

RECORD, Volume 31, No. 1*

Life Spring Meeting
New Orleans, LA
May 23–24, 2005

Session 66OF Asset-Liability Modeling Issues: An Open Forum

Track: Investment
Moderator: ELLEN COOPER
Panelists: TRACEY J. POLSGROVE

Summary: In this session a panel of experienced ALM modelers and C-suite professionals facilitate a discussion around ALM modeling issues and concerns. How can we utilize our modeling results into actionable activities within our organizations? What are the obstacles and how can they be overcome?

MS. ELLEN COOPER: I think this is the first time we've ever run a session like this. I'm going to run through some introductions and then talk to you about what we're hoping that we can all talk about today as a group. I'm Ellen Cooper, and I work with Ernst & Young in our New York office. I wear a lot of hats there. I think they feel like they get their money's worth if I have about four different jobs. One of them is that I lead our risk- and capital-management efforts. That includes our research and development efforts and our modeling efforts as well, which is part of what we're going to talk to you about today. I also manage the New York office, where there's a staff of 26. I've been in the industry for 20 years. This is my former colleague, Tracey Polsgrove. Tracey is with the Hartford. Her title is AVP and actuary. Her responsibilities include corporate pricing review, embedded value, M&A and other stuff.

We're going to talk to you about some of the issues that the Hartford has been dealing with and working to accomplish. As a consultant, I have been working with quite a few companies to help them think through their modeling processes, where

*Copyright © 2005, Society of Actuaries

they are today, where they want to go in the future and how they're going to get there. It really is a journey. You need to have a lot of perseverance and a lot of discipline to go from current state to future state.

Many of the company representatives I have talked to and surveyed feel today as if they've gotten to a place where their models are a bit of a mess. It sounds like we can relate to this. I think part of what happened is back in the late 1980s and 1990s (for those of you who were in the industry then) is we started building these models for cash-flow testing. At the time, these were models that we were basically building for compliance. We were building these models to run once a year, and then senior management started realizing, "Wait a minute, maybe we can get more information out of these models. We're producing cash flows; we have cash-flow projections." Then this whole idea about stochastic scenarios came through—stochastic pricing, embedded value, hedging, ALMs and economic capital. We were trying to run these models that were originally built to run once a year for compliance to monthly in some cases, and we weren't able. We didn't have the foresight at that time to think through where we were going.

Now we have a lot of duct tape and a lot of different pieces in a lot of different areas that we're trying to pull together. Typically, what I've seen in companies is there's a lot of manual intervention and a lot of room for error. I see a lot of companies struggling with very long cycle times. They have data as of September 30, and they're producing results in March. Senior management says, "This is great, but can you tell me what's going on now?"

Part of what we're going to do is talk about what some of our model uses are today and what we want to do in terms of model design. Let's start with a clean piece of paper. Let's pretend that we don't have all these models and think about what we need to do. What does the design look like? For those of you who have redone your kitchens lately or redone a room in your house, you need to have a design first. You need to build a vision and build what it's going to look like and think through all those pieces or you're going to wind up with something you don't expect. Even when you do design it, you're going to wind up with something you didn't expect.

Once you have the design, we're going to talk about putting together the framework. One of the keys I think actuaries have traditionally not thought through is infrastructure controls, roles and responsibilities. You need to have very clearly defined roles and responsibilities. Who are the owners? Who is responsible? How do we implement it now that we have all the design, now that we're actually going to get all the work done? How are we going to do it and what do we mean when we talk about model transformation? With that, I'm going to turn it over to my former colleague, Tracey Polsgrove, who is going to talk about other stuff.

MS. TRACEY J. POLSGROVE: As Ellen said, we really wanted this to be interactive. What I thought I could do today is share some of the things that really worked for us at the Hartford. We've taken on very large modeling projects, and we tend to

build some pretty complex models. Also, I wanted to share some issues we're still facing and see if everybody else out there is facing the same issues, and what types of things have they thought of to get around some of these. I have three important themes that you'll hear me talk about today. One of them is communication. As you go through this process, there's a lot of communication that has to happen. It's pretty important to make sure all the right stakeholders are in the room and that everyone is approaching things systematically. The second thing that I think has really been helpful for us at the Hartford when we've built various models is building a relationship with our IT folks. We couldn't have done much without having a pretty top-notch IT team that really was interested in helping us solve a business problem. I'm going to share a couple of stories about where we were and where we went. The other thing, as Ellen said, is controls. That's something I spend a lot of time thinking about.

Why the heck are we doing this? Why are we building these huge, complex stochastic models, and in some cases nested stochastic models? One reason is to look at things like variable annuities under the new C3 Phase II rules that are coming out. When you try to do a financial projection and figure out what the reserve is, you have to do a nested count, but that's pretty complex. So, why are we doing this? What types of information do these models really provide? A question you have to ask yourself as you go throughout this is: What am I going to change with this information? If you always keep that in the back of your head, it helps you sometimes when you start to go down. We can do this, this and this, but is there something simpler we can do that will get us to the same spot? Do we really need that much complexity? The funny thing is, usually management says yes.

