RECORD, Volume 26, No. 3*

Chicago Annual Meeting
October 15-18, 2000

Session 94PD
How is Behavioral Finance Behaving?

Track: Investment

Moderator: JOSEPHINE ELIZABETH MARKS
Panelists: DAVID L. NEAVEN
BORIS BRIZELI
Recorder: JOSEPHINE ELIZABETH MARKS

Summary: Behavioral finance incorporates psychological phenomena into our models of financial markets. While behavioral finance was originally deemed controversial, many respected financial economists are now working in the field. What progress is being made to integrate behavioral finance into models of efficient markets?

Ms. Josephine Elizabeth Marks: David Neaven is assistant vice president of pricing and product management at ManuLife. He's responsible for the fixed interest product lines and segregated funds. He's going to be talking about behavioral risks in segregated funds. Many of these funds have guarantees associated with them, so behavior becomes quite important. Segregated funds with guarantees are the Canadian equivalent of guaranteed minimum death benefits (GMDB) in the U.S., for those who aren't familiar with the Canadian terminology.

Mr. David L. Neaven: When I read the abstract for this session, I thought it would be interesting to present a case study on a product that I spent a lot of time working on that's very rich in options. I will look at client behavior and utilizations of these options after the product has been purchased, and what will affect the long-term profitability and risk profile of the product.

The product I'm going to be discussing today is commonly known as guaranteed investment fund (GIF). I'll start with an overview of that product and I'll discuss the embedded options and the risks within the product. I'll give you a bit of data on utilization experience and then briefly discuss how all of that links into model building. I know Boris will have some interesting comments on model building during his presentation.

Now, what is GIF? As Josephine mentioned, it's a segregated fund. Risks are quite similar to GMDB variable annuities. In its simplest form the GIF offers a 100% maturity benefit guarantee at a 10-year point. The contract doesn't actually terminate after 10 years. It continues on, but there's a top up at the 10-year point.

Also, it offers a 100% death benefit guarantee. There are some more complex versions of the product with annual ratchets and so on, but for the purpose of this discussion, I'll just focus on the basic 100% death benefit guarantee.

About 50 brand-name mutual funds are offered under the product. We're basically providing a wrap guarantee on these underlying mutual funds for companies such as Fidelity. The two obvious options within the contract are the reset option, where clients at their discretion can lock in their gains, and fund switching, where clients can switch funds anytime they want.

The reset option gets a lot of attention within the profession. It's obviously a risky option and very difficult to quantify. The one that doesn't get a lot of attention is the fund switching. I'll also discuss that and, hopefully, show you why that's also a very significant option.

I'll just spend a few minutes on the basic product. We charge a management expense ratio. This is common in the industry at about 2.25-3.25% per year. If someone deposits $1,000 during the year, we'll take off this fee but the full $1,000 is still guaranteed. If you compound that effect over a 10-year horizon, deducting 3% a year, it's a fairly rich guarantee.

The other very interesting and important product feature is that we guarantee the deposit level regardless of where it is allocated. If the client puts 10% into a money-market fund, 20% into a balanced fund, and the rest into an equity fund, we guarantee it in aggregate. If the equities drop, but the balanced fund goes up, it's the combination of the two that's guaranteed. Obviously, that is favorable from the writing company's perspective if there are benefits from diversification.

The GIF is the standard product being offered in Canada. Incidentally, it's a closed block of business. The particular contract for ManuLife has been closed for about two years now. Some companies are still selling it in Canada and we are expecting these to be closed soon enough as the regulators come out with fairly high capital requirements for this type of product which definitely makes it questionable to sell going forward.

Now for some information on demographics. For ManuLife (and this may not be typical of the industry), we would have an average age of about 58 by contract size and about 54 by count. The average contract size is about $47,000. The mix is nicely diversified with 5% bond or money market, 28% balanced fund, 29% Canadian equity, 18% U.S. equity, and 19% international equity. A heavy Canadian equity weighting is there because within Canada there's a foreign content limitation on registered business and the bulk of the business is registered, so there's a bias toward Canadian equities.

Now I'll get into a little discussion on the risks embedded in the option. The profit/loss distribution on this sort of product is an asymmetric distribution. If all goes well, the writing company stands to make a lot of money on this product. Even at the median of the distribution, it's still doing very well. All is well until you get into the tail of the distribution and this is where things, like behavior with the resets, can really swing things.

When we do the analysis, we're looking not only on the investment return path, but how client behavior can shift that graph. You would have a whole series of graphs based on what your clients or customers would do going forward, and it becomes quite complicated.

The client has the option to lock in the gains, to switch funds and to surrender. Surrenders are not normally viewed as an option, but we view it as an option. People can surrender when they like and they can surrender if they're in the money. Their guarantees are in the money or out of the money and it definitely affects that risk profile of who's surrendering.

Now, advisor influence is a huge question mark. These products are very complicated and we wonder if the people who buy them really understand what they've purchased. Some do but a lot of people are really dependent on their advisor to influence them in fund selection, fund switching, resets, and issues like that.

