RECORD, Volume 26, No. 1*

Las Vegas Spring Meeting
May 22–24, 2000

Session 44L
Computer Models for Retirement Policy

Track: Pension

Moderator: CHRISTOPHER M. BONE
Panelist: JOSEPH M. ANDERSON‡
Recorder: CHRISTOPHER M. BONE

Summary: Policymakers often turn to elaborate computer models to assess the effect of proposed changes in Social Security or pension laws and regulations. The SOA has sponsored research into the most important models used by public policy analysts looking at retirement benefit issues. At this session, descriptions of these models and their methods are presented, including a review of the models' capabilities and limitations. The potential for use by actuaries in connection with these models is also discussed.

Mr. Christopher M. Bone: This session grew out of a project to look at the different types of retirement models that are used primarily by policymakers. It was an initiative from the Pension Research Committee, and we've had a distinguished researcher working on categorizing models for some period of time. I'd like to turn this session over to Joseph M. Anderson, who has a Ph.D. in Economics from Harvard University.

Dr. Joseph M. Anderson: I'm president of Capital Research Associates, which is a small economic research and consulting firm in Washington, D.C. We do work largely on human-resources-related expenditures, usually relating to public policy: large-scale database development concerning retirement income systems, retirement savings, and health-care finance and model development in those areas. Over the past few years, we have become involved in many of the pension reform efforts going on around the globe, particularly in former Soviet republics in Asia and Europe.

Right now, the actuarial profession is engaged in some of the most exciting and significant work globally, because it's part of privatization and the trend toward market economies, both in the former Soviet republics and in emerging market countries in general. Pension reform and reform of insurance systems and health-care finance are some of the most important issues. And, in most of those countries, particularly the countries that do not have market economies, there virtually is no actuarial science, or there wasn't, at least 10 years ago, and almost no actuarial profession. Actuaries from the western countries are playing an exciting and vital role there.

†Dr. Anderson, not a member of the sponsoring organizations, is president of Capital Research Associates in Chevy Chase, MD.

Today I'll talk about fairly advanced computer models that have been developed to analyze various aspects of retirement income policy in the U.S. and Canada. These are large empirical simulation models that have been implemented on computers because of the number of calculations required and the size of the databases. I'm going to discuss several specific models, the ones that I think are the most important, currently being used in the U.S., and one Canadian model, and I'll describe some of their specific characteristics.

As a matter of background, I want to describe a little bit about modeling in general. This talk is based on a study that we are doing for the SOA. I'll talk a little bit about the background of the study, its basis in retirement income policy issues, and then give an overview of some types of modeling approaches, along two dimensions in particular: kinds of alternative paradigms for modeling in general, and, second, the way uncertainty is treated in these various models. I'll talk about one particular type of model that is perhaps the most important, or the one most widely used currently in policy analysis, called microsimulation models. Then I'll describe eight of the most important models currently being developed and used.

For several years, leaders in a number of actuarial organizations have felt that actuaries should have a greater capability to use the analytical tools that you already have and to contribute to public policy debates concerning particular issues for which actuaries have expertise and a particular interest. In a sense this comes from a feeling that the federal government takes actions without understanding the full consequences and, particularly, the effects of those actions on the response of retirement benefit plan providers and retirement benefit plans.

I think that there was a sense that in the health-care-reform debate in 1994 actuaries may have missed opportunities to contribute their expertise and perspective to that ongoing debate. And the reason was lack of a central analytical capability. Currently, of course, Social Security reform and private-pension reform are being widely discussed, and actuaries may seek an opportunity or a framework in which they can bring together their expertise to contribute to that debate.

Here's a little background about the retirement-income-modeling project. There was a particular interest from a number of leaders in the profession whether actuaries and their various organizations should acquire or develop a retirement income policy model to provide a framework and a capability to provide the actuarial perspective on policy issues. The issue was whether a new model should be developed, and if so, what should be the characteristics of that model? And, importantly, how much would it cost? The first step in this consideration was to review existing models to see if some existed that could be used or tailored to be used by the actuarial profession to provide this framework.

A project oversight group was established with representatives of a number of interested organizations. The SOA Retirement Systems Research Committee, the AAA, and the Conference of Consulting Actuaries had input into the project.

The first question before the project was to develop a framework in which the issues that might be of concern could be addressed, emanating from the view that issues determine what is the appropriate model. The Project Oversight Group reviewed types of effects of potential policy initiatives on the pension system and the economy, grouping them into six major groups: employer pensions, effects on employees themselves, effects on retirees, effects on industries, and effects on the government’s fiscal situation and on the aggregate economy.

To organize the discussion of issues and the nexus between models and issues, we developed a policy matrix to systematize the examination of how various models address these policy issues. The rows of the policy matrix were what we might call the policy inputs—the laws, regulations, and system features that would be affected by specific policy measures with specific changes that might be proposed. The columns would be the aspects in each of these six overall issue areas that I identified previously that could be affected by specific policy inputs, identified in the rows. There was one matrix for each of these six areas.

What issues do policymakers and actuaries want to address? We listed potential policy inputs in categories.

- Tax Policy
  (a) Pertaining to pensions in particular
  (b) General tax
- Social Security
  (a) Retirement age
  (b) Benefit structure
  (c) Indexation
  (d) Payroll tax
  (e) Trust fund investment
  (f) Individual accounts
- Funding and Guarantees
  (a) Pension Benefit Guaranty Corporation (PBGC)
  (b) Premium and funding rules for defined-benefit plans
- Pension Regulation and Policy
  (a) Employee Retirement Income Security Act of 1974 (ERISA) and IRS
  (b) Employer plans
  (c) Pension and saving incentives/mandates

In addition to the policy matrixes, providing a crosswalk between specific policy proposals and various areas of the retirement income system and the economy that would be affected, we also looked at seven different policy benchmarks for each of the models analyzed. I'll briefly describe those hypothetical policy proposals. And we looked at each model to see whether or not, and how well, it would address those proposals.