Some of the information that models provide needs pricing results and benefit costs. Having started my career in the early 1990s, and having started in a pricing role, my average return is 15 or so. But wasn't it nice that it ranged between 12 and 17 and we were all happy? One of the things that the 1990s and the early 2000s taught us was that an average doesn't help us when the range is financial ruin or tail risk and makes a whole bunch of money. Stochastic pricing has become ever more important as we build these products, especially ones that have guarantees and are very past-dependent. To me as an actuary, having taken exams and having grown up under this average, we really have to shift our thinking a bit.

The second piece of information is a "what-if" analysis. Maybe we call that a sensitivity analysis. If you have a model that can do a lot of things, you can look at all different scenarios. This will help you make decisions about which actions you can take from a management perspective to protect yourself against certain scenarios. There are some new uses of these models in hedging analysis. This is not only a systematic hedging program, but also, again, looking at "what-if" and perhaps entering into some hedges to protect yourself against certain potential outcomes. You want to make sure your statutory income doesn't fall below a certain level because it's marked. You might be able to enter into a hedge. You get there by

deciding whether that would be worth it by looking at the results of the models you build.

On balancing between economic income and statutory capital requirements, I think if you have a stochastic model on a lot of products, you can see what product features you can tweak that will really give you huge gains. For the statutory capital requirements you have on that product, sometimes there are some pretty small product designs and changes that will help you get there.

The old stand-by is the asset adequacy analysis. We still have to get the donuts out the door, so to speak, and we still have to do that analysis every year. If any of your companies are variable annuity carriers, use of C3 Phase II results and regulations will cause us to have to build nice stochastic models. If you, in turn, hedge it, it will have to be a nested model.

Who uses the information? Again, I think there are some new folks on the list who use it and there are some of the standard users. Obviously, actuaries, pricing and financial risk-management and the investments area are all going to look at these results and make decisions based on that. As for regulators and rating agencies, I went to a couple of presentations yesterday about economic capital where the rating agencies talked about looking at economic capital models and internal models. Again, most of those are going to be stochastic ALM-style models that can get quite complex. I put senior management, the board of directors and investors in the new users group. To some extent senior management has always looked at the work that actuaries have done, but I think it's with more frequency and in more detail than maybe ever before. I know it's true at my company that they really want to see the details, the stuff that you thought, "Wow, I just want to do a nice, glossy PowerPoint and put that in front of him. They're not going to want to know the real nuts and bolts." But we are finding that as we provide more information, they want even more. They're excited about it. They want to see it, and they ask a lot of questions.

The main point I wanted to make, and this is, again, one of my themes on the controls, is the stakes are higher than ever. There are complexities on some of the models that we're building and the amount of data that goes in (and then also the high profile on some of the uses). Some have always been there, but on some of the newer ones, the stakes are higher than ever. Peer review, model validation and model controls are critical in this new environment.

In the variable annuity models we typically run, we're not really modeling the general assets. For our fixed account, we have done that for cash-flow testing after we model the assets, but on most variable annuities that are offered in the market today, they have some sort of guarantee against market downturns like a return of premium no matter what happens. In that sense, you'd want to look at that tail.

MS. COOPER: You'd be looking at a range of results if you were looking at any item

that was either stochastic or dynamic. So if the actuary were doing a model and let's say he always had from the investment department some kind of fixed earned rate, no matter what, but if they had dynamic lapse in there, there would still be a range. There's not a range to the price, but there's a range to the outcome as to whether that put is in the money or not.

MS. POLSGROVE: And whether you as a company make money on it.

MS. COOPER: You look at a range of economic outcomes to figure out under which economic scenarios that put winds up in the money versus is it valueless.

FROM THE FLOOR: That's why I was wondering about this, because one big problem with my firm is it's not the assets, it's always the liability. That's something I have to incorporate.

MS. POLSGROVE: This is looking at the liabilities over a large number of scenarios?

FROM THE FLOOR: Over whatever.

MS. COOPER: Some companies feel like they have significant risk on the asset side, as well in terms of what kind of earned rates they can achieve. They have convexity around any kind of securitized assets.

FROM THE FLOOR: The investment committee will decide and the actuary pricing of the product would have no input ...

MS. COOPER: Right, and that goes to roles and responsibilities. Who's responsible for providing that information and are you modeling actual assets or not? That comes in to some of the design decisions and understanding within each company.

FROM THE FLOOR: The investment area is going to move into corporate bonds and they don't need your approval; they're not going to necessarily tell you what they did with your money, so don't worry about it.

MS. POLSGROVE: Yes, but we also have a lot of fairly long-tailed products. We have things like terminal funding and maturity funding. If you're taking over a retirement fund and there are still active employees, we might get the constant earned rate from the investment team, but in looking at it stochastically, especially in today's low-interest-rate environment, you can look at what future reinvestment would be. You could also play around with the investment strategies to see which strategy might be best for a particular liability, and typically that's how we use our stochastic models.

In corporate we have done work with our investment team to look at different investment strategies by line. A line, 401(k) or something, might say we have this

general account, we have this much money coming in and this much going out. You would take a look by cases and things and do some sensitivities around that. Do a stochastic model where you bury the investment strategy and say, gee, can I go a little longer here, would that be okay? What additional risk does it introduce? Those are some of the things where we're looking at how you use these stochastic models. Again, the Hartford's a fairly large variable annuity writer, so we tend to do a lot of stochastic risk management and pricing around variable annuity guarantees.