Another option, not an obvious one, is the option to become more aware of your options. Right now we would have three or four years of experience on this product. It's recently been drawing quite a bit of media attention, because the regulators are coming out with high requirements. Going forward people can become more aware of their options. It's not like a life insurance contract or GIC product. Once it's sold, people can do a lot of things. Of course, if they never reset the product or they never fund switch, because they don't know what they've bought, that's generally not in their favor. If they want to use what they're paying for and increase their awareness, they can do that and that can have quite an effect on the distribution.

What are the factors in this reset decisions? If somebody gives us $1,000 and the next day it spikes up to $1,500, they can call us up and we'll lock in the $1,500 to the point of maturity. What happens is that they reset the entire amount that they've given us. If one equity fund has done very well and another hasn't done so well and they reset, they have to reset the whole thing. They can't selectively reset pieces of it. Also, if they reset, it extends the term to maturity. There's a tradeoff there and it becomes a question. When is it good to reset? When is it not good to reset?

For example, if someone's had a contract for eight years and so they're two years away from the point at which we would be obligated to top them up and if they decide to reset that contract at that point in time, that pushes the top up out another 10 years. We've just collected fees for eight years and now we get to collect fees for another 10 years, so we're collecting fees for 18 years. When people are making that decision they have to keep this in mind.

We might see a lot of resets in the first year, because they're not giving up much in terms of when they'll actually benefit from it. Of course, the death benefit is still in place, but it's usually the maturity benefit they're looking at here. The age at which they reset is another issue. As I mentioned, the contract holders on average are about 57. There are some complications concerning the death benefit. Age is part of the evaluation equation.

Comprehension of the option is an issue as well. Once again, when we've done focus groups, a lot of people don't know they have this option. We describe it to them and they say that's a great option. Where do I buy that? Well, guess what, you've had it for several years now and thank you very much for not electing it.

The bulk of these reset elections, I would say in excess of 95%, come in through an advisor, though the advisor doesn't get paid for the election. If someone has given us $1,000, they might not know that their fund has gone up unless they're checking the unit values. But if the advisor is monitoring the client's portfolios they may advise them that it's a good time to lock in. Boris and I were having a discussion on advisors and potential clustering of these resets. Because you've got key advisors who have sold a lot of business, you might see some clustering of the resets by advisor groups.

Any experience we've had over the last three or four years has been along certain investment return paths. The markets have, generally, been quite favorable, so we've seen utilization of these options. Maybe along other paths this wouldn't have happened or in the future awareness could increase and there would be more utilization of these options.

What about reset experience by policy size? We have approximately 92,000 contracts with an average size of $46,000, for $4.2 billion of assets, so it's a good-sized block. Some people have never used their reset. Those are zero reset people. On a policy count basis, about 60% of the people have never used it, even though they've had it for three or four years.

Their average size is less than the overall average size, so there are a few things going on here. They have less money. Perhaps they are less well advised by their advisor. Maybe their average size is lower because their funds haven't done that well. Also, if you look at the people who've used it more often, their average size is increasing with increased use. There's an obvious relationship here as shown in Table 1.

TABLE 1
RESET EXPERIENCE BY SIZE

# Resets	# Contracts	Average Size
0	55,442	$39,145
1	18,469	50,129
2	10,586	58,630
3	4,817	72,419
4	1,747	90,624
5	517	115,155
6+	108	123,376
Total	91,686	$46,865

Now, this reset option is usually viewed as something very evil by regulators and the media. It's very risky, but most of the time it's very profitable for the issuing company, because most of the time things are going well. People are resetting. They're pushing out the option. It's really just in the tail of the distribution, where a few years before the block matures, suddenly, there's a downturn in the markets and people have just locked in immediately before that happened. Then things could turn quite ugly.

One other thing to mention is that this distribution is based on the business that we've had on the books. Now, as the business surrenders and people leave us, the ones who persist will be most likely the ones who value what they're paying for. If you look at the management expense ratios (MERs) this is not cheap. MERs are often 3-3.5% on equity funds with guarantees, so people are willing to pay for it and learning about it. We would hypothesize that over time we'll be left with a block of option utilizers.

Table 2 shows some data on reset gains. You can see that in 2000 so far the average gain that was locked in was 15%. You might think, that's not too exciting, but still if the person can lock in 15% growth, it's not too bad.

TABLE 2
RESET GAINS

% Increase in GV due to Reset			Average Reset Size
Jan	2000	12.9%	$7,198
Feb	2000	18.9%	7,650
Mar	2000	16.9%	10,126
Apr.	2000	17.2%	9,382
May	2000	14.7%	8,728
June	2000	12.1%	7,726
July	2000	12.8%	7,374
2000	YTD	15.0%	$8,261

Table 3 shows the utilization of this option by market value growth. In this particular month, about 33% of those resetting fell into the growth category of 101-110%. That was the growth of their market value to their guaranteed value since either inception or the last time they reset. Then there were also about 33% in the next category. You can see as the growth increases up to the 140-150% category, there are simply fewer clients in that mix, so they're accounting for a lower proportion of the resets.

