1999 Valuation Actuary Symposium Proceedings

September 23-24, 1999
Los Angeles, California

Session 38 TS
Applying Complex Adaptive Systems to Actuarial Problems

Instructors: H. Michael Shumrak
Vince Darley+

This session provides an introduction into how complex adaptive systems (CAS) can be applied to solve actuarial problems. Readers can:

- *Obtain a general overview of CAS, including genetic algorithms, neural network, nonlinear modeling, and agent-based modeling,*
- *Learn about recent examples of applications of CAS to practice business problems,*
- *Understand how CAS approaches can be used to address actuarial pricing and valuation issues, such as modeling customer behavior.*

MR. H. MICHAEL SHUMRAK: I'm a Hartford-based partner at Ernst & Young (E&Y) where I lead the New England life actuarial practice area. I am also involved in leading our effort to leverage technology and software tools to enhance our services. One of these technology development efforts is with a famous complexity scientist by the name of Stuart Kauffman, a member of the Santa Fe Institute. E&Y and Stuart have a strategic alliance run through a software/consulting firm called Bios.

Here is how we began to consider applying complexity science to actuarial problems. In our ALM practice we have been assisting clients with pricing and managing risk in variable annuity

+ Vince Darley, not a member of the sponsoring organization, is Senior Scientist at the BIOS Group, LP in Santa Fe, NM.

products, with particular focus on the so-called “guarantee” wrappers such as minimum guaranteed death, accumulation, and income benefits. We developed some very useful pricing and risk analysis modeling techniques. However, for lack of relevant historical data, our models applied a bimodal “savvy/naive” switch with respect to the assumption for policyholder behavior. Unfortunately, the range in price and/or risk margin between these two extremes is tremendous. We wondered if we might devise a new approach to modeling policyholder behavior that might help clients to better hone in on the expected behavior of their customers within the wide range between “savvy” and “naive.” A bulb went off, and we said, “Maybe this is a situation where we might be able to marshal certain complexity science modeling techniques to help us solve this problem.” We will elaborate more on this later in our presentation.

Our plan is to cover three topics. First we’re going to do an overview of complexity science in general and define and describe the various techniques in particular. Then we’ll move into the second phase of our discussions. After giving you a little bit of show and tell, in terms of what this technology is about, we’ll highlight a number of real life situations where the science has been successfully applied. In the final stage of our presentation, we’ll get into some of the details of the initial work we’ve been doing on annuity customer lapse behavior. We’ll also talk about some other actuarial problems that might benefit from applying complexity-science-based modeling techniques.

We are going to start off by introducing Vince Darley. Vince is a senior research associate with Stu Kauffman’s Bios Group. Vince started his academic life with a mathematics degree from Cambridge University in England, and then moved through theoretical physics and computer science degrees before being seduced into the complexity world at the Santa Fe Institute. His doctoral dissertation is concerned with understanding the natural dynamics of large systems with autonomous optimizing agents and how to design local interactions in order to bring about particular global goals. At Bios Group, Vince is applying these insights, understanding the behavior of stock markets and customer behavior in the recent work with us. He’s using agent-based simulation and developing more robust dynamic scheduling algorithms to leverage operational efficiencies.

MR. VINCE DARLEY: We'll start with a general introduction to complexity science. What's it all about? How does it relate to some scientific areas you might have heard of like chaos theory, genetic algorithms, and artificial intelligence? We'll look at examples where this science has been applied to businesses successfully. Finally, we'll talk about the work that Mike and I did together looking at policyholder behavior in the insurance industry.

The first section is Introduction to Complexity Science or Complex Adaptive Systems, which is another name given to the same area. Some of you may have noticed articles in the press over the last few years. In one, the Federal Aviation Authority decided that maybe it'd be a good idea to give pilots a little bit more control over exactly what routes they pick and where they want to go. A question arises. Why are they doing this? Isn't it surely much more effective just to control everything centrally and send out instructions to all the pilots and let them get on with it? The answer is that people are beginning to understand that it's not always better to do that and that we can let the tens of thousands of pilots have some control over what they do. They don't need to have complete control but some control. It can actually lighten the load on the central scheduling, the Federal Aviation Authority, air traffic control, and it can make the whole system work a lot more effectively.

One of the things that complexity science is about is determining under what circumstances central control is just a bad idea. How can we distribute that control or distribute the understanding of the system into a lot of the different parts it is made of? The systems that are usually studied by complexity science tend to have four properties. First, they're usually very non-linear and very dynamic. We're usually interested in the interactions between a large number of agents, which is the second property. In the previous example, it would have been the pilots flying the planes. We're interested in the emergent properties that the interactions of all these individuals creates. In the case of the air traffic control example, there's a belief that perhaps you can get a more effective scheduling, more effective routing, and more planes in on time by giving the pilots a bit more control and autonomy.

One of the difficulties lies in the third point. These properties tend to be extremely difficult to predict, and this is where computer simulation techniques and science lies in this. It is in carrying out this third step and trying to understand the individual parts and the interactions that creates these interesting global properties and how you can make use of them. The fourth point is that you often have rather sensitive dependence on initial conditions. Some of you may be familiar with chaos theory. There have been a few good books written for the layman. Some of those four points certainly are pertinent to chaos theory. However, there are some definite differences which I'll highlight here.

Complexity is often called anti-chaos, and the reason for this really is that superficially it seems to have the opposite property of what people tend to say about chaos theory. Chaos theory is really about how very simple, small, low dimensional systems, like a dripping tap, can have incredibly complex dynamics, and they can appear to be random. On the other hand, complexity theory looks at incredibly large, complex systems with millions of interacting parts, and it tries to find when and where patterns emerge despite all of that crazy complexity. Sometimes there are simple patterns that come out of that, which we can begin to understand. However, it's not really such a dichotomy. There's a lot of overlap, and if people are interested, I can get into that in more detail or you can talk to me afterwards.