Policy Benchmarks

1. An increase in the Social Security full-retirement age, or what used to be called normal-retirement age
2. Means testing of Social Security benefits
3. A hypothetical proposal for a mandatory minimum employer pension
4. Expansion of individual retirement account eligibility
5. The effects of a value-added tax
6. The construction industry benefit accrual rates
7. (Although it's not a policy issue) how each model could simulate and analyze the effects of alternative future macroeconomic scenarios

The project began with the premise that each model should be firmly rooted in recognition of those issues, which it is best suited to address. It comes from a recognition that hasn't always been widely shared in the past, that no model can address all issues. I think past failures and the disappointment with some of the early, very ambitious modeling approaches have arisen from attempts to develop or expand the model to address too many issues. We started with the notion that it is important to choose the most appropriate model for each issue, to select or develop the type of model most suitable for the policy issues to be addressed, given the time, data, and budget constraints.

Before we talk about some specific models, I think it would be useful to provide, by way of background and discussion, some types of modeling approaches. I'll talk in particular along two dimensions. First, what I call overall modeling paradigms, and second, overall general approaches. I thought it was useful to identify four paradigms in which most of the big economic models, and most of the models that we talked about, have some kind of economic content, and in which most, social policy models can fall. First there are the aggregative, or macromodels. These are primarily the large macroeconomic models.

The second category I call inter-industry models, and these emanate primarily from what I call input-output models. The work of an economist named Wassily Leontief, focused on the productive sector of the economy and the flow of resources among firms in the sector in forms of inputs and outputs. These types of models, this paradigm, have rather limited applications to retirement income policy, so I won't talk about it very much. In studies that we have done in the past, we have used straightforward inter-industry models. The largest currently in use was developed by the Bureau of Economic Analysis. Some of the private economic analysis firms, like McGraw Hill Data Resources Inc. (DRI) and Wharton Econometric Forecasting Associates (WEFA Primark) also have inter-industry models. We have used these models to analyze pension policy, because, as you know, pension and retirement income policies can affect labor costs. This was the primary analysis and then to see the effects through the economy, as inter-industry models look at the way that costs of certain inputs affect the price of outputs in different industrial sectors.

The third type of modeling approach I call cell-based or transition matrix models that focus on cells where whatever is being analyzed, ordinarily individuals and households, can be identified, dividing up a population over some well-defined characteristics, for example, age and sex groups, and then looking at what factors cause changes in the average behavior or average states of those cells, and how individuals move from one state to another over time. The fourth paradigm, or general approach, I'll call micro-analytic simulation.

The second dimension, along which I think it is useful to categorize models, is the way they treat uncertainty, in particular, whether they are deterministic or stochastic. All of the models deal with future outcomes, which of course are uncertain. But in a deterministic model, if the assumptions and parameters of the model are known, then the model outcome itself is determined. We may not know what it is until we run the model, but there is no uncertainty inherent in the model. Uncertainty can be represented in such models by simulating alternative scenarios by selecting assumptions that represent a range of possible outcomes that we think usefully spans uncertainty and doing alternative simulations. This is the approach that the U.S. Social Security Office of the Actuary uses. The modeling approach of the Office of the Actuary is a deterministic model. Once the assumptions are specified, the answer comes out. But uncertainty is treated by choosing three different scenarios, a low-, mid-, and high-cost, and hoping that that categorizes in a heuristic way the degree of uncertainty.

In stochastic models the model outcome itself is uncertain. Their randomness is inherently built into the model, so that there are various events that can take place, and the outcome is uncertain. For example, we'll get back to discussing micro-simulation models, where the specific events might include birth, death, and retirement. In other models, there could be uncertain environmental conditions, such as what future interest rates will be or labor market conditions. In these models, the outcome of these events themselves, in the process of running the model, are determined by stochastic processes—in general, by specifying probability distributions and drawing random numbers.

Returning to what I call the modeling paradigms, the major modeling approaches, whether aggregate or macro-models or model aggregates or averages, such as total consumer expenditures, or total private employer defined-benefit payments, are usually deterministic. The most well known examples, as I said, are the widely used macroeconomic forecasting models. These can be classified into short-term models and long-term models. Short-term macroeconomic models are usually what economists call income expenditure models. They focus on the determination of aggregate demand, or aggregate expenditures—the employment of the economy's resources, given the economy's capacity, which doesn't change very much in the short run and, therefore, we assume to be fixed.

This type of paradigm was first pioneered by an economist named Jan Tinbergen and Lawrence Klein in the U.S. Klein and a colleague, Arthur Goldberger, developed the Klein-Goldberger model, which was the first well-known U.S. macroeconomic model. Both Tinbergen and Klein, by the way, won Nobel prizes for their work. These models have been in widespread use for 25-30 years. The early developmental work by Tinbergen and Klein took place in the 1940s and the 1950s.

The Federal Reserve now maintains a large model, which is heavily used for monetary policy. DRI, McGraw Hill, and Standard & Poor's, it's all one company, has a large-scale, short-term, macroeconomic forecasting model, and the WEFA Group, now a part of Primark, has another commercially used macroeconomic model.

An alternative aggregative model is one we might call long-term models, which focus on the evolution of the economy's productive capacity in the long term and the determinants of potential output. These often have a macroeconomic growth framework that builds on the theoretical work of Robert Solow, who also won a Nobel Prize. Dale Jorgenson, an economist who has done a lot of empirical implementation of long-term growth modeling, applied to the U.S. economy and a number of economies around the world, developed the framework. Probably, the most full-blown application of long-term growth modeling to retirement income policy was supported by the National Institute of Aging in the development of a macroeconomic demographic model of the U.S. Retirement Income System, which we'll talk about later.

