University of Wisconsin-Madison

Institute for Research on Poverty

Special Report Series

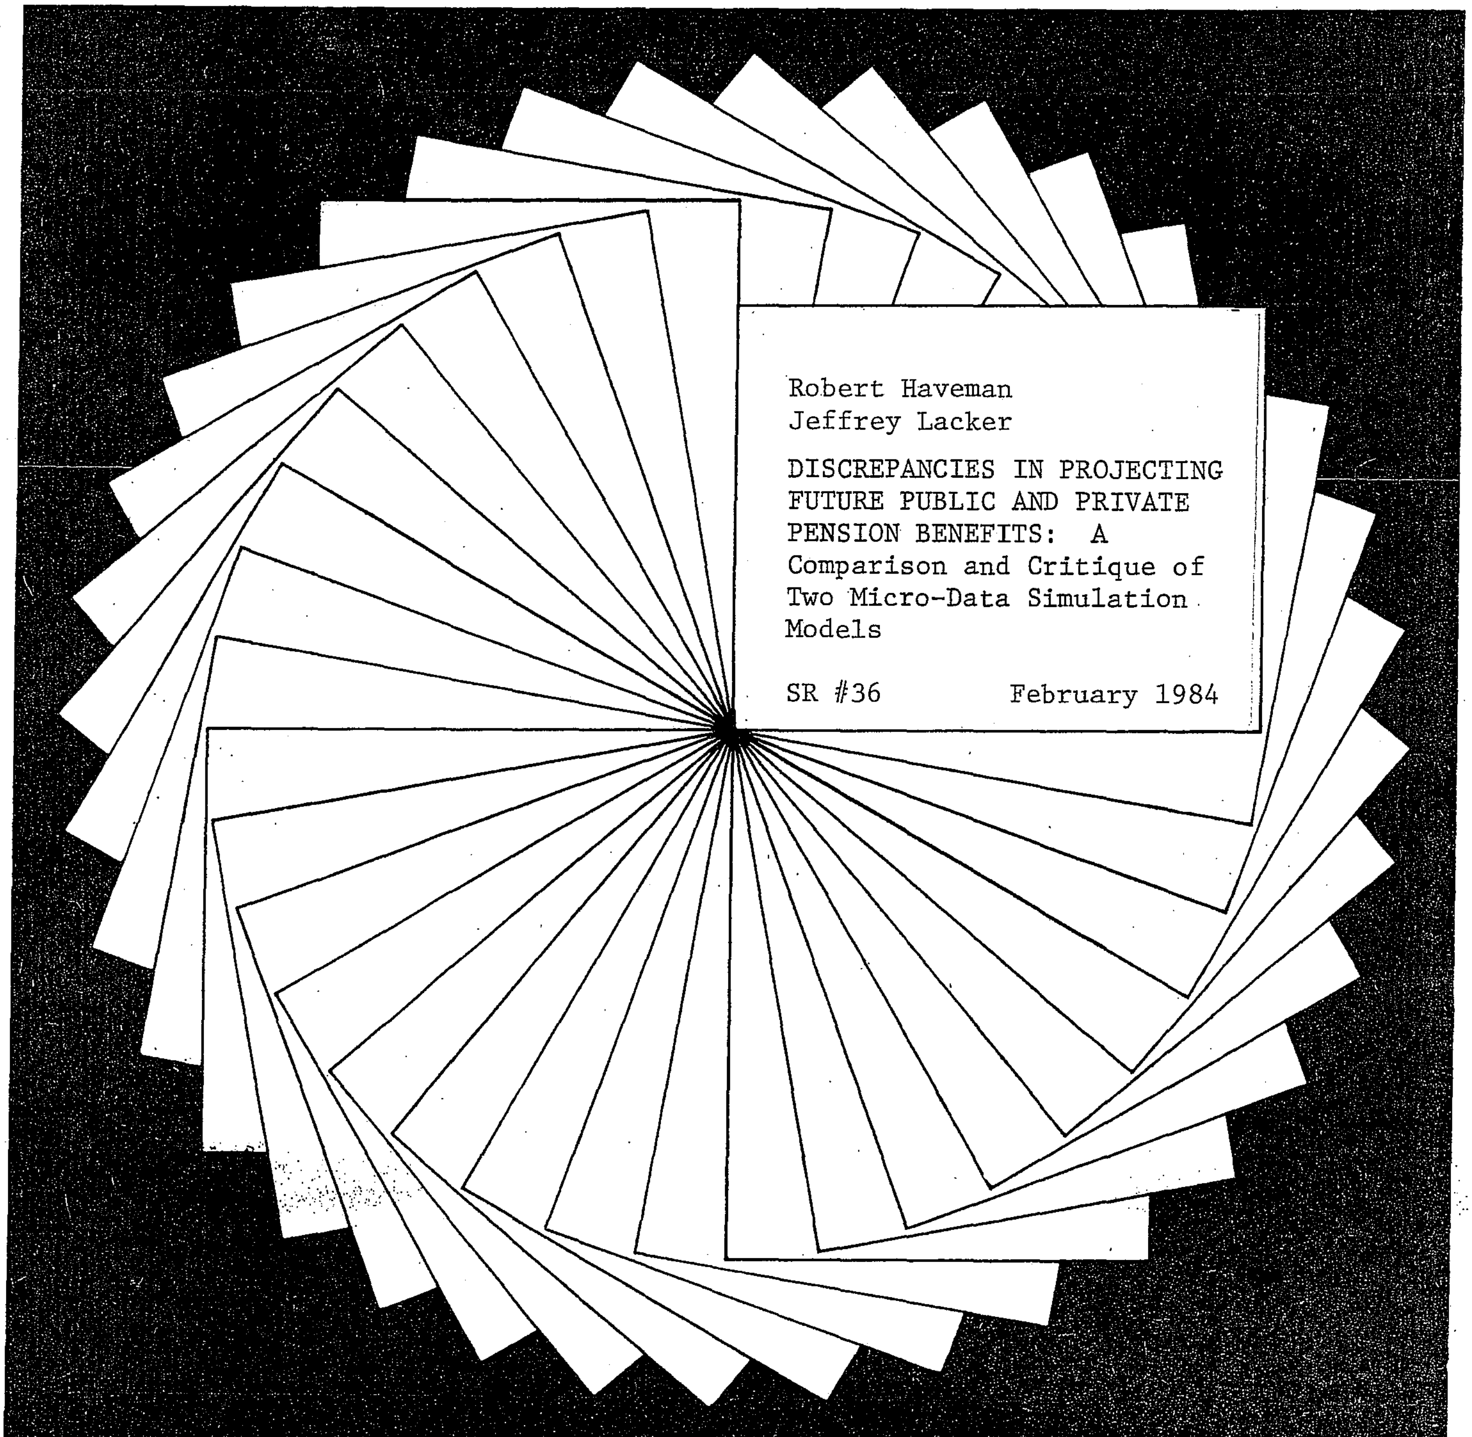

Discrepancies in Projecting Future Public and Private Pension Benefits: A Comparison and Critique of Two Micro-data Simulation Models

Robert H. Haveman
Professor of Economics
University of Wisconsin-Madison

Jeffrey M. Lacker
Research Associate
University of Wisconsin-Madison

February 1984

This research was supported by the Brookings Institution, through a grant provided by the U.S. Department of Health and Human Services, Office of the Assistant Secretary of Planning and Evaluation.

Table of Contents

List of Tables

1. INTRODUCTION

The need to accurately forecast the aggregate flows of future retirement income from various sources and their distribution among retired citizens has become increasingly important. The last three decades have seen tremendous changes in the sources and patterns of income for retired persons. Income derived from both private pensions and from Social Security benefits has grown dramatically, both absolutely and as a percent of the income of retired people. Policy decisions on the design of the Social Security system and on the regulations surrounding the private pension system implicitly rest on an assessment of the future adequacy and distribution of retirement income from other sources--primarily private pensions.