Another area of science that became very popular in the 1980s is artificial intelligence, and it's since fallen from favor because it wasn't able to deliver on a lot of the early promises. The main difference between a complexity type approach and what went on in the 1980s with artificial intelligence is that most artificial intelligence approaches were very much top down, always looking at the very highest level, concepts, semantic categories, natural language understanding, and expert systems. It tried to extract high-level knowledge, facts and information and rules from people and to embed that in the computer. The complexity approach is quite the opposite. It says start at the very bottom. Start at the simplest level and try to build up to the complex interactions of the higher level.

Genetic algorithms also became very popular in the 1980s, and they are still used today. Genetic algorithms are a kind of evolutionary technique for solving problems. The idea is that you have a very difficult problem to solve, and it might seem very hard to find a good solution. Let's just create a whole population of solutions and throw them in a computer and then let those solutions evolve over time. Let them interact, mutate, intermingle, and create new solutions. You'll gradually get better and better solutions. It has had some successes. It tends to be rather brittle, especially with regard to a lot of real-world problems where the problem is changing all the time. Many of these techniques, such as genetic algorithms, simulated annealing, and evolutionary strategies don't really cope very well when your problem is changing on you every minute, every day, every hour. Many of the real-world problems tend to be of that nature.

Much of what we'll be talking about are agent-based models. Agent-based models are really one of the main ways in which people try to understand complex adaptive systems. The main tool in the toolbox of complexity science is an agent-based model. There are others, but this is probably the main one. What is an agent-based model?

FROM THE FLOOR: What is an agent?

MR. DARLEY: An agent is the individual unit of interaction. In the case of the pilots flying their planes, the agent would be the pilot or the plane. It's the individual unit that you want to treat — the thing that's displaying some behavior and making decisions. In the examples of policyholder behavior, there will be two different types of agents in the models. There are the individual policyholders. Each of the million policyholders is an individual agent. The salesmen, more traditionally called "insurance agents," are also agents in the model, and there are interactions between those two agents in terms of sales decisions and advice. The agent is really just the unit that you're looking at.

An agent-based model always has a population of agents. There are often millions of agents and there are some rules of behavior. There are the rules of interaction. We often find that it's the interactions that are most important, although the behaviors of the individuals are important. The behavior of the system, as a whole, is often more subtly dependent on the interactions between the individuals than on the specifics of any one behavior. There have been some interesting studies.

I have given a few examples of where this kind of technology has been applied, and one area is in traffic modeling. There are now at Los Alamos National Lab some very big models of the traffic flow in Albuquerque and Dallas. In this case, all the agents are the individual cars, and there are rules for how people tend to operate. There is a range of different rules because people tend to drive in slightly different ways. The lab looks at what kinds of traffic patterns emerge: The lab asks, "How can we adjust these rules? If we change the way in which people behave a bit, can we improve the situation? If we change our road network a little bit, can that improve things?" They found that changing the way people behave doesn't really have too much effect. As long as people have a few simple characteristics for how they drive, you're going to get these same traffic jams in the same places in your road network all the time. The solution is to then redesign the road network. By using these kinds of models, they can then test out different road networks and try to work out what's going to improve the situation the most.

We now come to the fourth point. What we're interested in are these aggregate statistics. In the previous example, it was the statistics of traffic jams, how often they occur, and how long they last. These phenomena occur at the very highest level. Even though we have a very low level description of individuals or interactions, the observations in which we're interested are these high-level things. When we talk about insurance models, the observations will be aggregate statistics of decisions that policyholders are going to make about their policies; are they going to lapse their policies or annuitize the policies? These aggregate statistics emerge out of the underlying interactions, the underlying behaviors, and those are the things we're always interested in.

The agent-based modeling approach has a few advantages over most other approaches. One of them is that the agents can be heterogeneous. Most modeling techniques will tend to assume there are a million policyholders. They're all basically the same. We'll just take an average policyholder, a representative policyholder, and sometimes that's an okay approximation, but often it won't be. The advantage of agent-based modeling is that everyone is treated individually. If you have data on a million individuals, you can feed in all of that data, and you can really understand whether the fact that they're different is important and whether perhaps particular segments of the population are of much more interest than other segments of the population, something that you'd never discover if you just had one representative behavior type and one representative agent.

Another advantage is that often we can make these agents more behaviorally realistic. Because we're using the computer to do all the hard work for us, to simulate the system, we don't need to make simplifying approximations quite so often. We don't need to assume, for mathematical convenience, that everyone is perfectly rational and then solve some equations. We can put into the computer a more reasonable behavior for the individuals, and then let the computer work out what's going to happen. The third advantage, which I guess I've already covered, is that the interactions between agents are treated very explicitly. Because we have this fine-grained model of every single individual, we also have a fine-grained model of the interactions between them all. In the case of traffic modeling and the air traffic control examples, you can see that that can be very important and that somehow having an average system in which there's just one representative agent and the interactions are all averaged might really lead to the wrong answers.

The fourth advantage, depending upon your persuasion, is either an advantage or a disadvantage. Because these tend to be simulation type models, the computer will generate the results for you, so you don't have to worry about mathematical problems or analytical difficulties. You have a whole host of new problems that arise. We just have to understand the results of the simulation, how to make sure that they're robust, how to convince yourself that they're correct in some way, and we'll come in due course to ways in which you can understand that. You can understand sensitivity of the results that you observe to the initial conditions, to the behaviors you put into the model. There are many ways in which you can verify what's going on. They're related to the traditional ways of verifying a more traditional mathematical statistical model, but they are slightly different. We'll come to those later.

Here are a few examples. I've already given a couple as I've gone along. Much of this area of science really came out of looking at the real world, particularly at biological systems. There are two biological systems that have really caused enormous improvement in our understanding of them through this kind of approach and this kind of theory. The two areas are our immune system and our understanding of colonies of ants, bees, and those kinds of insects. The real problem with the immune system is that the number of possible antigens of diseases or whatever that can come into our body is extremely vast. There has been some enormous range of possible cell structures that are foreign to our body, which our body needs to be able to recognize and then attack and destroy. It's something like a trillion possible surface shapes of cells that you need to be able to recognize. There's no way your body can contain enough white cells to be able to recognize all of those individually. Your body has to take an evolutionary approach wherein when the foreign, unrecognizable particles are observed, your body actually constructs something on the fly that will recognize that specific thing very, very accurately. Then it replicates it billions of times so that the new white cells can go throughout your body and destroy the antigens.