The second important type of modeling approach for retirement income policy is what I call cell-based, or transition matrix models. I think that most actuarial models would fall into this category. These models can be the deterministic or stochastic; an example of a deterministic, or what I call cell-based approach, are the modeling systems used by the U.S. Census Bureau and the Social Security Office of the Actuary to project population, using its own assumptions, dividing the population in this case into cells of a single year of age and both sexes and, in the case of the Census Bureau, two racial groups and two ethnic groups, altogether three race-ethnic mixes. The

Social Security Administration (SSA) divides the population only by age and gender, and then applies fertility rates to move those cells through time and apply mortality rates to move the survivors of those cells through time and fertility rates to determine the number of newborns, based on the population of women in each cell of childbearing age.

The Canadian Pension Fund Actuary has a similar model called ACTUCAN. But also there is stochastic implementation in our cell-based models, where the assumptions used to move the population of cells through time are themselves stochastic, and I think the best example of that—I'll talk about it later—is the Social Security Stochastic Simulation (SSASIM) Model. A similar stochastic simulation model of Social Security has also been developed by Mountain View Research, a firm in California, with heavy input from economists from the University of California.

The third important paradigm I'll call microeconomic simulation and this approach actually simulates the behavior of micro-units. It can focus on individuals, families, and households as the units of analysis, but it can also equally be used to apply to firms or even specific governmental agencies. These models can be divided into two classes, according to what has come to be referred to as the aging methodology.

Aging methodology is used to adjust the model's database to simulate the changes in these individuals, and consequently an entire population from one period to the next. And the two types of aging, generally, that are used are referred to as static aging and dynamic aging. In static aging, the model's database, if you will, is adjusted directly to correspond with some future period and some future set of conditions. This is done primarily by reweighting the records—and I'll explain a little more in a minute what I mean by that—and by directly imputing values onto the records, so that a database collected in 1996 could be adjusted to correspond to the U.S. population in 2005, by changing the weights on each record to correspond to the number of persons in 2005, and by changing the variables of interest, such as income and employment on each record to correspond to projections of aggregates in the future or target year.

Two prominent examples of microsimulation models with a static aging methodology are: (1) the Transfer Income Model (TRIM), which was first developed at the Urban Institute in the 1970s and updated since then. It continues to be used heavily for welfare reform and other government acts and transfer income analysis, And (2) a similar model the Micro-Analysis of Transfers to Households (MATH) Model, developed by Mathematical Policy Research. Another example of a static aging microsimulation model is the Household Income and Tax Simulation Model (HITSM), developed at ICF Inc., also to analyze government transfer income programs and tax programs. The Department of Health and Human Services in-house has also used heavily static microsimulation techniques.

An alternative approach to aging, to adjusting the microdatabase in a microsimulation model, is sometimes referred to as dynamic aging. In dynamic aging the actual events that determine the characteristics of a population – such as births and deaths, becoming disabled, getting married, finishing one more year of school, specific events — that altogether, determine how an individual and a family evolves, and consequently, when aggregated up, how the population evolves. These specific events are simulated (themselves) so that given the number of births and deaths, the size of the population can be aggregated, and given marriage rates, the overall marital status of the population and of each individual can be determined in each year, year after year, so that the evolution of the population through time can be simulated.

Essentially, dynamic microsimulation builds on a base-year database a series of simulated databases, perhaps through a historic period, and then simulated future cross-section databases, period-by-period, so that the overall evolution of the population is simulated. In addition, longitudinal data for each of the individuals or families in the database are created.

Dynamic microsimulation models are inherently stochastic. Almost all of the events or characteristics referred to in a dynamic model are discrete events—A woman either has a child or not; individuals, if they're unmarried, get married, or if they're married, become either widowed or divorced. Most, not all, but most of the events in dynamic microsimulation models are discrete events that are being simulated for each individual, so the simulation process itself involves uncertainty, because not every individual can be simulated to have the event. Dynamic models have an inherent stochastic element built in.

In addition, there are models that involve mixed, static, and dynamic aging, and, in fact, most models have some elements of both. And now, a new kind of hybrid microsimulation model has been developed, which is essentially a static microsimulation model, but that directly simulates whole processes or whole longitudinal characteristics of the individuals. This is a model being developed under the auspices of the SSA, called the Modeling Income in the Near Term (MINT) model, and I'll talk a little bit more about the MINT model after I describe the existing dynamic microsimulation models.

The Department of Labor currently has under development what I would call a kind of experimental cell-based dynamic microsimulation model, again, specifically for retirement pension analysis, called PENSEN, where the units used in the analysis in the model are not actual individuals from an actual database, but pseudo-individuals. The model itself does not aggregate up to represent an entire population, but rather types of individuals. I think it's fair to think of it as a dynamic microsimulation cell-based model.

There's a lot that's started to happen in the last two to three years in the area of modeling, and particularly microsimulation modeling after a hiatus, if you will, or a lull in model development, for about the past 10 years. There can be micro-simulation modeling that applies to units other than individuals and households, although those are the most common. An example of that is a model of pension plan sponsors that has been developed at the PBGC, called the Pension Insurance Modeling System (PIMS).