It's very important to look at what the basis point cost of a return of premium death benefit is, because it can vary wildly depending on what happens in the market. It could cost you virtually nothing or it could cost you a few hundred basis points. Those are some of the ranges of pricing results. In the past, when I first got into looking at variable annuities, you'd hear folks say average cost of that is three basis points, and we really found out that that wasn't a really good way to look at it. The three is really great if it's average cost, but if you got there through various market movements, you either lost a gazillion dollars or made a gazillion dollars.

MS. COOPER: Let's talk about design. We really want to draw a road map. Whenever I'm working with clients, we always do one succession and come up with a single page. What we'll do is we'll set up three tiers, including modeling and then output. If you really think through the issues of things like where does a liability in force come from, where does it go, where do I want it to go? Whatever software I use inside here is really going to affect my calculations. The input is so critical. This all comes from the same place. Do I have little pieces of it? Are there things that I need that I currently don't have where I have someone in a room doing some sort of calculation? I recommend coming up with some kind of standardized data structure for what needs to go into the input. It sounds so simple, but think of really doing that and really having control around it.

Now when you pull the data in, you want to first pull it in as an aggregation. Do we not have access to it, and if so, what is that level of aggregation? I have issues where I have renewal premiums and I have it based on number of years. Therefore if I'm rolling it all into this little inforce record, we need to sit and look through it and redesign the process. We cannot do this alone. We may need to have some rules around compression. How are we going to compress for a model? And by the way, depending on our profit, here you have different levels of compression. So I'll use the model as one of my sort of calculations. You're probably going to have more sales than a financial plan.

The other thing that we don't do well is we don't have any way to control it, so we want liability implicated. We need to have names, and typically no matter what type of engine you're using to work with IT, whether it's Oracle or something else, there are a lot of databases. There are tricks, but the idea is there have to be rules. The same thing goes for assets. If you're trying to do all kinds of things and the

modeling engine is doing the asset reduction, sometimes they're not exactly what you need. The other thing is this whole area around economic scenarios. I've seen all kinds of things. We all know we struggled with this: Which scenario generator did you use? I don't know. You need to have better control around how to calibrate and what kind of scenario process you use when you do ALMs. Is it consistent with what we've done. What's the investment process using it? They have duration numbers using some different processes, and then somebody's making decisions based on that. Step back and think through all of your plans. I can go on and on, but I'm not going to do that to you. I'm going to move on, but that's the basic idea.

There are really two sets of processes. There's production where things that happen every month of every quarter of every year-end, and we know the people in IT should have control of that. Then there's what's called ad hoc. I typically recommend that production is production and that it's automated and as controlled as possible. The separate models, the separate processes, the separate line, whatever it might be, but you have rules around things, like whether variables have consistency.

I didn't talk about substance, but that's the other thing that's really critical over here. How am I controlling my input and putting it into a model? If I have a model, do I have any idea? Again, I recommend that it's all documented and that if you can, you build some kind of IT control. We're recognizing that you need some kind of a hardware infrastructure to run this. What do I need to do to optimize my system? In terms of output, again, I recommend that everything get put to the test. I know they have some kind of standardized support and standardized analytic, and we have some ways to audit. We have auditable assets that we can go back and look at. And finally, underneath the whole process and around the whole process you have to have control. We're going to talk more about that. I'm going to hand it over to Tracey, but does anybody have any comments? I know I threw a lot of stuff out there, but I've been thinking about this for a while and working with companies for a while.

FROM THE FLOOR: One question that I have right off the bat is IT resources are fairly scarce and for something like this, I question whether the resources that we have currently are appropriate to do some of these things. How do you get the resources allocated, how do you make sure you have the right resources to be able to do something like that?

MS. COOPER: That is an excellent question and I'll open it up to the group, but I can't help it, I have to answer it first. You need to build the case internally and you need to be pretty vigilant about it. I have one example of a client who I'm working with now. The person, the owner of that project, the executive sponsor of the project, he worked on this. He finally got it. Seeing more and more activity like this in the industry, I think it's going to become the new way. Companies are going to start to understand that they see the value. This is the value that you're going to get from it. I don't know how many of you have issues where senior management

says I want to see some “what-if” analysis and you go back and scratch your head and figure out. If you go back to them and say if you let me do this, I can have an answer for you and you demonstrate to them that you really can do it, then they start listening.

This variable annuity hedging production thing is really transforming the industry, because for the first time we understand a little bit what it’s like to be a real valuation shop, and in real time.