Another typical example is ants. Ants, termites, and bees all have astonishing properties when you look at them as a colony and as a whole, but the individual ant is really rather stupid. It has very simple behaviors, and it's astonishing that it's able to construct anything at all, yet termites are able to construct very complex colonies with sort of air conditioning systems built in so that the temperature in the colony is very constant. It only fluctuates by a degree or two year-round. This is all done by construction of peculiar air flow through the colony. You wonder how on earth this can take place when these termites seem to have incredibly simple brains and very simple rules by which they operate. There's another fantastic advantage of this kind of distributed intelligence; if half an ant colony gets destroyed, the other half still carries on, and it functions perfectly well. There are very few human-created systems that can do that.

What has happened in this science is that we've started up at the top, and we've looked at these systems in the real world in terms of nature and biology. We've begun to understand that we can apply the concepts in the created world to the same kind of ideas, and we can understand them a lot more effectively. A few examples include the Internet, companies, and the supply chain.

Let's discuss some typical sort of coordinated behavior by some ants in the rain forest. Here's a brief history of this kind of scientific approach. I guess it started way back at the end of the last century with Poincaré who made some incredibly foresightful statements about weather forecasting and about how it could be that systems could be unpredictable even though it looked like you could understand every little part. More recently, in the 1940s and 1950s, there was greater and greater understanding of these kinds of phenomena in which very simple rules can give rise to very complex behavior. Lorenz was one of the first who really computationally noted these problems in weather prediction. It is one of the main examples often quoted of chaos.

Then, through the 1970s and 1980s chaos and practical mathematics became better and better understood. In the 1980s, the computer began to be used much more. We're now pretty comfortable with the computer and the way in which it can help us solve real-world problems. I listed a few examples of where we've taken this area of science and we've pushed it sort of to the last stage where it's being applied to businesses all over the world. We'll go through a few examples of this afterwards. The areas in which I've been particularly involved have been in dynamic scheduling, decision support, and educational tools.

This technique is very new and interesting and exciting, and it certainly can produce wonderful results under some circumstances, but it's not going to solve all the problems. There are many cases in which it's not appropriate to use a complexity-based approach. There are some distinctions between when you should think about using a complexity-based approach and when it's probably not going to give you much of an advantage. I mentioned earlier that when your problem keeps changing, that's a case where many traditional techniques just can't keep up anymore. Many real-world problems change on a minute-by-minute basis nowadays or at least on a day-by-day basis. In those situations, this kind of technique is certainly something worth considering.

For the purposes of today's focus on the insurance industry, one thing I could add to this list is that when you have a large amount of real data about the behavior of customers, then a traditional approach is something that probably will work reasonably well. When you don't have a lot of data or any data because it's a new kind of policy (you're introducing with new kinds of guaranteed benefits), then it becomes very hard to use traditional statistical techniques. In this case, as we'll see later on, this kind of approach can really give you an enormous advantage.

For example, let's say all of our policyholders have some understanding of how they behave under different circumstances. Perhaps we've calibrated that using real-world data that we have from some past policies people have held. We understand how they're going to behave under some kinds of different market conditions. We know how they're going to make their decisions and if they're getting a certain return on their policy right now. If they could get a better return somewhere else, how many of them are going to switch their policy?

If we have some understanding of basic decision rules like that, then we can create an agent-based model of these customers, which is what Mike and I have done. We can then ask questions about how these customers are going to behave under rather different circumstances for which we didn't have any data in the first place? How will their decisions be applied to the new input, to the new environment in which they're faced? For instance, the stock market might drop a lot, and we really have no real data about how customers are going to behave effectively. We can understand how to value our policies in these cases, how much risk is really held in those policies under a circumstance for which we didn't have any data. That's the real advantage for the insurance policyholder and we can begin to answer questions for which we don't have any data and for which traditional techniques just aren't going to work.

I'll move on to the second section of the three to go over a few examples of businesses where these techniques have been applied and where they've been very useful. For those of you who are interested in learning more about this area, there's an excellent book by Kevin Kelly called *Out Of Control* which is very entertaining and interesting. He talks about all sorts of new things in the world, and at least half of the book covers these kinds of topics. The whole book is just very entertaining and well worth reading.

In this business application section, I'll cover two main areas in which these techniques have been applied very successfully. One of them is in scheduling, and the other area in which I've been involved is cargo scheduling for Southwest Airlines. They route an enormous amount of cargo around the states for the post office, and they found they had all sorts of horrible bottlenecks in their system. They didn't really understand why, and, in fact, by using these techniques we were able to improve their performance throughout the entire system by an enormous margin. I'll show a demonstration of that in a moment.

First, I'll go over another example, which was a paint booth scheduling problem, and this is one that really shows how it could be that these techniques can really work very effectively — much more effectively than traditional techniques.

General Motors makes a lot of trucks and cars, and its system for painting these cars was really quite simple. There is an assembly line, and when the trucks come off the assembly line, they move to a whole line of paint booths. The trucks have to go into a paint booth where they get painted with a particular color, and then off they go. General Motors had a central scheduling system for this entire system. There are many thousands of cars coming off the manufacturing lines every day, and all of this is planned in advance. They calculate a good schedule, and they know that Truck 1 is going to go to Paint Booth 3, and it's going to be painted red, and then Truck 2 is going to go to Paint Booth 1 and be painted white. It's all calculated so that the number of paint changeovers that a paint booth has to make to change the color of the paint it contains is very small because the paint is incredibly expensive. It's about $400 a gallon. When they have to change the color of paint in the machine, they must throw away many thousands of gallons. The problem was that this schedule, in advance, looked wonderful and very efficient, but then about a half an hour into the day, one of the paint booths might get clogged. Workers had to go in there and clean it out, which took another half-hour. If one of the manufacturing lines breaks down for a few minutes, and before half an hour or an hour's up, the whole system is thrown off. The schedule that they produced became useless, and they had to manage things on the fly all day long. It just didn't work at all.