What are the basic features of microanalytic simulation? As I said, the key aspect is that it operates on or simulates the behavior of individual units, which can be, most commonly, households, families, or individuals, but also firms, pension plan sponsors, or even government agencies. These models are structurally similar to models used by actuarial consulting firms to analyze pension plans, but differ in important respects. Microsimulation initially was developed largely to apply to government transfer income policy and tax policy, initially the static microsimulation models. But when retirement income and Social Security policy issues became important, toward the end of the 1970s, it became obvious that static aging approaches were inadequate when looking at retirement income policy. This really was one of the impetuses that led to the development of dynamic microsimulation: because of the importance of simulating life courses and the events that occur in individuals throughout their lives and determining the effects of pension plans and the effects of policies to change pension plans and Social Security.

The first feature is the focus on individual units. The second key aspect is that the models, for the most part, operate on actual population sample databases, usually a representative sample based on an actual survey.

The third key feature is the basic operation of the model. As I said, you simulate the specific behavior of each individual in a survey database, and sometimes behaviors or events that occur in families, groups of individuals, or individuals grouped into families; such as: whether a woman of childbearing age will actually have a child; whether each person lives or dies during the period being simulated; whether marital status changes; whether individuals of school age stay in school one more year or graduate or return to school; whether a person gets a job or loses a job or experiences wage growth. So, a model would typically simulate about 30 vital events or demographic and economic states or events that can occur to an individual.

What has happened to that individual, grouped into a family, in a given period or a given year—are simulated, and at the end of the (annual) period, we have a new database consisting of numbers of individuals of various ages, various educational attainments, marital statuses, employment statuses, wage levels, and we can aggregate over that database and analyze it, as we can analyze an actual database. This is repeated throughout the simulation period of the model to develop a database at some target year. For example, if we want to simulate between here and 2020, we would simulate each year between now and 2020. Then we would have a simulated 2020 population of the U.S. and also a history of what happened to each individual family each year, from the base year, perhaps 1996, through 2020.

Given a base simulation of the population, then we can change the environment. If we have certain macroeconomic assumptions, we can assume more robust growth, and we can also change the policy, the environment in which the population is living; that is, Social Security rules, or private pension regulations, or the behavior of pension providers, and do an alternative simulation to see how the population would have evolved under an alternative set of policies, and then compare the two simulated databases to analyze the effects of the policies.

I mentioned very briefly in characterizing the models the two basic types, static aging and dynamic aging. I think it can be useful to describe exactly what that means. In general, we're using in microsimulation an actual survey database, such as a sample of the census, or very typically, one of the databases collected in the current population survey, which is essentially a survey of the population conducted by the Census Bureau each month, primarily to look at labor market conditions. It's a survey from which the monthly unemployment rate is calculated. Annually in March there's a detail that's called the Demographic Supplement that collects a lot of information on the past year's labor market activities and income from a large number of sources in the population. That will be a typical database that will be used for microsimulation. To be specific, in that current population database, there are about 60,000 households, so it's far fewer than the total population. To draw inferences about the whole population based on that database, each person in the sample database and each household has a weight. So the database can be aggregated adding up the weights, which are generally stratified over the characteristics of interest in the current population survey — that would be by age, sex, and race, generally for about 30 states — and ethnicity, in about six different groups. For those characteristics, and a few others, we can aggregate the population and get totals from the sample and characteristics for each of those demographic groups. So the weight on each person would be the number of people in the total population that that individual in the sample survey represents.

We can change the total population and the population of each individual simply by changing that weight. And that's what we do in static microsimulation. If an aggregate population forecasting model forecasts that the number of white women age 35 will increase by 25%, between our base year of 1990 and our future year of 2005, then we can increase the weight on each white woman age 35 in the sample by 25%, and it will align to that aggregate proportion.

The other way we would age or adjust in a static aging microsimulation model is by actually imputing changes to the characteristics on each record. If we think that the unemployment rate is going to increase by 10%, we can randomly select employed individuals in the base year to become unemployed in the target year, so that the employment rate when we aggregate it up is increased by 10%.

If we believe that wages of a specific demographic group will be increased by 15%, we can increase the wages of each person in that group by 15%. If we believe that a particular policy will change employer pension coverage of a specific demographic group— let's say a change in vesting requirements will increase the proportion of persons with a certain employment experience, and a certain age by a certain percentage — then we can modify, if we have data on the pension coverage and vesting in the base-year database, that characteristic to correspond to the estimated effects of the policy to create a new database in which we are simulating what the outcomes would be for individuals that live under that new policy regime. We can modify each record directly, and this has come to be known as static aging. Static aging does not really simulate the actual behavior through time that has led to the change in these data, the change in the size of the particular demographic group.

In the second aging methodology called dynamic aging, rather than directly changing the weights, the numbers of people each record represents or their characteristics, we would simulate the actual events that occur— births, deaths, marriage, economic events, and the behaviors of each individual, period by period — so that we create life histories for each individual and a cross-section database for the whole population for each year. At the end of the operation of a dynamic microsimulation model, in addition to a series of cross-section databases, we also have a longitudinal database, one that follows the life course for each individual in the database.

And then, as I've noted, there are mixed models that often simulate the key events that we are focusing on; such as the demographic events of birth and death, and key economic events, such as whether a person is employed, or whether individuals are employed by employers who offer a pension plan, for example. And then other characteristics may be imputed directly for particular intervals.

Let's say we're simulating a dynamic microsimulation model over a 20-year period, but we don't really have behavioral studies that give us a way of modeling IRA accumulations, so that we might impute the IRA balances, based on the characteristics that we've simulated for the individuals in the population every five years. Or if we want to analyze the population at a future target year, such as 2010, we might simply impute, in this example, IRA balances in 2010, and we would use a static methodology to impute the IRA balances in what was otherwise perhaps a dynamic model. Many models are mixed in this sense.