MS. POLSGROVE: We run our hedging every night. I’ll talk a little bit later about one of the things that I think has been really key for the Hartford in building all this stuff. We started with the hedging implementation. We really found some partners in the IT folks, and it is rare and I understand that. It was funny how it started. I was building a stochastic deferred acquisition cost (DAC) on mock model, just ran a bunch of scenarios and my student needed a new computer. Every corporation has its computer procurement process. You get a computer and you might add some more memory to it, some more RAM because they’re going to be running some nice, big actuarial models. The young man came to deliver it, and I asked a question because we had gone to a new model and I could no longer fit three RAM cards in the back, so I could do half a gigabyte each at \$90. It only had two spots, so I had to put one big RAM chip in there, and they cost \$1,000 each because they were brand new. I was like well, why is that? We need really good machines. It was the first time I had gotten an IT procurement guy excited. There are these high-end workstations. We ended up with four, dedicated, dual-processor, high-speed workstations. This was in 2002, and we tend to roll those out and give them to our students when we replace them if we need those. We’ve also gone to a server, and they’re springing up like fungi all over the building. That started with our hedging implementation and, again, this was IT.

My counterpart is an IT guy and who is an actuary, so maybe he understands us and knows how to play with us, I don’t know. But when we had to do our hedging implementation, his team came together and said we need to solve this business problem. They love working with the actuaries because we’re doing all the cool stuff and we’re building this partnership. I have on my staff a C++ programmer who is an ASA, and he helps bridge the gap between the two worlds, and they’re very different worlds, but if you can bridge that gap you can have a lot of success in doing this. We started our hedging implementation and we had servers that were being run in a test environment, so you’d need a mock done really fast. I asked if I could get some time on their tech servers because I have to run this cash-flow testing model or whatever it was. Instead of high-end PCs, we started buying banks of servers. So now, like for cash-flow testing, it was like almost 100 percent run on servers. That actually has some issues as well, because servers don’t run like your PC at your desk. I don’t even think I really know how they work, but there’s a little more upkeep associated with them. There’s a little more in getting them done. They were proficient in this, and yet there’s a guy on my staff who spends quite a bit of his time as an internal help desk, helping people with model issues, programming

issues on the software and also with getting them to run if they get a little unstable.

I've heard a couple of comments on scenarios. I've heard a lot at some of the sessions I've been at. You need to run more scenarios and now we have the technology and with grids and all this other good stuff. It's not as easy as it sounds, and I don't know if anybody else has experience in this, but I loved the last couple of sessions I've been at when they roll off how many scenarios you can run. We really found that it depends on your purpose. I mean, if you're running a stochastic run with your variable annuity block, 250/250 has always been able to do and it gets those servers humming. I used to joke that the lights will dim when we run this model. Seriously, last February the electric bill went way up, so I quit joking about that. People throw out these scenarios and run 1,000 scenarios. What can you run overnight and in what kind of time?

I found that grid is not as easy as it sounds, either. When I first started hearing about grid you just kind of hit the button and then all the computers in the building are working on your actuarial model. Isn't it beautiful and grand? We really found that the software is not as easy as we thought. There are a lot of output programs and input programs you have to do to manipulate it beforehand. Also, not all computers can run your model. I finally got our IT folks to realize that actuaries need something a little bit fancier than the call center people would get when they get a new PC. Likewise, because half the building has PCs that don't have much memory or RAM and all the other things you need to run these models, you can't necessarily deploy your grid across all the machines in the building.

We have different banks of servers, and we have different banks that we've built to plug and plan, but it's not really like, oh, you have a bank of 15, 15 and 15, I want to run B3 and those four and just push the button. It's a little more complicated than that. There's a little more planning that needs to go into that and that's where this partnership with IT has worked out where we're working on projects to figure out what models do we want to grid, because sometimes it is very model-dependent as to what you can run on the grid because you have to do a little work to get things ready to do that. We're still working on that and hoping we make some more progress to make it a little easier.

I don't know what folks are doing out there, how many scenarios they're running for stochastic testing or other things. One of the things we've been struggling with is how do you choose. If you want to choose a representative sample of them, and I don't know if anyone has any experience out there of things that they've been trying, but that's one of those nuts that we just keep trying to crack. We choose them randomly and put them through some sort of process. I don't know if you've had a lot of experience in that, but that's something we're still working on.

Also, on the liability inforce data, some of the things we've been looking at, and Ellen touched on this really well, is what the right fraction is. For our variable annuity line, in which we have quite a few policies, we tend to have about

1,400,000. We have to compress that to various levels for various things. When we first started doing some work, obviously run time and all this good stuff, we tried to compress down to about 2,500 cells, which if you think about it, is pretty small. We actually run that in Oracle. We started out with Access, but we ended up going to Oracle. We've been doing more testing now, as we're going to C3 Phase II to try to figure out what the compression should look like. At what level of cells does it start to jump around? We found we're getting up in the 20,000 to 30,000 range before it starts to stabilize a little bit. You expand a little more and you have a huge change in your reserve or your capital requirement. You have to make sure how to time that, what you're comfortable with, how much that jumps around and how accurate it is.

We have not really changed our compression rules in a while. Now we are because we are looking at implementing the C3 Phase II on our variable annuities, and we have a big project going on. This is where they're doing the work. We started a couple of years ago and then what we've been playing the most with is something we call small cell compression. In the old days, when I was a consultant, you ran the compression. You'd do a nice click check to make sure there weren't cells with three policies, and you'd reassign them and move on. We run through our compression once and then we go back and look for small cells. Continually, because of model run-time issues, we go back and say nothing under 5 million. Okay, that's too many. It really is a trial-and-error thing to try to figure that out, and, unfortunately, there's no easy answer.