The solution was to throw out the old computer system that had millions of lines of complex codes. It was very difficult to maintain, and a good thing to get rid of. They replaced it with a very simple system. The new, simpler system had an individual paint booth with an agent. When a truck comes off the assembly line, each paint booth makes a bid for painting that truck. This bid depends upon how easy the job is, whether it's the same color paint, how important the job is, if there's some priority attached to this particular truck, how long it's going to take to actually get the job done because the agent might be in the middle of another job. The paint booth with the highest bid would win, and the truck would roll on over to that paint booth and be painted. It was a very simple system. Almost nothing can go wrong. If a paint booth breaks down, that's fine. It just stops bidding, it gets fixed, and the other paint booths take over. There's really no problem. The system kind of dynamically adapts to whatever's going on. By using this system they were able to save millions of dollars in just paint costs, let alone everything else, in the first year. This is an example that shows how throwing out the central control really improved things.

Here are a few other examples. Credit scoring has been one big area in which this kind of approach has been used a lot. Nowadays, a lot of big credit card companies have data on millions or tens of millions of customers. There's a lot that can be done with that. The other example is reasonably similar to the paint booth. It has to do with dynamic routing of cement trucks in Mexico based upon how the demand changes all the time.

I'll discuss a model just to give you an idea of what some of these models look like. Let's discuss a model of Southwest Airlines' entire flight network. In this agent-based model, the agents are both the planes and the individual pieces of cargo that are being shipped. There are about a million pieces of cargo being flown around. The routes of the planes are pretty much fixed, but there are many choices of how to route a piece of cargo between L.A. and New York. There are obviously many different routes you could take to get it over there, and the decisions have to be made on the fly.

The interesting thing about Southwest Airlines is they are very technology averse. Their system is designed to be as simple as possible. They want the system to be as simple as possible without any overhead attached to computers and central control or anything like that. They need a solution that tells them how they should ship a package from L.A. Should it fly down to Phoenix or Las Vegas? The rule that they were using was a very simple one. We call it the "hot potato rule." If it's going east, put it on the first plane going in the that direction, which meant that most packages went to Phoenix.

FROM THE FLOOR: How does that differ from the way the Internet works?

MR. DARLEY: That's a good question actually. The Internet uses a slightly more clever routing system than that. Each hub or each location on the Internet has a routing table that basically says, if something's coming in, and it's go to x, look up in the routing table where it should be sent. Those routing tables are re-calibrated over time. In Southwest's case, the routing table would have one entry: the next plane. The Internet is a little bit more sophisticated. To some extent, that's what we gave to Southwest; we gave them a routing table that was just a little bit more complex than what they had. It took account of things like, a plane that is actually going from L.A. to New York via two or three stops. Then you shouldn't put the package on the next one going to Phoenix. Wait half an hour and put it on the plane that's going to go directly to New York.

By implementing really simple rules like that, they were able to cut down on the number of transfers of their cargo by 20% during the day and by 80% on overnight transfers. The overnight transfers are the really bad ones because they have to cart everything off to their warehouse and back again. This had an enormous effect on their system. The way in which we studied this wasn't particularly sophisticated. We created this agent-based model of the system, plugged it in the rules they're currently using, which you find out by just interviewing the cargo personnel on the ramp, and then determine if there are some other rules that could work better. It's just really a hands-on, what-if tool. If I change a rule, what happens? We were able to cut down not only on these transfers but also on the number of delays (which is very important) because the post office pays Southwest based on its frequency of delays. If it can manage an 80% on-time arrival, then there is a bonus of so many million dollars. Southwest was very happy with the solution.

I'll go back to the presentation now. As you've no doubt gathered, the computer is the enabling technology here. Without the computer, we wouldn't be able to do any of this. With computers becoming cheaper and ever more accessible we can now easily build models in which there are millions of agents. Five years ago we would have been dealing with thousands. That's something that's of particular interest to the financial sector or to the insurance sector area where you do tend to have millions of individuals under consideration.

Let's move to another example, which is customer behavior. It's a bit closer to where we're heading. I'll share some work that was done for Disney. Disneyland is a very complicated system. There are all these rides. There are lots of people queuing up all over the place. It's extremely important that everyone is happy. Disney goes to incredible lengths to make sure that the parks function effectively. If there's congestion in one part of the park, they'll try and arrange for some parade in some other part to kind of pull people away. If they think the park has been extremely congested one day, they'll leave it open a little bit longer to make sure that everyone sees enough of what they want and goes away happy. There are all sorts of decisions that need to be made on the fly. Currently, these decisions are all made by one guy who has been doing this for years. He just knows the park well, and he knows how it works. He knows if there's a problem in one section, you just need to do something in another section. It's fantastic. Unfortunately, he's getting quite old now.

Disney had been investigating using some of these complexity-based or agent-based simulation techniques. We produced a prototype model for them to show how we could look at the flow of individuals through the park and at any individual ride and to see how full it's getting. We can look at waiting times for rides, how many rides individuals go on, and how full the park is. Is the relationship between the number of people in the park and the actual capacity of the rides a nice linear thing, or are there somehow bottlenecks in the system so that people aren't somehow getting to where they should be getting quickly enough?

How much people spend is very important. In fact, there are all sorts of interesting decisions that are made depending upon what kind of person is visiting the park. If there are a large number of Japanese tourists, they tend to spend more time in shops. In that case, Disney will have to lay on extra staff in the shops. There are all these kinds of on-the-fly decisions that have to be made.

