



SOCIETY OF
ACTUARIES

TECHNOLOGY
SECTION

CompAct

ISSUE 63 • OCTOBER 2019



Spreadsheet Controls Add Risk Resilience—Part Two

By Diane Robinette and Leslie Martin

Page 6

3 Letter From the Editors

*By Hugh Lakshman and
Bill Cember*

6 Spreadsheet Controls Add Risk Resilience—Part Two

*By Diane Robinette and
Leslie Martin*

8 Your Moment of Zen: An Article on Writing Articles

By Mitchell Stephenson

9 Techniques for Taming Model Risk

*By Bill Cember, Scott Houghton
and Dylan Strother*

14 Blockchain—Is the Newest Kid on the Block Going to Be All Right?

By Helen Duzhou and Jeff Guo

18 Predictiveness vs. Interpretability

*By Kimberly Steiner and
Boyang Meng*

24 Risk Management in the Digital Age

*By Rob Ceske, Kelly Combs,
Nadim Hraibi and Jamie Hooten*

29 TCARE: Unlocking the Power of Psychosocial Data

By Michael Mings

CompAct

2019 SECTION LEADERSHIP

Issue 63 • October 2019

Published twice a year by the
Technology Section of the
Society of Actuaries.

475 N. Martingale Road, Suite 600
Schaumburg, Ill 60173-2226
Phone: 847.706.3500 Fax: 847.706.3599
www.soa.org

This newsletter is free to section mem-
bers. Current issues are available on the
SOA website (www.soa.org).

To join the section, SOA members
and non-members can locate
a membership form on the
Technology Section webpage at
[http://www.soa.org/sections/
/technology/technology-landing/](http://www.soa.org/sections/technology/technology-landing/).

This publication is provided for informa-
tional and educational purposes only.
Neither the Society of Actuaries nor the
respective authors' employers make any
endorsement, representation or guar-
antee with regard to any content, and
disclaim any liability in connection with
the use or misuse of any information
provided herein. This publication should
not be construed as professional or
financial advice. Statements of fact and
opinions expressed herein are those of
the individual authors and are not neces-
sarily those of the Society of Actuaries or
the respective authors' employers.

Copyright © 2019 Society of Actuaries.
All rights reserved.

Publication Schedule

Publication Month: April 2020
Article Due: Feb. Feb. 3, 2020

Officers

Sean Hayward, FSA, MAAA, Chairperson
Ravi Bhagat, FSA, MAAA, Vice Chairperson
Robert Maxwell, ASA, MAAA, Secretary
Paul San Valentin, FSA, CERA, MAAA, Treasurer

Council Members

William Cember, FSA, MAAA
Han (Henry) Chen, FSA, ACIA, MAAA
Shyam Kolli, FSA, MAAA
Frank Reynolds, FSA, FCIA, MAAA
Garrett Shafer, ASA, MAAA

Newsletter Editors

William Cember, FSA, MAAA
william.cember@prudential.com

Hugh Lakshman, FSA, MAAA
hugh.lakshman@ibx.com

Program Committee Coordinators

Paul San Valentin, FSA, CERA, MAAA
2019 Valuation Actuary Symposium Coordinator

Han (Henry) Chen, FSA, ACIA, MAAA
2019 Life & Annuity Symposium Coordinator

Shyam Kolli, FSA, MAAA
2019 Health Meeting Coordinators

Han (Henry) Chen, FSA, ACIA, MAAA
2019 SOA Annual Meeting & Exhibit Coordinator

SOA Staff

David Schraub, FSA, CERA, AQ, MAAA, Staff Partner
dschraub@soa.org

Jane Lesch, Section Specialist
jlesch@soa.org

Julia Anderson Bauer, Publications Manager
jandersonbauer@soa.org

Sam Phillips, Staff Editor
sphillips@soa.org

Erin Pierce, Senior Graphic Designer
epierce@soa.org

Letter From the Editors

By Hugh Lakshman and Bill Cember

Welcome to the second issue of *CompAct* for 2019! First we wanted to thank you all for all the feedback on the spring issue of our newsletter. We continue to use your feedback to add to and improve *CompAct*.

We would also like to continue to encourage you, our readers, to submit articles or topics of interest for future publications. We really appreciate the contributions and feedback. We can be reached at hugh.lakshman@ibx.com and william.cember@prudential.com.

In this issue of *CompAct*, we have seven articles that cover a diverse range of topics. We even have an article on writing articles!

“SPREADSHEET CONTROLS ADD RISK RESILIENCE—PART 2”

Every actuary uses spreadsheets, but how do you measure and manage the size and complexity of your spreadsheets and the formulas they contain? In part two of this two-part series, Diane Robinette, CEO of Incisive Software, and Leslie Martin of Safety National offer additional insights on how to leverage technology to help solve some of these common data integrity issues in spreadsheets.

“YOUR MOMENT OF ZEN: AN ARTICLE ON WRITING ARTICLES”

Has this ever happened to you? You have a topic that you know would be of interest to our readers, but you are unsure how to go about turning it into something you can submit for publication. Guest writer Mitchell Stephenson of the Leadership & Development Section provides practical advice on writing your next article.

“MODEL STEWARDSHIP AND TECHNOLOGY”

Our actuarial models continue to evolve in increasingly complex directions. Scott Houghton of Valani Global, Dylan Strother of



Oliver Wyman, and our own editor Bill Cember of Prudential walk through techniques that can help manage model risk and complexity.

“BLOCKCHAIN—IS THE NEWEST KID ON THE BLOCK GOING TO BE ALL RIGHT?”

So, what exactly is Blockchain and how is it applicable to insurance? Helen Duzou of Oliver Wyman explains what blockchain is and provides examples of how it can be applied in an insurance setting. She also includes a very helpful decision tree for when to consider using Blockchain.

“PREDICTIVENESS VS. INTERPRETABILITY”

Kimberly Steiner and Boyang Meng of Willis Towers Watson discuss the tradeoff between predictiveness and interpretability in model selection. This article won Best Eligible Paper in the call for essays on the theme of “Risks Posed by Predictive Models” that was jointly sponsored by our section, in conjunction with the Predictive Analytics and Futurism sections of the Society of Actuaries (SOA).

“RISK MANAGEMENT IN THE DIGITAL AGE”

Advances in technology are allowing organizations to automate processes that previously required human intervention, but have

our risk management strategies adapted to this? Rob Ceske, Kelly Combs, Nadim Hraibi, and Jamie Hooten of KPMG discuss the implications of automation on risk management and ways that organizations can start thinking about adapting to these changes.

“UNLOCKING THE POWER OF PSYCHOSOCIAL DATA”

As our population continues to age, more people find themselves in the role of caregiver. Michael Mings of Tailored Care discusses a model for evaluating caregiver burnout risk and discusses potential solutions.



Bill Cember, FSA, MAAA, is a director and actuary with Prudential Financial. He can be reached at william.cember@prudential.com.



Hugh Lakshman, FSA, MAAA, is a director and actuary with Independence Blue Cross. He can be reached at hugh.lakshman@ibx.com.



LIVING to 100

SOCIETY OF ACTUARIES
INTERNATIONAL SYMPOSIUM

Jan. 13–15, 2020
Orlando, Florida


Register Today

Registration for the 2020 Living to 100 Symposium is now open. This prestigious event brings together thought leaders from around the world to share ideas and knowledge on increasing lifespans. Expert presenters will explore the latest longevity trends, share research results and discuss implications of a growing senior population.

New this year are teaching sessions that will provide practical pointers to help actuaries measure and forecast mortality at advanced ages.

Symposium speakers include:

- Steve Horvath, Professor of Human Genetics and Biostatistics for the David Geffen School of Medicine at University of California, Los Angeles
- Jacquelyn B. James, Director of the Boston College Center on Aging & Work and the Sloan Research Network on Aging & Work
- Ronnie Klein, FSA, MAAA, Director of the Global Ageing program at The Geneva Association



Visit LivingTo100.SOA.org for more information

Spreadsheet Controls Add Risk Resilience— Part Two

By Diane Robinette and Leslie Martin

In part one of this two-part series on spreadsheet risk resilience, we discussed how Excel spreadsheets, despite years of rumors predicting their demise, continue to be the go-to tool that actuaries use to get the job done. Common causes of data integrity issues relative to spreadsheets were spotlighted, along with insight into controls to help actuaries solve these issues. Part two of this series offers a more detailed look at these issues, using a real-world problem/solution example related to spreadsheet risk, with a special focus on the insurance industry.

MODEL SIZE AND COMPLEXITY

Actuaries are frequently tasked with creating sophisticated models to assist with risk calculations, valuations and pricing. The number of workbooks, spreadsheets, formulas and coding vary depending on the complexity of the model. Actuaries generally think of model size and complexity in terms of the number of worksheets or file size, primarily because these metrics are the most readily available. As the actuarial pricing team at Safety National learned from researching best practices in model governance and solutions for spreadsheet version control, there are many additional metrics for evaluating the size and complexity of a model. Examples include:

- Number of populated cells
- Number of formulas
- Number of unique formulas
- Number of external references
- Number of formulas with nested IF statements

Having these metrics changes one's perspective on how to manage spreadsheet risk effectively and helps explain the practical limitations of trying to review changes manually.

