RECORD, Volume 25, No. 2[*]

Seattle Spring Meeting
June 16–18, 1999

Session 83CS
Year 2000 and Beyond: SOA New Course 7
Applied Modeling Demonstration

Track: Actuary of the Future, Education & Research, Investment
Key Words: Education, Futurism, Investment Products, Research

Moderator: WARREN L. LUCKNER
Panelists: STUART KLUGMAN
                 WARREN L. LUCKNER
                 JOHN I. MANGE
Recorder: WARREN L. LUCKNER

Summary: This session is a preview of an actual Course 7. Attendees refresh their memories of common actuarial models in a nonthreatening setting.

Attendees participate in a demonstration of a case study from the modeling Course 7. They obtain some hands-on experience by looking at a problem, discussing the various possible models and analyzing the results. Two experienced researchers lead the discussion.

Current Fellows leave with an idea of the new skills that their employees can bring to the job after attending a Course 7 modeling seminar. Future Fellows have an idea of what to expect in the Course 7 seminar when they go through the Year 2000 transition.

Mr. Warren R. Luckner: I currently provide the staff support for Course 7. The other presenters are Stuart Klugman and Jim Mange. Stuart is an FSA and Ph.D. He is the principal financial group professor of actuarial science in the College of Business at Drake University. He is on the Course 7 Working Group and is Chairman of the Common Core Subgroup. Stuart will be one of the common core presenters beginning next year when we start the seminars. Jim is an FSA and managing director at Health Reinsurance Management Partnership. Jim is also on the Course 7 Working Group and is the SOA Education and Examination (E&E) System General Officer for Course 7. Both Stuart and Jim have served as faculty for the current intensive seminar on risk theory, which gives them a good background for their responsibilities for Course 7.

We will provide a general overview of Course 7 using slides that were first developed by Jeff Beckley, who is the Chairperson of the Year 2000 E&E System Design Team and also the Chairperson of the Course 7 Working Group. After I give that overview, Stuart will discuss the case study that he has prepared for the common core portion of the seminar, and Jim will present some information and an example from the general or practice-area-specific day-long case study component of the seminar.

A concise theory of how individuals learn is, "Tell me, I forget. Show me, I remember. Involve me, I understand." I think it summarizes well one of the key motivations for having a Course 7 and the importance of having interactive exercises. The participants in this seminar will understand better if they're involved in the learning process and interacting through exercises.

Course 7 is part of the new education system that will be effective in January. Throughout the new system, there is an emphasis on modeling. In particular, Courses 3 and 4 cover the topics of actuarial models and actuarial modeling. However, they use a traditional approach of studying the topics and taking an exam to obtain credit. Course 7 adds to that a different way to get credit. It adds an intensive seminar and interactive component to the new system. The Course 7 Working Group Report, which some of you may have already seen, is on the SOA Web site. It gives more detail about what Course 7 is intended to do.

The five learning objectives we identified for Course 7 closely parallel what we consider to be the important components of the modeling process: the context of modeling. That is, what is modeling and when do you model? There is design, selection, set up of the model, input data, selection analysis, output data analysis, or analysis of results and communications of the modeling process and results. Meeting the fifth objective is essential. If you meet all of the other objectives but don't communicate the process or the results effectively, it has essentially been a waste of time and can even be counterproductive. There is an emphasis in Course 7 on the communication of the process and the results.

The Course 7 format is similar to the current intensive seminar format. Some of you may have taken those elective courses in applied statistics or risk theory. The educational format has been found to be very effective in terms of developing skills different from the skills that you develop through self study or just studying and taking an exam. One of the motivations for developing Course 7 is the success of the current intensive seminars. Obviously, one of the differences between the current intensive seminars and Course 7 is that Course 7 is going to be required of everybody who goes through the system, while the current intensive seminars are elective.

Other differences include the fact that there is a substantial amount of pre-reading for Course 7, as well as a pre-test. There will be a set of common core case studies that's going to be the same for all seminars, but there will also be a general or practice-area-specific day-long case study that is presented at the seminar. The difference between the general and practice-area-specific case studies is that the practice-area-specific one presents an application that is specific to a particular single practice area. On the other hand, the general case study presents an application that may cross practice areas, such as social insurance, or may not be specific to any particular practice area at this time, such as public school funding. One of the criterion in developing the general case study, because these case studies will be different for different seminars, is that it should be designed so that no particular practice area background is significantly advantaged or disadvantaged for that case study. It's a challenge to do that, but using those kinds of case studies allows us some flexibility.

Let's discuss a bit about the pre-reading for the pre-test. The pre-reading focuses on the modeling process, and is more general than specific with respect to actuarial modeling. Much of what you have in Courses 1 through 6 gives you specifics about actuarial modeling. There may be some seminar-specific pre-reading, depending upon the content of the general or practice-area-specific case study. In the pre-reading there is some communication skills material, particularly some material from the Actuarial Standards Board that relates to communication.

What about the pre-test? What's the purpose? The purpose of the pre-test is to determine if the students are prepared for the seminar and prepared in such a way that the seminar does not have to review topics covered in previous courses, such as Courses 3 and 4. The pre-test also allows us to have a shorter seminar time. Three-and-a-half days are what is scheduled. Hopefully, the pre-test will also result in students being better prepared, which would allow us to perhaps have higher pass marks, and it would also eliminate or at least reduce significantly, the possibility that somebody would come into the seminar ill-prepared and try to take it without much background.

For the pre-test, the intent is to have a bank of questions and to offer the pre-test monthly on an open-book basis. The format of the pre-test will include both multiple-choice questions and written-answer questions. There is a sample pre-test being developed. It's not quite ready yet, because we're working on finalizing the illustrative solutions. The sample pre-test should be available soon on the SOA Web site.