TABLE 3
RESET TO APPRECIATION EXAMPLE

Growth	Reset %
0-101%	6.1%
101-110%	32.6%
110-120%	33.3%
120-130%	14.9%
130-140%	6.9%
140-150%	3.3%
>150%	2.9%

One very interesting thing was that we looked at utilization within these groups. If 20% of the book of business is in the 100-110% group, do they use this reset more than people at a higher growth or a lower growth? You would expect that the people with the highest growth would use it more. But for the month I looked at, that wasn't the case. In fact, it was the opposite that was happening. They can only use it twice a year, so maybe they're waiting another month.

We've also studied the correlation of resets to fund value changes and there's an obvious correlation when the markets are high. When the fund values are high, people are resetting. When they're down, they're not resetting. There's definitely an opportunity loss in that. There's a lag. This is similar to buying a stock. Nobody's resetting at the peak and people don't necessarily buy a stock at its lowest value. It's not optimal by any means, but there's an obvious correlation there as there should be.

Now, some very interesting things went on at the end of 1999 with Y2K coming up. As I mentioned, people can lock in their gains on this product and they can also move their money around. In December 1999 the phones were ringing off the hook in our customer service area. People were very busy, because we had a flood of resets coming in, as people wanted to lock in their gains. This was a tremendous product for people who had Y2K fears. We just locked it in and said you're not going to lose it, Manulife will guarantee that gain.

Now, of course, what it will also mean is that 10 years from December 1999 there are going to be a clustering of maturities coming up, so that also will be interesting. We saw about 13% of the block reset in that month which is a lot, about double the normal level.

Now, there's another time of reset spike in Canada. February is registered retirement savings plan time and that's when advisors are in contact with their clients and we normally see the resets happening then.

The other very interesting risk within this product is fund switching. We saw a lot of activity late in 1999, when people moved out of equities and into lower risk profile products. In 2000 they all moved back again. We're wondering if this is going to happen a lot. It's a little bit disconcerting. You're trying to manage risk and everybody's moving out of equities and back, but we've seen it return to normalized levels since then.

Now, you might think, well, what's a big shift and what does it really matter? But, for example, if just 3% of the fund value moves from equities into fixed income instrument, that's actually a 6% total change. You've now got 3% less equities and 3% more fixed income, for a total of 6%. When we evaluate such things as income at risk or various other capacity measures for the business, shifts of 6% are quite significant. Trying to predict what people will do in the future in modeling this business out, you get into issues of whether people will be switching a lot. If there are price changes, what will people do with their products? Will they move it around?

What we're foreseeing is that over time we'll have a block of people that use their options, they will be resetters to the point at which it's no longer sensible to them, because they just keep pushing it out forever. That's no good. They'll basically move into the most cost-effective funds, the ones where companies would be charging the least for the guarantee.

Now, surrenders are very important for the risk analysis on this product. Typically, we're seeing that segregated funds have about half the surrender rate of mutual funds. Normally, there's a spike at the end of the surrender charge period. We do have some older generations of this product and we can see what happens at the end of the surrender charge period. When the agent's compensation is tapering off, you see that the surrenders are going up.

What will people do if the performance is really poor? We have no idea. There's no experience. Will they say, "Well, I've got two years left to go, I'm going to hang on?" Or, "I'm now in the money by 20%. That's pretty good. I'm risk free. I'll just hang onto it for two years." Or, "if it's four years should I take my money out and just put it into a guaranteed certificate where I'm guaranteed to do better than that 100% that the company is guaranteeing to me?" This is key in our evaluation of the risk.

Also, how will the advisors influence them? Just going back for a minute to that group of people who have never reset, are they being advised? Have their advisors abandoned them? What kind of advice will they have if the funds go down? Will people even know that they have this wonderful guarantee? You know, they get a contract. Maybe they've read it, maybe they know. The larger contracts will definitely know, but you never know with the smaller ones. Certainly, for long-term prolonged equity market correction there is simply no data available.

I have a bit of information here on sales trends. There's obviously an annual sales spike in what we call peak season. That's at the registered retirement savings plan season. We see lumpiness in the money coming and we won't have a nice, perfect diversification of your maturity end dates. Ten years out you'll have concentrations in February, March, and then another concentration around October.

Sales are highest when funds are on the rise and lowest following the correction. If we look back at August 1998, when our funds were off about 15% or 20%, there was a long lag. The sales just dropped about 30% for four or five months after that and then slowly came back onboard. Obviously, from the company's perspective, that would be the best time to sell this product following a drop in the market. It will be guaranteeing less, but people don't buy low, sell high. They buy high.

Advisor influence in fund selection is another interesting one. Advisors tend to gravitate toward the hottest funds, the best funds. We'll see some clustering of the business in certain funds, which we also have to monitor, and you may have to cap capacity in certain funds.