In the case of Safety National, the actuarial team created a model for underwriters to use in pricing business. Given the large size of the model and small size of the pricing division, the team needed an efficient way to review and test model changes. While the model has a relatively small file size, learning the number of

populated cells (about 1.5 million) and the number of formulas (about 65,000) highlighted the challenges of trying to review all of the revisions manually.

INSIDE THE NUMBERS

Actuaries either create or work on hundreds of spreadsheets in a single year. While this number may not seem overwhelming, ponder the following. A single Excel worksheet can have 1 million rows and more than 16,000 columns. Again, that is worksheet, not workbook. In addition, the size of a workbook is only limited by memory and system resources.

Another jaw-dropping Excel spreadsheet stat is that the length of a formula can be as high as 8,129 characters. That is about the same as the average length of a paragraph. Editing and troubleshooting these long formulas can quickly become a challenge because they will wrap a few times. As a work-around, users typically copy and paste the large formula into a notepad, apply indentation for readability, and then copy and paste it back into Excel. Beyond the sheer amount of time it takes to manually cut and paste each formula, it is very easy for errors to be introduced during this process.

Other interesting stats include such things as cross-worksheet dependency, in which 64,000 worksheets can refer to other sheets. Keeping track of which worksheets link to the active worksheet and whether that worksheet contains formulas that reference other sheets can quickly become an overwhelming, if not impossible, task when relying on manual methods. Do not forget that four billion formulas can depend on a single cell. That is a lot of information.

There is great value in deploying software solutions that allows users to easily drill down into complex Excel formulas.

LEVERAGING TECHNOLOGY FOR EFFICIENCY

While the numbers cited above may seem extreme, users do not need to get anywhere near the limits for spreadsheets to become unwieldy. In an effort to manage risks, many actuaries turn to the troubleshooting capabilities within Microsoft Excel. Unfortunately, while adequate for the casual user, these tools are not enough for power users—like actuaries—who often encounter a single formula that is more than 8,000 characters.

Actuaries require advanced troubleshooting solutions that make the processes of monitoring and managing spreadsheets significantly easier. As the team at Safety National quickly realized,



there is great value in deploying software solutions like Incisive's Xcellerator that allows users to easily drill down into complex Excel formulas, visually see when formulas are different among contiguous cells, and run a comparison between two workbooks. Necessary capabilities for Excel power users include performing quick scans of a spreadsheet for cell references and properties that are not in the working area, cells having different formulas pointing to them, and the ability to highlight potential errors nested within complex formulas. This is particularly helpful when working on large and complex models like the one created at Safety National. In addition to reducing risk exposure, the accuracy achieved using automated spreadsheet risk management software helps lay the foundation for good business decisions.

Ensuring the integrity of spreadsheet data is an arduous task. Oftentimes a search for a better solution does not begin until an issue occurs or an internal audit identifies the need for more stringent controls around the actual versioning of a spreadsheet. In the meantime, actuaries spend countless hours trying to detect hidden errors that might exist. Rather than implementing

these tools because it is required, actuaries should take a proactive approach. In addition to greatly reducing risk and exposure, spreadsheet risk management technology enables actuaries to do their job more efficiently and with a higher degree of accuracy. Spreadsheet controls also add risk resilience, a state in which actuaries are able to quickly iterate processes in a way that boosts flexibility and agility, no matter what changes occur. ■



Diane Robinette is CEO of Incisive Software. She can be contacted at diane@incisive.com.



Leslie Martin, ACAS, MAAA, is actuary-LPT and WC Reinsurance Pricing at Safety National Casualty Corp. She can be contacted at leslie.martin@safetynational.com.

Your Moment of Zen: An Article on Writing Articles

By Mitchell Stephenson

Have you ever been interested in writing articles for an actuarial publication but not sure where to start? Do you think it's not something you have the skills to do? Well, it may not be as hard as you think. Many articles written on forums such as LinkedIn, or in publications like *CompAct* and *The Stepping Stone*, follow a basic cadence, which, if you see it enough, can help guide you to publish your own work. Once you've got the cadence, you only need a good idea to get you started.

If you're on the fence about contributing to a future publication, consider this quote from Sir Arthur Conan Doyle, the creator of Sherlock Holmes, regarding writing: "Anything is better than stagnation." Or perhaps this gem from W. Somerset Maugham, a British playwright, novelist and short story writer, who said: "There are three rules for writing a novel. Unfortunately, no one knows what they are." Fortunately, there are some guidelines you can follow if you want to write a short article. Here are some of them:

- Capture the reader's attention in the opening of the piece through a question, or statement they may find relevant. Many of today's news and publications are read online, and it is helpful to create an opening sentence that will draw people in, or make them want to click on it.
- Follow your opening with a fact, some statistics, or a quote from a relevant source to add credibility to your piece. It is always good, especially for an audience of actuaries, to provide some background and evidence about how you are presenting your case and drawing conclusions.
- Give some clear, tangible steps for addressing the problem you identified in the beginning of the piece, in the form of three to five bullets or short statements. These should each stand alone as a separate piece of advice, guidance or supporting evidence of the main conclusion you are looking to draw. If you create too many bullets, the reader may lose interest, and if you include too few, he or she may decide the piece is not substantial enough.



- Finish with a solid conclusion that briefly summarizes the context, ties it together, and gives the reader confidence that he or she will walk away having read something that addressed the initial opening statement. Make the reader feel optimistic that the problem you identified at the onset can be addressed through the outlined steps.

If you follow this approach, you can take almost any topic of interest and turn it in to a brief article. Keep it simple, focused and on point. You may find that the words flow more easily than you suspected or that you already knew what you wanted to say and only needed to organize it. Most importantly—just try. In the words of Harry Potter author J. K. Rowling, "You, yourself, will never rest until you've tried!"

To submit an article for *CompAct* or *The Stepping Stone*, contact Jane Lesch at jlescb@soa.org. ■



Mitchell Stephenson, FSA, MAAA, is the chair of the Leadership & Development Section Council and can be reached at mitchell.stephenson@prudential.com.

Techniques for Taming Model Risk

By Bill Cember, Scott Houghton and Dylan Strother

By the end of 2022, many major accounting regimes commonly used by actuaries will have significantly changed. Most of these changes are material and will increase the complexity of models needed to calculate actuarial-related balances. New regimes like Principle-Based Reserving (PBR), GAAP LDTI, and IFRS 17 require detailed cash flow projection models, often with multiple assumption sets that need to be updated frequently, increasing model complexity and risk.

As actuaries begin to use more-complex models, it's not enough just to have the right risk governance in place. The right infrastructure will also help reduce risk by making models easier to maintain and manage. The following techniques will help build a reliable infrastructure and increase model management:

- Modular design
- Model consolidation
- Model documentation
- Modeling roles & model change management

BUT AREN'T OUR MODELS ALREADY REALLY COMPLICATED?

For products like variable annuities, this type of modeling and the associated model risk has been part of reserve calculations for years, and models supporting these products are the focus of existing model governance framework. For many other products, especially traditional life and health products, reserve calculations have traditionally been classified as low-risk models, as the calculations are generally formulaic, and assumptions are prescribed or locked in. While more-complex models for these products typically exist within insurance companies, they are generally used for applications such as pricing, forecasting and pass/fail-type tests such as cash flow testing or loss recognition testing. Using more-complex models to calculate reserves directly extends the model risk inherent in these models to the financial statements.



MODULAR DESIGN

Actuarial models are made of many parts. There are inputs, such as liability in-force extracts, asset investment accounting extracts, scenario data, and assumptions. Once the model has its inputs, the next step is coding the methodology for calculations such as reserve calculations, cash flow projections, and capital calculations and projections. A model run processes these inputs and calculation methodologies and produces outputs, which are then taken by an end user or utilized as input by another model (See Table 1).

Table 1
Model Components by Type

Component	Component type
Liability in force	Input
Asset portfolio	Input
Liability assumptions	Input
Asset assumptions	Input
Statutory & tax reserves	Calculation Engine
Embedded value	Calculation Engine
GAAP / IFRS	Calculation Engine
Cash flow projections	Calculation Engine
Capital	Calculation Engine
Reports	Output

Designing models as a set of model components, also known as modular design, offers several advantages:

- **Reusability:** Components can be reused across models and can be developed and tested once rather than multiple times.
- **Change Management:** Management of models is easier when model components are modular. Having distinct and well-defined components streamlines development and testing, allowing model changes to be done once and then leveraged in multiple ways.
- **Model Releases:** It's easier to show progress to users of the model when they are designed out of smaller components, which can be changed and released more quickly. More frequent releases allow the user of the model to more quickly use it and provide feedback and also decrease the probability of projects going overtime and over budget.

As an example, Figure 1 contains modules needed for a cash flow testing model. A model needed for VM20 deterministic/stochastic projections may use the same modules (with different assumptions) but not require a formulaic reserve projection.

In addition to the advantages listed above, modular design can increase the ability to consolidate, document and manage changes within models.

MODEL CONSOLIDATION

Using components across models leads to the idea of model consolidation. While it's tempting to consolidate models as much as possible—after all, who doesn't want to minimize work

and maximize sharing?—there can be challenges with sharing components. These are summarized in Table 2 (Pg. 11).

