2002 Valuation Actuary Symposium
September 19–20, 2002
Lake Buena Vista, Florida

Session 16PD
General Modeling Techniques

Instructor: Rebecca L. Burton

Panelists: Rebecca L. Burton
Scott Kane
David A. Ricci

Summary: Session topics include model hierarchy; statistical, deterministic and stochastic models; real-life situations; how to organize a model effectively; structuring the model to fit the project or report; data sources/integrity validation techniques; documentation; and constraints and limitations.

MS. REBECCA L. BURTON: This panel discussion will be on general modeling techniques for actuaries with moderate experience. We'd like to tell you a little bit about ourselves. I have been with Tillinghast a little over five years now. Before that I started my career at Equitable in New York where I worked as a student rotating through several departments for a couple of years. I've been involved with many large modeling projects, which include development of actuarial appraisal values, embedded value calculation and review, and development of the closed block and actuarial contributions to surplus for demutualizations. For the past couple of years I've assumed the role of project manager for these types of projects. I hope to share some useful information regarding general modeling techniques. My co-presenters are Dave Ricci and Scott Kane, and I'll let them take a moment to introduce themselves.

MR. DAVID A. RICCI: For the past 12 years I've been involved mostly with corporate financial work, asset/liability management, and statutory GAAP reserves. Before that, I developed models of distribution systems on a consulting basis to large mutual companies, and since there aren't anymore left, I guess my market dried up.

MR. SCOTT KANE: I work at AON Consulting, known to most people as the Avon Consulting Group in Avon, Connecticut. This is my second year. Before that, I worked at Ernst & Young in the Hartford office. I worked at the Chalke/SS&C organization in Virginia, and prior to that I was at National Life in Vermont, where I started my career. I've done quite a bit of modeling using commercial actuarial modeling platforms and home-grown platforms. I've run into most of the different issues you can run into in modeling.

MS. BURTON: We, as actuaries, use models in almost every area for which we are responsible: cash-flow testing, risk assessment, asset/liability analysis, financial planning, GAAP unlocking, and other regulatory requirements. These models need to address the issues we are facing, and they need to provide a meaningful answer or a range of probabilities. In addition, the models we develop need to be independently verifiable. We hope that this panel discussion will help you gain a greater understanding of the process of successful model development. Dave will begin by telling you a bit about project management.

MR. RICCI: Every modeling opportunity can generally be divided into three types. You have standard repetitive modeling that generally falls along like mortality models for XXX, standard product development models and general planning. Then you have longer-term developmental models that are basically product updates, sensitivity analysis, substitution for the current process and the like, and then you have short-term, reactionary, want-it-yesterday type of models that are usually the result of some high profile issue. I'm going to deal with the latter two types because, in general, you'll have plenty of time to associate yourself with the first, and I'm going to talk about them in a project management atmosphere.

For standard project management, you take a look at the issues that generate the need for the modeling in the first place, but they're mostly economic, improving profitability. There are many sub-issues. Distribution systems may be out of whack. You may need to promote better

relationships with the regulators, quality sales, and that kind of thing. Then you're going to define the objectives from the issues and be sure not to exceed the boundaries.

A large part of issues relate to profitability. Do you need to generate a model that will enhance the increase of productivity, or will it decrease the cost? Do you first need a model of the current situation? If you don't have a model already prepared for that, you may need to do it anyway to get into the new model in order to make those calibrate correctly. Then you need to have a model that's sufficiently flexible that you can generate alternatives.

What needs to be changed for current model assumptions? Do a sensitivity analysis. In the new model you will probably verify your results against experiences as best as you can. You may want to modify the performance grid of your distribution system so that you're paying more money for more quality results. It might include a distribution bonus of some kind, or you may be evaluating one distribution system versus another.

The model structure has to be such that you stay within the limitations of computational results, and you have to be able to eventually produce results that you can sell to whomever is requesting the model—senior management or whatever. You must generate answers that will be approved by management. In consulting, one of the overriding principles we use with the task force is that senior management will rubber stamp it, and the reason we did this is to make everybody believe that they had—and generally this was true—a responsibility to get things right, so that question would never come up.

I guess you need to know from your model or before generating the model what kind of risk profile you have in the current product or the situation, whatever it is. Do you have a narrowly defined risk around the mean, does it spread out like a normal distribution, or does it have fat tails? Whatever it is, you must be able to look under a wide range of scenarios, generate what-if scenarios of various types so you completely cover the waterfront, be able to migrate to alternatives, and find out whether there is empirical proof now that you have generated an acceptable alternative that will bring you closer within the profitability constraints.

You must use the model as a test of financial results. Does it trend right? Are the results reasonable in light of the model? Many models are generated just to basically understand that the valuation system or the financial system or whatever's coming out of the ledger appears to be correct. The ultimate users have to be comfortable with the results. They have to be believable. You have to reduce the possible impact of negative results, kind of maximizing by minimizing. In many situations these days you're faced with cutting your losses, and models can help significantly in doing that. Of course, a lot of changes in products and incentives will subsequently change behavior. They may not change it the way you like it, but they'll change behavior. So you have to be prepared to understand the consequences of that.