If we look at a map of Disneyland, we would see all the various rides. Individuals will come in. They flow around between the different rides. You can pin down any one ride and look at how long the queues are and what's going on. You can look at which parts of the park are congested. Then you can start making decisions about placing other attractions elsewhere, and you can look at how the people will flow in the park based upon that. There's actually a secondary benefit of a model like this. Not only can you study the existing park, but you can actually study what happens if Disney wants to add a new ride to the park? Where should it be so as to best prevent congestion and maximize the public's benefit? Again, you can use it as a what-if tool for the future for situations in which there aren't any data. There's obviously no data if you're going to construct a new ride. How are you going to deal with that? That's an issue we'll come back to later in the customer behavior modeling.

The second business application is, in some ways, a little bit more closely related to the policyholder behavior. It is in portfolio management. There are two ways in which the complexity-based approaches can add some new benefit to portfolio management. One is portfolios consist of a large number of parts that make up the portfolio. Sometimes the interactions between those parts are nonlinear. There's often very strong positive or negative correlations between the different parts. Are they portfolios of stocks or portfolios of projects? There are investment and research decisions. Some of the greatest interest in these kinds of tools has actually come from pharmaceutical companies who are interested in making research decisions. They'd like to know how they should divide their research amongst a whole number of different goals when they know full well that if a particular goal succeeds, it's going to have enormous repercussions on other projects. Any kind of conventional portfolio modeling technique to try and understand the different research projects that they're considering just doesn't tend to handle these kind of interactions between the different parts.

The other way in which the complexity-based tools can help is in this first example — group decision-making tools. If you have a large number of experts trying to make decisions about which of these research projects they should invest in, everyone often has their own agenda and their own beliefs. All of this information from these individuals can easily be combined in a model like this so you can understand how the different people's perspectives influence your overall view of how good your portfolio is. In particular, you can understand whether perhaps there are misunderstandings amongst the group.

One of the ways in which these tools have been used is to take a large group of executives from the company and make them go through a portfolio decision-making process by answering a bunch of questions about various research proposals in terms of looking at potential long-term success, risk, and potential revenue streams. By using this kind of tool, you can analyze how effective different projects are and what the different characteristics are. You can also look at situations in which there's kind of contradictions in your data in which different people actually seem to have contradictory viewpoints about long-term potential success in a project. It can actually help you to drill down on why these contradictions exist. You can then go back to these people and create some discussion. Why is it that they have completely opposing views? You can then change the dates or update based upon the results of those discussions. You can really understand the trade-off between the kinds of things that you'd like, net present value, cost, number of employees, how novel the project is, and various risk estimates. This project really allows you to balance all of these tasks. One of the very nice things it does is it lets people see the impact of their own individual decision. You can put a company's future plans into this tool. An individual in a particular branch of the company can say, "This is my pet project. I'd like to throw that into the mix." That person can instantly see how much impact that has on the sum total of all these other projects. He begins to realize that maybe the pet project wasn't so important after all. At least he can make more balanced decisions. That's another rather different kind of tool from what I discussed earlier. It's a sort of new area in which these techniques are being applied from the basis of helping groups make decisions more effectively. In decisions about portfolios, there's a large number of parts amongst which we're trying to choose.

Here's another summary of some business applications. Boeing is a typical example. Boeing obviously designed planes, and there's this very complex design process. A plane has an enormous number of parts. The goals for each of the parts often come into conflict. You'd like to have the fuel tanks as large as possible so you can have a long range, but then you have a much higher take-off weight, so you need a longer runway. You can carry less cargo. There are all these trade-offs. If the wings are longer, that's fantastic, but that makes the plane more expensive. Maybe you need a bit more fuel. There are all these trade-offs between these different characteristics.

One thing that Boeing learned from the design of the 777 was that the design process lies in the interactions among all of these conflicting goals. They estimated it cost them about $7 billion to design the 777, and they see no reason why it should cost more than $1 billion. The $6 billion in between shows that no one can come up with designs that interact effectively with all the other designs. They constantly have to redo designs because they all come into conflict. Ways in which you can examine trade-offs among the different designs more effectively to get separate groups and separate design teams interacting together to create one unified plane are incredibly important.

I guess these are typical examples of these kinds of techniques. The U.S. Army is interested in dynamic scheduling of drones for information gathering purposes. You can send little drones over a battlefield, and they'll arrange themselves to go to those sites of interest that have been identified as being of great importance so they can take pictures and gather data. Of course, they're being shot down all the time as well. They need to have a system that rearranges itself on the fly to sort of maximize the potential intelligence-gathering capability.

At this point we'll move on to the third part of the talk, which is the actuarial applications; in particular, we'll discuss the work on policyholder behavior that Mike and I have been doing. At this point I'll switch back to Mike for an introduction on this topic.

MR. SHUMRAK: We'll now turn to actuarial applications. First, we'll talk a lot more about some of the early work that Vince and I have done with a couple of our colleagues. Marshall Greenbaum is a senior consulting actuary in our capital markets practice. Our fourth contributor to our work was Rob Axtell who's a fairly well-known economist at Brookings who's also a Bios fellow. After we discuss our initial work on applying agent-based models to model policyholder behavior, I will describe a number of other ways we might solve actuarial problems using these techniques.

As you recall from my introductory remarks, we have set the scene. In our capital markets practice, we've developed certain models to price and value the risk and the guarantees around variable annuities. One of the most challenging right now would be the so-called guaranteed minimum income benefits (GMIBs). Not only is this a new product with historical experience, based upon the definition of the product benefit, there will not be any historical experience for some time to come since the decisions policyholders face relate to how they might behave in reaction to your product and capital market movements several years after they purchase the product. We were pleased with our actuarial financial approach to pricing this product feature until we came to the assumption for policyholder behavior. In the past, actuaries may have run a set of deterministic scenarios to get a feel for the impact of variations in customer behavior and attempt to set an expected "case" that would represent the so-called "pricing" assumption. More recently, actuaries would consider logistics equations tied to various factors such as what competitors and capital markets were doing to "dynamically" model expected customer behavior. However, even this is simply curve-fitting based upon informed opinion from the actuary's view of the underlying situation as to product design, distribution channel, etc.