Many different forecasting procedures have been employed to project retirement income from various sources. These run the gamut from simplistic extrapolations to judgmental forecasts based on estimates of demographic changes and expected patterns of income growth to inferences from the economic forecasts yielded by intermediate-term macroeconometric models. However, none of these approaches is able to adequately capture: a) the amount and distribution of retirement income from various sources; b) the flow and distribution of non-pension income, both earned and unearned; c) the underlying evolution of the size and the demographic characteristics of the population; or, d) the endogenous interaction of pension policy and individual labor market, migration, and work-leisure choice. In order to achieve these modelling objectives, large scale microdata simulation models have been developed and adapted to projecting individual behavior and retirement and non-pension income flows. The two leading models which have been developed for this purpose are DYNASIM and PRISM.

The current version of DYNASIM is a descendant of a model originally developed at the Urban Institute between 1969 and 1976.[1] The original model was designed to analyze the impacts of a wide range of public assistance and income transfer programs. It did not include specification of private and public retirement income flows and the behavioral responses which determine and which are determined by expected pension flows. The latest version of the model contains sufficient specification of retirement income flows and endogenous behavior to enable estimation and prediction of individual accumulation and receipt of private pension benefits.

Other modifications have been made since then as well, primarily involving reestimation of behavioral relationships and the enriching of the various components of the model. Nonetheless, the basic structure of the model is intact and many of the original component models have been retained.

PRISM, was developed at ICF Inc., a private research firm, in 1979 and 1980 specifically for the analysis of pension policy. The model was initially developed for the U.S. Department of Labor and the President's Commission on Pension Policy, but was then extensively modified for projects undertaken for the American Council on Life Insurance in 1981 and the Employee Benefit Research Institute in 1982. Because the sole purpose of the model has been the analysis of pension policy, fewer individual characteristics are analyzed than are specified in DYNASIM. It is a more streamlined model, focused directly on the work-retirement decision and the pattern and level of post-retirement income flows.

Although the microdata simulation procedures on which these models rest are marked improvements over previous methodologies, the ability to project retirement income with accuracy has yet to be demonstrated. The accuracy of a given model's projections becomes an important issue to the extent that different models yield different forecasts. For the two models under consideration here, the baseline projections do diverge.

Table 1 reports the most comparable forecasts from each model. Focusing on the figures for individuals reaching age 65 in about the year 2000, we can see that PRISM projects slightly higher Old Age and Survivors Income and significantly higher private pension income than does DYNASIM. In addition, PRISM projects a substantially larger percentage of females and a smaller percentage of males to receive private pension income when they retire.[2] The comparisons of predictions for the early 1980s and for the period around 2020 reveal a similar pattern, with the divergent trends for males in the 1980-2000 period continued until 2020.

Why is there such a significant discrepancy? There are numerous potential answers:

1. The initial population samples on which the two models operate was different.

2. Different specifications were employed for the endogenous relationships of the models.

3. Relationships were estimated on different data sets.

4. Different judgments were made in situations where no data existed.

5. Different exogenous parameter values and assumptions were specified.

Table 1

Comparison of Projections from DYNASIM and PRISM
(average annual benefits in constant 1978 dollars)

	Males		Females	
	DYNASIM 65-67 Year Olds	PRISM 65 Year Olds	DYNASIM 65-67 Year Olds	PRISM 65 Year Olds
	1982	1985	1982	1985
OASI	$5,084	$4,401	$3,115	$3,002
Private pension	$1,876	$3,903	$ 846	$2,321
Percent receiving private pension	31.1%	29.3%	11.2%	11.7%
	2000	1995-2005	2000	1995-2005
OASI	$5,573	$5,733	$3,452	$3,992
Private pension	$3,509	$6,160	$1,584	$2,287
Percent receiving private pension	54.1%	48.5%	24.3%	30.2%
	2020	2015	2020	2015
OASI	$7,865	$7,875	$4,808	$5,532
Private pension	$4,521	$7,438	$1,897	$3,756
Percent receiving private pension	60.5%	49.3%	40.6%	46.5%