The pre-test grade will be pass/fail with a predetermined pass mark, which is different from what we've done before. The period of validity for a pass on the pre-test is six months for registration purposes and 12 months for attendance

purposes. That is, you have to register for a seminar within six months after having passed the pre-test, and you have to take the seminar within 12 months of having passed the pre-test. This is partly to make sure that the pre-reading you have studied is still relevant to the seminar you attend, since the pre-reading can change over time.

Other Requirements
As I noted, you have to have passed each of Courses 1 to 6 before you can take the seminar. You have to bring a laptop, and you have to have familiarity with a spreadsheet program and a word processing program.

Format
As I mentioned, the format is similar to the current intensive seminar format. The format is a little bit shorter since it is three-and-a-half days, compared to the current five-day seminar. The first day-and-a-half will be for the presentation of several case studies that are common across all seminars. These case studies will cover various aspects of the modeling process. The presentation will be interactive in the sense that there will be some lecturing, but there will also be some assignments and some give and take in terms of the students doing some computer work and presenting results. The students will be asked to work in teams of no less than three when they work on these assignments. The faculty for the common core part will typically be academic faculty, but not always.

After the first day-and-a-half, the next full day is a presentation of a single general or practice-area-specific case study that covers all aspects of the modeling process. The instructor will be from business or government. It will be a practical application. Again, the presentation will be interactive. The last day is the tough day. The student will be given a project assignment and will be asked to complete that assignment and prepare a report. The assignment will be to analyze a practical problem. Instead of working in teams, the seminar attendees will be working individually to complete their assignments. The project assignment is the only means of evaluation for the seminar.

The criteria for evaluating the written report, which is the end result of the assignment, is as follows. There must be use of an appropriate modeling process, the report must be presented in a clear and concise way, and there must be accurate and appropriate analysis of the problem.

Our goal is to provide results of the seminar within eight weeks after the seminar is completed. We expect that, partly because of the pre-test, 80–90% of the students will pass the seminar on their first try.

We expect that we may have to offer at least 15–20 seminars a year. We'll be offering seminars throughout the year at various locations and with various general or practice-area-specific case studies.

Two major concerns have been identified. How do we accommodate students outside of Canada and the U.S. who may have difficulty getting to the seminar for financial reasons? How do we keep the fee in a reasonable range? With respect to students outside of Canada and the U.S., we are going to offer an alternative for students who meet certain criteria related to financial hardship. There will be two levels of criteria: the income level of the country in which they're living, and the income level of the individual.

The current alternative includes use of a videotape presentation, within the same timeframe as an actual seminar, with the same project as the actual seminar, and with the instructor being available via e-mail. We don't anticipate that we'll have a lot of these cases, but we do feel an obligation to make sure that we're ready for any particular problems due to financial hardship. With respect to disability, there is already a procedure in place within the current E&E system guidelines.

We estimate that the fee for the seminar will range between $700 and $1,200. This is a wide range, but even the endpoints of that range seem reasonable when compared to commercial three-and-a-half-day seminars. We understand some of the financial constraints on seminar participants, but we feel fairly comfortable that this is a reasonable fee to charge.

Now Stuart will take us through his common core case study.

Mr. Stuart Klugman: The common core case study, of which there'll probably be about five in use at the seminar, has three phases. Phase one is information that will be handed out to the students before they come to the seminar. That information is basically in two pieces. One is just to tell the students very briefly what their case study is about. This one tells them that they're going to be doing a case study on the pricing of warranties on compact disc players and that they'll be comparing different warranty options. It gives them a small flavor of what they're in for. More importantly, we wanted to make sure that the students are prepared for the case study and, in particular, we wanted to re-emphasize items from earlier courses that they might want to review as they do this case study.

For this particular case study, the student's spreadsheet program—and we don't care which one it is—had better have a maximization routine. They would also have as part of their package an instruction manual that would tell them how to run the Excel Solver. I'll eventually make one for the Lotus equivalent. The students can then practice on that particular spreadsheet tool before coming to the seminar. There is a reminder that for statistical technique they're going to

need to remember what maximum likelihood estimation is all about. Then they get a write-up on a new item, a distribution called the mixture of exponentials. We'll be asking the students to have looked at that before they come to the seminar. For each of the common core case studies, as well as the business case study, each instructor will prepare appropriate materials so that the students are ready to roll when they show up.

At this point we would be at the seminar, and we would hand out to the students the remaining pages. You have a memo to "actuarial consulting trainee and modeling wizard" from your manager asking you to do something. In the one-hour time period allocated either in the first morning or afternoon of the seminar, the instructor would first take the students through that handout and explain the problem.

Our problem is that we have a client named WorstBuy Electronics. We currently do their pension plan. Recently they began selling warranties on their CD players. They think they are taking on some sort of risk and that actuaries may know how to analyze those risks. Then they go to their pension consultants and ask if we can analyze the risk. We can do that, and we want to know exactly what the warranty situation is. There is a reminder that says a written report is due in 90 minutes. Even though it's not going to be a continuous 90 minutes, that's basically the timeframe allocated for each common core case study: one hour on the first day and about one-half hour on the morning of the second day.