Now we're going the other way because before we wanted to take 1.4 million policies to 2,500 because we're trying to run something; we're trying to run the entire inforce. If we try to run it nested, we have to go even smaller, and we can only do like 250 scenarios. Now we're going the other way and saying now we want to base our capital on this, so we're expanding back out. We're doing that today. We're not there yet, and this is how do you get there. You envision you have a compression process and you choose for whatever purpose. If it's a financial plan then you want more detail, and you run it through something again that keeps compressing it down to the various purposes. I'm not sure we're at that utopia yet, but we're trying.

FROM THE FLOOR: What about the actuarial journal on scenario selection that only dealt with interest rates?

MS. POLSGROVE: We used to use a different approach. We still use it in some instances. It's Fowler. There was a particular estimation thing where for cash-flow testing we would take the 1,000 to bring them down. We've since gone to just running the straight 1,000 regenerate. But on other things where we need to bring it down to even less than that, we've had some issues in that kind of the consistency that Ellen hinted at. We want to be consistent across the operation. There could be reasons why you would have different ones for different purposes, but let's approach it systematically and make sure we know where that is instead of

everybody calls up the investment guy and says, hey, I need some scenarios today. We're working on bringing that in line.

What happens when we try to choose a representative 250? We want to be consistent across the organization, and we have other countries. We've become too sophisticated to be simple. We're struggling with how do you choose those if you choose a representative 250? How do you get your foreign product to have scenarios that are consistent so when you go to look at everything together scenario one lines up with scenario one and scenario two is consistent? We've learned the hard way a few times where we've put things together and thought, "That's probably not a good idea." You need interest rates consistent with equity markets.

FROM THE FLOOR: When you do the last compression, how much is driven by computer speed?

MS. POLSGROVE: For cash-flow testing and various purposes, I think we've always had to compress into representative cells. I think now it just becomes more and more difficult as you build these bigger models. We had an issue with the structure of our model where physically the computer wouldn't run if it was over a certain number of cells. You just got an error; it was kind of nebulous, it didn't really tell you.

When we help a company implement these variable annuity hedging programs, typically, there are a number of reasons for it. I spoke yesterday and Marshall Greenbaum talked about this, that there are some correlation offsets to some of the separate accounts you lose if you do too much compression, so that's one issue. The second thing is that some of these complex products, like guaranteed minimum withdrawal benefits (GMWBs), for example, as people start electing, the compression exercise there becomes so complex. What bucket are they in if they've elected one benefit versus two? Are they going to elect again annually? Know your inforce. It does give you insights. Know your inforce and know it well.

If you're doing single premium immediate annuities (SPIAs) for example, you may say what's the point of running these things? I don't want to tie up my production computers for 42 hours this weekend on this stuff because I'm not going to get any value out of that; I'm going to compress the whole thing. If you really are going to get value out of it, there's cost benefit.

FROM THE FLOOR: Well, this compression is not easy when you model. When you start compressing, you have to estimate assumptions...

MS. POLSGROVE: Yes, definitely, when you have death benefits and living benefits on the same policy. There's also a memory issue. I guess for us considering how large our block is, that's so far off.

MS. COOPER: Some of what I've seen, for example, and I've observed this multiple times, the pattern is a client will come to us when he's dealing with the issue of cell compression in these modeling transformations. We'll first define what the hardware infrastructure is going to look like. So once they define that, given the number of scenarios that I have, given the number of time steps that I have, if it's a variable, how many funds, whatever, how many cells can I run, given that I want my run to be overnight. Or I want my run to be over a weekend and that's how they come to the level of cell compression that they need to have. They back into it for now.

MS. POLSGROVE: We did that in our hedging. We had a limit because we wanted it done overnight on our hedging program, and we recently rewrote that and took something that took overnight to taking a matter of minutes. The first thing is what else we can run. We'll get back to overnight. I don't think we're there yet, but I have been out of the office a couple days, so we might be back to overnight.

We found that even with all the hardware we have, even on a pricing basis we have a bunch of different cells that we want to look at. You have to figure out what's going to impact your results. Even on a pricing basis, maybe if you were pricing life you'd grab a few different premium payment patterns if it was a flexible pay and you'd do a few different things like that. Now as the benefits get more complicated and there are more things that can impact that, you have finer and finer distinctions so that even on a pricing basis you end up with 40 cells or something like that.

MS. COOPER: Again, these are things that we'll always be working on. We're going to talk about controls and some of the insights that I now personally am having relating to control. Again I'm going to give credit to IT. We started working with some IT folks and they gave me insights that as an actuary I really liked. Going in, making the change, calculating the numbers and then moving on and then if somebody needed me to figure out how I got back to that answer, I'd figure it out on the back of an envelope somehow. This is just intended to be an example of some of the controls that we probably should be thinking about—putting them into production, having them so that we really build an infrastructure around what we're doing.

Most of us do these today in terms of validation. Most companies are doing static and dynamic validations, some kind of analytics when we're first building models. But are we documenting them? I know from my example as a consultant, if I go into a company and I say can I see your documentation on a model and I get a couple of e-mails that I'm putting them together. We're just bad at this. We need to document what we want to validate on a static or dynamic basis and where we are going to put those validations, besides off in a file some place or in an e-mail; let's have a process around it. Give it to IT; they seem to be more organized than we are.

