Interview with Mark R. Shapland and Jeffrey A. Courchene, authors of the monograph *The Actuary and Enterprise Risk Management: Integrating Reserve Variability*

A new volume in the CAS monograph series, *The Actuary and Enterprise Risk Management: Integrating Reserve Variability* by Mark R. Shapland, and Jeffrey A. Courchene, proposes moving beyond reserve variability quantification to allow for full integration of key reserve risk metrics into the larger enterprise risk management framework.

Janice Young, of the Monograph Editorial Board, discusses stochastic reserving models with the authors.

Janice Young: What motivated you to write this monograph?

Jeff Courchene: After our successful IAA Monograph on stochastic modeling ("the textbook") and working in an environment (Europe) where Solvency 2 best estimate was defined in Article 76 of the Solvency 2 Framework Directive to be a "probability weighted average of future cash flows," it was abundantly clear that actuaries needed to increase the attention given to stochastic reserving models, complementing their existing deterministic reserving methods.

From my perspective, calibrating a distribution of possible outcomes for individual and aggregate reserve estimates was a precondition for complying with the Solvency 2 best estimate definition. Observations from the market, however, indicated that the understanding and use of such models was underwhelming and often disjointed.

The process often focused on downstream activities (e.g., the calibration of solvency capital and/or the pricing of portfolio transfers). My initial motivation was to develop a framework for (and to incentivize) implementing stochastic reserving models in a way that complemented the well-established and time-tested actuarial practice of deterministic reserving.

As the framework developed, it became abundantly clear that a number of powerful metrics were unlocked which supported the understanding and management of reserving risk and most importantly positioned the actuary to proactively engage other insurance professionals in the reserve uncertainty discussion.

So that's the European take, but I think it's pretty accurate, right?

Mark Shapland: Yes.

Jeff Courchene: I mean at that time we, you know, had amassed a significant amount of knowledge and developed this pretty significant tool that was superior to what most

people were doing in the market.

And I think that we were both motivated to see more actuaries do more work in the space.

Mark Shapland: I agree. I recall one of the other things that triggered this was Jeff and I had both been giving a lot of talks at GIRO in the early 2000s and everyone was talking about Pillar One (of Solvency 2), which is I think a lot of what Jeff was talking about, the technical details of coming up with this stuff and a lot of discussion about how we do these calculations to get to the solvency calculations using a one-year time horizon.

But no one was talking about Pillar Two and Pillar Three, which is about governance and back testing and validating models, and so maybe, Jeff, you thought about this before I did.

But as we were going to those sessions and thinking about that and realized that no one was talking about Pillar Two and Pillar Three, we thought maybe that's something we should be exploring more deeply.

And to me, that was one of the motivating factors, as well is the desire to expand the discussion beyond just calculations into actual real-world use of the models, and to help manage a business. Because if you're not, if you're just going through the motions of calculating solvency or risk margins and your reserves, then you're kind of only doing half the job, right?

Understanding how to manage risk, and understanding the governance and validation aspects of that process as well, are important. In the deterministic world, actuaries are used to that, so they'll go back and say, OK, here's how I selected last year; here are my picks from last year; did I do a good job of estimating the numbers? But they don't really validate the models all the time. The validation process isn't the same under stochastic reserving and the dimensions change quite interestingly.

So, I think that that's kind of what Jeff was getting at, too. And his response that the metrics we talk about in the monograph are moving from “how far off is my point estimate” to whether it’s significantly off, is a very different answer and different mindset.

That does that jibe with your recollection, Jeff, in terms of part of the genesis for the monograph?

Jeff Courchene: I definitely think you're right.

And I omitted the discussion about Pillar Two of Solvency 2 and the integration, which essentially is the risk management framework that surrounds Solvency 2 and includes a number of things involved in managing and governing various risks.

So that's for sure an area where we saw limited research and limited input from actuaries in a space where actuaries should be highly involved.

I think you're right that over time from a deterministic reserving perspective, actuaries are very happy with their own methods, maybe even validating their prior assumptions before they carry on with their current assumptions.

But perhaps they’re kind of locked behind a closed door providing a number that management ends up using the way that management prefers to use the numbers.