I'm going to talk about variable annuities later in the study. The problem I think that exists with a lot of companies is, can there be an acceptable premium developed for the benefit or is it just not possible? Is there a reasonable intersection between risk and profitability? Improvements can create disadvantages for the current policyholder to the extent that the profitability on that business is reduced to a greater extent than if you pick it up on the new product, of course. There has to be a perceived desire for change. What you come up with has to be seen as being beneficial going forward. It may be a rough sell—maybe something that people won't initially like to do, but it may be very important for you to bring them to it. As I mentioned before, there's profitability, and then there's risk, and we hope there's some area in there where they intersect, but the model may come up with the fact that there isn't. Then you have an entirely different problem on your hands.

What kind of profitability measurement do you want? Do you want ROI or ROE or a number of different measurements? A lot of that is important in today's interest environment. You're not going to generate sales when the product that you developed is too complicated, when there's not enough perceived compensation, or when it's perceived as not being sellable.

MS. BURTON: Our next topic is on model organization or perhaps how the pieces fit together. I'll point out the components of a generic model, followed by an example of a hierarchical structure. Then I'll follow that with considerations that you should bear in mind when developing your model. My discussion on model structure assumes you will be using a generic asset/liability model. Of course you might just be working with an asset-only or a liability-only

model. In that case, you would want to make the proper adaptations to the model I describe. You might also have special considerations, depending on the type of project you are completing. I would like to move onto the next slide containing a hierarchical structure and discuss each of the components.

Your model will have, as its backbone, what I call a master or control file. Perhaps this is the software you are using or perhaps it is actually a single file. Nevertheless, it will be the source that links all of the pieces of the model together. It will be the engine that will complete all of the processing. You will need some liability data that you will put into a structure accepted by your model, and I will discuss data in more detail in the next section. Generally, liability data used by your model will be in the form of cells or records. Your model will know how to project these data forward based on the product assumptions and definitions you input into your model. If working with an asset/liability model, you will also need your existing assets backing the liabilities just described.

As with product assumptions and definitions, you will also need to include asset definitions, which will define for the model how it should project the assets forward. These asset definitions might include investment return, default rates, and investment expenses. You will also need to give your model a reinvestment strategy or a negative cash-flow strategy. As assets mature and the monies need to be reinvested, the reinvestment will be done according to the strategy defined in the investment data file.

If the model has negative cash flows, the model will need to know the strategy to deal with negative cash flows, perhaps borrowing at an assumed rate or selling off assets in a particular order. If you're running a stochastic model, you will need to attach a set of stochastic economic scenarios. At a higher level you would define aggregate-level corporate data, perhaps apply tax assumptions, or if you wanted to alter some aspect of the model at a higher level, apply a multiple. Perhaps all expenses should be increased by 10% or mortality in aggregate should be reduced by 5%. These might more appropriately be called global assumptions.

Particular to each model, you need to provide calculation specifications. Perhaps you want to run your model for 30 years with a valuation date of December 31, 2001. What discount rates do you

want to use to determine present value of profits? These assumptions are particular to the model as a whole. Finally your master or control file, after you complete the projections, will also provide you with the output that you need in some organized format. Perhaps it will give you an income statement and balance sheet or exhibits relative to the work that you are doing, and perhaps you can customize the reports you receive from your model.

As I mentioned before, there are some additional considerations you should bear in mind while developing your model. Really think what it is you are trying to do. If, in your mind, you can break your model down into these components just described, you can develop an organized plan to set up your model component by component. Having this component-like structure will allow others to help you with your model. Be sure to automate as many processes as possible. From your data work, to table making, to running, to generating reports, you will have to do these tasks over and over again. Manually repeating processes introduces a definite source of error. Automation of as many of these processes as possible will certainly reduce this error. Be sure to allow for more than one set of hands. Have someone work on the data perhaps while others create tables, code specifications, and assumptions, and work on spreadsheets for static and dynamic validation or report compilation.

Maintain your model in such a way that it is easy to expand it if needed, to add more products, plans of the same type of product, or asset or liability data. We have a tendency to try to own our model, thinking we can do the work the best. Be able to divorce yourself from it. What if you want to go on vacation or you need to take an emergency leave? Won't you feel much more at ease if some others could take care of the model while you were away? Be sure not to hard code. For example, if your software allows you to alter the code and say you want to do one thing if the face amount is greater than 500,000 and something else if it's less than 500,000, put that 500,000 as a variable that is part of your set of assumptions or set of inputs that will be called on by this additional bit of code rather than hard coding it directly in the code that you write.

Use efficiently structured output compilation spreadsheets in which you compile your results. Set these up so that your process can be automated. I work with Excel or ACCESS, using a file that reads in output from the model at the punch of one button. This provides me with results

summarized in a way that is acceptable for my purposes at the punch of another button. It's not necessarily that simple. It takes a little bit of initial setup, but it's a process that can be used over and over again for future and additional runs of the same model.