Around data control, and this goes back to the data inputs here that I drew earlier, we need controls around things like bad data. In terms of putting out bad policies, do we have account values that are negative, that kind of don't make sense? Do we

have variable annuities? Do we have guaranteed minimum death benefits (GMDBs) that are lower than our current account values, or does the sum of our sub-accounts not equal our account values? Again, I've seen all kinds of things like this that don't make sense. A policy that's coded with some kind of a ratchet roll-up but then it has a product code that shows that that rider is impossible. We have bad data in there. We need to have controls around being able to detect it or clean it. I talked earlier about consistent data definitions. We need to understand what goes in what field in our input files and what it means.

Infrastructure environment is also really important in terms of if I have a production environment, then I should have certain modeling runs that go to production on some kind of periodic basis and anything else that I do, any kind of ad hoc development, building new models, testing, should be someplace else. It shouldn't be on the same infrastructure. I should have two infrastructures, one that does my production stuff and one that does all my other stuff. There are some other examples around controls, things like my documentation, having things like official testing quantification of differences. In terms of model maintenance, and again this is really important, when I'm going to make a change to a model, I recommend having somebody own that process. Let's say, for example, my company's going to start doing economic capital and I'm going to build in the ability to calculate the value of my embedded options and guarantees. At some future point, what are my business requirements and how am I going to do it within each of my models? It needs to be documented. Somebody needs to own that process.

Once I deal with the requirements of what's going to go in and what's not going to go in, that I have documentation around, what am I going to do in terms of my model design, and basically to build specifications? Okay, we're going to calculate economic capital. That means we're going to do nested stochastic. We're going to figure out the future value of my embedded options and guarantees at some future point. What do I now need to do to my model to get that in there? Then you start the model development. In terms of the model development, there needs to be a whole control process around any kind of coding activity, any kind of validation, checking of codes. I need to have a stage process in terms of a unit testing, testing individual cells, integrating into my formal processes, having some kind of back-testing approach and then finally having somebody else, the person who wrote the business requirements, accept my process and sign off. I call that user-effective testing.

Then we need to go back and quantify any changes in our models. I made a change to my models. How does that impact my previous results? Let me quantify it and document it and put it away someplace so that we understand if we're getting any differences. Once we do this, the entire model needs to be deployed and some kind of a very controlled process started again. Now it's ready, and it goes into production.

Along with that I'm going to talk quickly about roles and responsibilities. As you're

viewing your new design and your new controls, you also need to think about what your roles are. This is going to be very company-specific, obviously. It's going to depend, but here's an example. In addition to writing things like business requirements, I need to have some kind of a governance structure. One thing that we see work very well is to have somebody like a rising star, like a new FSA, who can be responsible for what we call the model steward. It's a little bit of a queer name. If somebody has a better name, let me know. The model steward is really the only person who is allowed to accept controls. He's responsible for ensuring that all documentation is in place. He's responsible for executing actual production runs. The person who writes the business specifications is the process owner, so that could even be potentially his boss. The person who runs the UL line of business, for example, may say, "I need to get my UL secondary guarantees more refined, I need a better way to map my premium levels or something like that, I'm going to send the specifications for that and the rules and requirements around that to my model steward." My model steward is then going to put all the codes in it. He may not actually do the work, but he's responsible for ensuring that the work is done, and then once it's done he sends it back to the process owner. The process owner then accepts that the changes are there and in good shape. Finally, there's IT, and IT can have a range of roles, including being responsible for data and structure, and certainly for hardware.

MS. POLSGROVE: We had a thing on single-platform versus multi-platform because I'm changing my thinking on that. I think I used to be a single-platform person and I'm not sure I'm there any more.

One of the things that got me concerned about controls a couple years ago is we were moving to Moses, which is more of an open-code environment and it just scared me. You have students in there and it gets a little nerve wracking. It makes you lose sleep at night, especially when you start being further and further away from being a hands-on guy who does stuff; it makes you even more nervous. You start thinking about controls, and I agree with Ellen, we've had a lot of success in our hedging implementation, and now we're working on C3 Phase II implementation for this year. As we go about those, we're really looking at what controls we build around it. Sarbanes-Oxley (SOX), forget about it. You have to have the controls there and signed off on, and even on things that are statutory, whereas right now SOX only applies to GAAP. We're still building in SOX-compliant controls around the processes that count, because we really think you need to do that, it's coming. Controls have been kind of a thing.

FROM THE FLOOR: It's almost impossible to implement what you're supposed to with that.

MS. POLSGROVE: It makes people nervous, and it should. I think this is where we really borrow from the IT world on that production environment. You have the things Ellen talked about, and I brought something that my IT guys gave me of the types of controls that we had built around our things. These were things that we did

for SOX, which included change request tracking, regression testing, archival and access controls.

Disaster recovery planning was another thing that we've had to formalize. Where I struggle is if you're using an open-code system for pricing and you're doing some pretty complex stuff, you can't really have production controls around that. But what we have now is this peer review informal, high-level and everybody does his own thing and checks his own models in his own way. That makes me a little nervous, too, because we're a company in which actuarial students rotate through. Are you making sure that the knowledge gets handed down to each and every one of them? They all do it differently.