Other things that can add risks to the company would be this concentration of maturity dates, where if the market suddenly spiked and we have a lot of resets happening, there's a concentration of dates. We have seen there is definitely a correlation between resets and the markets spiking.

Fund manager behavior is a very interesting one, another one of Boris's favorites, I know. This refers to the risk of the fund manager doing something that is counter to our best interest. Say, for example, the fund has been around and you're in your ninth year of the fund. Suddenly, a rookie manager is put on to manage that fund and the fund tanks for some reason. Well, the insurance company is guaranteeing that block of business. We have a lot more at risk than the mutual fund company. It's very tricky.

Another very interesting thing here is that we charge the clients most for this when their guarantees are most out of the money, so that's got to be a bad thing. That's very odd when you think about it. When they're hugely out of the money you keep charging the same spread on their fund to cover these guarantees. If it's 100 basis points and they're way out of the money, the dollar amount you're charging them becomes quite high.

There is model risk. We're getting into quite complex models here. I'll just spend a few minutes on model building. The Canadian Institute of Actuaries has recently set out some parameters and some rather tight guidelines for model building. At Manulife we were already doing these things, but we think it's helping some of the other companies along. But the one thing to keep in mind for a model is that, of course, they'll never be absolutely correct. You're getting into tremendously complex models.

For example, in our pricing area for this product, we have an actuary with a Ph.D. in mathematics and a professor at the University of Western Ontario doing this modeling. He has another post doctorate from MacMaster University in option-related work helping him out. The expertise going into these models and the complexity is quite tricky.

In terms of model building, rational behavior is always mentioned. I think we're all coming to realize now that people do what's in their best interests, but they can never be totally rational. They can't behave optimally. It's just not possible.

Just to make the model building even more complicated, within our basic model we'll have a Brownian motion stochastic investment return path generation model. Then we've got to incorporate into that model the contract holder behavior, for which we do the best we can, but it's fairly deterministic. You're trying to surrender people who are most out of the money and reset those people who are most in the money. But building this into the model, the stochastic investment return path generation can become quite tricky. Then there are things like uneven distribution of sales, how do you factor that into your model building? How do you factor advisor influence in? It's almost like we need two models, one for the high-end clients that are monitored closely and advised very well by their advisors and one for everybody else.

Then there are trigger levels. What will initiate a decision to reset a fund? We have to build these trigger levels into the models.

I'm going off the topic a little bit here, but there are also state dependent assumptions for volatility and fund switching. The last one is projecting increasing awareness of these options and utilization with time for the persisting blocks. We don't use past history as an indicator of what people are going to be doing five or six years from now when the block is downsizing, so building that into a model is a challenge, also.

As I mentioned, we almost need two models: one for those people using these options and one for the ones not using it. This one model has to capture all of these things, these investment return paths, client behavior, fund switching, and resets. A model like this takes a couple years to develop. It requires a lot of peer review and is very complex. You run into huge run time issues. If you're running 10,000 simulations and then building behavior into that, a little laptop isn't going to get us anywhere.

Ms. Marks: Boris Brizeli is a principal at Insource Limited, which is a consulting firm. He advises clients on product design and the development of investment risk management strategies.

Mr. Boris Brizeli: I'm going to talk about behavioral finance from the perspective of complex adaptive systems. First, I'm going to talk to you about what is a complex adaptive system. Then I'm going to mention very briefly some nonactuarial applications of these systems, some existing actuarial applications, and then some potential actuarial applications.

When I was about 80% through the building of this presentation, I realized that the audience is actually investment people. The slant I was taking was more toward pricing. However, I will make several remarks about use of these models in investment modeling.

First, let's summarize how actuaries typically approach modeling of behavior. The way they usually do it is top down. What I mean by top down is essentially saying we're going to take this whole system that we're looking at and we're going to look at it through the telescope. We're going to look at patterns of aggregate behavior.

The key thing we're concerned with is macro-dynamics of aggregate populations. We look at how the aggregate behaves. What are the usual tools that we use? We use trends. We use summation, because we say if we take a bunch of things that behave linearly when you sum them up, the sum behaves linearly. We use averages. This is the typical actuarial toolbox.

For example, when we model lapses, we say X% of the people exposed at Chem Labs will lapse at time T. For those resets that Dave has described, we can say X% of people who are eligible to reset will. That's a consequence of modeling things in the aggregate.

For some of the more sophisticated approaches, we do some things that are more complex. Some of the things actuaries do is they model optimal behavior. We look at a particular option, be it guaranteed minimum income benefit or a reset. If it's a financial option, we can price it. We can see what it means to optimally exercise this option. Let's also assume that most of the people don't really know that it's optimal, but some of the people do. We model the aggregate and say part of the aggregate behaves like this, another part one behaves like that.