MODEL DOCUMENTATION

Documentation is a very effective tool to manage risk from models as not all team members are involved with the technical aspects of a model. Model documentation helps stakeholders and other business partners understand what the model does, what it doesn't do, and what the input, output and calculations are.

Key items to include in model documentation are shown in Table 3 (Pg. 11).

In addition to items in the Model Documentation chart, it is also helpful to have a "Day 2 list" of potential future improvements to the model. What goes on this list? Everything that someone might want that isn't there now. This list can include:

- Functionality desired at the time the model was built that is not currently present, perhaps due to complexity or software limitations
- New business requirements due to new regulations and new products
- New experience studies to improve model assumptions
- New and improved data elements and data feeds
- Approximations that are in the model now that could be removed

Figure 1
Example Cash Flow Testing Model as Viewed as a Set of Components




Table 2
Advantages and Disadvantages of Sharing Model Components

Design Consideration	Advantages of Model Consolidation	Disadvantages of Model Consolidation
One source of truth	One model or model component can force consistency	Getting buy-in from multiple stakeholders Consistency not always desirable Technological limitations
Avoid repeating work	One model avoids situation where multiple teams doing redundant work	Getting buy-in from multiple stakeholders
Flexibility	Fewer models	Different use cases have different requirements for flexibility (e.g., pricing vs. valuation)
Controls	Fewer controls	Different use cases have different requirements for controls (e.g., pricing vs. valuation)

Table 3
Model Documentation—Key Items

Key item	Description	Comments
Model business requirements	<ul style="list-style-type: none"> What the model is supposed to do What the model is used for; intended uses Functional and other requirements Uses of model output Other downstream processes supported by the model 	<ul style="list-style-type: none"> Examples: <ul style="list-style-type: none"> Produce reserves on business Day 4 Produce asset and liability cash flows to support VM20
Key elements of model design	<ul style="list-style-type: none"> Choices and trade-offs made when model was being constructed Rationale for choices Limitations of model due to design decisions 	<ul style="list-style-type: none"> Include rationale Cost savings, time savings
Model input	<ul style="list-style-type: none"> Where input comes from and how it's used 	<ul style="list-style-type: none"> Include any limitations of data or other input
Model output	<ul style="list-style-type: none"> What the output is What the output means What the limitations of the output are 	<ul style="list-style-type: none"> Documenting output and limitations of output limits the risks that the model output is misused or misunderstood
Procedures for running the model	<ul style="list-style-type: none"> The “nuts and bolts” process for updating the model and running projections 	<ul style="list-style-type: none"> Include enough detail so that another person with correct qualifications could run the model
Model approximations	<ul style="list-style-type: none"> Tracking of current approximations in the model 	<ul style="list-style-type: none"> Documenting these reduces the risk that approximations impact results inappropriately Approximations with financial statement impact may need to be tracked separately
Model specifications	<ul style="list-style-type: none"> Model technical specifications 	<ul style="list-style-type: none"> Include in appendix or reference another document
Roles	<ul style="list-style-type: none"> Model owner Model steward Stakeholders 	<ul style="list-style-type: none"> Include descriptions of owners of model and users of output

MY DAY 2 LIST IS AN EXCEL FILE ...

The Day 2 list is a helpful tool to set priorities at a company level, not the level of a single stakeholder—it ensures the company's priorities for the model are set correctly.

To help ensure there is a single source of truth—i.e., one Day 2 list—enterprise project management software such as Jira or Trello are much better tools than multiple Excel files being emailed back and forth.

MODELING ROLES AND MODEL CHANGE MANAGEMENT

Model development, including design, documentation and consolidation, takes a lot of effort. It is not practical for a single actuary to design, implement and validate a model. Generally, a model development and maintenance process involve roles borrowed from IT application development. Roles can include:

- **Developer**—Responsible for coding the model according to specifications
- **Tester**—Unique from the developer and responsible for testing model development
- **User**—Specifies model requirements and uses model results
- **Steward**—Responsible for governance and change management of the model

Many actuaries may be working on a model at a given time, which can make the model steward function challenging. Many IT application developers use programs such as GitHub and Subversion to manage changes made to source code. These programs allow control and documentation over the model development process to help reduce model risk. For example, a developer or tester can check out a copy of the code, simultaneously make changes or perform testing, and check the model in

with documentation. The management section of the program allows the model steward to review sequential changes to the model and assess whether the proper change management steps were followed and then decide whether to accept changes into a master version.

This type of change management functionality is starting to gain traction in the actuarial world but has not yet gained widespread adoption. Actuaries can learn from tools typically used in IT settings and advocate for integration of similar tools to their model development process.

CONCLUSION

In this article we walked through techniques that can help manage model risk and complexity. As regulatory change increases the complexity of our models, utilizing best practices and tools from software development such as modular design and software-assisted change management reduces the risk of making complicated changes to our model. Our models don't have to take us to the moon (at least not yet)—but let's build them as if they should. ■



Bill Cember, FSA, MAAA, is a director and actuary at Prudential. He can be reached at william.cember@prudential.com.



Scott Houghton is a vice president at Valani Global. He can be reached at shoughton@valaniglobal.com.



Dylan Strother is a senior consultant at the Actuarial Practice of Oliver Wyman. He can be reached at dylan.strother@oliverwyman.com.




CPD Tracker

A Free and Convenient Way to Track Your CPD Credits

- Track multiple Professional Development standards
- Download data conveniently to Microsoft Excel
- Load credits automatically from SOA orders
- Offers catalog of Professional Development offerings

SOA


Current Cycle
5.78 credits



Add new
activities

+

Track now at [SOA.org/CPDTracker](https://www.soa.org/CPDTracker)



Blockchain—Is the Newest Kid on the Block Going to Be All Right?

By Helen Duzhou and Jeff Guo

In the world of digital excitement, Blockchain is the newest kid on the technology block. Pull up any search engine, punch in “Blockchain” and find yourself inundated with articles praising improvements over current centralized network systems that many insurers use. But is this praise warranted, and does Blockchain truly cure all ails? In this article, we pose commonly asked questions as well as provide answers from the perspective of an insurance company.


MECHANICS OF BLOCKCHAIN

Definition: Blockchain is a **decentralized** and **distributed** database (also known as a ledger) from which accurate and secure information from any point in the past can be efficiently retrieved. A Blockchain is a connection of a set of blocks in historical order, and each block is a collection of data in a pre-defined structure.

“Decentralized” and “distributed” are not the same! In distributed databases the data is spread out over many



Figure 1
Simplified Visualization of Blockchain



computers, also known as nodes. For a decentralized database the ability to edit is not centralized and each node can edit the database directly.

Each block contains a cryptographic hash which serves two primary functions: (1) summary of data contained in the block, and (2) marker for the block’s location within the chain.

Figure 1 provides a simplified visualization of a Blockchain. The Blockchain begins with a single block called the **genesis block**, which defines the Blockchain’s initial parameters. Data (such as transactions or records) are validated by the network through a cryptographic hash algorithm before combined in a block.¹ The subsequent blocks are appended to the previous block, which can be traced all the way back to the genesis block.

Consider the following example Blockchain:

Beatrice and Anthony live together and split expenses evenly (the “genesis block” that defines the parameters). In January Beatrice and Anthony combine their receipts, check correctness, and create a summary of the transactions. All monthly transactions (for January) would be one block (“block 1”).

Thereafter, the February block (“block 2”) would be appended to the January block, the March block (“block 3”) would be appended to the February block, and so forth.

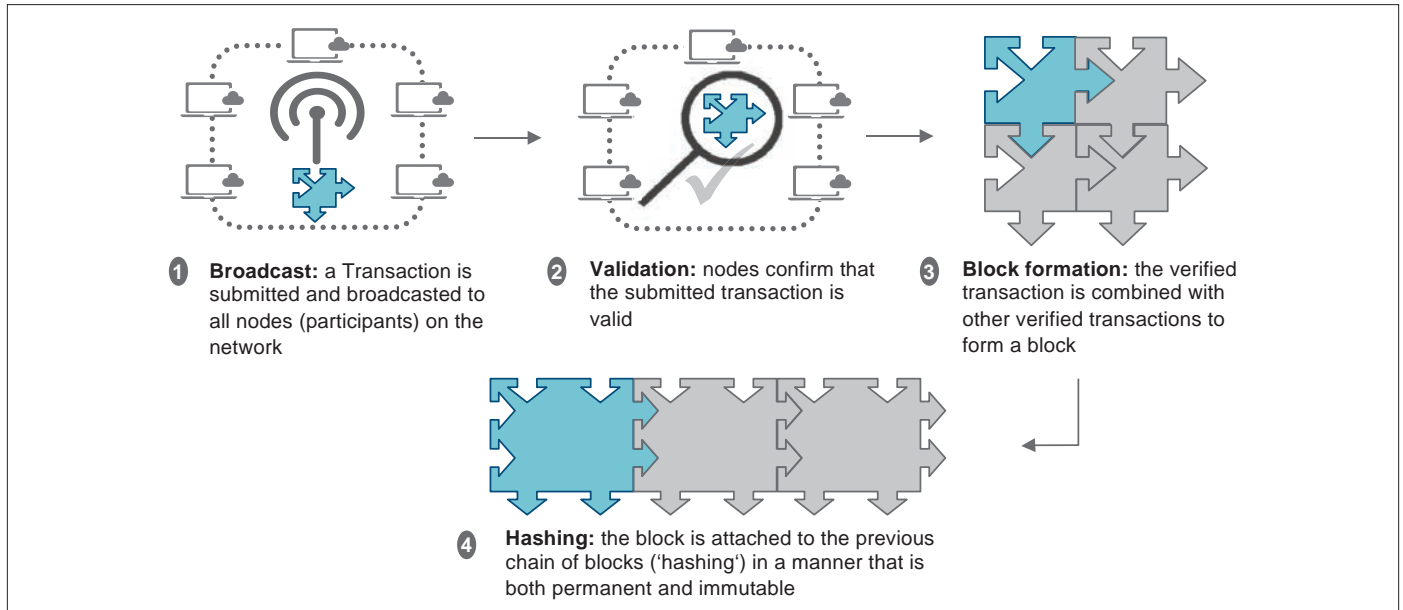
To figure out how much Beatrice or Anthony have spent, they would be able to review back to January and add up all of their receipts, similar to how information can be retrieved from and added to a Blockchain.