Here are some details about the warranty. It covers all damages. I originally had set this up with a much more complicated problem, but then that 90-minute thing got in the way. So instead of having options about repairing (the original problem had the dealer being given the option of repairing or replacing), we're just going to replace the CD players. This is based on a real story, as they say on TV. My son bought one of these warranties from an electronics store, and I was rather surprised that when it broke he walked back in and said it's broken, and they handed him another one. They didn't even bother to check to see if it worked or not. I think they had decided the cheapest thing to do is hand him another one rather than waste anybody's time figuring out whether it worked or not. One unsettled issue that the actuary might help the manufacturer decide is: should the warranty cover the replacement player or should it be a one-time thing and would you then have to buy another warranty on your second player? The cost for the warranty will be proportional to the selling price. We note that there'll be a 15% commission to the salesperson for selling one of these warranties. There's an additional fixed cost to us of $2 per sale, and it'll cost us $1 to process any replacement.

Table 1 gives the data that we have available from cash register records. We find there's a problem. We don't have much information on the frequency with

which these CD players break down, but we think it's somewhere in the 15–30% range. One of the things you'll note is the data were collected in a group data format rather than by the individual numbers. One of the reasons we're bringing the students the data presented in a table like that is to spend some time at the seminar discussing advantages and disadvantages of having data aggregated rather than working with the individual numbers. Our time constraints today won't allow us to do that, but I wanted to let you know about some of the issues we would talk about with the students at the seminar to help them be better modelers and data analysts.

TABLE 1
SALES DATA

Price Range	Number of Units Sold
50-75	3,343
75-100	2,055
100-125	1,262
125-150	806
150-175	597
175-200	446
200-250	515
250-300	305
Over 300	581

Then we're going to talk a little bit about the mixture of exponentials distribution. The paper has just been accepted for publication in the proceedings of the Casualty Actuarial Society. The students will have read the key details in their handout, and then one of the questions will be: is this a good distribution? Even though Mr. Keatinge says it's a great distribution, students must determine how it works for their problem? Are there some advantages or disadvantages to having a sort of all-purpose model? One of Keatinge's claims in his article, is that if you have this model, you don't need many others.

Let's move on to our second question. We're going to sell 10,000 of these warranties, so what's our risk? What's the risk that the retailer is taking on by doing this? We also want to help them estimate what they're going to charge for this warranty. If they're going to buy insurance to cover it (that is, if they're going to go out and find someone to insure their warranty business), what sort of insurance scheme would we recommend that WorstBuy go out and purchase? Remember that there was the question of the effect of the warranty continuing on the replacement players. What are the consequences of each of those decisions? We're still going through the set-up phase. Then there is one more open-ended question. When you're all done, what don't you feel good about? If you had to give the report, what are all of the caveats you'd include? What other information do you wish they had given you so you could have done a better job? We're asking the students to, when they go back and, and based on what they

have here, do the best they can. What should they ask from the client that might help them do a better job? We won't ask them to take that and go anywhere with it, but we kind of like to leave these case studies with the students understanding that these things don't end in 90 minutes. In the pre-test, there's a good bit about the feedback loop: you do something, you learn something, and then you go back and do it better. As long as your clients are willing to keep paying you to improve it, you want to keep cycling through and improving what you're doing.

The formula below gives you the mixture of exponentials. For those of you who've studied for the Risk Theory Exam 151, you've seen this before. This is that same distribution— the only one for which you could get the exact probability of ruin when claims followed a Poisson process. When you took that exam, your fervent hope was that they would not ask you a question about that calculation because you remembered that it took about two pages in the book to get it done, and the chances that you could reproduce it under exam pressure were small. Keatinge gives some other rationales for liking this distribution, and the fact that it can give you ruin probabilities wasn't one of them. We will not be asking our students to do that. We have the cumulative distribution function as well.

<div align="center">

FORMULA 1
MIXTURE OF EXPONENTIALS

</div>

The probability density function is:

$$f(x) = \sum_{j=1}^{K} a_j q_j e^{-q_j x}, \ \sum_{j=1}^{K} a_j = 1$$

The cumulative distribution function is:

$$F(x) = 1 - \sum_{j=1}^{K} a_j e^{-q_j x}$$

Note that the number of exponentials to mix is itself a parameter. If we ask how many parameters are in this mixture of exponential distributions, there is no answer because the number of exponential distributions (the value of K) is itself a parameter. Therefore, it's different from some of the other distributions we tend to encounter. But one of the things Keatinge proves in his paper is that if you keep adding more exponentials, at some point, it doesn't do any good to add another one. There's a point at which the likelihood function can't be made larger by throwing in another exponential into the distribution. That's something you might not have expected, but it does mean there'll be an end to the process as you keep adding exponentials. If you look back at the data, no CD players sold for less than $50. Because the mixture of exponential distributions puts probability from zero to infinity, we'll actually be building a model for the cost in

excess of $50 of one of these players. We would work the students through that a little bit.

At this point they would have done some work. They would have taken the data and their spreadsheet and would have worked out the likelihood function. We would ask them to do a little work at this point and produce the following results. Some summary statistics for fitting (cost minus 50) with various values of K are:

K	-ln Likelihood	Chi-square	df
1	19,172.4	162.04	7
2	19,006.3	6.02	5
3	19,005.3	4.02	3

Note: Increasing K further did not change $-\ln L$. A likelihood ratio test indicates that $K = 2$ is appropriate.

They would find out that once you have three exponentials, you don't need any more. That would turn out to be a five-parameter distribution: the three exponential parameters and the two weights to allocate how much of each you want to use. We would also talk a little bit about the likelihood ratio test and about the Chi-squared goodness-of-fit test. The students studied some of those things in their Courses 3 and 4. Here we're using it to help decide what to do. Based on that, we might come out with a recommendation to go with $K=2$, mixing two exponentials, a slightly simpler three-parameter model.

The other thing we want our students to be able to do is make a picture. Chart 1 is a graph pasted out of Excel, and this would give someone a good visual confirmation. The mixture of exponentials density function matches the histogram of the data quite well. If they don't understand or appreciate the Chi-squared goodness-of-fit test, we want the students to know that by making a graph they might have another argument that'll persuade the people they're reporting to that they have found a model that does a very good job of describing the distribution of the selling price of these CDs.