For some housekeeping considerations, keep models on the local area network, not multiple copies on the hard drives of the people working on the project. Maintain and organize directory structure. I often harp on the importance of organization. Things will probably go wrong with your modeling project, but keeping yourself organized will get you back on track and will keep you progressing successfully to the end of your project. Maintain separate directories for project management files, documentation of all your work, spreadsheets where you're compiling results or doing additional analyses, and another directory where you have your data work. Try to keep your models in an entirely separate directory.

For example, I will have separate directories for data work, and then I have a directory for models and for output, separate directories for spreadsheets for value compilation, miscellaneous analyses, or documentation, or if you need to produce a report.

Archive or zip-up models and the spreadsheets that were used to compile the results are of importance, perhaps because they back the final submission of embedded value results. If I communicate results three times to a client, with the last communication being the final results, I like to archive the models coinciding with the final results and specially label spreadsheets containing interim results that have been communicated to my client. Lastly, you should be completing your documentation every step of the way. Now, if you do this, I'd like to see a show of hands. It seems really elementary, but I think we're all guilty of perhaps waiting until the end of the project or even weeks or months, in some cases six months, after the close of a project. Ideally on day one of your project, you should set up some kind of filing system, a box or an expandable file folder, where you will place all key elements of your work as you're going along.

Our next topic is data. We'll discuss data sources, structure, and integrity. Your data work process will commence by starting with data and information requests from several people. You will want a seriatim data extract with information from your valuation and administration systems. The fields that you will need include those required by your model, as well as

additional fields that will assist with grouping your seriatim data into model points. In addition to the data extract, you will also need statutory annual statements and exhibits to which you will reconcile and reserve other items from your data extract. You'll also want these items in order to validate your assets or to look up other important information. Actuarial memorandums might also contain other information to which you can reconcile, and experience studies come in handy when developing or validating assumptions. Bloomberg might be one source of asset type information that you could use. The data and information list are certainly not all-inclusive. Depending on your project, how much additional or new modeling is required, and how many changes do you need to make to your existing models? You might require additional bits of information.

You must decide how much grouping you would like to do. The first thing I always do with data that I've determined to be good is to look at the distribution of the business across the plan type and to get a feel for where the natural splits of the data are. I might have ACCESS open alongside Excel and may run a series of queries to try to split the data up in multiple ways. Maybe I'll copy some of these queries of information over to Excel to preserve them. I'm generally just trying to get a feel for the various sizes of the business in relation to each other, the size of one model plan in relation to the rest. I do this in order to develop for myself a set of rules, how we'll go about grouping and building the model points. I might make some rules for myself (for example, no more than five ages for a plan). If it's a super large plan, I'll throw in a sixth age. If I have plans that are less than 250,000 face value, they'll get two ages, and that's it. It's not always that easy, but these are some kind of general rules that you can make for yourself.

The trick is that you are trying to maximize accuracy under the constraints of budget and time. This type of work is really more of an art than a science. We have a tendency to split model points out into additional plans, issue ages, issue years, and even more. Through the data work and model-building process, we really need to consider the materiality of the various pieces of the business. If we include more model points, which we often try to do to achieve more accuracy and I think we all have this tendency, we run into greatly increased runtime and a model that is much more difficult to maintain and much more prone to error. Smaller, more immaterial

and extraneous business should either be grouped with larger model points or dealt with approximately in spreadsheets. Be sure to model smart, not big, and remember that more often than not you can get away with the more compact model that will still produce reasonably sufficient results.

Now, for new business you'll want to obtain forecasts of future sales production and distribution. You'll also want to be sure to consider past trends in sales while bearing in mind changes to the current or recent sales environment. Using model points is usually typical for modeling new business, as it would simply be guesswork to try to create seriatim forecast policy data points. Seriatim data is typical for your assets. Some of the data fields that you might want to obtain would be book value, participating (PAR) value, market value, coupon rates, maturity dates, and information regarding options embedded in the assets or explicit options.

Just after you receive your data extract, you must be aware of problems that could cause you to have a meltdown. For example, perhaps the data that you receive is not exactly what you asked for. You might have duplicate records or blank fields or null fields that should be populated with important information. You might have inappropriate text instead of numbers. You might not understand what you have. You've submitted your data request. You've received back the exact number of fields that you've submitted with these cryptic headings above them that don't make any sense.

Don't just assume that you know what those cryptic headings mean. Perhaps it's simply the incorrect elements are missing. Maybe there's bias in the data, or it's not consistent or reasonable. Say, you do a quick check and find you have 80% female smokers. The problem generally can be taken back to a communication breakdown between you and the programmer. You've submitted your request, and you think that the data are perfectly correct because you've taken the time to put together this very specific data request. However, it is this assumption that will get you into trouble.