Adding a block is more complicated than two roommates reviewing their monthly receipts together. Figure 2 represents how a new block is verified by the network before being added to the chain.

Cryptographic Hash Algorithm and Mining

Integrity, validity and anonymity associated with a Blockchain are centered around a “magical” **cryptographic hash**

Figure 2
Adding a New Block to the Blockchain



algorithm. The algorithm combines a set of inputs (the hash of the previous block; the hash of the transactions of the current block; and a **nonce** (a **number used once**) to create a **hash**, which is a unique string of fixed length.

Through the cryptographic hash, the following key properties are realized in a Blockchain:

1. **Immutability.** Transactions cannot be altered since each transaction contains a digital fingerprint (the “hash”) that summarizes all prior transactions, preventing manipulation.
2. **Anonymity.** Each participant has a codename, allowing concealment of identity and only the owner can approve transactions.
3. **Integrity.** The hash ensures that the blocks are in proper order and that related information can be reconciled to the block.

A Blockchain is susceptible to a **collision issue** where modern-day computers can create thousands of potential blocks that could be appended to the Blockchain at the same time. To address the collision issue, the participants undergo a **mining** process to ensure only verified blocks are added to the Blockchain at a preset time interval.

During mining, a node intending to update the Blockchain will repeatedly guess a nonce until, by luck, the produced hash meets a network-defined criterion. Benefits of mining include deterring

network attacks and distributing the chance of updating the Blockchain across the network based on computation power.

Mining ensures the veracity of the Blockchain. If the majority of the participants are honest and only append to the correct blocks, then an incorrect block will be discarded by the network, and fewer people will append blocks to it because there is greater computational power in the network, and so a greater probability of the network recognizing the right block.

Together, the properties of mining create a **memory of the network** that prevents fraudulent or erroneous transactions.

During mining, a node intending to update the Blockchain will repeatedly guess a nonce until, by luck, the produced hash meets a network-defined criterion.

APPLICATIONS IN INSURANCE

The immutable, synchronous nature and integrity of the Blockchain allow for applications in optimizing loss management, streamlining existing systems, and penetrating new markets.

Several startups and initiatives are taking advantage of these opportunities, noted in Figure 3.

Figure 3
Use Cases

Example	Description	Use Case
Crowd funded industry initiative	Insurance companies merge their efforts in leveraging Blockchain to optimize processes, realize cost savings, and offer network users a variety of integrated applications.	B3i was founded by 16 market participants in 2017 and is now an initiative of 40+ company participants
Fraud detection	Improper transactions are detected by validating the “fingerprint” of encrypted and immutable Blockchain transactions.	Blockverify reduces fraud by tracking product’s Blockchain tags along a supply chain Everledger (e.g., <i>Diamond time-lapse</i>) verifies asset authenticity through Blockchain fingerprint
Smart contracts ²	Smart contracts are written as code in the Blockchain, which self-execute once a triggering event is met, without the need for third-party intervention.	Etherisc provides real-time insurance quotes based on flight delays ³ Edgelogic uses sensors to detect an event that triggers payment for repairs ⁴
Real-time insurance quotes	Blockchain can self-regulate the appropriate insurance premium continuously, which allows for development of tailored products addressing each customer’s unique concerns.	Safeshare Insurance facilitates short term commercial insurance contracts

POTENTIAL COMPLICATIONS

Although many insurers are interested in Blockchain technology, few have embraced it as a business proposition. According to one observer, “insurers do not necessarily need a current Blockchain strategy to remain competitive.”⁵

Potential complications for a Blockchain implementation for an insurer are as follows:

1. It is unclear how Blockchain will adapt to new regulation, since its structure is immutable. Since Blockchain does not allow tampering with existing blocks, removal of certain data is hindered, which limits compliance with regulation such as General Data Protection Regulation.
2. Future technology such as quantum computing⁶ can overpower Blockchain mechanisms, can be used to disrupt hashing as the validation mechanism and can expose private data to the public.
3. Blockchain is entirely dependent on the core source code, which can leave the entire ledger at risk if it is hackable.
4. Since the Blockchain is dependent on the memory of the network, collusion can occur if a party with majority of the computation power commits fraud.

In addition to the points already mentioned, other costs and considerations when implementing Blockchain include:


- First movers who invest in Blockchain will incur considerable up-front costs, due to lack of existing technical standards and expertise.
- Blockchain (currently) is not capable of scaling to handle large amounts of transactions.
- The process of mining is data-intensive and has high storage requirements relative to central databases that are currently used.

WHEN BLOCKCHAIN MAY BE NECESSARY

The main advantage of Blockchain is that it solves trust issues between firms and individuals. A common situation for which Blockchain would be very useful is if multiple parties are involved and their interests are not aligned. However, most of an insurer’s data needs can be addressed with existing technology. For more information, the flowchart in Figure 4 illustrates when it is appropriate to use Blockchain.

While Blockchain is not necessary today, it could become a business requirement in the future. Recent allegations around data breaches and privacy invasion are troubling, as insurers who hold sensitive policyholder information find themselves increasingly at risk of an attack. As Stephen Mildenhall pointed out, the internet has created a **trust vacuum**, which highlights the requirement for verification.⁷ The trust vacuum will become increasingly apparent as the world grows more interconnected. With Blockchain’s anonymous nature, data can only be approved by the owner.

Figure 4
Do you Even Need a Blockchain?



Source: B. Suichies, revised by Oliver Wyman

CONCLUSION

Although Blockchain offers significant advantages such as improved level of data integrity, anonymity of users, and inability for tampering (immutability), the insurance industry has been slow to adopt the technology. The heavy initial investment and expertise required can be prohibitive to entry by smaller insurers. Even now, most start-ups utilizing Blockchain are in their infancy, with heavy emphasis toward R&D. The general approach is to wait for a more applicable use case to be developed, or to participate in a crowd-funded industry initiative such as B3i.

Even with successful implementation, operational costs and regulatory and technology risks are higher than for traditional databases, with limited potential remediation methods. Insurers may be better off investing in a well-managed relational database.

Like all new kids on the block, Blockchain will develop and resolve many of its initial issues by becoming more scalable and efficient. And with its crucial data privacy benefit, as policyholders seek more control over the use of their data, Blockchain may become necessary in the future as it helps resolve the trust issue between insurer and policyholders.

Growing up is tough, but Blockchain might just be all right. ■

The views expressed are the authors' own and may not represent the views of Oliver Wyman.



Helen Duzhou, FSA, CERA, is a senior consultant at the Financial Services Practice of Oliver Wyman. She can be reached at helen.duzhou@oliverwyman.com.



Jeff Guo is a consultant at the Actuarial Practice of Oliver Wyman. He can be reached at jeff.guo@oliverwyman.com.

ENDNOTES

- <https://blockgeeks.com/what-is-hashing-digital-signature-in-the-Blockchain/>
- <https://blockgeeks.com/guides/smart-contracts/>
- <https://www.ccn.com/Blockchain-disrupt-air-travel-insurance-flightdelay/>
- <https://phys.org/news/2015-12-Blockchain-bitcoin.html>
- <http://novarica.com/Bitcoin-and-insurance-overview-and-key-issues/>
- <https://medium.com/the-quantum-resistant-ledger/no-ibms-quantum-computer-won-t-break-bitcoin-but-we-should-be-prepared-for-one-that-can-cc3e178ebff0>
- <http://www.aon.com/reinsurance/gimo/20180711-gimo-Blockchain>

Predictiveness vs. Interpretability

By Kimberly Steiner and Boyang Meng

A common criterion for the selection of predictive models is predictiveness: one model is considered better than another if it gives more accurate predictions of the outcomes of unknown events. Apart from making intuitive sense, this criterion is attractive because there are measures available (e.g., Gini coefficient, R^2) that allow us to easily rank models by predictiveness. This paper demonstrates that relying on predictiveness alone can result in choosing a model that exhibits behavior that may not be intuitive. It also demonstrates that this unintuitive behavior may not be immediately obvious.