We now have a little discussion about why having one all-purpose model is a good thing. The pre-reading talks about overfitting. It's one of the topics that's covered fairly heavily in some of the pre-reading. The idea is that if you try 20 different models, there's a good chance one of them will match your data, not because your data really come from that distribution, but because if you try enough models, sooner or later one of them will wind up looking good. And the fewer models you have to begin with, as long as you can get one that works, you'll have a lot more confidence that it's the real thing, that your data actually comes from that distribution. We'll talk a little bit about parsimony, about overfitting, and again reinforcing some of the concepts that were in the student's pre-reading.

We'll then build a little model that describes the random outcome facing the seller of the warranties. The cost of providing the warranty depends on two random variables. One of them is the random R, which will be the number of replacements out of the 10,000 each year. The other is the random X, the cost of the CD player. The broken CD player that comes in our door is also random because we don't know whether the next one we'll have to replace will be a cheap one or an expensive one, because there are fixed costs associated with it. Our price is proportional and that randomness has an impact on our answers. If R is the number of replacements on a single sale, it will be zero or one if we don't continue the warranty, but it could be any number if we do. This would be the loss we expect to pay on a single player, and then we just have to add 10,000 of those random variables to see about our entire collection. Again, as mentioned, if the warranty doesn't carry over, R is just zero or one, and it's a Bernoulli distribution, otherwise it'll have a geometric distribution because sooner or later the customer will get one that works, that never breaks down. We're just counting the number of replacements until they finally get one that outlasts their patience to keep coming back to get a new one.

The next calculation is a little bit involved, but we're basically just invoking the breakeven principle, the net-premium principle. The charge is a function of the unknown *p,* which is the probability of a CD player malfunctioning. Remember all we knew about it going in is that it's between 15% and 30%. We do some calculations that we would go through with the students or have them go through and then we'd review how to do it afterwards. We haven't quite got the time management down for these case studies on what we'll ask the students to do separately and what we'll work through with the students all working together. We'll note that if 15% of the players break down, the break-even premiums are 20% of the cost or 23% depending on whether or not we replace the CD player. If it's a 30% breakdown rate, we'll need to charge 37–53% of the cost. We would have all the calculations worked through with the students.

Table 2 covers some of the intermediate cases. They will know from their actuarial mathematics work and elsewhere that one way to set premium that has a risk margin in it would be to set the premium at which you'd have a 97.5% chance of not having to pay more than that number. By adding 10,000 random variables, we can invoke the central limit theorem. That means we're just adding two standard deviations. We go through the standard deviation calculation with the students as well.

TABLE 2
PREMIUM THAT WILL COVER LOSSES 97.5% OF THE TIME
WHEN 10,000 POLICIES ARE SOLD IS AVAILABLE

| | No carry over | | Carry over | |
| | | 97.5% | | 97.5% |
P	Break even	Premium	Break even	Premium
0.15	0.1964	0.2000	0.2277	0.2323
0.18	0.2319	0.2359	0.2788	0.2840
0.21	0.2675	0.2718	0.3337	0.3396
0.24	0.3031	0.3076	0.3929	0.3996
0.27	0.3386	0.3435	0.4570	0.4645
0.30	0.3742	0.3793	0.5267	0.5349

A question we'd like students to think about is: now that we have some numbers, we've looked at the model, and we've had it tell us what it can possibly tell us and what we can recommend? What are some things we should think about? We observe that there's almost no process risk. That is, if you insure 10,000 of these CD players, they're going to behave as they ought to. The law of large numbers is on our side. The standard deviation is relatively small compared to the mean. The risk of a few extra CD players breaking down is small. The margin needed to bring us from breakeven to the 97.5% percentile is adding 1% to the premium. Collect an additional 1% of the selling price, and we'll have the process risk taken care of. There's no reason that they should insure against the process risk. On the other hand, the uncertainty about the value of *p* is very large. Remember, if *p* is wrong, it's wrong on every player. It's not something you can overcome by insuring a lot of them. We want the students to recognize that this is the source of the difficulty, and, therefore, we're thinking about an insurance arrangement. We'd want to do some kind of aggregate insurance rather than individual insurance. The real risk they're facing is not knowing the probability of the breakdown of the CD players. Again, the pre-reading and our discussion here will talk about all the various sources of risk and sources of error in a modeling exercise. We'd want to think about which ones might be more important here than others.

While difficult to quantify, we can make a feeble attempt at measuring the risk of not knowing *p*. Assign probabilities of 0.1, 0.2, 0.2, 0.2, 0.2, and 0.1 to the six values. For each of the *p* values and a fixed *w*, we compute the probability of losing money. Then, the weighted average of these probabilities reflects the probability of losing money with uncertain *p*. To make this probability 0.025, we must essentially set *w* at the worst case value (based on a *p* of 0.3). It's essentially maxi-mini. Basically you have almost no choice but to assume *p*=3 and take this worst case and work from that one.

At this point we would go back with the students and ask them what they assumed. Up to this point they have been following instructions and believing

everything the instructor has told them. We will intentionally not go into any detail about things we've been assuming, and so we would want to work through with the students at this point the various assumptions that they had to make to get those numbers. A complication is something called moral hazard, and so we'd talk a little bit about the ever-present problem that as soon as you insure something, it changes people's behavior. I'm sure my son was much more willing to toss his CD player around, maybe throw it across the room if he was mad, knowing that he could always get a free one by just trekking over to the store.