Therefore, to determine that data are good and in order to move onto the model-building process, you must do the following: Make sure that you understand it. Ask yourself, where the data originated? What was it previously used for or has it been solely created for my purposes? Be

sure to inquire about the definition of those fields with headings that make no sense to you. Now, this one is very important. After you get your data, you want to tie the reserves from this data extract back to your annual statement, to your figures that have been published. Perhaps you are higher, and in that case it might let you know that you're double counting some of the data, and if it's lower, then you definitely need to dig in and understand why. Make sure you get ahold of the pieces that are missing.

You'll want to perform some reasonableness checks. Print out and examine several records and check for duplicates and blanks by using simple access queries. Have expectations for the distribution of your data. What do you think that the male versus female ratio should be? What do you think the smoker versus nonsmoker ratio should be? How should it look across the ages? Run some queries in ACCESS or some other database facility and check these. After you've gone through all of these analyses, write out your list of questions and go meet with the individual who collected this data for you face-to-face, just to make sure you are both on the same page with what you've asked for and what you've received.

I can't stress enough that, at this stage, there's never enough time spent. Maybe at the start of your project you budgeted a week to do all of the data work—to ask for the data, review the data, and get it into a form that you can use. I'm not saying that you need to sit there and look at it longer than you really need to. I'm saying that three weeks down the road, you'll still be working with and looking at the data because you've had to go back three times for a new request, because you've found these problems. Just expect to spend some time with your data work. Our next topic is on validation, and we move to Scott.

MR. KANE: You've got your model, you've got your data, and you've run it. You want to validate it. What's the purpose of model validation? You want to know not only that the model produces reasonable output for reasonable input, but that it is reasonable input. That's back to what Rebecca was talking about on your data. The entire thing—the data, the calculations, the output—all needs to hang together. Even if you have a perfect calculation engine, if you put garbage in, you get garbage out. That's been drilled forever.

The first thing we want to do is determine the right answer, and you'll want to do that before you run your model. You cannot validate your model unless you've got something to validate to. This has to be something published that would be acceptable to your management and to other people. Therefore, you need to establish up front the sources of your information and your tolerances. You can't look at it after and say, that's pretty close; that's good enough. Before you even run the model, you should know what your tolerances are that you're going to require to be acceptable. For instance, you may want policy counts to be exact, or face amounts to be exact. You want your premiums to be within 1% of what you project on a different model. Therefore, we do static validation. This isn't just the validation of the initial counts and balances and stuff; it's validation of the results of your model with static assumptions. For example, policy charges, coupon payments, benefits, investment income, and so on.

Dynamic validation is a validation of how changes in the assumptions flow through the model. Sometimes one assumption depends on another. If you have a dynamic lapse, changing your interest rate could change your dynamic lapse, but even without a dynamic lapse, you want to know how a change in assumptions is going to affect your model. We're not just looking at the results after we've changed an assumption for the level. We're looking at both the direction of the change and the magnitude of the change.

Back-testing always sounds like a great idea. It's not often done. One of the problems with actuarial models is that we are making projections of the future. When we produce model results, we don't know the future. When the future is here, we don't want to go back and revisit last year's model. We might have made improvements to the calculations or refined the data-mapping process. Most likely the economic conditions or sales volume were quite different than the assumptions used in last year's projections. There are a lot of reasons why the results produced by last year's model aren't accurate, and we're generally interested in our current model and changes we've made.

If we don't want to actually rerun last year's model by going back and putting in the actual economic conditions or the actual sales volume because our model may have changed, we could take this year's model and go back, put in last year's data and last year's assumptions, run it, and then compare it so that we feel that we're testing our current model. This is assuming that it's

probably changed since last year. If anything is out of line with what it should be, then you would want to know why. Even if your model is working as you expect, the results it produces may be significantly different than actual. There may be things that you would want to look at.

What if your model doesn't validate? Don't use it, or don't run the results and give them to your management. By not validating, you produce results that you cannot explain. You've done some checking, and it doesn't make sense. So what do you do? You have to dig into the problem. Part of any validation should include some detailed checks at the most basic level. Even if you've used the model many times before, it is worth the extra time to do some review at the most detailed level. The more complicated the model is, the more important this step is. Also, actuaries are very good at rationalizing results, especially with complicated models. I've seen many times when the results were not what was expected, but after thinking about it for a while somebody comes up with a plausible reason why the results make sense. Don't fall for that trap without actually going back and checking. If you've come up with a reason, you should still check and make sure that it's right.

An audit report is a report of the entire model. It ties with what Rebecca was also covering, but it should cover the process you went through in building and validating your model. It should be of sufficient detail that someone else could start over and duplicate your results, and it should cover the entire model from data gathering to the final reports.

Next we have model control issues. How do you deal with model components outside your area of expertise? There was a proposed standard of practice for this. There has been an adopted standard of practice for property and casualty actuaries but none for life and health actuaries. Some examples are economic scenario generators or mortgage prepayment models. If it's just a black box, and you have no clue about what it does, then that's insufficient. Maybe you're better off using something simpler that you would understand, so at least you can explain the results.