We have one guy who has recreated stuff in the spreadsheet, and we test everything out there. I hired a guy who is non-rotating and does programming. He's an ASA, so he's the guy who bridges the gap. One of the things I told him is we really need to figure out what we want to do with controls around things like pricing because you have to find some medium ground there. You can't do all the production controls because in our environment the speed-to-market issues would just kill us; nobody would stand for that. I know there are some companies out there that have the people who do the coding, and you brought it up, the actuaries don't. Do they do it in pricing? When you guys do pricing, do the actuaries code the liabilities into it when they're running like a pricing model?

FROM THE FLOOR: You can price the liabilities without.

MS. POLSGROVE: I'm from the aspects on just coding the liabilities with open code systems and you're making changes. What we're working toward is to introduce some test banks so we can do some regression testing, even have these types of things you always run through your models to test them. You should be able to choose where things changed. We're building that into one of our processes and we hope to roll it out for other folks to use.

We have actuaries who make the changes. Where we have partnered with IT is on the hardware requirements. We don't do this, and we don't have particular models. In our production environments we do, and it's made up of actuaries and IT. Like our hedging process, for example, we have to have controls for the SOX requirements. We have built an entire control system around that. Where I struggle with it is I still think on some of the other things you do, static and dynamic validations might not be everything you need.

FROM THE FLOOR: Do you partner with your IT or is it a large company?

MS. COOPER: We try just like everybody else. We struggle with this stuff, too. In our shop we do have IT that is responsible for ensuring that the hardware is ready for us to run when we're going to do production runs for a client. But we have actuaries who are actually submitting those runs, and we do not yet have any sense

of this. It's a little bit different because we are typically working with clients, but we're helping clients to set this up. The model steward is an actuary. The model steward is not intended to be IT. Our feeling is that the IT people are not comfortable with actuarial calculations, and they don't want to be, but they're willing to deal with data issues typically, depending on the firm. They're willing to deal with cleanup of all the data infrastructure, and they're certainly willing to deal with all the hardware stuff and typically the output as well.

FROM THE FLOOR: There's no manual. The interface is hardly something you'd want to advertise to people.

MS. POLSGROVE: We just aren't good at that stuff, I think. I don't know. This is where IT has helped us with discipline as far as writing the documentation. We haven't used it everywhere, but on places where we've used it, it's been pretty successful.

MS. COOPER: My sense is that as a general stereotype of who we are, I think we're very good at being quantitative. I mentioned in the beginning that I have a staff of 26. I see actuaries really struggle with needing to write sentences.

I force our people to think it through first before jumping in and actually doing it. They need to actually do some planning, and the planning requires documentation. I don't just want it in their heads. The process does go a lot smoother. It's painful, but it's in our heads. A lot of the time I think it's in our heads, but it's really not well thought out yet.

MS. POLSGROVE: Even with FSAs, when people leave, there's always that, oh, that's what that person did and I don't know what to do. We're famous for calling people at companies, saying, "Can you help me out here?" I don't think it's just actuaries, I think it's human nature in general. I've seen other people do that and there is a point I will make of something we found that helps with this.

That's like with the controls as well. I have meetings once a month with my counterpart on the IT staff, and he gets along well with me because his dad's an actuary, but we sit down and talk about it and he's like controls, controls, controls. If you're in a pricing, how do we find that in-between state that we can feel relatively confident about the results? I think we don't have change control and an online Web-based system to say you made a change to a model. That doesn't work necessarily for all uses, and he and I talked about how you have to sit down and say for this model here's the level that I think I need. I think it can be lesser stuff.

We partnered with this particular IT team, and it's one of a couple in the building. But these are the guys who have helped us get the servers, that have helped us rewrite some of our production programs into other languages, working with an actuarial student. They really worked together to get us at on-time speed and all this other good stuff, but one of the things in partnering with them and building that

relationship, is that as actuaries we never had project managers. We'd never built an actuarial model with a project manager. When we had a large modeling job and we wanted to build a large corporate risk model, we had a project manager. They did things like put together a charter, got executive sponsors; they did all the organization of identifying the key resources. It helped us figure out that we promised this on four critical projects and there just aren't enough hours in the day, so resource management. It's silly in a way. But going through the actual process of writing a charter and saying here's the goal of what we want to do, here's the inputs we have, the outputs we have and they take care of that housekeeping, documenting even the decisions you make along the way. Because we're famous for getting actuaries together saying that's a good idea and then we all walk away and two months later something similar comes up and we wonder what did we decide and why did we decide that? Just making sure that we keep that discipline of keeping the project going, making people accountable so that things get done in the time frame is important. That's something that has been phenomenal for these types of huge model undertakings, because it really keeps it going. Bring the executive sponsors in once a month to present results, and you'll get the importance of that in a minute. Those are things I think we've really had a lot of success with. There's still a lot of work to do.

I think it doesn't work for every purpose. That's when we do it. That's when we bring it in. I think the next step is: What can we learn from those to use in our every day and then how do we make that the norm operation, best practice throughout the organization when somebody builds a pricing model or builds a cash-flow testing, an embedded value model?