If we look back at all of the assumptions we had to make, we realize we probably haven't accomplished a whole lot that's very useful. Our best advice to WorstBuy is to do a pilot study, get the moral hazard effect, and get the probability of p estimated before embarking on a large-scale warranty program because that's the real risk. We're not sure we can find someone who'll cover that risk when they don't know the value of p. That's where we're likely to leave the students on this one.

Mr. Klugman: If there are any questions about this common core case study or how we plan to run common core case studies in general, we could do that now before Jim talks about the business case studies.

Mr. R. Allan Ireland: In this sample case, the consulting actuary could have probably come up with just applying the probabilities. You could do it in a deterministic manner just applying the probability of 30% and coming up with some sort of value in a sort of a back-of-the-envelope fashion. Would it be your intention to sort of compare the kind of prices that you come up using ad-hoc methods with this technique?

Mr. Klugman: In fact, you would have done pretty well, again, because the process risk is negligible. Working with expected values would do a pretty good job, and good actuarial intuition would have told you where the risk was. The message was go out and learn more about that risk. But that would be a good thing to add in; to say here's the old way. As it turned out, that would do OK, but we wouldn't want to deprive the students. They've all gone through the very challenging Courses 3 and 4 that enabled them to do all these things. I'm sure they'd be disappointed.

Mr. Ireland: The solution drives the problem?

Mr. Klugman: In this case that's fair to say. I think Jim's case will show you that there are times when having the probabilities is very beneficial to your analysis. Jim, I don't know if you were going to talk about some of the discussions we had. He was pointing out that most of the time, the customary ways do just fine, but

every once in a while you get a problem where the probabilities can really help you do a better job, and his case is one of those times. In this one, it's not as obvious.

Ms. Jennifer A. Vandeleest: I have a question about one of the requirements for this course. You listed that you had to pass 1 through 6 before you could take this one. I thought I read on the Web site that you could have 5 out of 6 and apply for it. My concern is, for example, if you fail Part 5, and you have 1 through 4 and 6, and you have to wait a whole year to take Part 5. Could someone not take Part 7? How are you going to handle that?

Mr. Klugman: Due to the timing of the seminars, if you are particularly interested in taking a seminar in your part of the country or on a special topic that's of interest to you, the seminar might only come along once a year. It may be that the next opportunity to take the seminar is shortly after you'll be taking Course 5 or Course 6. We're allowing people to do the pre-test and get ready so that if the seminar's the week after the results come out, and if you pass, you can go right to the seminar the next week. If you don't pass the exam, then you'll have to wait. Just as we do now with the current intensive seminars, you can sign up for them even though you haven't passed the qualifying exam, and if you then pass the qualifying exam before the seminar, you can go.

Ms. Vandeleest: If you're waiting to take Part 5 for a whole year, you can't take this pre-test and take Part 7? You wouldn't be taking any exams for a year? Is that what you're saying?

Mr. Klugman: It could work out that way.

Ms. Mary P. Kirby: If this course is kind of an extension of Courses 3 and 4, what does it matter if you have 5 and 6? If I have Course 3 and 4, and I have all the tools I need to take Course 7, why can't I do that?

Mr. Klugman: Tools, yes. Context, no. Maybe this case wasn't the best example. We'll expect you to be comfortable with the insurance world and with the actuarial world, including the practical items in Courses 5 and 6. In particular, I know we have some finance common core case studies, and I think there'll probably be some things from Course 6. Some of the finance material on Course 6 will be in some of the common core case studies, but we wanted our students to have a little bit of the practical side as well so they can think about these problems as more than just probability and statistical exercises. I think, if nothing else, it's a matter of maturity. We want the students to be a little older, to be able to do the writing, to understand the context, and to think about the practical issues as well as the technical issues. There is the technical Part 3 and 4 aspect, but we hope there's more to it than that.

Mr. Luckner: I think that's right. I apologize if I might have misled you on that one by emphasizing Courses 3 and 4 in my part of the presentation.

From the Floor: I have a question for you on the exam. What I heard Warren say was that to pass this course you really do need to be focusing in on the delivery, the communication, and the results that you have. You'll have Courses 3 and 4 already behind you. When you come in here, you should have the tools available to you. How are you going to teach the development of the communication necessary to pass the course?

Mr. Luckner: My view is that there are a number of things that'll happen. Of course in the pre-reading we have some reading material that'll help. Also, there will be examples of reports during the seminar and in material that's distributed in advance of the seminar.

Mr. Klugman: Jim, is it part of your day as business instructor to specifically talk about that as well, like it used to be?

Mr. John I. Mange: I used to be the business instructor for the 152 seminar, and part of the presentation that I made was to actually go through how to write a good report. You spend 30 to 45 minutes, and the students do get that as part of their seminar study materials.

Mr. Luckner: Also, one of the common core case studies focuses on the communication aspect.

Mr. Klugman: I think you're on, Jim.

Mr. Mange: I'd like to cover some additional background on what we'll cover in the practice-area-specific presentation of Course 7 and then give you an abbreviated illustration of a practice-area-specific presentation. Finally, I will explain why we think it's important to include some practice-area-specific content.

The practice-area-specific presentation would be done in about eight hours of classroom time plus homework. We think it's important to have practice-area-specific content because it makes the seminars more relevant to the candidates. When I was taking the exams, there were certainly portions of exams, and I really wondered: why am I bothering? What's important about this? Practice-area-specific content will help make the seminars matter to individual careers.

We also expect to find that many candidates have not identified a practice area. In my case, I was working in the group pension area when I completed the exams, and I continued to do that for three or four years. I then moved to the

group health side, and I've been there ever since. My practice specialty is health, but when I would have taken this exam I would have been on the pension side. Recognizing this issue, about half of the seminars will have no practice area specifically identified. Half of the seminars will be specific to a practice area, and the students will be expected to have some knowledge of that practice area from their business life. The other half of the seminars will be much more general.