And so, in the same way that actuaries have expanded beyond that closed door in the past decades for their point estimates, we identified that this as something that actuaries need to do with respect to the discussion of uncertainty around those estimates as well.

With the advent of Solvency 2, that was really an opportunity for actuaries because that's when those discussions were expanding with risk functions and with management boards within large organizations.

We were motivated to make sure that actuaries had a seat at that table as experts in risk to lead those discussions.

Janice Young: Can you go over what Pillars One, Two and Three are?

Jeff Courchene: Pillar One is, as Mark mentioned, the quantitative calculations of an economic balance sheet and the required capital to support that economic balance sheet.

Within Pillar One, an important element on the economic balance sheet is technical provisions and an important element of technical provisions is the claim provision, which is analogous to a GAAP loss reserve estimate.

There's also a premium provision which is analogous to an unearned premium reserve, although it's based on an economic value as opposed to a GAAP balance sheet.

Pillar Two deals with governance and talks about data quality standards that have to be upheld. It also talks about various opinions that the actuarial function has to render upon a company which include not only an assessment and a certification of reserves, but also an opinion on reinsurance arrangements and an opinion on underwriting processes.

There is a significant amount of validation and assessments with respect to the risk management framework that a company is employing, how they're managing those risks and whether they're managing those risks effectively.

Pillar Three has to do with reporting and disclosure, and so Pillar Three is essentially

the defining of a number of reporting templates that the various regulators in Europe require. It's similar to the Yellow Book or the NAIC annual statement, but on an economic balance sheet basis, not on a statutory accounting basis.

Mark Shapland: That's the European perspective. It's interesting that our motivation started in Europe, but I'm not sure that it really made its way too much to the US, although a few years later the ORSA requirements came in for companies in the US, usually the bigger companies that were doing, capital modeling and so forth. To me it's a little disappointing that the US didn't adopt some of the things that the Europeans were doing in that area.

And as I think about it, I mentioned earlier actuaries are comfortable back testing, but I don't really have a feel for how many actuaries actually do that In their day-to-day life.

So, I would say hopefully one of the motivators for people reading this monograph is maybe it will get them thinking about doing more back-testing.

In the deterministic world of back-testing and validating methods, I suspect actuaries may just rationalize using chain ladder, Bornhuetter-Ferguson and some other common methods and note that they're not perfect, but they're what I have and they've stood the test of time, so back-testing and validating probably don't come into their minds all the time.

I'm just guessing, but I would say that I would hope that maybe that would be something that actuaries reading this monograph would maybe think to themselves: "Maybe I should do more of that."

I don't know. We'll see.

Janice Young: The back-testing is similar to a runoff analysis?

Mark Shapland: Yeah, but not just to make sure that the models are working right, so are the models performing the way they should be?

And if you're projecting a risk metric for management of reserves, how good was that estimate? And there are things that you could do to improve that estimate over time?

Jeff Courchene: And if you've made a number of assumptions in order to come up with that estimate, whether it's a point estimate or an uncertainty estimate, are you able to demonstrate that those assumptions are consistent and reasonable?

And over time, either continue to be reasonable, or, if not validated, change.

In Europe, we're required to do that. The actuarial function report is one of the outputs of Solvency 2 that addresses this explicitly.

The actuarial function report requires actuaries to do essentially a comparison of the

performance of their past assumptions and disclose this to the board of the company as well as to the regulator upon request.

Mark Shapland: I haven't done an opinion in the US for quite a while, so I don't really know what the current requirements are. I don't think they're anywhere close to as stringent as the European requirements are, but I could be wrong.

Jeff Courchene: That's my understanding as well. I mean the demands from our clients are such that they are motivated and obviously impacted by what they see in their sister companies in Europe or what they hear about in various professional conferences and such.

So, I think the demands are high, but what ends up falling into the US opinion is extremely limited, whereas in our case, what falls into our reports, or what is mandated in our reports, is quite significant.

Janice Young: What would you say are the key takeaways in the monograph?