Because we were not satisfied to tune the curve fitting between "savvy" and "naive" customers, we wanted a policyholder modeling methodology that might capture the real dynamics so that we might narrow the range between these two extremes. The approach we used was a combination of behavioral economics and the agent-based approach that Vince has already described to you.

Behavioral economics is the science of determining how people might respond to various situations, such as policyholders reacting to advice form their salesperson or to moves in the capital markets. The data can either be historical behavior or data developed through various market research techniques. Those of you who attended the Investment Section's guest speaker luncheon at last year's SOA Annual Meeting in New York heard from one of most famous behavioral economists, Richard Thaler. We all know from our growing involvement in finance applied to actuarial science, markets are assumed to be efficient, and people playing in these markets are assumed to be perfectly rational — everything works like a machine. However, when you really step aside and think about it with respect to markets such as the retail annuity market, efficient markets with rational players do not seem to happen. We also know many insurance buyers are risk averse. Behavioral economic research has shown that, depending upon how you frame a question to a person, you may get two different answers, even if the true economics are identical in both cases. The analogy to "framing a question" with respect to our business would be situations like how a product is offered in terms of design and pricing points, how a conservation campaign is positioned to annuity customers beyond the surrender charge period and so on. For example, I can say I'll sell you two suits for $200 each and get one result in terms of sales penetration in a market and then go into an identical market selling buy one for $400 and get one free and get a different sales penetration result. This example may seem trivial and based in hucksterism, but "framing" clearly makes a difference, even though the mathematics of the game theory would tell you otherwise.

A second example is as follows: You are advising a CEO of an oil company. One of his competitors has gone under and is going to auction off some of the potential plots where there might be oil. The CEO tells you he figures there will be three companies bidding. The CEO decides to bid $2 million for the plot. At the last minute, there are eight bidders. The CEO asks you how he should change his bidding strategy. Perfect rationality would suggest that with so many more bidders, you would expect that each of them is making a guess at an offering price that would result in a "good buy." However, with more bidders, the likelihood that the winning bidder may be one who overvalues the plot is increased. Therefore, you should advise the CEO to lower his bid. The research shows that most would increase their bid to increase their chances of the winning bid. In fact Richard Whaler's amusing and informative book on behavior and economic phenomena is titled *The Winner's Curse,* and he had this example in mind. Having established the importance of behavioral economics, how do you develop behavioral economic data? The obvious starting point is if you already have individual-specific information on the matter at hand, you simply use it. In many cases, such as with the new GMIBs, you won't have this information. Here you will need to develop the behavioral data through interviews and surveys. These methodologies will produce interesting but not so hard and statistical data. If you ask the wrong questions or get the wrong group on the wrong day, the answers might only be directionally useful. However, we're trying to integrate the behavioral data into a quantitative model.

Another more formal approach is to do a focus group. A focus group is a little bit more organized. We really get experts to stage and ask the questions and pick the right representative participants. The results of the focus group, however, may still only be directional and not so quantitative. If more quantitative edge is required, some companies will apply conjoint analysis, which is a market research technique also known as trade-off analysis. This can be quite powerful. Basically what you do is use computer analysis to analyze policyholder and sales force responses to trade off questions with respect to one or more factors that you believe impact their decision to act (i.e., for the salesperson — give advice to a policyholder, for a policyholder — decide whether or not to lapse). For example, if you had two or three guaranteed benefits you're thinking of surrounding around your variable annuity, you would combine them in different ways to ask the value the customers might place on each benefit or benefit combination. You would then compare their values to the economic costs. The benefit with the greatest excess of "perceived value" over "economic cost" would provide the greatest potential.

In our work, both policyholders and salespeople were examined with respect to their behavioral economic behavior. In Vince's parlance, these were the two agents that would be studied in our agent-based modeling approach to policyholder behavior. We wanted to capture what an agent might do, what an end buyer might do and then how they interact to produce policyholder behavior. Anecdotally, most of us agree that if you have an aggressive, active broker who can write for lots of companies, he may well, just like the perfectly rational "quant" (financial engineer) on Wall Street, keep moving a piece of annuity business around to the best current deal. Conversely, data have shown that annuity business sold by captive insurance agents and/or rural banks tends to stick on the books despite the better opportunities that may be out there. These are the savvy and naive extremes. The strength of behavior economics combined with agent-based policyholder behavior modeling causes you to be able to start moving toward an individual-specific paradigm and away from the everyone-fits-some-average paradigm.

MR. JAMES VAN ETTEN: Are there relations that you lay in off the data or is this a statistical analysis only?

MR. DARLEY: You are really developing a bottom up characterization of the way salespeople and annuity policyholders might behave and interact against the back drop of stochastic variations in the economic environment.

MR. SHUMRAK: Right. We're going to infuse a lot of if-then statements around a customer and around an insurance agent, and those would be the behavioral economic aspects of each of the agents, and then we would take that and do the agent-based modeling to let that sort of synthetic computer-based world play out and see if we can learn something from it. It's not creating *the* right prediction, but it's giving you a sense, just like the ant colony, where any one ant's actions seem random, but collectively the ants' progress is very compelling.

MR. DARLEY: I'll go over the model that Mike and I put together which is basically just an elaboration of what Mike has just described in more detail. I'll try and show you how we go about creating an agent-based model of this kind of system, what kind of behaviors we put into it, and then what kind of analysis and calibration you can do. What kind of results do you get out of it in the end?

In this case, we have two types of agents in the model. There are all the policyholders which are called customers, and then there are the policybrokers and the salespeople, which we'll call sellers. There also are the current market conditions. How well is the market performing right now?