The practice-area-specific content enters the seminar in two different ways. First is eight hours of classroom time devoted to a practice-area-specific case study. It's a single case study presented over that length of time. It will begin on the afternoon of the second day of the seminar. The case study will be set up for the students during the course of that afternoon. Then the students will be left with an assignment to work through part of the problem that evening. There will probably be some consultation time available in the evening for the students. The next morning they will solve the remainder of the problem. What I will try to do for this discussion is illustrate the flow of such a day.

The second way in which practice-area-specific content enters the seminar is the candidates will have 24 hours to analyze, solve, and present their solutions to a practice-area-specific business problem. I think we found in the Intensive Risk Theory seminar that this was a very challenging task for most students.

Here's how we are carving up the seminars that we expect to deliver in the year 2000. We anticipate that we will deliver 16 seminars over the course of the year, nine of which will be general. The general practice area will be a subject area that does not advantage or disadvantage any student who is a specialist in any particular area. We'll expect two seminars devoted to life topics, two devoted to pension, one to investment, one to finance, and one to health.

I'll take you through a practice-area-specific area illustration. The illustration is from Appendix 3 of the Report of the Working Group on Course 7 Applied Modeling. It's available on the SOA Web site if you don't otherwise have a copy. It's a case study that I originally developed for the Intensive Risk Theory seminar, and it was based on an actual business problem that I faced some years ago.

How many of you have been in a large company's pricing department before? A few of you have been there, but most have not. It's a place where you're very much separated from the underwriting department and the sales department. You're insulated to a large degree. At some companies, it is more so, and at some companies, it is less so. The players in this particular scene include somebody from the underwriting department, your manager and you. Underwriting has a product variation, aggregating specific stop loss, that it has

never seen before, and your manager said you will help price this product. You've probably never seen this product before either.

You have several objectives. First, you know your price for traditional specific stop loss covers. How do you adjust that price for this variation? Second, you must do it with pretty severe time constraints. You might want to do this again. In fact it might be a great product idea for the future. Your third objective is to solve the problem such that similar problems may be solved more easily. Finally, you need to identify any effects on medical rate filings. The ideal solution, then, doesn't just solve the business problem; it lays the groundwork for a more general solution.

If this was presented as a general business problem, then the business instructor would have to spend a lot of time describing the health insurance area. He would describe the sales process, the underwriting process, and define some of the group health insurance financing arrangements, talk about specific and aggregate medical stop loss, and how they fit into the whole medical financing structure. If this example was delivered as part of a health practice area seminar, then you could assume the audience is familiar with much of the basic background. You could delve more deeply into the business problem itself and perhaps provide a more interesting solution. You could also probably assume some general knowledge of the regulatory environment. You probably wouldn't ask questions about the regulatory environment unless you were in a health practice session.

Looking at the data, we have some frequency information, but it's not frequency data. It's frequency assumptions, and they tell you about the likelihood of a claim and also about employee and dependent relationships. Medical stop loss is usually sold per employee rather than per member, a common HMO pricing basis. To price medical stop loss, you need to know about employee and dependent relationships. The key idea here is that we're actually starting with the results of an earlier study. We're not starting with data. The typical problem that you would have solved on Course 3 or 4 would present you with data that you'd then turn into a result. Here you're starting with a result that you must apply in the context of a unique business problem. The other problem with the frequency information is that the study was conducted three years ago. If you know anything about medical, you know that it changes quickly. You have to think about whether this information is really useful or not.

On the severity side, you have paid claim data, and not incurred claim data. It was claims paid during the 12-month period ending six months prior to the effective date. Now, a lot of medical stop loss covers are sold on a paid basis. The paid aspect is not an issue in this context, but the data will need to be trended forward to the effective period. We've provided a trend assumption to

the students, but students must determine how to apply that trend assumption to the data that they've been given. Some expense information has been provided, but the expense and profit information is stated with respect to their traditional medical stop-loss product and doesn't necessarily apply to this aggregating specific stop-loss problem. Since the expense assumptions have been stated as a percentage of specific stop-loss premium, you need to ask how the assumptions should be adjusted to reflect the nature of this product.

We come to this point. We need to build a model of aggregate losses in excess of $100,000 per person. What strategies might we apply to developing that aggregate loss model?

I'll throw out a few of the ideas I had and see if there are other ideas from the floor. There are some different ways you could go about developing an aggregate loss model. One would be to use a Monte Carlo approach. That's what a lot of actuaries I know apply to these kinds of problems. Another one would be the compound Poisson model. That's such a familiar model that a student might use the method of moments averages to drive the parameters of the compound Poisson model. The Intensive Risk Theory seminar in which this was originally delivered studied parametric frequency and severity models. The students in that seminar leaped right to that solution. For them, it was just a question of applying that technique.

Are there other approaches that people would take to solve this problem?

Mr. David E. Olsho: We just set this problem up at our Howard Johnson office. We're taking a Monte Carlo approach.

Mr. Mange: Howard Johnson is going to use the Monte Carlo approach. Any other ideas?

From the Floor: I guess I would have used the Monte Carlo also.

Mr. Mange: Other techniques? Anybody else have another idea of how you'd approach this problem? Yes.

From the Floor: You'd buy somebody else's manual.

Mr. Mange: But you've got 24 hours to do it.

From the Floor: We did have one good idea, I thought, at this table. Just cut the premium 33% and go for it.