In this article, we compare two kinds of predictive models, built using the same data, on the criteria of predictiveness and interpretability, in the context of life insurance mortality. The two types of models compared are generalized linear models (GLMs) and gradient boosting machines (GBMs). We demonstrate, using a double lift chart on holdout data, that a GBM can give better predictions than a GLM. We also demonstrate that while GLMs are easy to interpret, GBMs can be difficult to interpret, in the sense that profiles that are similar can have very different, and sometimes unintuitive, behaviors.



In conclusion, we emphasize that the desired attributes of a predictive model must be taken into account when determining what type to use, and we discuss some implications for the wider use of machine learning techniques in the insurance industry. We do not dispute the importance of predictiveness. However, we do argue that depending on the context, interpretability is an important consideration, and that, in some contexts, interpretability should not be sacrificed for predictiveness.

This article is organized into the following sections:

- **Predictive Models Considered:** General remarks on GLMs and GBMs
- **Data Used:** Details of the data used for this study
- **Details of the Models:** Details of the actual models' fit
- **Predictiveness:** A comparison of the predictiveness of the models
- **Interpretability:** Discussion of the interpretability of results
- **Conclusion:** Discussion of these results and some consequences in the context of life insurance, as well as some possible directions for further study

PREDICTIVE MODELS CONSIDERED

This section includes a high-level description of GLMs and GBMs. Further details can be found in the predictive analytics literature.

The types of models we chose to compare in this study were generalized linear models and gradient boosting machines. GLMs have been widely used in property and casualty insurance for decades for pricing purposes and have been increasingly used in recent years in life insurance for experience studies. GBMs are a trendy machine-learning technique becoming more widely used in many sectors. Models involving the use of GBMs are frequent winners of predictive analytics contests such as Kaggle (www.kaggle.com), which determines winners based solely on the Gini coefficient (i.e., a measure of predictiveness is the only consideration).

Generalized Linear Models

GLMs are a generalization of ordinary least squares regression. They are characterized by the selection of an error structure, which comes from the exponential family of distributions (this includes normal, Poisson, Gamma and binomial distributions), and a link function, the inverse of which relates the linear predictor (the linear combination of features included in the model) to



the response or independent variable. Common link functions are the identity, log and logit functions. Features are selected using a combination of statistics, heuristics and judgment. Each feature has a parameter associated with it, and model-fitted values are calculated by summing parameters of the appropriate features and applying the inverse of the link function.

Gradient Boosting Machines

Gradient boosting involves fitting a model on a randomly selected subset of the data, calculating the ratio between some proportion of the predictions of the previous model and the response on another random subset, fitting another model of that ratio and continuing the process unless some convergence criterion is reached. The model is selected by determining combinations of parameters such as the proportion of data included in each sample, the proportion of predictors available in each model and the proportion of the previous model predictions used at each step (the learning rate), as well as the characteristics of the underlying model. The underlying model is often a classification or regression tree. In this case, the final model is a weighted sum of a (potentially large) number of trees.

DATA USED

This study used single life mortality experience data provided by 23 companies for Willis Towers Watson's TOAMS4. The data include \$25 trillion face amount of exposure over the four-year study period (calendar years 2011–2015), representing more than 123 million policy years of exposure. More than 1.5 million death claims, corresponding to \$82 billion, are included in the

data. The data were split randomly into training and testing data. Both models were trained on the same training data and compared on the same testing data.

DETAILS OF THE MODELS

Generalized Linear Model

The GLM used a log link function and Poisson error structure. Attained age, issue age and duration were included as polynomials. The model included many interactions, including between categorical variables and polynomials (e.g., smoking status and duration or attained age and gender) and between combinations of polynomials (e.g., between duration and issue age). Categorical variables were grouped as necessary.

Gradient Boosting Machine


The response GBM was assumed to be distributed Poisson. Attained age, issue age and duration were included as continuous variables. Different groupings of categorical variables were experimented with. Hyperparameters were optimized using a grid search and cross-validation on a random split of the training data with four levels.

PREDICTIVENESS

Double lift charts are commonly used to compare predictiveness of two different models. A double lift chart is created as follows:

- For each observation in the testing data, predictions according to each model are calculated.

Figure 1
Double Lift Rescaled



- The ratio of predictions is calculated for each observation, and the observations are ranked according to this ratio from low to high and segmented into a number of bands (we used 50) of approximately equal exposure.
- In each band, each average model prediction is calculated and divided by the observed (i.e., actual) mortality in that band.

A double lift chart is effectively an actual vs. expected analysis by discrepancies between predictions in a pair of models. Where the model predictions are different, meaning where the ratio is high or low (i.e., in the extreme left and right of the graph), the model that gives better predictions is that for which the actual vs. expected is closer to 1.

To compare the predictiveness of the GLM and GBM, we used a double lift chart on the testing data as shown in Figure 1.

According to the double lift chart, the GBM was clearly more predictive than the GLM.

INTERPRETABILITY

As stated earlier, for a GLM, predicted values are determined by calculating a sum of parameters of the appropriate features and applying the inverse of the link function. In the case of a log link function, this is equivalent to multiplying the exponentials of the model parameters; that is, the model is multiplicative. This allows us to have a complete and interpretable understanding of the variables and combinations of variables driving estimates of mortality and the quantitative impact of each. It also allows us

to make statements like, “In segment x, mortality is y percent higher than in segment z.”

As previously stated, a GBM is a weighted sum of (an often-large number of often tree-based) models. There is no practical way to extract an interpretable characterization of the model predictions. Techniques (e.g., partial dependency plots) do exist that allow a general understanding of drivers of the model, but because of the nature of the model, it is possible for predictions associated with sets of observations to differ in unexpected ways. We illustrate this using several examples. The examples were created by:

- preparing profiles corresponding to different combinations of policy characteristics, including sex, smoking status, underwriting class, face amount, product and issue age;
- for each profile, creating observations corresponding to different durations; and
- calculating the GBM prediction on each observation for each profile.

Mortality by Duration for Selected Profile

In this example, we used male, nonsmoker, residual standard, face amount band \$500,000–\$600,000, current assumption universal life with level risk amount (ULNG). We compare the q_x by duration for selected issue ages (Figure 2).

We note that the q_x pattern for issue age 35 is monotonic and might be considered reasonable for all durations, whereas for higher issue ages the pattern breaks down (mortality decreases in certain

Figure 2
Qx for Selected Profile





Figure 3
Smoker Relative to Nonsmoker for Selected Profile



durations compared to the prior duration) at higher attained ages that lack credibility. While this is not surprising, the duration at which the pattern breaks down will vary by profile, and the only way to determine the point at which it breaks down is to evaluate the curve for all required profiles, of which there may be a very large number. While GLMs also struggle where credibility is lacking, we can identify and understand exactly how they are lacking.


Smoker Relative to Nonsmoker Mortality by Duration for Selected Profile

In this example, we used male, residual standard, face amount band of \$500,000–\$600,000, male universal life (level net amount

at risk), ULNG. We compare the ratio of smoker to nonsmoker mortality by duration for selected issue ages (Figure 3).

We note that even for combinations of issue age and duration where exposure is high, the ratio between smoker and nonsmoker q_x can exhibit patterns, including zigzags, for which there is no obvious explanation. We also note that these patterns can be different for all possible profiles. By way of contrast, GLMs allow a complete understanding of patterns describing relative levels of predictions (i.e., the relationship between smokers and nonsmokers is straightforward to determine with a GLM).

Figure 4
Best Preferred Relative to Residual Standard for Selected Profile



Best Preferred Relative to Residual Standard by Duration for Selected Profile

In this example, we used male, nonsmoker, face amount band of \$500,000–\$600,000, male universal life (level net amount at risk), ULNG. We compare the ratio of best preferred to residual standard mortality by duration for selected issue ages (Figure 4).

The patterns can contain unexpected “jumps” for which there is no obvious explanation. As explained in previous examples, detecting such behavior inherent in the model requires significant analysis of model results.

CONCLUSIONS

We do not suggest that machine learning techniques have no place in experience studies or other applications in life insurance. We do want to emphasize that the characteristics of the model (including interpretability) are considerations that in some contexts are as important as predictiveness. There are serious consequences of not fully understanding the relationships inherent in your assumptions:

- Since virtually no data sets are homogeneous through all durations and ages in life insurance, you may end up with assumptions that are inappropriate for your new business and it will be difficult to evaluate this since relationships are not immediately obvious.
- It will be difficult to set charges such as cost of insurance without knowing all of the patterns inherent in the mortality assumption.

- Modifying the assumption in places where little credibility exists in the data will be difficult given that relationships are not easily identified. With that said, further areas of research that could help limit these consequences include the following options:
- Exploring ways to detect unintuitive behavior (such as that illustrated in the examples) in GBM predictions
- Exploring ways to limit the GBM (or other machine-learning methods) so that results are more likely to be intuitive (e.g., to guarantee that mortality increases with duration)
- Extracting value from the GBM in ways that can result in an improved GLM (e.g., finding more sophisticated features that can be used to improve the predictiveness of a GLM) ■



Kimberly Steiner, FSA, MAAA, is senior director at Willis Towers Watson. She can be reached at kim.steiner@willistowerswatson.com.



Boyang Meng, ASA, is consultant and senior actuarial analyst at Willis Towers Watson. He can be reached at boyang.meng@willistowerswatson.com.