Jeff Courchene: I think the three key takeaways are: (1) powerful KPIs for managing reserving risk are made available through integration of reserve risk measurement within an ERM framework; (2) KPIs providing the direction and significance of deviation from expectation are much more powerful compared to KPIs providing the direction and magnitude; and (3) proactive engagement with insurance professionals outside of the actuarial silo at the front end of a reserve analysis provides better results.

I think that number 2, the one that Mark mentioned already, is probably the most compelling and that has to do with when you're looking at how well your models have performed.

Simply lining up what you expected with what actually happened gives you a directional indication of how well you did, either higher or lower, and gives you a magnitude indication.

If you were a lot higher or a lot lower than what actually happened, it becomes a question of significance and there's a big difference between magnitude and direction and significance and direction.

And so it is exactly the inclusion of the stochastic reserving element which needs to be closely aligned with the processes in your deterministic methods.

That enables you to think about whether the recent performance is close to what would have been expected, meaning mean or median expectation, or whether it falls outside of what your own uncertainty analysis would have concluded as a one-in-100 year event or anything in between. Another thing that we found is actual performance is potentially outside of the distribution that you've calibrated for that line of business, which has

implications for the uncertainty distribution that you previously had come up with.

Janice Young: It's wrong?

Jeff Courchene: Not necessarily wrong, but not wide enough to handle what we've just experienced in the past year.

And these observations are available not only in total, but based on the framework that we put together, drilling down by line of business, drilling down further by accident year and depending on your appetite for complexity could be drilled down even further.

Mark Shapland: That's a great list. The only thing we could add to that list is the ability to help manage actuarial resources as well.

So, having these KPIs, you can fairly quickly see where the problems are and maybe you need to reallocate some of your actuarial resources to the more problem areas.

One of the things that wasn't part of my thought process during the creation of the monograph, but now in retrospect I've added, is hoping that this might motivate US actuaries to do more.

I presented the monograph about six months ago or so at the Central States Actuarial Forum and one of the things I said when I was introducing this is, to me this is kind of in line with Jessica Leong's theme of the actuary of the future that while a lot of actuaries in the US are not doing this—not that I know of anyway—I think this could be something that that could be part of standard actuarial practice in the US.

I certainly won't be around to see it, or at least not unretired long enough to see it, but I think it's something actuaries could do here. And I hope it does motivate people to do more and do some of the things that the Europeans are doing.

Whether that happens or not, I don't know. It probably depends on whether the accounting standards change. As long as our accounting standards are point estimates and that's all we need, it may not happen, but I would throw that in there, anyway.

Jeff Courchene: Yes, I think that the point about resources shouldn't be underappreciated.

I did not include that, I had forgotten about that, but we talked about that on a number of occasions as kind of an unintentional consequence of the work that we did: that actuaries tend to build up an actuarial team and then allocate various individuals to various parts of the portfolio that the company writes.

Those individuals grow up being the workers comp specialist actuary out of a team of 10, dealing with the reserving issues that deal with workers comp and they kind of stay in that silo. There's value to building up experience in that silo.

But what our framework offers is an early indication of which of the methods and models have performed particularly well and particularly poorly at any point in time.

So, if a team has resources that have, let's say, a technical skill set, which is more advanced than others on their team, then an indication that there's a problem in the methods and models could motivate the allocation of that resource on a targeted basis to exactly those parts of the business which have essentially performed differently than what was previously expected.

And so you end up as an actuary managing a department. You end up in a very powerful position to not only rely on individuals gaining experience in various lines of business, which is beneficial to the process, but also being able to target technical skills in areas where significant technical skills are needed in order to refine or to think about how the model supporting that part of the business are being run.

So, it's kind of another dimension to the management of the team that's unlocked.

Mark Shapland: Yeah. It's management of the enterprise risks within the firm, but also management of personnel within the actuarial function.

Janice Young: What percentage of insurance companies (including smaller ones) do you think have an enterprise risk management structure in place?

Jeff Courchene: According to Article 269 of the Solvency 2 Delegated Acts, all insurance and reinsurance companies in Europe must have a risk management function that, among other things, “manages the risk management process” and “cooperates with the actuarial function.” The sophistication of approach varies widely across insurance and reinsurance companies with the most advanced companies employing an ERM framework across all risks, both quantifiable and unquantifiable. Reserve risk, which is the subject of the monograph, is one of the more material risks faced by insurance and reinsurance companies.