Who's responsible for passing judgment and understanding all components of the model? You should be able to pass judgment based on your knowledge and reliance on experts, understanding at a detailed enough level what constitutes reasonable outputs and reasonable inputs. You must also know that it's an appropriate model. Who's responsible? There are generally many people

involved. It's not necessarily the actuarial student running the model. They should try to understand all they can, and they should be looking at results and trying to understand them, but it's really the responsibility of the actuary who's in charge of the model and who is responsible for sending the results up to management. For example, for cash-flow testing, it's ultimately the responsibility of the appointed actuary who's going to sign it, but in most companies, there are going to be people to whom authority is delegated to understand the model. Another point is disclosure. If you use something outside your area of expertise, even if you've done all this, disclose it.

With model documentation, I'm changing directions a little bit. By a model in this case, I'm talking about your modeling platform, whether it's your Excel spreadsheet, your expensive actuarial software you purchased, or something that somebody else has developed. I'm not really talking about the data here. For this part of it, there should be sufficient documentation of the modeling platform you used to prevent errors and, for the most part, so that people can follow it. That includes everything from how the assumptions should be input, data file formats, and actuarial calculations. There should be algorithms that you can follow and understand.

How good is good enough? If you have six weeks to do a project, you don't want to still be running your model the day before results are due. You should have time to gather the data, run the model, validate the results, and produce both the documentation of what was done and the final report. This should be done with enough time for someone else to review the results and, if there are questions, go back and do some additional analysis. It is better to produce results that you know are valid and are a couple days late, if that's an option, than to produce results on time that you're not sure of. However, you're better off still producing results on time that you are sure of, and if that means using a simpler model to get it done, then you should do that.

We discussed handling the data and validating the model, and we could have been through all of that. You also need to make sure that nothing screws up your model, and it doesn't even have to be something somebody else did, like your actuarial student. It could be you. You might have done a sensitivity test, just changed an assumption, run it, and then forgot to change it back.

Good practices should be defined ahead of time, and Rebecca's example with her model structure, all her directories on the network, would be good practice. You have to know who is responsible for what, how procedures for certain things like sensitivity testing will be done, backup procedures, naming conventions, and documentation of changes made. Now we're onto Dave again.

MR. RICCI: I'm going to briefly go through this part of the session. This basically is how you wrap up the process and the report. Oftentimes, you don't really allow enough time to do this, and it becomes very critical, particularly when you're creating a model for an extensive process that has to be implemented. The three final stages pertain to developing your alternatives. First, there may be a very small stage if you're really just trying to verify a given situation, given a very large number of constraints, but it could be something completely different when you have a task force situation and people who are constantly suggesting changes to the approach, process, distribution, or whatever. Then you have a resolution phase when everybody needs to be on board. Not everybody needs to be happy, but at least they have to be on board and understand why you're going in the direction that you're going. Then, once everybody has resolved that senior management has approved it, you're in the application phase.

There are a number of considerations for an alternative development. I'll go through these quickly. Reinsurance may be very significant this day and age. You may have policy limitations with regard to reserves and underwriting classes, and asset returns may have some asset/liability modeling situations to address. There will be a transition if you're developing something new that is basically evolving out of something that you had, of course. With new products, there should invariably be comparisons with the competition and limitations that are structured because of that.

You have to have a fairly good concept of what the market share's going to be. If you're a Johnny-come-lately, you can't expect to bull in and produce a large market share. There has to be something special, or you have to have a distribution system that's dedicated to what you're trying to do, or both.

This is in a context of our case study, the variable annuities. You may need to have auxiliary models that test different levels of guarantees, and going back to Rebecca's statement concerning data integrity, you might have a situation where you're producing a lot of data based upon your current structure. However, at the time it was developed the product people were rushing it through the house; therefore, there were some pieces left out of the administrative process. They really care about it initially because maybe these guarantees didn't click in until the end of the surrender period. Now they're getting tough, and there are certain holes, and I guess it really goes along your lines. Scott, it's very important to come up with the right answer. No one ever remembers that you came up with the wrong answer on time. They never remember that. What they remember is that you screwed up, basically, and it's taken me 30 years to figure that out. Maybe you all can beat me.

Interest bonus refers to any kind of bonus really. Is it lapse supported? Resolution is basically that you have now developed the appropriate profitability threshold. You determined the attractiveness to the market. You've set up a firewall with new business. You've made sure that there are no complications with regard to regulators, which is a big consideration, and it would be even a bigger consideration going forward. You have to have sufficient time to develop and launch it. Just getting it done is like the hot potato. The next person who gets it has to develop and watch it, and he has a half a day.

This is the application phase. You basically have to decide what kind of void you are filling and how quickly you're going to do it. How do you get the word out to both your customers and your distribution system, or if your distribution system is your customers, then how do you get the word out to them? It has to be an all-out team effort. You have to monitor and control the process so it doesn't get out of hand and depart severely from what you intended it to be. You need to update the models frequently as the results roll out to make sure that you're on the right track. Everything that's developed always can be refined, so you have to go through the process of making those refinements and making sure you're on track.