I noted there's a new type of model that has been developed over the last couple of years, although this approach has been talked about for a number of years, the MINT model, which is not really a dynamic microsimulation model because it doesn't simulate year-by-year events. But in a static sense it imputes behavioral patterns and roughly imputes lifetime earnings patterns. It doesn't simulate the earnings based on demographic and economic characteristics and the individual in the environment year-by-year — but assigns individuals to one of usually 27 different lifetime earnings patterns or earnings paradigms that have been developed by analyzing past-earnings histories, and assigns them econometrically, based on the characteristics of those individuals. It also will assign to the individuals in the base-year database the year in which they die. It assigns to them an age of death, a retirement age, and various retirement savings and IRA accumulations at that retirement age. It's an intermediate approach between static and dynamic microsimulation.

Let's talk about the stochastic features of dynamic models. I think there are two key features of dynamic microsimulation models: event probabilities, and then a way of assigning events, all of which is often referred to as Monte Carlo simulation. The actual aspects of the individuals in the database that are being modeled; that is, that are being solved for, usually using statistical equations or econometric equations —the actual variables that are solved for ordinarily are event probabilities, so that for a woman of childbearing age, the probability that she will have a child in a given year could be solved for by an equation with reference to her age, race, ethnicity, family income, whether she is in the labor force, her educational attainment, and a number of other variables.

For a 35-year-old woman, the solution to the equation may be 0.001, the probability that she would have a child in that year. And for every woman in the childbearing ages in the model, a probability would be assigned, and then drawing a random number, the occurrence or nonoccurrence of the actual event would be assigned to the woman. And this would be similarly done for mortality. The probability of death can be assigned based on the age, gender, race, income, educational attainment, and a number of other causal variables.

For each individual the actual event is assigned, based on random numbers. Then that individual probability is solved for each person, and sometimes for each family in a large sample, to determine the characteristics of all the individuals in the sample. This could be done numerous times for each individual in the sample to reduce the variance, and essentially if the process is replicated, it's similar to using a larger sample. And then we would aggregate over all the individuals in the sample, to estimate the characteristics of the population.

Two other features of dynamic microsimulation models that are important are alignment and controlling methodology and variance reduction. The dynamic microsimulation models are being used to simulate events for each individual, and if we had perfect research, we could probably build up aggregates with a great deal of confidence, based on this microeconomic simulation, or microanalytic simulation alone. But our research, in general, is not that good. Typically, dynamic microsimulation models are actually aligned to external assumptions or projections of aggregate outcomes.

Let me explain what that means. We're forecasting the future population by forecasting the mortality of each individual and the fertility experience of women of childbearing age and immigration. In concept, we could forecast the entire population that way. Typically, though, these models are aligned to other forecasts of the population, perhaps those done by the Census Bureau, the Office of the Actuary, or the Social Security System. After an initial population is simulated by forecasting whether each individual lives or dies, the simulated population is compared to the alignment population; the external assumptions or the external projections and the probabilities, not the outcomes, for each individual are adjusted, and a second iteration or second simulation of the model is undertaken to adjust the aggregate outcomes, so that they come closer or match the aggregate alignment totals. This is usually done for population by the cells for which we have an accurate population forecast: age, sex, and race, usually. We can do this to hit a target for aggregate unemployment by adjusting the probability that each individual is employed, and then reassigning a job status for each individual.

Many dynamic microsimulation models are often used to study the effects on individuals and the distributional effects, and the specific policy interactions of policy changes within a set of assumptions and forecasts established by another modeling system. For example, the SSA may look at alternative Social Security policies and use a microsimulation model to do so, but it would want that model to be consistent with the projections of the population and the projections of the total expenditures and benefit payments of the Social Security System, done by the Office of the Actuary, so that the microsimulation model would be controlled or aligned to correspond to the Office of the Actuary's projections.

Similarly, Canada has a dynamic microsimulation model, DYNACAN, that simulates the individual characteristics of the Canadian population, but it is always aligned and controlled to the actuarial projection model for the Canada Pension Plan, ACTUCAN. To do so, the rates and the event probabilities, which are actually solved for by the equations in the model, are adjusted after an initial solution for these events and probabilities, in order to come closer to the external targets. It is important to recognize that this is what is being adjusted, not the events themselves. The events are still simulated stochastically, so that, in general, the actual outcomes of the dynamic microsimulation model do not ever align precisely with the external controls, but the solution for each equation is adjusted so that the expected values are aligned to the external projections or the external assumptions.

For example, if a population is simulated over time, and it varies from the external population, then the fertility rates for each woman can be adjusted, and a second simulation can be done. In general, the population totals would be expected to move closer to the aggregate controls.

The second key characteristic of dynamic microsimulation models in this regard is what is sometimes referred to as variance reduction. As we said, these models are inherently stochastic, so that the actual outcome, the actual effect of any policy proposal on the population, is inherently uncertain. But, we want to take steps to reduce or control the expected dispersion around the expected values, so that when we are comparing two different environments, or the effects of two different policies, we can have a high degree of confidence that the differences in the population that we observe under Policy Regime A, versus Policy Regime B, are due to the different policies, and not to the simple random, stochastic nature of the model itself. Various techniques are used to reduce the variation in each event, or to control the amount of variation as the model is simulating over the population, so that the expected value of the event in the given year corresponds to the expected value of the control population.