Elevate Your Analysts with the Certified Actuarial Analyst (CAA) Qualification

Support your staff with this professional growth opportunity:

- **Enhance** technical and analytical abilities
- **Expand** your data analyst team's actuarial skill set
- **Strengthen** data set analysis and communication of financial concepts
- **Build** employee loyalty and demonstrate your commitment

Empower your analysts to take the first step today at [CAA-Global.org](https://www.caa-global.org)

CAA Global is a joint venture of the Institute and Faculty of Actuaries (IFoA) and the Society of Actuaries (SOA).



Risk Management in the Digital Age

By Rob Ceske, Kelly Combs, Nadim Hraibi and Jamie Hooten¹

This article first appeared on www.theclearinghouse.org. It is reprinted here with permission.

Numerous technological advancements are now available to financial institutions that allow them to increase efficiency and keep up with changing consumer demands. Intelligent automation (IA) can range from simple algorithms to cognitive technologies, which have the ability to “learn” and adapt. (See page 27 for an overview of IA) Each type of automation can drive efficiency and effectiveness but also can introduce unique new risks. Traditional risk management techniques, which attempt to detect bad decisions or “rogue” employees and ensure appropriate lines of defense, must be adapted to address these new risks. With fewer human touchpoints throughout IA processes, the importance of design and appropriate usage, anticipating potential unusual circumstances, testing, and monitoring becomes paramount. Risk management teams will need to adapt their thinking and approaches to these new technologies and be proactive in reducing design risks and detecting unintended consequences of the new digital landscape.

FRAMEWORK

Leveraging IA can help financial services firms to automate processes, increase efficiency and consistency, and allow existing human labor to focus on more strategic activities while also improving customer experiences. However, once IA’s strategic mandate is defined, its adoption should include strategic, design, and operating considerations as part of an IA risk management framework (see Figure 1). The risk management of IA is a key part of these considerations. Although existing risk management approaches and disciplines (such as those for model risk management) can be leveraged for IA solutions, machine learning and cognitive tools, in particular, require adaptation to traditional risk management approaches that were developed for human decision makers. Many risk techniques were designed to mitigate and/or detect isolated bad decisions or “rogue” behaviors. However, as Class 2 and 3 IA technologies speed operations and decision making, risk approaches need to adapt to focus on design and monitoring activities.

Because Class 2 and 3 IA solutions have less involvement from humans in their operation, risk strategy will require a heightened focus on design components—for example, technology “fit for purpose” reviews, appropriateness and comprehensiveness of calibration approaches, and planned monitoring mechanisms. Solutions that have not incorporated risk monitoring into design are likely to be more challenging to monitor because decision processes in machine learning and cognitive solutions are much harder to discern than processes from operations run by people (e.g., poor or biased underwriting decisions).

Risk management in design should have much more focus on the technology component versus process and people components than would be the case with more traditional activities. Depending on the class of IA that is being used, there will be relatively more emphasis on programming and scripting (with Class 1 IA), tool configuration, and data used for calibration/machine learning (as with Class 2 and 3 IA tools). Class 1 IA tools require much more careful design planning, particularly regarding technology interaction, than traditional solutions. For example, changes to the color of an “OK” box, email formatting, or latency (for computer response time) can all affect Class 1 IA, but none of these changes would affect processes managed by people.

Leveraging IA can help financial services firms to automate processes, increase efficiency and consistency, and allow existing human labor to focus on more strategic activities.

Risk management of operations for processes that leverage IA solutions requires a heightened focus on business continuity and contingency planning. This may be uniquely challenging if IA has been used to displace humans or process-driven activities because sufficient staff and contingency processes will be harder to implement on a short-term basis. Depending on how critical the IA process is, if unanticipated outcomes are experienced (e.g., underwriting anomalies), the overall provision of affected services could be impacted because cognitive tool decision making will need to be investigated and retrained.

The training process for machine learning algorithms often is not easy to understand or back solve, giving rise to the added risk that these decision processes need to be more actively inferred from outputs (and with challenger models). In addition, risk management oversight professionals will benefit from more active data/analytic techniques as well as traditional monitoring

(e.g., volume, throughput, heat maps, etc.). In addition, programming and process enhancements should include a clear audit trail generated by the machine, exception-handling logic, processes to address exceptions that are “kicked out” of the IA process in a timely manner, and a feedback loop to incorporate these exceptions into the algorithms as adaptations over time.

Figure 1
Intelligent Automation Risk Management Framework



Monitoring approaches may also be candidates for intelligent automation (e.g., leveraging IA tools to perform independent validation, file reviews, and oversight of call center activities). Financial services firms that have not implemented Class 2 and 3 IA may find that existing review activities for second and third lines of defense are good candidates for “training” machine learning and cognitive tools because they are likely to have relatively greater amounts of data and be less time sensitive because the reviews and decisions already will have been made by humans (allowing the “right” responses to be known).

CAPABILITY CONSIDERATIONS

In adopting IA capabilities, executives should consider the current operational environment, governance, change management, resourcing, and integration with existing technologies. Adopting an evolutionary approach will lessen the risk inherent in technological disruption. Companies may wish to start IA implementation in lower-risk areas where results are easier to observe and verify. Firms may also wish to run IA and traditional processes in parallel and slowly transition from human- to IA-based

processing. After a company gains experience with IA, implementation can progress to higher-risk, higher-reward areas.

For an organization to advance to Class 2 or enhanced process automation, technology teams should have the ability to analyze structured and unstructured data. Intelligent automation technology required for Class 2 should support a built-in knowledge repository, from which it can perform some elements of machine learning.

In adopting IA capabilities, executives should consider the current operational environment, governance, change management, resourcing, and integration with existing technologies.

Regardless of complexity, all IA technologies consume data to complete tasks in a more efficient manner. As organizations progress through the classes of automation and data becomes increasingly more important, so does the need for effective data management and governance. A model is only as good as the underlying data. It is important that roles, responsibilities, and ownership are clearly established related to data.

IA implementation should follow traditional model validation processes. The model must be clearly documented and independently reviewed and tested. Model documentation becomes more important as human touchpoints are removed from the process. A monitoring function is required to review IA results and ensure that the model is operating as intended (in addition to the exception-handling process referenced previously). For example, the monitoring function could analyze input data to evaluate whether new patterns or conditions are prevalent in this data that was not anticipated in model development or training. In addition, an automation Center of Excellence can serve as a central point of contact for organizations to share knowledge and best practices.

In adopting IA capabilities, executives should consider the current operational environment, governance, change management, resourcing, and integration with existing technologies.

Companies must be equipped with data scientists who will have the ability to train and evaluate the model, as well as transform the data as the model evolves. In instances where anomalies

are removed or data is modified to enhance the outcome and accuracy of the model, documented approval and justification for why this has occurred should be in place.

Risk managers of the future will need to use more sophisticated data analytics to monitor artificial intelligence and have direct involvement with process owners to do root-cause analyses of issues. Risk managers will need to understand the implications of their models and be agile enough to respond to model corrections, understand the output, and evaluate risk of the model as it evolves over time.

ETHICAL RISK IMPLICATIONS

Modelers have been building statistical models used for predicting outcomes for decades. So why is artificial intelligence different? What are the ethical implications that need to be considered?

Artificial intelligence models need to consider the availability of historical and current data, be able to identify and correlate patterns in data, and be able to predict complex outcomes based on the same indicators as the human brain. Humans have inherent biases, however, so how is it possible to build a model that thinks like a human without the societal bias? And how does the model determine what bias is considered good within the appropriate context?

The projections out of algorithms are only as good as the data entered into the system. If the data is skewed or biased, then a destructive feedback loop can ensue, only worsening with time.

Because machine learning in itself is theoretically unbiased, the designers of the model need to be explicit and thoughtful about the design to help ensure that unintended bias is not created from unanticipated sources (e.g., data or flawed logic in the algorithm design). Think of machine learning in the context of a parent: Did you raise the child (build your model) well enough to ensure he or she has good morals (i.e., a low propensity for bias)? Poorly designed or managed machine learning models can have detrimental effects on individual stakeholders (e.g., through credit scoring or mortgage/loan decisions) as well as enterprises.

As machines continue to learn, they alter and develop their own algorithms so complex that the engineers who designed the

system may not be able to identify the reasoning behind a single output. Therefore, the disconnection between humans and artificial intelligence opens up risks for predicting when failures might happen.² A model that is transparent—when the design of the model can be understood and the factors that attribute the outcome are known—allows the user of the model to understand what influences the outcome of the model.

To help improve the accuracy and integrity of IA-driven decisions or predictions, organizations may want to consider implementing feedback loops. This process allows for better monitoring of conclusions reached by the algorithm against factual data sets (expected outputs) to identify degradation of the model, which, in some cases, may require model retraining.

DATA IS THE NEW OIL

When companies use cognitive solutions, they will also need to recognize that “data is the new oil”—that is, data will be the most vital component of a cognitive model—and companies will need to evaluate whether the company has appropriate historical data to feed the cognitive algorithms. Organizations will be challenged with evaluating whether competitors have better data or more accurate data sources than they do.

The projections out of algorithms are only as good as the data entered into the system. If the data is skewed or biased, then a destructive feedback loop can ensue, only worsening with time. Because cognitive systems learn from patterns,³ it is detrimental if they do not identify errors early. Therefore, when exposing a system to data, there must be a balance between the overfitting and underfitting of data. Data that can be directly attributed to the model outcome should be used where possible. Where proxy data is utilized, or data that indirectly is correlated to the outcome, this data should be understood and evaluated for its influence on the outcome over time.