Source: Memorandum, Sheila Zedlewski, Urban Institute, to Gary Burtless, Brookings Institute, January, 1983, courtesy of Sheila Zedlewski; PRISM results supplied to her by David Kennell, ICF, Inc.

Ideally, these factors could be isolated and evaluated by running each model so as to test the sensitivity of projections to varying assumptions employed, initial data, or the estimated relationships, for example. This would be a costly and time-consuming procedure due to the computer-intensive nature of the models. It was beyond the range of the budget for this project.

Instead, we have qualitatively assessed and compared the structures of the two models. We take as our sources the detailed documentation of the workings of each model, and compare the two, sector-by-sector. In each area we describe both models, emphasizing the differences in the crucial assumptions and specifications which are incorporated in each. Then, in each area we assess the differences between them, and offer a judgment on the implications of the model differences for retirement income forecasts. Clearly, certain differences can be expected to make large contributions to the observed discrepancy in the forecasts, and we attempt to identify them. Similarily, we attempt to indicate those identifiable differences between the models which can be expected to have only small or negligible effects on the observed discrepancy. There are other differences between the models, however, whose contribution to the direction or size of the observed discrepancy we cannot assess, given the documentation available. From this effort, we cautiously suggest which of the alternative assumptions or procedures appear most reliable, both overall and in any given sector.

While our procedure is able to identify some sources of the difference between the projections yielded by the models, a major question remains unanswered. Which of the projections are likely to most accurately track actual developments in the structure, level, and distribution of pension benefits; which model offers the most reliable future projections? This question cannot be answered with confidence without an extensive study of the conformance of the data used to an ideal data base, and the conformance of the estimated underlying relationships to state-of-the-art procedures for econometrically estimating complex relationships. Again, this was beyond the resources available for this project. However, scattered through our evaluation is our judgment of the degree to which the models conform to the ideal.

In section 2, we present a description of the overall structure of each model. This discussion is accompanied by a schematic description of each model and an extensive table summarizing the primary characteristics of each. Section 3 addresses the demographic modules incorporated into each model designed to simulate death, marriage, divorce, childbearing, education, location, and disability patterns in the population over time.

Section 4 describes and assesses a crucial component of each of the models--the procedures for simulating the labor force activity of each of the observations. These patterns are important determinants of expected pension benefits, insofar as most retirement benefit plans are conditioned on prior work and earnings histories. The actual pension modelling procedures incorporated into each of the models is assessed in section 5. This discussion distinguishes the separate procedures for simulating private pension coverage and benefit levels, social security, disability, and survivors benefits, supplemental security income (SSI) benefits, and retirement income available from individual retirement accounts (IRA's).

In section 6, the procedures used by the two models to estimate the pattern of retirement and benefit acceptance decisions is described and assessed.

Finally, in section 7, we describe how the two models are linked to independent macroeconomic time series forecasts. Section 8 concludes.

2. OVERALL STRUCTURE OF EACH MODEL

DYNASIM and PRISM are both microdata simulation models, and hence share a distinctive methodological strategy. In both cases, the goal is to project the detailed demographic and economic characteristics of a population over time, starting with a very detailed baseline data base describing a population sample at some recent date. The method is to simulate the changes which occur in the characteristics of the population sample by modeling a variety of behavioral patterns and the occurance of certain exogenous events as they pertain to individuals in the population sample. In such models, individuals are probabilistically selected each year, on the basis of complex behavioral rules, for marriage, divorce, childbearing, entrance to and exit from the labor force, job change, retirement, and eligibility for pension benefits. The probabilities used to make these selections are interrelated in very complex ways, and are often derived from analyses of behavioral patterns or data on the recent historical occurance of particular events. In this way, the evolution of the characteristics of a population over time is built up from the stochastic simulation of events occuring to individuals in the population.