The basic structure of dynamic microsimulation models starts with an input database, which is an actual sample population, and those data are operated on by what I call the core simulation model. Sometimes these models are modulated, so there will be two or three modules. I can talk about that when we talk about some specific models, so that a record will be selected, with the characteristics of an individual. And then, if it's a woman, the equation for fertility, mortality, employment, educational change, marital status could be solved in a module. And then each of those events would be simulated, using random numbers, and that would be done for all the variables on the records and all the records in the database. With inputs of user commands and assumptions — and the user commands could also specify which particular events are being simulated — and the alignment factors and data, the projections, for example, of an external population forecasting model or macroeconomic model as inputs to the core simulation model— that model produces an output database, which would be the events that occurred to each individual, that record for the year. Then it would create a cross-sectional database for the population at the end of the first year. And that could be done for a number of specific events and behaviors: birth, death, marital status, earnings, disability typically, educational attainment, and employment change, whether the person is working for an employer who offers a pension plan, whether the individual retires — so a database is created for the end of the first year. Then that becomes the input database for the second year and is run through the core simulation model and solved to create the output database for the second year, and similarly for each year in the forecasting period until the end of the simulation horizon.

The various report-producing packages would be used to analyze each of those databases or selected databases, and also to sort them and to create longitudinal databases, life histories for each individual. And then reports, summary tables, and graphs are produced.

These dynamic microsimulation models have two key requirements. The first requirement is very severe data requirements. We have to have generally a large-scale database with all of the attributes in which we are interested. If we are interested in retirement income policy, we generally want a number of demographic characteristics — information on labor market activity, and some characteristics about the employers, generally firm size, industry, etc., — and then some characteristics about the pension plan offerings of the employers, including the characteristics of pension plans themselves, including participation requirements, vesting, and benefit formulas.

An actual single database with all of that information is a pretty severe requirement. Usually a single database doesn't exist that has all of the information that we need. To create the database several different databases are combined or merged, either directly — where we have information in more than one database on the same person that we can link using some ID number—or more commonly, statistically—we can impute data, to create the large database by using data on similar, but not the same, individuals in more than one database.

The second severe requirement is the computational requirements. We're storing, as you can imagine, a very large amount of data, and generally inputting and outputting a lot of data, so that to make modeling simulation feasible, computers with large storage capacities and rapid input-output capabilities are required. Ten to 20 years ago, those computational requirements were severe limitations on the application of microsimulation. Dramatic changes in computer capabilities and computer costs over the past 20 years have significantly changed the microsimulation environment and costs. Twenty years ago, a huge part of the budget for doing a study using microsimulation models would be for computer time, and it would cost sometimes $5,000–10,000 for one run on a huge mainframe computer. As you can imagine, people were very careful, and the number of simulations that could be performed was quite limited. Of course, that's not true anymore. Very large simulations can be done now on microcomputers, with essentially zero incremental costs. There's been a huge change in the computer environment. And, only recently is the analytical technology of microsimulation catching up to the change in computer capabilities.

Briefly, what are the key strengths of microsimulation models? The first key strength is the ability to analyze individual behavior, so that we can analyze policies that affect individuals differently, depending upon their specific characteristics. Second, we can do distributional analysis. In aggregative models we're talking about the whole population. Or if they're aggregative models with a cell-based structure, we may be talking about the population of males of a given age group, say, 40–45, 46–50; but in general, we're assimilating only the average characteristics, the average income, or the average rates of pension plan coverage, for individuals within these cells. We can't really do distributional analysis. We can't say how low-income versus high-income persons are affected, or how many people are moved out of poverty in old age by a particular change in Social Security rules. That can be done with microsimulation models that simulate individuals.

Third, we can analyze interactions among programs, which are important for most federal transfer income programs and are certainly important for retirement income policy. Social Security changes are going to have effects on the employer pensions, and regulations of employer pensions are going to change retirement ages. These interactions among programs, and in particular interactions among program features and individual characteristics, can be analyzed using microsimulation models. For example, we can look at how a change in Social Security rules may affect women in old age who have different marital histories during their careers.
And fourth, we can create longitudinal histories: life courses for individuals, lifetime employment experience, lifetime earnings, periods of disability, marital histories, and in particular, pension accruals for which we need to analyze the effects of different policies, actual employment experience and pension coverage, and accruals that an individual may experience in each year of his or her career.

Important limitations to microsimulation models include, first, the data requirements. Because large, accurate and complete databases are required, that imposes a heavy cost and limits the environments in which dynamic-simulation models can be undertaken. As a result, models so far have pretty exclusively been developed in North America, the U.S., Canada, and the countries of Western Europe — countries that have a large number of survey databases and a lot of data about the populations.

A second limitation is the models' complexity, which makes the operations of models difficult to communicate. In general this makes the outcomes of the models nontransparent, which is, a severe drawback. In any model, as you know, there will inevitably be bugs, and most of us discover bugs when the model produces implausible results. But truly implausible results produced by a microsimulation model indicate that there is certainly something wrong in the model, but the results are the effect, as in the real world, of the interaction of so many events that it is often very difficult to track down the validity of each operation.

More than that, even when the model is performing correctly, it's not always obvious how a particular result or a particular simulated result has been achieved, again because of the complexity of the interactions being simulated. It's the mirror image of the strength of the model that is actually seeking to replicate a lot of real-world processes, which are, in fact, very complex.

A third important limitation is that because of the size and complexity, all existing current microsimulation models have a very incomplete representation of most of the social and economic systems that they are actually modeling. Most models represent only one side of the labor market, for example, the supply side, the household sector that is providing workers, not the demand side, because in most household microsimulation models, employer behavior is not represented endogenously. They simulate only one side of the capital market, again the supply side: savings but not the demand for savings, the interaction of which should determine rates of return in some kind of general equilibrium framework, for example — only one side of the health-care market.

Microsimulation models were applied heavily during the health-care-reform debate, but primarily focusing on the reactions of consumers to various changes that were proposed in discussing health-care reform, and in particular, only one side of the employee benefit system: again, the reaction of individuals, or the effects on individuals of regulations that will change the characteristics or the behavior of pension providers. These models don't really simulate that behavior, so they don't simulate how providers may react to regulations or respond to the reactions of their employees.