Another approach is principal component analysis. Principal components are used. For example, one area is interest rate scenario generators. You look at the principal components, essentially, the key triggers for people to behave in a certain way. Dave mentioned some seasonal effects in sales. We have a time in Canada when it's the last day possible to get a tax deduction if you contribute to a tax-advantaged vehicle. You're going to get a spike of activity around there. That would be one principal component driving behavior.

However, principal components look at behavior of the aggregate. All you're really doing is parameterizing that behavior. You're saying there's a bunch of parameters that can describe the aggregate, but still working with the aggregate.

Another approach that actuaries use is to introduce market condition thresholds. Dave talked about triggers. He showed you several distributions that allude to that. There are some well-known models like single premium deferred annuity (SPDA) lapses, SPDA disintermediation, policy loan utilization, and mortgage-backed securities (MBS) prepayment. They all utilize some or all of these kinds of approaches.

What are the major features of the kinds of things we as actuaries do? The models of behavior we use are either static or deterministically dynamic. They're homogenous in that we assume that there's some kind of average behavior with which we can approximate the aggregate. Usually they're analytic models. That is, we seek the compression of information into a formula.

We assume linear aggregation. This is a key component here. One of the big implicit assumptions we make is that if you take a group of elements that behave a certain way, that once we create a sum of the individual behaviors, the aggregate can behave linearly. That's a very big assumption and often violated.

Let's talk about complex adaptive systems (CAS). This approach considers dynamical systems. What's a dynamical system? It is a system that evolves through time. Essentially, it's a time-indexed function. It's nonlinear. It looks at behavior of agents. What is an agent? Agent is a very abstract being and that can apply to either insurance agents or the economic agent I'm talking about.

This approach assumes people are different and it doesn't assume that I can aggregate everybody. What it assumes is that people behave in very different ways. It is a rule-based behavior. There are people involved and they interact. There's feedback in this dynamical system from time to time, determined by how the people interact among themselves.

The key feature is that it doesn't look at what the system is doing from the top. It's saying that the key properties of the system should emerge as a consequence of the behavior of the individual agents.

I'll give you an example. Look at any capital market. You can look at it from the top down and say the way capital markets behave is, let's say, Brownian motion. The other way to look at it is that it's a bunch of people using a bid/ask process. We can model the rules based on how the information is communicated through the system and based on which people make decisions to buy and sell. Then I can simulate trades between different people, and as I look at it from the top down it looks like Brownian motion. But that's an emergent property of the system. I never programmed into the system that's its Brownian motion. All I said is that there are people buying and selling things, but when I looked at it from the top, certain properties emerged.

The key feature of CAS as well is that it looks at things via simulation. It is not doing it through analytics. I'm contrasting this to actually putting a formula in. All you're really doing is saying that there is a bunch of agents. They have rules for behavior. These rules can evolve. They interact with each other and interesting properties come up.

The key question that CAS answers is the following. You have local interactions and decisions by different agents. How do they lead to emergent and self-organization phenomena? What I mean by self-organization is organization that just happens by itself and not by design of the individual agents. If you go join the refreshment break, you can see self-organization in action. There's no random distribution of people in a hall. You see a bunch of groups forming. That's an example of what self-organization is.

Why do we want to assimilate? Why is assimilation an important part of complex adaptive systems? First and foremost, is because some things you can't really perform experiments of. You cannot develop an experiment that tells you what will happen if we pollute this planet to see if the ecology survives or not. Another impossible experiment is to crash the market and see how the economy behaves. You really cannot do these things.

Simulation allows you to have a lab. Essentially, it's a behavioral lab where you can perform experiments without having to do any of the things I mentioned. We need to simulate to see how emergent phenomenon comes about, because we have no analytics. We have to actually let the things run in a computer, because we don't have an analytic formula. We don't have a function.

We need to simulate to come up with the nonlinear expression of the interaction. We cannot come up with the appropriate result just by summing. We need to simulate, because there's no linearity.

The key purpose of simulation is to allow experimentation. For example, you can have new regimes that are externally introduced or new policies. Two researchers from the Brookings Institute did a study on how population adapts to an externally introduced required retirement age. At one time the U.S. Social Security Administration introduced a mandatory retirement age and it studied how quickly the population would adapt and change the distribution of the retirement age. Social Security was able, with very simple dynamics and very simple rules of how people make decisions, to simulate and very realistically approximate the new retirement age patterns. You can well imagine that that would be something very useful for Social Security planning or for insurance companies.

Another place where it can allow experimentation is to test new interaction structures. Let me give you an idea. Look at dissemination of information pre-Internet and post-Internet. Those are two very different information structures and allow very different levels of interaction. If I introduce new possibilities for interactions what they will do? Is this something I can easily experiment with? Not really. But with CAS, which means outside the computer, I definitely can.

Let's talk a little bit more about the features of complex adaptive systems. What kind of agent behavior do we expect? This is where behavioral finance comes in. We assume that agents don't behave like optimizers. The big assumption in economics and of efficient markets is that everybody's being optimal all the time. It doesn't work like that. There's considerable evidence that people are adaptive. They really look to fit themselves to the environment that is to optimize short-term advantages and change their behavior locally. They don't look at some kind of infinite horizon and discount forever their cash flows. You know it would be very funny if people did. Everybody would be running around with a calculator.