You've gone through and you've said we want to build a model for a particular purpose; maybe we want to retool our entire cash-flow testing. So how do you get that done? That's why I talked about us having project managers who really helped us with all of these questions. It also helps you see where systematically you maybe have some holes in the types of people you want to hire. Again, we've gone to hiring some corporate modeling resources that we loan out. I'm trying to think if it's Allstate that has a corporate modeling staff where they do control all the model changes. We haven't gone that far. We've gone more to where we have resources that we can loan out periodically. Lately they've been doing things like implementing C3 Phase II, but that means somebody else doesn't have to do that.

Do you have the right resources? Again, I think we have open-code systems, so we've tended to find out that we really want to start generating people who are more programming than actuarial. I never was. I came into the world in a task world where you didn't have to program much and I think it's kind of interesting in a way. Folks who maybe started out in the ATL world where you knew ... I did use ATL for a while and I could never figure out what all those fun little symbols meant, and I actually have a claim to fame at my office for breaking PPF. In certain areas we're still using it and the pricing areas and individual life are still using it. You can't import tables or something. I think it's me. It always worked before I showed up. I

think it's interesting now with the newer tools that we're kind of going back to being more of the programming with these open systems and the ability to go in there and code and C++ directly. Some skills that haven't been important faded out, and now they're coming back into play.

How do you make sure everything gets done and how do you get senior management? Most importantly, how do you make sure you rein the budget in?

MS. COOPER: We basically tried to think about the schematics from day one to sort of putting everything that you ever wanted into production. Obviously, this is extremely high-level, but typically the first step is to do a real self-assessment: How bad is it? What's our current state? That means you have to take inventory of every part of the company that has a model and what it's using it for and what the strengths and weaknesses are. Typically, you'll find that the people who are actually running these models know what they love and know what they hate and really know what needs to be improved. They know what secret sorts of manual manipulations they're doing to try to get at the right answers, and they know the questions that are coming down that they can't answer.

We've had some companies that have said to us, "Can you go around and we're going to identify 25, 30 people within our organization that we want you to interview?" We'll come up with appropriate questions and basically pull it together. The feedback is quite helpful if you go through all different organizations. You don't have to hire a consultant to do this, by the way. You can give this to somebody internally and just say let's really assess this from the most junior level to senior management and figure out what's broken. Do we need to blow the whole thing up or are there just pockets that if we just fix them we'll make the whole process run smoother? The whole time that you're doing that, you're really building your business case, because you're really learning inside the organization. You're going to be able to recognize pretty early on where the low-hanging fruit is. You want to be able to map out, based on where we are and based on what our weaknesses are, to get some kind of consensus.

You're going to get common themes across the entire organization. You're going to get the same things—cycle times are too long; ad hoc is too hard; I want to be able to do more what-if analysis and I can't. Those are some of the common themes. Run times are too slow. How can we fix them? Where is the low-hanging fruit? Where do I start? That will help you to start getting buy-in. The most important thing is you need to identify who the owners are, who has the budget internally and what you need to do to get them convinced that this needs to be fixed. Are there competing needs out there? What else is going on within my company that might take higher priority? That's also really important. Then from there, as Tracey was talking about, build the plan, build the private charter, figure out what the key deliverables, the milestones and the budget are over a long time horizon to give senior management some perspective. How is this going to help us not just with this effort, but with all of our other efforts?

Then you start to get into the actual project. There are a couple of important points here. One is that this is going to take some time and so senior management or whoever the sponsor of the project is needs to be told, "You have to be patient, we're going to do this right, and we've never done this before. We're going to document this time, we're going to really think through the issues, be patient with us, here's how much time we need." Come up with an initial private study or a quick hit or something to demonstrate success early on, something small that can be accomplished within three months. That will then help to build the rest of the business case. If the five-year plan or the three-year plan is implementation, and it just seems too overwhelming, give them the three-month plan for a quick hit. Our theory is that that quick hit could be very powerful. Then you could go to all these other things around production and refinement and those various other things.

MS. POLSGROVE: With the changes in the actuarial software that's available in the market, you might need to decide whether you would want to choose a new platform or undertake building new models.

I came from a consulting firm. A lot of companies said we want one actuarial piece of software that does every single product we have; that's the only thing we're going to use and we want consistency. I know our company did that in about 1998. I think things have changed a lot and I was probably in that camp that said, "That's a great idea." I think it's a kind of utopia. We never got there and the utopia says, gee, the pricing actuaries build it and then when the time comes to do cash-flow testing, they can test those models on stochastic cash-flow testing and then everything will be wonderful and great. You still ended up with people for whom the tool wasn't quite right, so they built something for a spreadsheet or they do ad hoc, off-line sort of things. One of the things that I'm at a turning point deciding is: Do I go with all one platform or do I use whatever the best software for a particular product or job is? Then if I do that, how do I make sure I'm comfortable with the results across the different lines?

MS. COOPER: Part of that single platform idea originally came from the fact that we didn't have databases. We were running these models that were producing files, and we didn't have a good way to aggregate information. But now that a lot of these newer actuarial models at least have some kind of a flat file that can be ported to something else, you really can aggregate output from different models from different places and put them together to be able to look at one big picture.