Similarly, microsimulation models have very limited feedbacks, so they may simulate a change in wages and a change in employment conditions, but they don't simulate the way a change in wages will affect the labor market activity of the individuals themselves. And, as I said, they don't simulate how a change in employees' labor market behavior, for example, choice of retirement age, in response to changes in pension policies might in turn affect the pension offerings of employers themselves. Because of their incomplete representations, for the most part, it's difficult to simulate these feedbacks endogenously in the model which can really be crucial in understanding the effects of a lot of the types of policies that these models are designed to analyze, in particular, pension and Social Security retirement income policy.

Let's talk about some of the specific characteristics of the most important dynamic microsimulation models; and the most widely used microsimulation models, in the areas of retirement income, retirement policy, and retirement income analysis are the models DYNASIM, PRISM, CORSIM, DYNACAN, and MINT.

We'll start with DYNASIM, because it is the grandmother of microsimulation models. It stands for Dynamic Simulation of Income Model. It was the first large-scale, microanalytic-simulation model, probably the first dynamic microsimulation model that was implemented anywhere. It was developed at the Urban Institute, roughly in the period 1969–75 under the auspices of an economist named Guy Orcott, who had essentially invented dynamic microsimulation, starting in the late 1950s in a series of conceptual papers. Then he actually gave birth to his ideas at the Urban Institute with heavy sponsorship from the Office of Economic Opportunity and the Department of Health, at that time the Department of Health and Education and Welfare. It was revised to include considerable employee pension capabilities in the late 1970s, and between 1979 and 1989, and then has been revised again in 1999.

The original DYNASIM model had a base-year data file comprised of a March 1973 current population survey, which provides a lot of information on individuals and families, considering their assumed demographic characteristics, economic characteristics and income, and particular labor market activity. This database was matched with their actual Social Security earnings records, going back to 1951, so that altogether the database consisted of 60,000 persons with current information from 1973, and their earnings, according to their Social Security records, going back to 1951.

In updating the model, which was completed in 1999, the database has been updated and now consists of a combination of four panels from the Survey and Income Program Participation, the 1990–93 panels. The Survey and Income Program Participation is what I call a semilongitudinal survey. Each 1990 panel is named for the first year that the households in the survey were in the survey, and they were interviewed at four-month intervals, usually eight or nine times, over a period of about three years. In each interview considerable information about demographic and economic characteristics is collected. And then for many of the interviews, there is a topical module on particular information, one of which is employer-pension plan coverage and participation and vesting. That provides a quite comprehensive, although smaller, database for those individuals.

Four of those panels have been combined to create the database for the updated DYNASIM model. But they have earnings histories that are imputed, not the actual earnings histories of these individuals, but past histories imputed using yet another longitudinal database called the Panel Study of Income Dynamics, which is a survey that began in 1968 with 5,000 households and has followed those households, reinterviewing them and all of the households that are subsidiary households, that is, that have broken off because of children growing up and leaving home, or divorce, for every year since 1968. The most recent data available for the development of DYNASIM were through 1992. They also used imputed earnings histories from Social Security earnings records for 1951–67.

The reason why, in the current update and in other microsimulation models, imputed data are used is because the Census Bureau no longer makes publicly available the exact match to Social Security earnings records. Even though the identity of each individual in these databases is carefully protected, to protect confidentiality, the Census Bureau has become hypersensitive to any possible release of information on particular individuals, so it no longer makes available the actual matched data with earnings records to the public. This is one of the things that held up the advances in microsimulation for a lot of models that we've developed in the 1980s until recent years—these rather crucial data, particularly for retirement income analysis, were no longer available. The Urban Institute has solved this problem by imputing these past histories.

The DYNASIM model of employer pension plans assigns individuals to one of five prototypical types of plans. There are four defined-benefit plans and one defined-contribution plan. And of the four defined-benefit plans, two are single employers and two are multi-employer plans, so that no actual plans are used in the database. This is based on 1979 and 1983 data and really reflects the original pension model developed for DYNASIM in the early 1980s. This was a strong component of the original model, but I would say it is now out-of-date, and I think the Urban Institute would agree, and it would be high on their priorities to update this part of the model.

As you all know, better than most, there has been a dramatic change in the nature of pension plan offerings, so that the very simplified structure used in the original DYNASIM pension model, is inadequate for current analysis. The model includes simulations of the Social Security System, employer pensions in the simplified manner that I have described, IRAs, and supplemental security income. It is a fairly complete depiction of most of the important sources of retirement income. It models earnings, but no retirement savings other than IRAs.

The original model was at the time the most sophisticated retirement decision model. For each individual it simulated a variable referred to as Social Security wealth and another variable, pension wealth, which is equal to the present value of future expected Social Security or pension benefits less contributions. They did that for each individual and each year, and they used this variable as well as other economic and demographic characteristics of the individual in a comprehensive two-stage model of the retirement decision. The individual in this model was essentially a perfect, rational, economic man and woman, if you will, and was evaluating the effect of working one more year or retiring on his or her current and future welfare as summarized in the Pension wealth and Social Security wealth variables. This is, I think, a powerful modeling approach, but this also needs to be updated.

DYNASIM has very limited feedbacks, and, in particular, the modular structure of DYNASIM that carries over into the current revision imposes inherent limitations in the simulation of feedbacks. In DYNASIM, a modular structure was developed in the original model, primarily to save on computer time. This was developed in the early 1970s, so that in one module most of the demographic and labor market activity variables are simulated, and after the population was simulated in that module, a second module operates on that longitudinal database. It simulates particular job characteristics, including pension coverage and retirement decisions, so that pension and Social Security policies that primarily would be implemented in the second module, have no way of directly, within the model, feeding back on overall employment and labor market activity and demographic behavior, which is simulated in the first module.