When we talk about agent behavior, we say that it's rule-based. Does this mean that we're actually saying people have a bunch of rules they follow like robots? No. What we're really saying is that if we think of the agents in our system as adaptive and that they follow rule-based behavior, we can reproduce real world. What it really allows you to do is to very realistically simulate what goes on.

The nice thing about modeling agent behavior is that it's common sense. It's much easier to believe that people look at the stock market and say it's going up, maybe I should put some money in and there's been a bull market for a while. It's easier to believe that people do that than people run around with calculators constantly optimizing. It's a very common-sense approach.

Another part of agent behavior is the behavioral changes. People have these rules that tell them how to behave, but also they interact between themselves and with the environment. They want to optimize some kind of utility function. How an agent changes his behavior is based, first of all, on how his rules evolve. There is methodology for how rules change based on how the agent interacts with the environment and the other agents.

Another nice thing about agents' behaviors is that it has stochastic elements. None of it is deterministic. Nobody does something with 100% probability. There's stochasticity in almost every rule that people follow in a real world.

Let me give you some ideas of what I'm talking about, what I mean by agent, systems, interaction, and emerging property. Some of the systems you can look at are the immune system, ant colony, Internet companies, and supply chain.

What are the agents in the immune system? They are the antibodies and antigens. In the ant colony we have queen ants, worker ants, and soldier ants. This is one way we can define agents. It's probably also the most natural way. There's no other way to define agents in an ant colony. There are other ways to categorize agents in some of these systems.

We deal a lot with supply chains. Insurance products and liabilities are really one big supply chain. What are the interactions? Well, in the immune system antibodies and antigens talk to each via chemistry. There's a whole structure of biochemical interaction. Queen ants and worker ants actually communicate via pheromones. There's pheromone trails all around ant colonies and this is how they know who is doing what and who should be doing what.

What are the emergent properties? In the immune system creation of the antibodies allows a huge leverage effect. When you vaccinate somebody, that's a huge example of how a complex adaptive system works. You introduce a tiny change into the environment of the immune system and what does it do? All of a sudden, it becomes totally immune to a particular virus, whereas, before it wasn't. The actual change in the system was tiny. In fact, the main goal of complex adaptive systems is to be able to identify these leverage points. This is a very new area and I can tell you they're nowhere near being able to identify the key leverage points in a system.

The key researchers working on this kind of stuff are part of the Santa Fe Institute people in the Biochemical group. The key thing they're after is how do you identify what will make a huge difference.

Let's look at the ant colony example. What does an ant colony do well? They forage and build nests. Now, each one of them can do any of these things, but as a group they follow very simple rules.

When do you use complex adaptive systems and traditional type systems? Traditional is by no means obsolete. It's very useful. It served us extremely well for a long time. All I'm trying to propose here is there are situations where a complex adaptive system can give us a much better way of looking at what goes on out there.

You would use a traditional model if you have a problem that's easy to define. Dynamics don't change very quickly. What you're really after is optimization. A lot of actuarial problems do that. But you would use complex adaptive modeling where, first of all, the problem keeps changing. It's never the same. The solutions are only relative. You can say that this is better than that, but that's all. Optimization may not even be possible, but what you're after is something better.

Dynamics change very quickly and the focus is on adapting to the environment. If you have a problem, the environment and the interactions may be so unstable that you really cannot use the classical tools. The best you can do is tactically optimize. This would be one such area.

By the way, I should give credit where credit is due. Even though Bruce Sawhill wasn't able to speak at this session, some of the ideas in this presentation come from him.

Let me talk about how complex adaptive systems relate to other nonlinear models. They really incorporate several other nonlinear models, a lot of which are actually becoming more and more useful in actuarial science. For example, chaos theory and chaos modeling incorporate the key feature of chaos as initial condition sensitivity. You can't really ever do long-term prediction, because you can never be precise enough as you're looking forward. You'll never know initial conditions, so you have this initial condition sensitivity. You never really know enough to predict a future state of a dynamic system. That's what chaos theory has to do with it.

However, complex adaptive modeling allows you to be much higher dimensional. We can look at millions of agents. Chaos theory really has to do with the deterministic system and its sensitivities to initial conditions.

What do genetic algorithms do? Genetic algorithms essentially solve different types of optimization problems. They use algorithms that are derived from Darwinian-type adaptation. It allows you to solve a bunch of problems using these algorithms. Genetic algorithms are subsumed by CAS modeling. However, complex adaptive modeling is also much more robust. A lot of these genetic algorithms are not very stable.

What about agent-based models? Agent based-models are actual practical implementation. Agent-based models allow us to say here is the world I would like to recreate and see what kind of emergent properties can be produced. An agent-based model is really a computer technique more than anything. It's an implementation device if you want to model a particular complex adaptive system. It's a way of looking at a complex adaptive system and to be able to run the simulation. You endow the agents with rules of interactions between themselves and with the environment. Then you let the model run and see the kind of patterns that can emerge.