So, in Europe we definitely have enterprise risk management structures firmly in place within insurance companies.

For insurance companies in the US markets it's a little bit different.

Mark Shapland: Yes. I don't really have a great feel for it, but my sense is that maybe some of the large companies that must do ORSA reports have that and maybe some of the medium-sized companies are striving toward that.

But I would be very surprised if any smaller companies were doing anything like this. It's a pretty large effort to get this in place now.

The value I think is there but maybe there's other things to tackle in terms of expense ratios and other issues that smaller companies have to worry about.

But like I said, until there's a regulatory mandate for it, I doubt that will change much in the US. The other thing we could address is outside of Europe and the US.

I think there's probably a shift starting to happen with IFRS 17 taking hold in a lot of other countries outside of Europe and the US. When I was in Dubai a few years ago they were early stages but I don't know how much further they've come along since then.

Jeff Courchene: I'm not an IFRS 17 expert so I can't really speak to that.

There certainly are companies that are complying with IFRS 17 and of course they have an economic balance sheet and of course they have a risk adjustment, which you're unable to calibrate without doing some significant uncertainty analysis.

So you're right that a lot of that is being picked up by companies who are complying with IFRS 17.

I think, independent of that development, the Chinese developed a new solvency regime called C-ROSS. The South Africans have their own version of Solvency 2 that that they have brought to bear. Australia as well.

And Canada has a totally different approach with their PFAB.

But the US, and I guess North America in general, is kind of the exception to the rule of global development in this space, in the space of risk management and economic balance sheet thinking for capital requirements.

Janice Young: Did you see the percentage of insurance companies with an enterprise risk management structure changing?

Jeff Courchene: In general, the level of sophistication naturally improves over time, benefiting from new technology and innovative actuarial talent.

With respect of the management and measurement and reserving risk, I've witnessed a number of companies consider, and a subset of those companies initiate, a reserve transformation project with the goal of streamlining processes at the front end of a reserve analysis in order to allow actuaries to spend more time engaging with professionals from other silos.

Unfortunately, not all of these projects included the integration of stochastic reserving models.

So more work needs to be done to decrease the resource cost required for the calibration of such models without sacrificing the knowledge gained through the calibration process.

Mark Shapland: Yes, I think that answer holds true in the US for the few companies

that have done it.

Hopefully more companies will think about the benefits from what we are describing in the monograph that they probably weren't thinking about. I suspect a lot of actuaries just kind of say, well, ERM is great thing we should do this too. They don't necessarily have the tools to do it or it's hard to convince management of the benefit of doing it.

Perhaps, I don't know.

Jeff Courchene: Yes. An early version of our monograph was a paper in the call paper program. Many, many years ago, I want to say a decade ago, but maybe it's less than that.

And at that time, I would have expected many more companies to have picked up the approach that we are recommending because there's really not a whole lot of difficulty to the implementation.

In theory, they are all doing the component parts. At least in Europe, they're doing the component parts already for other reasons.

It's simply a matter of aligning all of those moving pieces within a single holistic framework.

But, unfortunately, my observation in Europe would be that only a small number of players have risen to the challenge, which is a bit disappointing considering how long ago we started talking about this.

But I think the number one reason why has to do with the resource cost. The uncertainty analysis that we're referring to takes time. The more time you invest in it, the more benefit you get out of it.

You learn and start to understand your data quite a bit more than you may have understood it before.

What I see instead is the application of models where they can firstly decide on a way forward and then essentially turn the crank for however number of segments they need to analyze within a quarter or within an annual review period.

And the reason why that is, there's just so many minutes in a day, so many actuaries on the team, and what we're talking about would require a significant investment of time and actuaries are expensive and that time is simply expensive.

So, I think a little bit more convincing needs to be somehow communicated or in a better case a client or a company needs to go down this path in a significant way, gain the significant advantages that we think they would gain from it.