In order to train algorithms, enough training data must be available. The more data variables that can be evaluated, the better the overall model. However, with every new dimension added to the model, the more computational power and storage is needed. As this computational volume increases, the available data to support the validity of the model decreases.

REGULATORY OVERSIGHT

Regulatory oversight of financial services firms (particularly oversight of risk management processes) will need to evolve with the increasing use of intelligent automation, particularly with Class 2 and 3 tools. A particular challenge will be in regulation and supervision that is designed to combat human bias in sales practices, extensions of credit, and similar financial decisioning for retail customers. Although it will not be acceptable for financial services executives to just say “the computer

made the decision,” supervisors will need to adapt oversight techniques and approaches to combat intentional (or directly embedded) bias in IA decision making—and not assume outcomes result from programmed bias. While correlation does not mean causation (or design, in this case) and supervisors should focus attention on intent/design, there will still be a need for risk managers to prevent unintended bias and to detect issues based on outcomes.

Intelligent Automation technologies create opportunities for improved efficiency and effectiveness in financial services firms, but they can also create risks that need to be managed. By expanding existing data and model risk management techniques as part of a comprehensive IA risk framework, companies may benefit greatly from these new technologies, while managing their risk. IA is here to stay. Let’s get the greatest net benefit from it!

OVERVIEW OF IA

Intelligent automation solutions can be broken down into three classes:

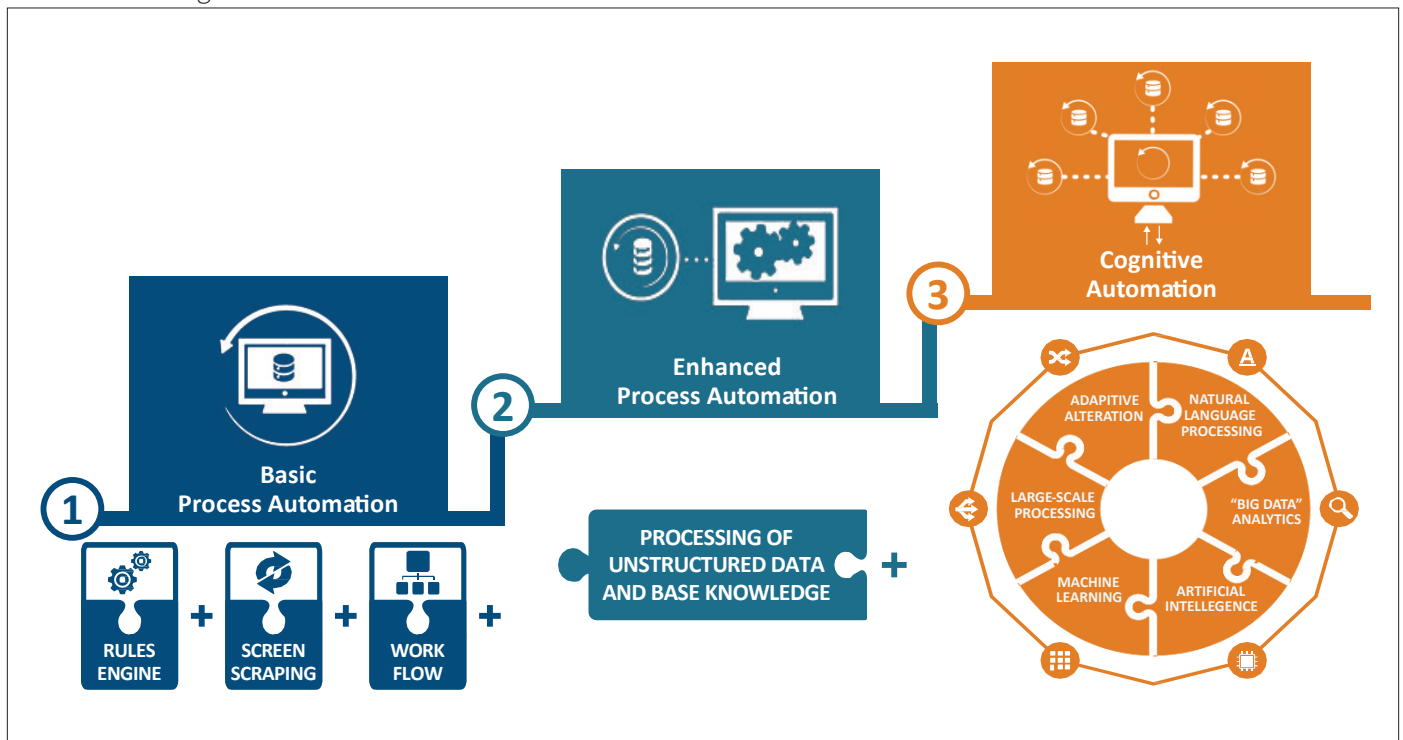
1. Basic process automation
2. Enhanced process automation
3. Machine learning/cognitive automation

Basic process automation (Class 1) addresses transactional work activities that are rules based and primarily repetitive in nature and typically completed in existing IT applications. This includes screen scraping, macros, incorporating workflows, and basic design capabilities. This is the simplest form of IA, where macro-based applets synthesize structured data to complete a noncomplex, limited judgment task or job function. Class 1 automation is used where there is no ambiguity in the processes and uses structured and standardized input data. Common types of basic process automation include robotic process automation (RPA) and screen scraping.

Example usage: *Systematic form population or bank account reconciliations*

Enhanced process automation (Class 2) enables the recognition of unstructured data and aids in adapting to the business environment. It builds upon basic process automation by incorporating a knowledge base and repository (RPA with the addition of a simple script/API add-on). The knowledge base is an important part of Class 2 automation, which allows the script and other capabilities to handle minor variations in input (e.g., date, address, business acronym). Such scripts can structure subprocesses or manual work that is not fully incorporated into the IT applications. Additionally, a key role of historical data

Figure 2
Classes of Intelligent Automation





includes use in performance evaluations. Class 2 automation requires moderate to heavy involvement from business users to structure requirements along with structuring rules to build computational algorithms and knowledge base.

Example usage: *Level 1 sanctions screening or cash flow forecasting*

Cognitive automation (Class 3) enables decision support with the help of advanced algorithms. The evolution of these tools is generally linked with advances in artificial intelligence, natural language processing, big data analytics, and evidence-based learning (machine learning). Machine learning is best defined as the ability of computer systems to learn and improve performance by exposure to data without explicit programming. Computer systems observe and recognize patterns, save the patterns in a knowledge repository, and later build on patterns to make predictions and offer solutions. Cognitive automation is the most advanced type of automation and can be used to automate tasks that require a relatively high level of human judgment. Cognitive technologies have the ability to mimic human reasoning and adapt as they self-learn. Cognitive solutions combine natural language processing, big data and predictive analytics, machine learning, and artificial intelligence. Class 3 automation is probabilistic and does not require business users to structure algorithms or logic; instead, models are typically “trained” by leveraging historical data. Additionally, key business users play a big role with evaluating model performance and enhancements. Further, historical data is used in model building and

performance evaluations. (Note: It is important to split training and testing data in order to avoid overfitting.)

Example usage: *Level 2 sanctions screening, email classification automation, or cash positioning and investments* ■

Disclaimer: Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates and related entities. The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.



Rob Ceske is a principal at KPMG LLP. He can be contacted at rceske@kpmg.com.



Kelly Combs is a director at KPMG LLP. She can be contacted at kcombs@kpmg.com.



Nadim Hraibi is a director at KPMG LLP. He can be contacted at nhraibi@kpmg.com.



Jamie Hooten is an associate at KPMG LLP. He can be contacted at jhooten@kpmg.com.

ENDNOTES

- 1 The authors would like to gratefully acknowledge Junghoon Woo, Steve Gaeta, and Mayuresh Kulkarni for their valuable assistance.
- 2 Will Knight, “The Dark Secret at the Heart of AI,” *MIT Technology Review*, April 11, 2017.
- 3 Cathy O’Neil, “When Not to Trust the Algorithm,” *Harvard Business Review*, Oct. 6, 2016.

TCARE: Unlocking the Power of Psychosocial Data

By Michael Mings

Editor's note: The Technology Section awarded to TCARE the first prize of the InsurTech Innovation Networking Event at the 2018 SOA Annual Meeting & Exhibit. As a prize, TCARE was to write an article for the Technology Section's newsletter, presenting themselves and the gap they are filling. Enjoy the reading!

The United States is facing a looming crisis: the baby boomer generation is getting old. The Population Reference Bureau has estimated that by 2060, Americans older than 65 will make up 24 percent of the population (compared to 15 percent today), which is an increase from 46M to 98M). Not only is the US population rapidly aging but the means to support this generation in its older years is proving a more difficult challenge than it has been in the past. This is further complicated in part by dementia and other neurodegenerative diseases; by 2050 Americans with Alzheimer's Disease could nearly triple to 14M (compared to 5M recorded in 2013). Insurance providers are increasingly challenged to respond to these demands as exhibited by long-term care (LTC) insurance. Traditional LTC insurance sales have decreased by more than 90 percent. As the population ages and available LTC insurance options decrease, family caregivers are becoming increasingly relied upon to provide services in the absence of other viable options.