Here are some examples of the actual agent-based models that have been implemented. There exists a model simulating every single car in Albuquerque. Don't ask me why Albuquerque. There is another model simulating financial markets. Agent-based models exist to simulate firm formation. Crowd behavior, airline cargo routing, and retirement decisions, all sorts of agent-based models exist. Some of these implementations are very small implementations. Others are into tens of millions of dollars. The traffic simulation was a very big effort.

What are the advantages of agent-based models? It allows us to look at a population as a bunch of heterogeneous agents. We don't have to assume that the aggregate all behaves the same. The interactions between different agents are treated explicitly. The beauty of this is that this is a computational model, so the analytics don't really matter. You'll never run into analytical difficulties. You'll never run into an integral that you can't solve.

I have heterogeneous agents allowing diversity of behaviors. I can find out rules. How can I find out the rules? Do I just make a guess of them? Well, there are ways. It's really data mining, but there are some very creative methods that allow you to get at the rules that people use. This allows the model to be nonlinear.

You can calibrate agent-based models against history and that allows you to replicate what has happened. There's no magic in that (it's like a fortune teller telling you what has happened in your life), but it's an important feature. It confirms that the model is real. It also allows you to do a "what if" analysis and distributional analysis of outcomes and some quantification of emergent properties. All these are nice features of an agent-based model.

I'll give you some examples of agent-based models. One is a traveling salesman problem. It's not a real problem. No salesman really needs to go through 30 cities selling encyclopedias. However, routing of the information flow is a very real problem. When you have huge networks with many different nodes and you want to effectively flow whatever it is through these networks, it can be physical goods or information. It's very important. It can mean savings of millions and millions of dollars if you properly route it.

This problem suffers from a very big flaw in that when you put enough nodes and enough connections in there, it suffers from combinatorial explosions. It's computationally extremely difficult. Trying to do it with dynamic programming is not easy. But what you can do is you can use the ant kind of thinking. You set a bunch of "ants" into these networks and ants start locally optimizing their roads. Not only that, but they leave pheromone trains, so that information can be passed to other ants. Another model using ants is "Sugarscape." Here you have ants hunting for sugar in a landscape full of sugar hills. The bottom line is, it allows you to come up with an optimal solution to this traveling salesman problem.

What does this really tell you? This is really what economic behavior is all about. You endow people with the right rules and see what kind of emergent phenomena we get. We get things like conflict. We get tribe formation. We get wars. We get environmental evolution. The ecology of the whole environment can go through phases. We might see a huge swing in the resource. If you introduce trading rules, you can get clearing markets, which is a key phenomenon in the efficient market.

From looking at a bunch of ants finding sugar and spice, you get all these emerging phenomena. To me this is a little bit depressing. You know, human beings think we're all so special as we ruin the environment, but I can replicate that kind of behavior with a very simple program.

Here's some work with more of an investment flavor. A group at biochemists did some work for National Association of Securities Dealers Automatic Quotation system (NASDAQ). NASDAQ at one point wanted to change "tick" size, but they needed to know in advance what the impact would be on trading. They needed to simulate the outcome. This was an expensive model to develop. They used real features we observe in the real world. You know, people panicking and selling stuff, crashes, and people rushing to buy particular stock, like initial public offerings and things like that. All of these things were observed as part of that model, bottom up. The tick size change was successfully implemented.

Another application is credit assessment. It seems natural, wouldn't you think? Here we have a bunch of people applying for money. You'd think you'd want to recognize some individuality and not think a lot about the aggregate. What this makes me think is that underwriting is a natural for this. What kind of emergent phenomena do we get of risk selecting one way versus risk selecting another way? I can easily test different underwriting schemes, whereas, doing it in the real world is not necessarily affordable. I'll never know the kind of risk selection that I'll get.

Another application on the actuarial side was pricing of variable annuities. Ernst and Young, who I think are pioneers in using agent-based modeling in pricing, did this. I haven't seen anybody apply agent-based models prior to them, so kudos to them. I think it was a very good application.

What that model did is model customer and broker behavior and interaction. Essentially, it explicitly looked at how customers communicated among themselves and, more importantly, how brokers communicate with the customers and how that leads to options. As part of their model they even introduced methodologies to deal with lack of data. How do I get at the data at the rules of behavior? You could use focus groups, but there's a more ingenious methodology. It's conjoint analysis. Through asking people "compare and contrast" questions and "trade off" questions, you can get at what kind of rules people use to create their perceived value as opposed to the economic value of a particular feature.

The output of their model is lapse behavior and it was found to be qualitatively consistent with history. In new situations that model can be used to be predictive. For example, we do not know how variable annuities behave in prolonged bear markets. What happens? Nobody can say. This kind of model can help you get an answer.