Table 2 displays a simple schematic overview of the structure of both DYNASIM and PRISM. For convenient reference, Table 3 summarizes the characteristics of each model's component modules.

DYNASIM

DYNASIM simulates its sample population in three stages (see Table 1). In the first stage, the Family and Earnings History (FEH) model simulates demographic behavior and certain labor force activity for each individual. The output of the FEH is a longitudinal record for each person in the sample of birth, death, marriage, divorce, childbearing, leaving home, geographic mobility, education, disabiilty, labor force participation, hours worked, unemployment, and annual earnings. The record extends for the entire simulation period, or until a person is simulated to die.

These histories are inputs into the second stage, the Jobs model. The Jobs model adds to the longitudinal record of each individual the number of jobs held, the industry of each job, job tenure in each job, and by what type of private pension, if any, the worker was covered by each job.

These augmented records are inputs for the third stage, the Benefits History model. This model computes private pension benefits, social security benefits, Supplemental Security Income, Individual Retirement Accounts, and disability benefits, and simulates the timing of the retirement decision.

The initial data input for the FEH model is a population sample at a specified point in time. The results that have been reported for pension benefit analysis have employed the "March 1973 CPS-SER Exact Match File" (see section 8 for a description). The output of the FEH model, and the other models as well, is a set of longitudinal records of families including records for the demographic and earnings experiences of each individual family member.

The results of the FEH, the Jobs, and the Benefits History models are adjusted during simulation to conform to given aggregate time series assumptions. The FEH model is constrained to replicate aggregate time series for birth rates, death rates, marriage rates, remarriage rates, divorce rates, labor force participation rates, unemployment rates, hours worked, and mean earnings. The Jobs and Benefit History models are capable of being adjusted to replicate user-chosen time series for labor force participation rates, number of working individuals, average hours worked, average earnings, employment and job entry by industry, pension coverage by industry, and participation in IRAs.

PRISM

PRISM simulates its sample population in two stages (see Table 1). The first stage--the Work History model--simulates demographic patterns, labor force histories, and pension plan coverage. The output for each person in the sample population is a longitudinal record of birth, death, marriage, divorce, disability, childbearing, hours worked annually, wage rates, industry of employment, job changes, pension coverage, pension plan assignment, and benefit acceptance. The second stage, the Retirement Benefit Simulation model, takes these longitudinal records as inputs. This model takes the characteristics of the pension benefits which each person is simulated to accept, along with the employment and earnings histories of the individuals, and calculates retirement income from private pensions, IRA's, Social Security, and Supplemental Security Income.

The PRISM simulations are recorded by family, and the data for each individual are retained. The input for the Work History model is a population sample at a given point in time. For the results examined here this is an ICF match of the May 1979 Current Population Survey with Social Security Administration records of each individual's actual covered earnings prior to 1979. In addition, the Work History model takes as input a data base on Pension Plan Provisions, derived from an ICF survey of over 300 representative pension plan sponsors.

The results of the of the Work History model are constrained to equal the results of the ICF Macroeconomic-Demographic model. This macroeconomic model is a version of the Hudson-Jorgenson input-output growth model with a disaggregated labor demand system, and produces time series projections of employment rates and average wage rates for 22 age-sex categories.[3] The simulation of individuals in the Work History model is adjusted to replicate these time series. In addition, pension coverage is adjusted to be consistent with specified forecasts of coverage rates by industry.