And then of course, have a competitive advantage for a short period of time where

everybody else catches up.

Mark Shapland: Yes, I think that's a good observation, Jeff.

I think the flip side of that might be that at some point maybe someone will come along that makes this much more cost effective that the amount of work that goes into the models and the presentation of data and all the KPIs, production process becomes much more affordable at some point.

Jeff and I had discussions about doing that years ago and we weren't quite to the point where we could do that, but maybe at some point somebody will, especially with the advent of the AI models that are coming out and the actuarial future may have a very different environment than we do now.

Jeff Courchene: Yep.

Mark Shapland: So I probably will be long retired by the time this becomes fully available, but it's exciting.

Janice Young: Did you learn anything surprising as you wrote this monograph?

Jeff Courchene: Not really surprising, but my appreciation for the limitations of the most popular stochastic reserving models (e.g., the ODP bootstrap and the closed form solution developed by Thomas Mack) motivated a number of important additions to the case study, which included stepping away from popular bootstrapping approaches, when necessary, in favor of a GLM approach on which bootstrapping is grounded.

Essential limitations reinforce the message that solutions for integrating stochastic reserving models need to be flexible enough to accommodate situations where most popular stochastic reserving models are inadequate for the purpose.

Mark Shapland: Yes, I totally agree with that.

I mean, like Jeff said, it wasn't necessarily surprising, but one of the things we hope we get out of this monograph, too, is to help people appreciate the fact that not only stochastic models, but deterministic models probably need to be more flexible than they typically are.

Now, having said that, a lot of companies have vanilla bread and butter kind of things where their triangles were just fine, right?

But even there, there can be trend changes that happen that are very difficult to pick up in a typical chain ladder method, right?

So having more tools can make the analysis process more robust and it can give you more insights into the data that Jeff was alluding to earlier.

Hopefully people will recognize that if they're limiting their tools to the few that they are comfortable with, then they may be missing something.

Jeff Courchene: I was a little surprised by Mark's follow-up research, where he looked at thousands of companies across many accident years and essentially back-tested the entire schedule P database over a number of calendar years.

And one of his conclusions there was that the models that we rely on underestimate the true uncertainty within reserves.

Mark Shapland: Yes.

Jeff Courchene: And so from my perspective, I've carried on some research regarding what we call ENIDs in Europe, which is events not in data, and this is based on the idea that the distributions that we attempt to produce are truncated and we expand them in order to estimate the unknown piece. The Australians produced some research where they were adding some systemic loads to try to increase the width of the distributions of possible outcomes that fell out of stochastic reserving models. But I think for your research I was actually quite surprised how conclusive that was.

Mark Shapland: Yes, I was a little bit, too! But it's amazing how consistently every model we tested underestimated the variance, meaning the width of the distribution.

Janice Young: So is that kind of like the unknown unknowns?

Mark Shapland: To some degree. I think adding systemic risk in or adding the ENIDs Jeff mentioned are certainly an appropriate part of it. I've seen other research where if you start looking at more years of data or quarterly data, you start to get closer.

So, part of it, I think, is a lot of the time we just have 10 years of data, which is what most actuaries use when the tail is shorter. But all the variability that could happen in 10 years doesn't necessarily happen. So, I think we just don't see enough variability in the data we use to fully measure it. Which is the kind of what the point of adding systemic risk and ENIDs and other things that to add to it.

Jeff Courchene: Yes, but it's tackling the unknown unknowns and the two approaches I mentioned add something, so the question is, have we picked up all the unknown unknowns and the answer is certainly no. Well, we know things are missing because Mark’s research shows that quite clearly. I think the only the way that you can try to get your hands around that challenge is to continually do the work, validate the assumptions based on your past expectations and refine your models over time to try to keep them at pace with the knowledge and information that you are able to collect over time.

Mark Shapland: And adjust assumptions accordingly, right? To me, part of that is going back to the models. For example, even if you were using the ODP bootstrap with a

Bornhuetter-Ferguson, I would want to use a much wider variance assumption that I might have otherwise for the a priori assumptions.

That helps broaden the models output. So, it's part of a cycle of knowing the models are underestimating because of my back testing and other sources, and then proactively and interactively start changing the assumptions accordingly.