Mr. Mange: That is actually one of the questions that I asked the students in the Intensive Risk Theory seminar. Having laid the problem out, I would then ask them to apply their actuarial judgment to this problem. We'd get different answers. Some would say yes, and some would say no. The particular solution that I worked out came to the answer that you really couldn't do it from an actuarial perspective, but it was close. It might be a case where your underwriter might want to, as you said, go for it. I think that's what makes it a more realistic business problem. Somebody has to make a business judgment given the limited information at hand. I think that helps to make the problem more real, and that's what we would try to do in the context of this practice-area-specific modeling portion of the seminar. Any other ideas?

Mr. Darrell D. Knapp: If you have to do it in 24 hours, then you'd probably have to do the Monte Carlo approach. I might use a parametric approach to develop a more general solution.

Mr. Mange: There's commercially available software that would enable you to do that relatively quickly. If your company owned that software, you can do it, and you could actually do it fairly quickly.

To continue the illustration of the practice-area-specific presentation, we've identified three or four different approaches that we might take to solve the problem. What are the advantages and disadvantages of each of those approaches? What criteria should we apply to choose among the possible approaches? I think Darrell Knapp suggests one criterion. If I'm going to have to do this again, I might want to adopt one approach versus another. Finally, we need to ask how we will address the data deficiencies, particularly the frequency assumptions.

You'd spend most of the first afternoon setting the problem up, helping the students understand what it is you really want them to do. Over the evening they would work through a good part of the problem with faculty support. The next morning you would complete the solution to the problem. Finally, the business presentation would conclude with a presentation of the business instructor's solution to the business problem. My solution to the problem we discussed today is part of the working group report. If you want to see it, you could pull it down from the SOA Web site.

From the Floor: Did you do that one in 24 hours?

Mr. Mange: I did it over several evenings. When I presented it to the students we did it in six hours, I would guess, because I typically started at around 8:30 and ended at approximately 3:30.

Mr. Klugman: But nobody actually wrote it up during that period.

Mr. Mange: That is true. We think that the seminar approach really offers us a more effective way to convey some of the critical modeling skills that we think the actuary of the future needs to have, and that's why we've moved towards the seminar approach for this Applied Modeling Course. Practice-area-specific opportunities will enable some candidates to focus on those things that are really important to them. I think it will make the experience much more relevant and will help students see how the examination process will not just be a stumbling block on their way to success but it will help them prepare for their careers.

I'll turn it over to Stuart to present a few other case study ideas.

Mr. Klugman: The last thing we wanted to do was to talk about some examples of other case studies. These aren't necessarily ones that are going to be done, and I wanted to sneak an extra one in. On the way out here on the airplane I was reading *Foundation and Chaos* by Greg Bear. It takes place in the 12,000th year of the Galactic Empire. That's some time in the future, I presume. They must track down a lost spaceship, and apparently it's a challenge because they stopped putting recorders in them. Here's why. It's not studied much anymore because most hyperships get where they're going, and the Empire's actuaries believe it's more trouble than it's worth to track lost ships since so few ships are lost. You might get a case study on tracking lost ships through hyperspace and you have to decide whether it's worth the effort to put in a tracking device.

Let's discuss two more down-to-earth case studies. One is a project on mortality differences. The set-up is that your company sells annuities in some other country. That country wants your company to use a particular mortality table. You know your annuitants have different mortality, and your task is to build a case that you don't have to use the country's standard mortality table. It's not enough just to prove that the Qs are different. You can't say, "Here's my Qs and they are different from yours." It's not even enough to say they're different by a certain number of standard deviations. We have to explain why. We have to say there's something truly different about the people we sell annuities to than the average annuitants countrywide. We have some demographic information about the annuitants and perhaps do a logistic regression to discover if there are some differences in the population of folks to whom we sell annuities versus the national average. That would be a non-practice-area-specific or general case study. The fact that there are annuities in the background isn't really relevant to this particular case.

The other part that might also enter into this particular one is the fact that the data you've managed to collect to build your mortality information have lots of mistakes in it. One of the issues is cleaning up data. We have a common core

case study on cleaning up data, and this business case could further some of those issues.

Another one is a fraud case with some statistical issues. This is based on something I came across several months ago, but most of the details are obscured here. Some of your agents are doing things they shouldn't. Is there any way to automate the process of picking them out, going through data, finding patterns in what's happening, and getting some formula we can apply to our cadre of agents that'll pop out the names of the ones we need to investigate further? This would be a discriminant analysis type problem and a chance to again build a model. This case is specific to a particular practice area to the extent that the bad behavior might be related to the particular kind of insurance being sold. If it's not that difficult to get that across as a general case study, we would do that. Otherwise it would be specific to that kind of insurance. If it seemed that knowledge of the type of insurance that led to the bad behavior was truly important and would take time to explain to those who don't understand it, then it becomes a practice-area-specific case study. These are a couple of examples of some other things that could, as opposed to will, become case studies. Jim is still in the process of recruiting business instructors and nailing down the cases they will each be doing.

Mr. Mange: I think at this point, most of the business instructors are identified. Warren's been helping a lot with that.

Mr. Luckner: We've been working on it. We have a number of business instructors on board. We're looking for some more. I think some have already thought about case studies, but we don't have any officially developed at this point. As long as Stuart alluded to that data clean-up case study there are some slides of another common core case study that we are not intending to cover here, but I thought that some of you might want to look at it. It is the one that he's referring to. It came from the research project that the SOA did a number of years ago, and it's an ongoing experience study now. It's an asset risk study, and it was an example of applying what I've characterized as traditional actuarial techniques, applying disability, frequency and severity concepts to assets. Just as people get sick, die, or recover, assets can do the same types of things. The issue in this case study is the data quality issue. Because of the complexity of the data, it was a real chore in that particular research project. It is now one of the common core case studies.