This study has been confined solely to looking at lapse rate. The decisions that individual customers are making are solely whether to lapse their policy or not. It wouldn't be a particularly difficult task to extend it to annuitization decisions, or other kinds of decisions, but in the first case we're looking at focuses only on lapse rates. You'll see that one very important thing that we can get out of this is a sensitivity analysis. We produce this model, and then we have to ask ourselves whether we believe this model? Does it seem okay? There are a variety of steps you can go through in trying to understand that and in also trying to understand whether you're going to start using the results for some other purpose. How much of an error is there going to be in these results, and how much uncertainty will there be in the future? If you want to estimate risk, estimate probabilities of particular future events. This kind of model can actually give you some understanding of those probabilities. The last really important point is that this model has potential to predict lapse rate for other conditions for which we just don't have any data at all.

I'll go through the different agent types and behaviors that we put into them. The agents split into two halves. There are the customers and the sellers. To gain a deep understanding of how customers behave and how sellers behave, we used the expertise of Mike and some of his colleagues. To apply this to a real block of policies, you'd want to have greater understanding of the individual customer behavior, and you'd want to carry out focus groups or surveys to get more detail; however, we view this as a good proof of principle. We've embedded reasonable knowledge of how customers behave based upon the experience of practitioners, which is a good first step in the right direction.

There is a basic decision process that a customer goes through. First of all, there's a small chance that they'll die (0.01). The numbers in parentheses are kind of estimate values for what these particular parameters should be, and later on we can calibrate the model and adjust those to fit a particular piece of data effectively. These will give you an idea of roughly what we're using as a first-instance guess. If they don't die, then some customers may just lapse for other reasons (0.02). They need the money to buy a house or whatever the reason happens to be. There are various liquidity reasons for which they'll just lapse independent of whether it is or isn't a very sensible, economic thing to do.

Assuming that they don't lapse for that reason, then some of the customers will have received advice from their brokers or from the person who sold them the policy in the first place. They might have received some advice that they should switch their policy to something else or keep it. Maybe the market is going up. They'll look at that advice and compare that advice with the current return they've been getting on their policy and obviously factor in any kind of surrender charges that happen to be applicable depending on the age of their policy. They'll decide whether to take the advice or not based upon that information.

If they don't receive advice, some of the customers may be your kind of more economically aware people who might think, maybe I should look at how good this thing is. They'll compare the returns they're getting with what the market is offering and with competitors' policies, and they'll make some decision based upon that. Some customers may just imitate the population at large. There's a vast amount of commonality of information sources throughout the media. If an individual hears that people are lapsing their policies these days, or the stock market is dropping like a stone, or whatever it happens to be, there is a chance that he or she will follow the trend and do what everyone else tells him he should do. So that's the basic decision rule in the customers.

Because of the way the model is designed, we can figure how many customers we want. You can run with a million customers on a simple notebook without any great difficulty. Each customer just has one policy, and we just assume they all started at age zero. The model doesn't dramatically evolve through time with people changing policies and having the policy age reset to zero. It's just kind of a one-shot model, although that's not something particularly difficult to change. We just uniformly distributed the customers' ages over a given range, and we used a kind of 80/20 distribution of account value so that 20% of the customers had about 80% of the values, which seemed like a good estimate of reality. For any particular case, where you had a block of real customer policy data, you just feed that data straight in, and that would be a very simple thing to do. Each year the agent decides whether to lapse the policy based upon the previous rules I outlined, and each year the agent obviously ages by one year.

So now we move over to the other side — the policybrokers. There's a user adjustable number, and we've used up to 10,000 of them. We distribute the customers amongst the policybrokers according to a kind of 80/20 type distribution so that 20% of the policybrokers have about 80% of the customers, which again seems roughly what goes on in reality.

There's one last point. At the moment we just use one distribution channel. All the sellers seem to be of the same kind of characteristic. Policies are, in reality, sold through a variety of distribution channels. For any given block of policies, you'd actually want to create a model like this wherein some of the policies were sold through a bank, and some were sold through policybrokers. Maybe some were sold over the Internet or whatever it happened to be. In those cases the interactions between the salesperson or bank and the customer would probably be rather different.

There is the behavior of the sellers in our one distribution channel. Each year the seller will pick some fraction of his customers, and rank its customers from high to low account value. He will tend to look at the higher account values first, and then for each of those customers, it is going to pass some advice on to the seller who will look at the current market conditions, the returns that a given policy is getting, and pass on some advice about whether the customer should lapse or not. This advice may be rather optimistic for those sellers who tend to try and get policies churning a bit more frequently than others. There's a sort of range of optimism throughout the population of sellers.

There are a few other aspects of the model. We generally used a linear surrender charge schedule, usually from 7% to 0% or 6% to 0% for one case where we were calibrating to get some Life Insurance Marketing and Research Association (LIMRA) data. Every year account values are updated. The accounts are split into a variable part and a fixed part, and those are updated according to the prevailing market conditions. We can drive the model with various assumptions about interest rates and equity returns and things like this so we can examine a strong, bull market or a bear market, and we can drive the model with these kinds of inputs.

What do we get out of the model? The first thing we get is this sequence of lapse decisions that the customers make over time. Because the model is stochastic in nature, there's a lot of probabilities embedded in the model. We can run the model multiple times and get some estimate of how much variance there is in the outcome. That would give us an estimate, if you wished to characterize risk. Of course, we can just examine the behavior of the model as a function of all of these input parameters and a function of the behavioral parameters. We can adjust how much attention individual policyholders pay to the advice that they receive, and we can examine how their lapse decisions vary over time. We produced a tool in which there are various parameters that we can input for the behavior of the individuals. How many are there? What are the different kinds of scenarios with which we drive the model? Is the market performing well or badly? What are the interest rates? We can choose how many iterations to run the model for. There are output statistics. The horizontal axis would show time varying from now to about 8–10 years in the future. The vertical axis would show a percentage of policies lapsed. One graph would show a percentage of total policies lapsed, and another graph is percentage by policy value, which is the more important case because that's really going to help you decide how to value the entire block of policies you have.