MS. BURTON: That concludes our remarks on the general portion of our presentation. We're going to move on to part 2, our case study. As some of you may be all too aware, in today's environment, the downside risk of variable annuity secondary guarantees has increased. By secondary guarantees I mean guaranteed minimum death, income, annuity, and withdrawal benefits. As fund values are tanking, more and more of these secondary guarantees are becoming in-the-money.

Let's think about a worst-case scenario where the variable annuities issued in 2000 experience an average 50% decrease over the next 20 years. For the guaranteed minimum death benefit (GMDB) alone, would asset charges cover this? I don't think so. Consider just a trivial GMDB tested with 100 randomly generated scenarios. There's a high probability that under some of these scenarios, this trivial GMDB would cause the variable annuity block to be insolvent. Imagine what this test would look like with non-trivial GMDBs that are coupled with other secondary guarantees. Now we have this risk. Who cares about it or that it is properly analyzed? All of those folks who were worried about the money and the solvency of the company care.

MR. KANE: I'm talking about modeling techniques for this risk that Rebecca just mentioned, the guaranteed benefits on variable annuities. Since we're trying to capture the risk of these guaranteed benefits, we need to make sure that our model accurately reflects the benefits. The risk is that we'll have to pay benefits above the account value, either when someone dies for a GMDB or when someone annuitizes for a guaranteed minimum income benefit (GMIB). We're trying to quantify how likely we will face the benefit and how bad it can be. In a perfect world, we wouldn't have to group our data at all. We could run every policy through using the exact information.

However, I assume we want our answer in a reasonable time. We also want to run at least 100 stochastic scenarios, and we have to group our data. We don't want our grouping to hide the risk. If only 10% of the policies are currently in-the-money, we don't want to group these policies in with the other 90% so that none of our model points are in-the-money. We may be better off grouping males and females together or doing larger groupings on issue age. We have to keep in mind the risk we're trying to measure and the way we group our data; we're not just losing it.

Since benefits are paid at the whole-policy level, and a GMIB and a GMDB cannot collect on both, I suppose that a GMAB and a GMDB could collect on both as long as you died after the AB, but we want to model our policies holistically. We don't want to try and slice out and look at GMDBs here and GMIBs there. We don't want to oversimplify that different benefits behave differently. In this case, we wouldn't want to group a three-way and a five-year rest together because of the way the account values would work. If we just called them all a four-year, we can miss some serious risk.

I'm assuming that the different fund types would have different risk characteristics. They have a different mean return and different volatility. If so, we want to take that into account. Generally, these are the AG 34 fund types that you have to model, and there may be within certain of your products another type of fund that has significant amounts of money that would be worth modeling as well by itself. Also, by fund allocation, policies that are 100% equity or 100% bond will get a different risk profile than policies that are 60% equity and 40% bond. I said that too much grouping can hide the risk exposure.

In this case, there's a guarantee and that relationship to the account value is the crucial part of the risk. It's worth more than your time to make your analysis more accurate. You want to group it by the amount of risk in each. Put cells of policies with similar amounts of risk together so that when you're running your scenarios, you'll be capturing 10%. The ones that were more than 10% out-of-the-money won't pay, but the ones that were within zero to 10% of being in-the-money will become into-the-money, and we want to capture that.

If we're trying to quantify risk, a single scenario doesn't make sense. The probability of any single scenario coming true is zero percent. We're trying to put a range on the risk probabilities. Using a number of stochastic scenarios, we hope to cover a range of likely possibilities to get a risk profile curve.

A good scenario generator is crucial. If we have taken the effort to model the different fund types, then we need a generator that will actually model the different returns in a realistic way. There's currently a lot of research going on in scenario generators. One of the key drivers is that

people trying to understand risks like this in the last two years have certainly driven that point home. The Variable Annuity Guaranty Living Benefits (VAGLB) Task Force has been at the forefront of this, and one of the major discussions on scenario generators is how well they measure the tails of a distribution. It's not generally the 25th to the 75th percentile where you're worried. It's the last 5% or the last 2%. To get a generator that realistically gets that, you want to know is it really the last 2%, or is it the 0.2%?

We want to know what information we're trying to get out of our model, so that we can write our reports and present this to management. We're trying to capture the risks of these guaranteed benefits. We want to be able to present it in the balance sheet information. This is currently how it looks. We also want to get the income statement in our projections. We want to be able to isolate the claims and the fees for these guaranteed benefits, including reinsurance premiums and reinsurance claims, so that we can see the part we're looking at is not just buried in a general income statement.

Are there analytics we might want to look at? Is a present value of distributable earnings a value at risk or a conditional tail expectation where you would actually look at the average of the last 10% or the last 5%? Rather than just looking at the 95th percentile, we're ignoring the bottom five. You would actually average the 90 conditional tail expectation. Even if the last scenario is just terrible, that gets in there. You don't ignore it. Now Dave will talk about maintaining results.