Mr. Ireland: Mr. Mange, can you clarify what you were saying about the business case study? With the whole seminar do you come to a consensus on what approach is going to be followed or do different groups get to use their particular approach and present what they think are the advantages and disadvantages of each?

Mr. Mange: I think it can be done either way. In the Intensive Risk Theory seminar that Stuart and I did together, we led them to a particular approach that we wanted them to take. There were some situations where we had half the room apply one model to the problem and had the other half of the room apply a different model so they could see some of the differences in the various models that were applied. I agree there's some value to be gained by that approach.

Mr. Ireland: We talked about two probabilistic kind of problems. Will the seminar get into the pitfalls of taking simplistic solutions? I alluded to the deterministic approach using expected values? It seems to me there would be value in maybe talking about the pitfalls and the dangers of some of those approaches if, indeed, there are any.

Mr. Luckner: I'll respond to both questions quickly. With respect to using alternative approaches I think there are two things. I'm not sure if this is what you're after, but first there is the actual presentation of the case study. Then there's the project itself. In the project, which is the evaluation tool for the seminar, we do say that the criterion is use of an appropriate approach. It doesn't necessarily say a single appropriate approach, but there may be other appropriate approaches to use. With respect to pointing out pitfalls, advantages, and disadvantages, I think, as we do case studies and make presentations, we believe that it is a natural thing we ought to do. Making it more explicit and requiring people to think about it and present it is another thing we might want to consider.

Mr. Klugman: Jim mentioned that we have divided the room up at other seminars. After we've agreed on perhaps more than one good idea, we'll say, this group huddle together and carry out that one, and the other half will do the other one. That gives us an instant sensitivity analysis as well so we can learn a little more by seeing the consequences of adopting two different strategies. When possible I encourage students to think in those terms, that there isn't a unique correct answer. Dividing up the room, that creates the time, and we can try out alternative approaches and compare them.

Ms. Kirby: I'm actually a little encouraged and a little discouraged. I'm someone who's looking at taking the seminar; I'm a veteran of 121 and 152. When you compare it to that I get a little concerned. One of the things I'm encouraged by is that the pre-course reading material will be referenced in the seminar because currently in 121 and 152, it's not. One of the things that concerns me is that when taking 152 over the course of five days, I only got about 10 hours of sleep.

Mr. Klugman: Can I ask when you took it?

Ms. Kirby: I took it a few years ago.

Mr. Klugman: It has changed somewhat, and we've recognized those issues. In that seminar we stopped requiring the homework. We cut out the pre-reading because it wasn't being covered. We've learned a few lessons. We're hoping that by supplying some specific preseminar items that the students will be a little more geared up for what they're getting into when they come to Course 7.

From the Floor: I have a few logistical questions. I noticed that only half of these are going to be general, and being a health actuary I see only one health. Is there going to be a limit on how many people you're going to let take the seminar? If I don't win the lottery, I'm going to have to wait a year or I'm going to have to take a general one.

Mr. Luckner: A couple reactions to that. One, for educational and logistical purposes, we're anticipating that we wouldn't have more than 40 and would prefer only 30 people at a seminar. The distribution of seminar topics is based on a projection of the people that would be eligible during 2000 based on their current status and the conversion. The reason you see one health seminar is because that's what they anticipated. The predominance of general ones gives the opportunity for health actuaries to take one of those as well. It may not be exactly what you want, but it would be one that wouldn't disadvantage you.

Mr. Klugman: Because of some of the locations, we feel compelled to offer general. Hong Kong is one of those places. The people who are going to Hong Kong probably don't have another choice so, to be fair, it's going to be general. The health actuaries in Hong Kong will be stuck unless they really want to travel a long way. The more remote locations are likely to be the place for the general topic, and the practice-area-specific ones will likely wind up in the larger cities, which should help accessibility. As Warren said, we'll adjust as supply and demand indicates. We would hope by the second or third year that we'll have it pretty well in sync. We're hoping one health is the right number, but our models aren't real good.

From the Floor: I understand. The $700–1,200 cost is quite expensive. You're likening it to the 121 and 152, but those are electives. I know when I took 121 and 152 my company said, "We'll pay the fee for the seminar, but you're paying for everything else because it's an elective, and you can take another exam." This is required now. Companies are going to be forced to pay that fee, and the question is, is that just the fee for the seminar? You're going to have to deal with lodging, and you're going to have to deal with traveling expenses as well. You could be looking at close to $2,000 for Course 7. Is that correct?

Mr. Luckner: It's correct that it doesn't include expenses other than the seminar fee itself. We tried to identify locations based on where people would be. That presumably would reduce the travel cost, but it could be expensive. It's a breakeven price. We haven't settled on a final price, but it looks like it's going to be around $1,000.

We do have a couple of quotes we want to share with you about modeling to put things in perspective. "Far better an approximate answer to the right question, which is often vague, than an exact answer to the wrong question, which can always be made precise." Tukey (1962). This first quote is a caution against spending too much of your time and effort trying to make a model precise and accurate or fining tune it. You should really make sure that you have a model that is addressing the correct question.

"Don't fall in love with your model because ultimately it will be unfaithful to you." Bill Schmidt, Virginia Tech, Simulation Seminar Instructor (1999). This second quote is probably true regardless of the type of model you're talking about. The point is that a model is not a once-and-for-all-time thing. Things change. The techniques you use and the tasks you're assigned to do change. Don't get locked into thinking that your model, once it's developed, is something that's going to last forever. Be flexible. Be ready to let it go and think about other ways to do things.

CHART 1
GRAPHICAL VIEW OF MODEL FIT

Fit of exp. mix with K = 2