Direct changes and assumptions can be made, based on one simulation, to resimulate both modules to mechanically represent these feedbacks, but the model has limited capabilities to endogenously represent such feedbacks. DYNASIM has no representation of pension plans themselves and no representation of plan sponsors or the behavior of plan sponsors. The original documentation of DYNASIM was quite good and complete, but it is now very out-of-date. New documentation is under development, but it is not available yet.

A second interesting and widely used dynamic microsimulation model, or retirement income model, is the Pension and Retirement Income Simulation Model (PRISM). It was initially developed to analyze retirement income policy for the President's Commission on Pension Policy in 1980 and 1981. The group at ICF that developed that model through various corporate evolutions is now at a consulting firm called the Lewin Group, where the PRISM model is still maintained.

The original base-year data file for PRISM was a powerful matched database that was created by matching the March 1978 and March 1979 CPS, between which there is a 50% overlap, that provided employment changes and demographic changes for the individuals over those two years, with the main 1979 employer pension supplement, that provided detailed information on employer pension coverage and participation. And, in turn, the persons in that matched database were matched with their actual Social Security earnings records to provide a database of 28,000 adults with extensive demographic and economic pension and earnings history information.

That database has been updated recently, to use the 1993–94 March CPS, to match with SSA data, but that particular implementation of the model can be used only with severe restrictions because of this change in policy at the Census Bureau to protect confidentiality, so that it can be used only under very limited circumstances for general application. This database has been augmented with extensive information from other databases on institutionalized persons. The March CPS collects only data about the noninstitutional population, and the PRISM model is used heavily to analyze long-term-care policy, for which the characteristics of the institutional population are important, so that considerable data on institutional persons have been added to that database.

The original PRISM model had a secondary database of employer pension plans, which included 325 actual plan sponsors and 475 separate plans, based on data collected by the Department of Labor. It had many, large pension sponsors and numerous pension plans, so it had detailed information about primary, secondary, and supplementary plans for a number of sponsors, along with detailed characteristics about those sponsors, including industry, firm size, and specific characteristics from 1984. That has been updated to use data from the Bureau of Labor Statistics on small, medium, and large government employer plans. This is a key difference between DYNASIM and other microsimulation models that generally use prototypical pension plans: the use of actual data on actual plan sponsors, and actual pension plans.

That model also simulates Social Security, employer pensions, IRAs, and SSI and has extensive retirement income assets information and, in particular, a detailed module on long-term-care financing and disability, and this has become one of the heaviest uses of the PRISM model. Much of the support for the updating and development of the PRISM model has come from interest within the government on long-term-care policy.

The PRISM model also has a modular structure that limits its feedback capability. Although it has extensive information about the plan sponsors, it also does not have any plan sponsor behavior, so it also does not simulate the effects of changes in pension policy from the offerings of plan sponsors themselves. Like the DYNASIM model, its documentation needs to be updated.

One model that is emerging as quite important for retirement income policy is a model called CORSIM. It's essentially the Cornell microsimulation model. It's the first one that was developed exclusively to be operated on PCs, and it has a single modular structure, which gives it a greater capability of simulating feedbacks. The key aspect of this model is that it has been adopted by the SSA to experience extensive augmentation development, both within the SSA, including extensive input from the Office of the Actuary, and at Cornell, where the model is maintained. So it has now a very detailed Social Security module, and one of the plans for this model is also to develop a detailed model of private pension behavior.

CORSIM serves as a basic framework for a similar dynamic microsimulation model developed in Canada called DYNACAN, which simply took the CORSIM modeling engine and essentially piece-by-piece replaced the specific behavior modules, for example, fertility, mortality, employment behavior, with modules developed using Canadian data. They refer to the Canadianization of CORSIM—so they have now developed a Canadian dynamic microsimulation model based on this framework. And there is extensive interaction between the Dynacan Development Group and CORSIM. DYNACAN is developed closely in conjunction with the actuarial model of the Canadian pension plan; the two are run always in parallel, and DYNACAN is always controlled to the assumptions of ACTUCAN.

The other important models that apply to retirement income policy include the SSA dynamic simulation module, which essentially is a stochastic rendition of the modeling system used by the Social Security Office of the Actuary. Instead of assuming future values for the key demographic and economic environmental variables, as the Office of the Actuary does, in developing its low-, medium-, and high-cost projections, in the SSASIM model, each of 13 key demographic and economic variables, such as fertility, mortality, disability, interest rates, and wage growth, are specified according to a probability distribution, estimated looking at historic data, and then for a particular run of the model, by using random numbers, the actual fertility experience of the population, mortality, and economic environment are simulated. The model then would be run a large number of times, typically 1,000 times, to generate, rather than a single variable, such as the Office of the Actuary's mid-range projection, a probability distribution of outcomes, so that we can actually see, based on the underlying probability distribution of each event, what the likely experience of the trust funds would be in future years.

Finally, a very interesting microsimulation model of pension plans has been developed at the PBGC, where the key economic, environmental, firm, and pension plan characteristics, first of the defined-benefit economic environment and then of defined-benefit plan providers, are simulated for about 400 plan providers that represent one-half of the defined-benefit plan liabilities in the country. The objective in the model is to simulate stochastically the probable future outcomes for PBGC liabilities.

This has been an overview of some of the major models that are currently being used to analyze pension and retirement income policy. There has been a rebirth of interest, like in some of the revisions of DYNACAN and the updating of PRISM and the advances that have been developed in the SSASIM dynamic microsimulation model, and two or three other interesting models have been developed in the last two or three years.