Here's the simplest thing we can do with this model. Once we've adjusted the various probabilities and the various parameters that the individual customers and sellers have, we can then change the market scenario and ask, what happens if the market tanks? We can see how the behaviors change over time. We can actually go a little bit further than that. We can really do some proper calibration. In this case, we took some data from a LIMRA study for lapse rates over time, and we ran them to try and find good behavioral characteristics for the customers and the sellers to match the data as best as possible. The final calibrated value is a little bit closer than the average data from our model. So that's great. Once we can do that, we can also calibrate the model. We can also hand-tune this because we designed these behavior rules. If the lapse rates are a little bit low in the sixth year, I should probably increase the way in which the customers accept advice from their brokers. You can do a mixture of hand-tuning and calibration. Once you've done that, the really interesting thing is you can now change the market conditions. You can say, if we have a very inflationary situation and a bad market, returns are very bad. What are customers going to do with their policies? Customers will tend to lapse a lot earlier in much greater numbers than otherwise.

On the other hand, if there is low inflation in a good market, the customers will tend to retain their policies longer, and the total number of policies lapsed is much smaller. The spike in the curve would be a lot lower than it is in the base case. That sort of shows you pretty much what you can do in a model like this.

You design the behavioral rules, you take in some input data for the current market conditions, and you run the model. You can understand the sensitivity of the model to various parameters. You can understand the variance in outcome. Those two factors are very important. You can actually calibrate the model if you have some real data, and you can change the market conditions and understand how those given policyholders would behave under a completely different set of market conditions for which you have no data at all. That's pretty much a summary of the model. At this point I think I'll switch back to Mike.

MR. SHUMRAK: As we gave our overview of some of those distributions and assumptions, I'm sure some of you were thinking, that's all very interesting, but is that right? Or, you might wonder, where do you get it? Clearly the next step, if this were to really be applied, is to really go through some full-blown market research. There's the real populating of the behavioral data, so that what you'd be looking at is particular companies, particular distribution channels, and a particular customer's situation. However, despite our abbreviated and simplified approach, we felt confident that at least we've proved the concept. Obviously there's no guarantee regarding the effectiveness of agent-based modeling in any given situation that it'll provide a silver bullet. The way I like to look at it is it gives you a *wind tunnel* to realistically understand how these things work as opposed to simply relying on logistical formulas that we used in the dynamics lapses. These dynamic formulae are much better than "simply running a few scenarios," but they are still very artificial. It is also very difficult to determine how much to turn the formulae's parameters up or down to characterize savvy and naive. Within the variable annuity situation, we'd also be applying this to other aspects of variable annuity product customer behavioral dynamics such as, fund switching from volatile subaccount funds to other less volatile subaccount funds such as a fixed account, utilization of the variable annuity guaranteed income benefits, and the propensity to annuitize (moving from accumulation to payout phase).

If we go beyond the variable annuity situation, it is quite easy to consider a number of other policyholder behavior factors like utilization of policy loans, periodic premium retention, and so on. The same applies to salesforce issues. For example, if we could be more selective, we know a smaller number of the agents working smarter and producing more profitable customers could produce more of the value. This could give us some sort of a genetic fingerprint for who is more likely to be a more productive sales agent.

We could also apply this modeling approach to product design. This is already used by some other folks in non-financial consumer products, where they're trying to figure our the next great hit product. What type of toy will be most wanted this season or how should I arrange the goods in my store so customers first see a certain category of merchandise? Instead of reaching product design and pricing structure decisions in smoke-filled rooms, those of us that are pricing actuaries, could actually do it by gauging how much impact a product feature or pricing structure might really have on sales and on risk.

MR. JOHN A. WOOLSEY: It seems like a real lucrative market would be Fidelity or Vanguard or something.

MR. SHUMRAK: I think you're right. I'm pretty excited about it because of the idea of moving toward individual-specific, actionable databases and customer information or agent information, and then being able to act on it is very compelling.

MR. WOOLSEY: Do you think Vanguard might actually have some behavior from the early 1980s that we could use?

MR. SHUMRAK: It could. That's part of it. Obviously, if companies don't do a particularly good job of starting to track individual-specific information, you're not stuck, but you'll be depending much more on the soft focus group type information. I've been around the marketing side a lot, and I think it is a heck of a lot better than nothing, but it's more directional. However, you can also expand on this through conjoint analysis. This may provide the strongest behavioral information by giving you a distinctive competitive advantage. For example, if three companies looked at the same product design with three different focus groups, the directional answers are going to be the same. You're not going to sell the wrong product, but your design or pricing structure may not be optimal or distinctive from the competition. You won't have a competitive advantage. If you have individual-specific models (even though you might be in generically similar markets with hot or cold money or selling through brokerage or agents), the ability to act on the exact "genetic fingerprint" of your best distributors and customers represents a sustainable competitive advantage. If nobody else has that, you can act on it time and time again.

MR. WOOLSEY: In our case, we're a good example of that. Is it a general purpose simulation engine or is the object-oriented compiler a synonym of that?

MR. DARLEY: Yes, it's a synonym.

MR. EDWARD M. HUIZINGA: Could you talk a little bit about available modeling software? Are there packages out there that would enable you to do some of this?

MR. DARLEY: In the context of insurance?

MR. HUIZINGA: I was thinking generally, but within insurance, too.

MR. SHUMRAK: Insurance-wise we don't know what some company is doing somewhere, but my sense is that this work that we've been prototyping is probably the first work in that area, and then I would defer to Vince to elaborate. Do some people package a basic agent-based algorithm that people use, if they understand all this?

MR. DARLEY: There is sort of one semi-general purpose tool kit called Swarm, but it's not well-suited for a lot of applications. I'd say the best solution is always pretty much to purposely build something, and that's basically what my company does.

MR. HUIZINGA: The second question is, how much modeling with lapse rates have you done? Have you uncovered any hitherto unsuspected properties?

MR. SHUMRAK: I would say we've validated the prototype. The next step would be to actually start to run it off in real data. I think we've convinced ourselves that if we took this and deployed this with real data from a company situation, we think we would discover some interesting and valuable insights regarding lapse behavior.

I know this was sort of an offbeat topic for the valuation actuary symposium, but if we can get somewhere on the behavior, it certainly would have an impact.