"In 1950, you had a one in thirty chance of becoming a family caregiver. Today, it's one in three."

—Theresa Harvath, founding director of the Family Caregiving Institute at the Betty Irene Moore School

Family caregivers juggle a variety of tasks depending on the specific circumstances of their care recipient, and some of these tasks are becoming increasingly more time consuming and complicated. Tasks range from providing transportation, coordinating doctor appointments, and preparing food to more complicated tasks such as managing catheters, operating home dialysis equipment and/or other medical duties. Given these challenges, caregiver burnout is a real and growing problem.



IDENTITY DISCREPANCY THEORY IN CAREGIVING

In response to the scale and criticality of this problem, we (TCARE) examined a new approach to tackle these challenges, using psychosocial data in addition to more traditional approaches. Beginning with the hypothesis that support hours (ADLs) alone was too simplistic and not the best indicator of caregiver burnout, other factors were considered beyond the physical aspects of caregiving. This led to the development and application of the Identity Discrepancy Theory in the identification and remediation of at-risk caregivers.

The general concept behind Identity Discrepancy Theory is that people have an internalized expectation/perception of their role and activities in the world and then there is the reality of what they actually engage in day to day. If there is a difference between these two it causes stress and the greater the difference the greater the stress. In a family caregiver situation as the responsibilities increase, the caregiver should experience increased emotional distress due to the differences between how the caregiver perceives their role in the relationship vs. the actualities of their role in the relationship with the care recipient. For example, a spouse who spent much of her life in a specific role with her husband might become distressed when her husband is now dependent on her for activities he previously performed in the relationship, for example, management of household finances, yard work/home maintenance, and so on. Throughout the caregiving journey the caregiver is in transition from their current role (spouse, sibling, child, etc.) to evermore increasingly that of caregiver.

In developing a model that could accurately evaluate caregiver well-being, a variety of questionnaires were created and evaluated for efficacy and fit. The stresses discussed previously as part of the identity discrepancy were further partitioned out into a set of multidimensional component question groups (discrete burden measures) as follows:

1. **Stress Burden**—A measurement of anxiety or depression the caregiver experiences.
2. **Relationship Burden**—The extent to which the caregiver perceives the care recipient to be manipulative or overly demanding.
3. **Objective Burden**—The degree to which caregiving imposes observable aspects onto a caregiver’s life, such as time for self and others.

Gender and length of caregiving were used as control variables, as these were known to have different effects among caregivers. The questionnaire was finalized around the burden measures listed above. In the study, the number of hours caregivers helped with activities of daily living (ADLs) was also measured. Through the collection of data from 358 spousal caregivers (caregiver and care receiver dyad data) and analyzing the hypothesized model (Figure 1) with structural equation modeling (SEM) for fit,

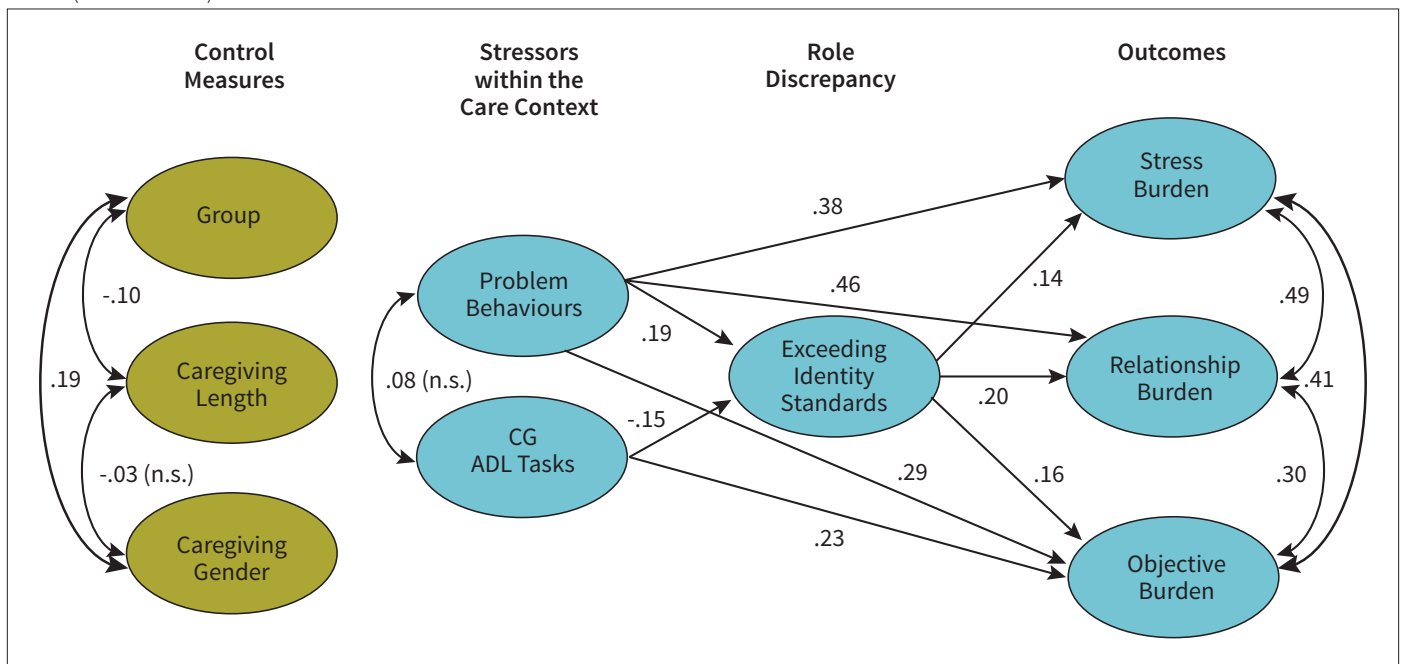
two important observations were established. It was discovered that the hours spent with ADLs are not inherently distressing; only when the performance of those ADLs goes beyond the perceived call of duty does it become stressful. A statistically significant relationship was supported by the multidimensional stress burdens (stress, relationship, and objective) and the caregiver’s identity discrepancy.

OPERATIONALIZING THE TCARE SOLUTION

With a functional model available, the buildout of the full protocol, inclusive of all the operational components (assessment, response, remediation, monitoring and prevention), commenced. Identification of caregiver burnout risk alone is not helpful without the ability to effect change for the caregiver through the application of targeted resources to help remediate the caregiver’s stress. To support this, a decision tree model was developed that identifies the type of caregiver burnout issue/s currently present within the dyad (issues can be different between a single caregiver with multiple care recipients). These fall into six major categories, which can be identified through their distinct fingerprints across the burden scores and the caregiver’s expressed intention to place the care recipient into a care facility.

A resource database was also constructed and populated with providers/solutions, mapped to the discrete problem drivers,

Figure 1
Final (Alternative) Model



Note: All parameter estimates are standardized and significant unless otherwise stated. Controlling for variables in SEM requires direct paths (not shown) from control measures to each latent construct in the model. ADL = activity of daily living; CG = caregiver

for the caregiver to perform the needed remediate for their issue(s). These components were then wrapped together in a software package with additional services supporting caregiver engagement (communications and monitoring) and workflow/case management. Since then the solution has been deployed to numerous customer groups.

TCARE SOLUTION RESULTS AND FUTURE POSSIBILITIES

Often if you ask a caregiver what they need, they will request some form of respite. Most of the time this is not the most impactful response and does not tackle the root cause of the caregiver's stressors. This approach provides for a more acutely focused solution to the caregiver's needs. Analysis of caregivers who have gone through the TCARE protocol vs. the control group reveals both a significant reduction in insurance service utilization (-20 percent in Medicaid groups) and a 21-month delay in nursing home placement, which results in millions of dollars in savings (health and LTC insurance) in addition to the benefits realized by both the caregiver and recipient due to their ability to age in place longer. Follow-up research is also being conducted to analyze potential increases in longevity in

the cohort, which would drive additional benefits, for example, increased revenue yields for associated life insurance providers.

Work continues to enhance the effectiveness of the existing solution, leveraging the ever-growing set of historical data augmented with third-party data and applying new, advanced and adaptive modeling techniques (supervised and unsupervised learning). The results, both historically and those from the new research efforts, are better than initially expected and suggest many other possible use cases in adjacent and unrelated domains. This includes research work already underway on risk rating populations of insurance policy holders and the creation of hybrid LTC insurance instruments with less risk and better yields. With the recent World Health Organization (WHO) designation of burnout as an official medical diagnosis (ICD-11), an employee risk/retention solution is showing promise. ■

Mike Mings is CIO at TCARE. He can be contacted at mike@tailoredcare.com.

SOA e-Courses

SOA's e-courses offer actuaries a broad range of forward-thinking topics. From decision making and communications to fundamentals of the actuarial practice, actuaries who enroll will gain a better understanding of relevant topics relating to the actuarial profession.

Enroll now at soa.org/ecourses



SOCIETY OF ACTUARIES®



475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
p: 847.706.3500 f: 847.706.3599
w: www.soa.org

NONPROFIT
ORGANIZATION
U.S. POSTAGE
PAID
SAINT JOSEPH, MI
PERMIT NO. 263