MR. RICCI: For this particular case study, some of the results will be counterintuitive, and people are going to need to be brought on board. For example, people may think what goes down comes up, but not necessarily. What goes down, gives you a certain distribution of net amount at risk among the various plans, but just because the equity markets rise doesn't mean your risk is going to be reduced proportionately.

If your Standard & Poor's rate went back up to 1,200 or wherever it was before all this nonsense occurred, you would not be absolved of all the risk. There'd still be significant risk there. Your model should indicate that, along with many other things, and determine what the cost is based

upon various policyholders' levels. If you're doing a GMIB, you need to talk a little bit about utilization rates. Policyholders loathe going through the annuitization process, but as the net amount at risk gets further and further apart, they have no choice. They could lose 50% of their guarantee.

This material is mostly Course 7 type stuff that I stole. Anyway, you need to convey whatever message you come up with from this modeling process and motivate action. You have to make it reasonably easy to read and understand, particularly for the people who are at the decision-making point. With the variable annuities, there's going to be a lot more people who understand the technicalities of it, but that doesn't absolve you from trying to get everybody on board. For example, the accountant may not understand it at all. It's good to have charts, particularly something like distribution of risk of scenarios and descriptions of the difference between the state of in-the-money and a kind of reversion to the mean assumptions. You need to take time to develop reasonably good pictures of the point you're trying to get across.

What's in report writing? Obviously I'm talking mainly about a written report that's going to go to people. Many times you don't do that. Perhaps you have a set of technical papers, and you submit something orally. For something like this, you probably will have a written report, an executive summary stating the reason for the report, general conclusions, and the methodology, generally used for people who have to make the decisions and need to understand, on a high level, what the process is. There's an introduction stating the basic cause of this particular modeling and what exactly precipitated it, what kind of information you used, and the data section. In the case of the variable annuities, there will probably be a fairly extensive data section. You need to explain a lot revolving around what data points you chose and what kind of sensitivity you used.

In the particular case of the model selection, even if you say stochastic models, there are all kinds of underlying distributions that you could employ. You may be trying these various distributions to determine whether the results are coming up correctly, or you may be discussing this as you go along with your supervisor, your boss, or with senior management to determine the best way to approach this to come up with a rational answer.

The analysis of results is probably a lot more straightforward, although you probably will have at least a variance by utilization if you're talking about the GMIB. If you're talking about the GMDB, there's always the lapse factor that is critical. Then you have conclusions and recommendations. Can you build a reasonable benefit with a reasonable premium that will sell in the marketplace? Do you have to cut and run? Do you need to get reinsurance? Do you need to buy hedges? This includes whatever you need to do in order to minimize your exposure. In this particular case, there are appendices and references.

The body of the report is the central part. Generally, only the actuarial people are interested in the executive summary, but I wouldn't say they're the only people that are interested. Everybody has an interest in the executive summary. How do you do the report? It is recommended in Core 7 that you need to plan the finish, and that's true. However, the finish might just ask another set of questions that are more refined. You have to determine what the finished product is supposed to be. You need to know how to get there. Do you need to just shock them and then come back and see how they developed it? Exactly what needs to be done? Peer review, as has already been mentioned before, is essential. You can't be out on a limb by yourself, particularly with something this sensitive.

How should you write it? They suggest building it from the bottom up just like a pyramid. Isn't that wonderful? I don't necessarily agree with that. I think you need to have your executive summary pretty much in mind before you lay it all out because what you want to convey will dictate, to some degree, the structure of the rest of the study. Much of it is done in this manner. Obviously, the end result has to be done in this manner because you could be doing something of the top part and then not have the appropriate file memorandum to support it.

As for professional standards, all major assumptions need to be documented. If you're not personally involved, you have to disclose who was personally involved and their responsibility in coming up with the numbers. It has to be detailed enough for another actuary to be able to explain it to another actuary. As I said, you have to acknowledge responsibility. I'm going to go quickly through this because it's stuff that's fairly obvious and meaningless in the context of our conversation.

An oral presentation will be somewhat different because, even if you supplement it with written information, you're really trying to get the message out in a much more clear, understandable manner. I would never go about something like the variable annuities with just an oral presentation. I would have at least part of a report and definitely tons of slides to be able to show where I'm trying to go. I think the most important thing with an oral presentation is that you have to be confident that what you're saying is important, that it has been researched thoroughly, and that you respect the judgment of the people you're reporting it to. Now we'd like to open it up to any questions that you have.

FROM THE FLOOR: This is for Scott. In your discussion of the nonactuarial component or model components outside your area of expertise, you said there's no standard practice for life actuaries. I think the one dealing with nonactuarial components, extends to life as well as property and casualty because they mention scenario generators and stuff.

MR. KANE: I remember when I was doing these slides that there had been a proposed Standard of Practice, so I tried to find the most current information. I found one for property/casualty, and I thought it excluded life actuaries.