With the modeling that we have done at Insource, we've observed some clustering of resets. Also, we observed some transfer clusters and some transfer apathy. That is the kind of cluttering we get because of the kind of interaction we introduced. There will always be people who will never utilize their options.

However, Dave said something interesting. He said people can behave optimally and I believe people actually can. It's all really a function of what information flow structure exists in the system. If an information flow drastically changes (imagine there are billboards everywhere), people can start behaving very optimally. We also observe diffusion of behavior.

What about potential actuarial applications? One thing you can do with this is product diffusion. This is nothing new. It's being done in other domains. This stuff is really used to analyze how a new product diffuses itself and how sales go.

You can model different levers in the product diffusion. There may be key agents that are connected to pretty much everybody. It's sort of like the theory of six degrees of separation of the actor, Kevin Bacon, from any other actor. You can model how that would impact the diffusion of your product. It can tell you how to concentrate your efforts. It can also let you model things like people imitating each other. You know, like your neighbor has the newest 4 X 4. It allows you to actually model these things. It allows you to test new distribution or products with a computer as opposed to introducing it and seeing what happens.

Can that help you analyze competitive dynamics? That is, you introduce a product, but there are other firms. Especially in commodity markets, everybody's trying to out price each other. You can actually model that. This would be actually a very simple implementation. In Canada we don't have that many companies. We could pretty much explicitly model every company and given that we know a lot of people in these companies, we would have a very good idea of rules.

Conjoint analysis gives more power to marketing in the micro-pricing algorithm. You can create different economic alternatives for prices and create these indifference pairs. That is, you create economically equivalent scenarios and let marketing pick the price. This kind of analysis can allow marketing to make much better decisions than guessing at what price they'll sell how much. It can actually allow them to model results.

There are other actuarial applications. You can model how people elect different resets. An interesting twist is that some options are not financial options. They are called real options. That is it's an option to do something, but it's not something you can actually approximate as a financial instrument. An example is a conversion option. How can you evaluate such an option? There's a methodology for a real option valuation that allows you to do that. It allows you to model things like option election.

Another thing that Dave alluded to is that the Canadian market will need to change price in segregated funds because of the new capital requirements. How can you test the impact of that price on the in force? Will people change the mix of their funds? This kind of stuff can easily allow you to test that.

There are a lot of applications in distribution. You introduce something into a new channel to see how a product diffuses itself through a particular channel, how the different compensation incentives relate to it.

Of course, there are the investment applications, such as MBS option values. I'm not aware of any attempts to do it based on an agent-based model of the prepayment option. It would be a natural.

What about hedging behavior? We know that some people manage their variable annuity guarantees through delta hedging. What you can do is come up with behavior-adjusted deltas. How do you adjust delta for behavior? You measure it through the agent-based models.

What about option trading and replication? We all know that Black-Scholes doesn't work. We've been told that it's just an approximation. Well, no wonder. It's based on very unrealistic assumptions. One thing you can do with agent-based models is reproduce some very realistic behaviors in how option traders and hedgers operate.

For example, one phenomenon observed in hedging options is what is called the liquidity holes. The way to think of liquidity holes is to imagine everybody rushes to trade. It's like people rushing to a fire exit, but they can't all fit in. You just can't execute all these trades fast enough. Too many people are trying to buy one thing. You can model things like liquidity holes with agent-based modeling. Is it easy to implement? For most of the stuff I talked about, a two- to four-person month can give you a decent start.

As far as lapses, resets, and repricing, if you want to model the stock market with an agent-based model, I don't even know if it's computationally worthwhile. Applications such as artificial stock markets tend to take a very long time to create. Some things are just computationally not possible. You may want to do them, but they aren't possible.

Mr. Howard M. Auerbach: Once you have a model that simulates well, what makes you think that it's actually the dynamics behind the system you're seeing? Suppose it does a good job of predicting and it has no market crashes. But then it turns out there are market crashes. What would make you see that there was more to the dynamics than you've already put in?

Mr. Brizeli: I think that what you're getting at here is a different issue. I think it's model risk. I may have a model that I'm parameterizing with a bunch of rules and interactions. It may well be that my parameterization may be good locally or over the history that I am analyzing and in some foreseeable future, but it doesn't capture another level of abstraction.

It may be that the rules are actually following some kind of genetic algorithm and you may not capture that in your model and nothing really tells you that. Model risk is huge in any of this. However, if I have a top-down model that I'm using for stochastic simulation and then I have a bottom-up model, independently created by someone different, and they validate each other, I've decreased my model risk.

Mr. Auerbach: It seems to me that you might easily find a way to simulate behavior, but there might be many other models that simulate the same behavior and they may behave differently. Could you tell us anything more about how you validate the model?

Mr. Brizeli: You're talking about the problem of calibrating the model. Keep challenging the rule calibration and it's only as good as how much data you have and the processes that you use to data-mine the data that you have. You can also look at feedback systems. What you can do is take the conclusions that you derived in your models and if you actually introduce it in real world that will allow you to do some rule refinement.