations provided by Mr. Frank.

Further, the Discipline Committee agreed that Mr. Frank violated ASOP #41 in his failure to disclose necessary information regarding the defined benefit plan in question. Mr. Frank was negligent in his professional obligations by allowing his principal to direct and limit his work to exclude disclosure and reporting responsibilities required under the ASOPs. Moreover, Mr. Frank admitted to not having read ASOP #41 prior to his ABCD hearing.

Conclusion

The Discipline Committee recognizes that Mr. Frank faced a difficult situation, particularly as a sole practitioner. The circumstances of this case highlight the need for clear and appropriate terms of engagement, strong processes, and a complete understanding of the law and regulations. This case also demonstrates the importance of obtaining peer review of one's work. Peer review helps to ensure relevant issues are addressed, work is completed in compliance with actuarial standards of practice, and supports maintaining the appropriate level of knowledge and understanding of the substantive and ethical obligations of the actuarial profession. These obligations exist for all credentialed actuaries.

Mr. Frank demonstrated a lack of appropriate professional judgment when he accepted an assignment that limited his ability to comply with federal regulations and ASOPs and for which he was not supplied sufficient, timely or accurate plan data. Mr. Frank also seems to have lacked the necessary tools and/or knowledge to complete his assignment, even if provided with sufficient and timely data. His willingness to allow his principal to direct his services in a manner which undermined his responsibilities under the *Code of Professional Conduct* and under the law materially impacted his professional actuarial obligations. Therefore, the Discipline Committee has determined that discipline is warranted, and that Mr. Frank be suspended from SOA membership for two and one-half years.

All members of the SOA are reminded of their responsibility to follow the *Code of Professional Conduct*. Members are also reminded that when they are faced with potential issues regarding professional conduct, the ABCD is available for counseling.

References

[1] **PRECEPT 1.** An Actuary shall act honestly, with integrity and competence, and in a manner to fulfill the profession's responsibility to the public and to uphold the reputation of the actuarial profession.

[2] **PRECEPT 3.** An Actuary shall ensure that Actuarial Services performed by or under the direction of the Actuary satisfy applicable standards of practice.

[3] **PRECEPT 4.** An Actuary who issues an Actuarial Communication shall take appropriate steps to ensure that the Actuarial Communication is clear and appropriate to the circumstances and its intended audience and satisfies applicable standards of practice.

Special Webcast Offer
for Groups Working Remotely

The SOA is committed to fulfilling your team's professional development needs during this unprecedented time. With most organizations working remotely, we want to make it easier than ever to access our webcasts. We have temporarily eliminated the need to purchase additional links for remote viewing.

How do I take advantage of this opportunity?

- Choose a webcast from our wide range of available topics
- Select the 3+ group purchasing option at checkout
- Use your organization's technology to broadcast SOA webcasts to those working remotely

This option will remain in place until social distancing guidelines are lifted by the CDC and other governmental authorities.

Check Out Available Webcasts at *SOA.org/Webcasts*

SOCIETY OF
ACTUARIES®