FROM THE FLOOR: I think now it does include it because they said the background was the hurricane models and stuff like that, but I think they've extended it to cover life because they mentioned scenario generators and things like that.

MR. KANE: I'll go back and look.

MR. KANE: When I originally was doing it, I was going to talk about the Standard of Practice and what we're all required to do. Then I thought it wasn't applicable anymore. So I changed what I was going to say.

FROM THE FLOOR: Scott, you had mentioned the increased interest in stochastic generators. I have started doing some research on that, and I've been kind of frustrated as far as finding stochastic generators or information on them. I've learned that it's a more complex topic than I

had originally thought, but I was curious whether you have any suggestions as far as reading material or sources of information you'd recommend on stochastic generators.

MR. KANE: There are a couple of people in our company who have been working on developing one. I could give you their names. I also heard that the Society was going to publish a paper on one, and I'm not really sure of the timing or whatever, but I think the Society itself was going to be putting out some good information.

MR. RICCI: You may just want to discuss this with the SOA. There has been a flurry of activity in the last couple months relating to this topic, particularly the value-at-risk, which is an associated topic. There have been quite a few people who are trying to come up with standards, and there are all kinds of stochastic generators. It's from the regime-switching model lognormal. So, I know it's a very complex situation.

FROM THE FLOOR: This is an opinion question for all of you. I find this work is very difficult for a number of reasons. You need somebody that has big-picture and little-picture skills, and they're pretty hard to find. My question is when you bring people in to help you or people to maybe take over parts of the project, are you better off taking a little-picture person and trying to turn them into a big-picture person, or vice versa? Each comes with pluses and minuses. I think we all know little-picture people sometimes don't see the forest for the trees, and big-picture people can be disorganized and overlook details and create data problems. If you had to pick, which would you do?

MR. RICCI: You need to identify the issues. You need to come up with a set bunch of issues. I don't necessarily think that somebody who is involved with the details will be able to do that, but you have to be able to get those people involved quickly because it's the only way you're going to come up with believable results. That doesn't answer it. It says yes, basically.

MS. BURTON: I found that any modeling project is going to have areas where you can have individuals who don't have the experience or don't need to know the big picture, but they can still be very instrumental and do key pieces of the project. It's generally a learning experience

more than anything for these people. But you'll also always want to have that person who knows the entire big picture and who can always remain in contact with the group of individuals working on the project.

MR. KANE: I think we have a tendency to ignore one end or the other. It's unfortunate. For something like this, there has to be a team effort. People have to understand the absolute necessity for both types of people. You can't do without both types of people in this kind of a situation.

FROM THE FLOOR: You mentioned in the course of the presentation that you would do at least 100 stochastic iterations. When I think of stochastic work, I think in terms of thousands of iterations. Can you comment on that and the statistical significance, and how you determine what the right number of iterations is?

MR. RICCI: In general, more is better. Much of the software takes a very long time to run. You have to run at least enough to get meaningful results. You'll start with 100, which seems to be a common number, but you would always want to do more if you could. If you're doing a 30-year projection of 200 model points, and you have monthly processing and all sorts of stuff going on, you generally can't run thousands, as much as we would like to. As computers get faster, we make our algorithms more complicated. If we're just doing something like some Monte Carlo option valuation, then, yes, you could do thousands, but for modeling an entire block of in-force business, there's just practicalities we have to deal with, but statistically more is better, of course.

MR. KANE: I would team up with the information technology (IT) guys on this. Maybe they can build the basic model that's sufficiently efficient to be able to produce a number of scenarios, and then you can take it and make sure you're involved in the process, and then start out with maybe 150, and see what 300 does to the results. Do they converge or don't they converge? Eventually get to the point where you feel comfortable about the number on a convergence basis to give you the kind of results that convince you you're going the right way.

FROM THE FLOOR: When you're talking about stochastic generators, are you talking about just interest rate generators, or do you have generators for equity changes?

MR. RICCI: No, equity as well.

FROM THE FLOOR: I've never dealt with any of that.

MR. RICCI: Generally you would want to have more than one class of equities, even with different parameters, and have a correlated matrix of coefficients. The different classes would have a mean and a standard deviation. Therefore, I think many generators would have two interest rates, a short and a long rate, so you get movement within a yield curve and something like a large cap fund, maybe an international fund or a small cap, and then a bond fund.

FROM THE FLOOR: What is the output for the equity part?

MR. RICCI: Monthly return, generally. You might have your short-term rate, your long-term rate, then there's some interpolation algorithm to fill in the rest of the pieces. Then you'd have a large-cap equity return, a small-cap equity return, and your bond return, which might just be an algorithm based on your interest rates, and you might have inflation as well.

MR. KANE: This is, of course, exclusive of the kind of modeling you would do for asset/ liability modeling, where you're talking about the assets that are backing up the liabilities in the general account. That's another kettle of fish.

MS. BURTON: We hope that you found that we spoke with passion from the heart, and although we did not touch you physically, we hope that we touched you mentally.