Janice Young: Anything else you would like to add?

Mark Shapland: I guess I would just hope that a lot of actuaries in the US take the time to read it and think about whether it can help their actuarial function or not.

Jeff Courchene: I think the one point I might make that we didn't speak to much is that actuaries like to work with their methods and models. And actuaries like to put their head down and do their work.

As part of the third key takeaway that I mentioned earlier, what are the things that actuaries need to do? I think there's broad agreement in various jurisdictions about the need to expand the set of individuals that actuaries speak to on a regular basis as they are tackling the challenges of calibrating a loss reserve.

Instead of focusing only on the numbers and only on their models, there's a wealth of information in claims departments, reinsurance departments, IT departments, and elsewhere.

There are market experts that are selling the business that understand different trends that might be occurring. However, they sometimes speak a different language than actuaries do, and so to the extent that you can increase the set of metrics being developed to try to be able to communicate with these other individuals using language that they understand, I think that that improves the analysis that actuaries do.

It complements the actuarial skill set with knowledge from other spaces and I think it's a missed opportunity when an actuary is behind a closed office door coming up with their numbers, without that communication to the wider community of professionals within their organization.

Mark Shapland: Well said. I think that's very appropriate. One thing I was thinking about is the case study we describe and when we presented this in the past, we talked about the fact that you have new information available before you even start an annual reserving analysis that says how far off you were last year and in the implications for what that means to your reserve numbers going forward, and that prospect can be scary for actuaries. It's like, holy cow, the whole world's going to know that I misestimated last year on day one.

And actuaries don't like that.

I don't like it, but I think if you embrace that process that Jeff was describing and get to the point where you can have risk management discussions with other managers and the company about the fact that we're using this as a way to know what's going on and to react to it as a management group ahead of time and not be surprised at the end of process before we close the books.

That changes the dynamic and it changes the way actuaries interact with other risk managers. So I think it builds on what Jeff was saying.

I hope that people look at that and don't get scared, but say OK, I need to think about doing this.

Jeff Courchene: Yes, Mark and I have given presentations on this topic for a number of years for the European Actuarial Academy seminar and this is one of the case studies that we discuss and that's the moment that gets the biggest gasp in the room. The idea that you've only just received the data and you're already telling your boss what the answers might be.

Given all the expectations from the prior year, you know what your new answer might be, based on your one-year time horizon view of what's happened. And that always gets a big gasp because people think I couldn't possibly talk to my manager or my CFO until I've thoroughly scoured through all of the new information and double-checked all of my work.

And I'm not sure that they're going to receive the message and not get upset if that number changes it anyway, which is a very different environment that we're proposing, which is you say this is where we are up front.

We're going to look through our models, see what we can find, and in the meantime, this is where lights are blinking red.

Who from claims, who from reinsurance, etc. should we engage with to try to get more information so that we have an all-hands-on-deck approach to making sure that the next set of models and methods is better than it was last year.

Mark Shapland: Just to clarify something Jeff said, you don't at day one know exactly what the results are going to be, you just have a really good idea based on the analysis you did before where the holes are going to be, where the big rocks are going to be.

Jeff Courchene: If we don't change any of our assumptions, we know the answer. Of course, we're going to change assumptions as we digest the newly available information.

Mark Shapland: Yes, but we want to know why. So that's the key. The key point is that once you have that, the question is why is that? Why did that happen? What caused

that? What changed that we didn't pick up on before? What assumptions did we make that didn't hold true?

And then it gets back into the resource allocation and risk management process and all the things that they come with it.

Janice Young: Anything else you want to say to wrap up?

Jeff Courchene: This has been fun. Thanks for leading us down this path.

Mark Shapland: Yes.

Jeff Courchene: After many years of writing and rewriting this, it is finally across the finish line. So, I'm delighted to see that it's finally there.

Mark Shapland: Yes, I think we presented the concepts several times before we even wrote the paper. So the genesis of this probably happened at least 15 years ago, didn't it? Maybe longer.

Jeff Courchene: Must be.

Janice Young: Thank you so much.