Table 2

Schematic Overview of the Structures of DYNASIM and PRISM

DYNASIM

INPUT:		Family and Earnings History Model		OUTPUT:		Jobs and Benefits Model
March 1973 CPS-SER Exact-Match File	→	Simulates, year by year and individual by individual: *birth *death *marriage *divorce *household informatin *disability *educational attainment *region of location *labor force participation *number of hours in labor force *wage rate *unemployment	→	Population history describing demographic events, and labor force behavior	→	Reprocessess historical record, year by year and individual by individual: Jobs submodel: *job change *industry of attachment *pension coverage and participation Employer Pension Submodel--calculates for each individual retirement income from: *Social Security *Individual Retirement Accounts *Private Pensions *Supplemental Security Income *Simulates Retirement Decision

PRISM

INPUT:		Work History Model		OUTPUT:		Retirement Benefit Simulation Model
May 1979 CPS-SSA Match ICF Pension Plans Provisions Data base	→ →	Simulates year by year and individual by individual: *birth *death *marriage *divorce *disability *childbearing *hours worked annually *wage rates *industry of attachment *job changes *pension coverage *pension plan assignment *benefit acceptance	→	Partial population history describing demographic events, labor force behavior and pension records	→	Calculates for each individual retirement income from: *Private Pensions *Social Security *Individual Retirement Accounts *Supplemental Security Income

Table 3

A Summary Description of DYNASIM and PRISM

Model Sector	DYNASIM	PRISM
Demographics		
Mortality	Individual mortality probabilities calculated using estimated regression equations which are functions of the age, race, sex, marital status and education.	Individual mortality probabilities depend on the age, sex and disability status of the individual. Probabilities are assumed to decline smoothly over time.
Childbearing	"Desire for a child" is an exponentially declining function of the woman's age. Probability of a women giving birth depends on "desire for a child" and "contraceptive efficiency" which depends on race and education.	Probability of a woman giving birth depends on age, marital status, employment status, and the number of children already present.
Marriage	Probability of an individual marrying is calculated using estimated regression equations in which marriage (0-1) depends on age, sex, race, education, and employment status. Matching depends on the differences in age and education between two individuals.	Probability of an individual marrying depends on age, sex and marital status. Matching depends on the age groups of individuals.
Divorce	Probability of divorce depends on the date and duration of the marriage, and an estimated regression equation which depends on the presence of offspring in various age categories, the age differential of the spouses, the age at first marriage, wife's education, weeks worked by each spouse, and the ratio of their wage rates	Probability of divorce depends on the age group of the husband and of the wife.

Table 3, (cont.)

Model Sector	DYNASIM	PRISM
Demographics (cont.)		
Disability	Defining the disabled as those reporting work impairment in the Michigan Panel Study of Income Dynamics (PSID), the probability of becoming disabled is calculated using an estimated logit function of age, race, sex, and marital status. Probability of exiting from disabled status depends on education and the factors listed above.	Defining disability as applying for social security disability benefits, the probabilities of becoming disabled and of exiting from disability status depend on age and sex.
Education	The probability of students' acheiving one more year of education depends variously on the year of education, the schooling of the parents, the age, race, and sex of the student.	Model only operates on a data set for which education up to at least age 25 has already been observed. Educational attainment thereafter is ignored.
Location	The probability of moving interregionally depends on age and sex of the head of the household the duration of the marriage (for marrieds), and education.	Initial data set contains state of residence and individuals are assumed to remain there.

Table 3, (cont.)

Model Sector	DYNASIM	PRISM
Labor Force Activity		
Participation	Probability of participation is calculated using estimated probit regression equations in which participation (0-1) depends on a set of individual characteristics (listed below) and an autoregressive term.	The number of hours worked by an individual, given that the individual worked a certain number of hours the year before, is determined by applying transition probability matrices. Separate matrices are applied for different categories distinguished by age, sex, marital status, education, hours worked in the three previous years, whether retirement income is received, and whether the individual is a woman who bears a child, bore a child recently, or became divorced or widowed.
Hours in the Labor Force	The number of hours spent in the labor force (both employed and unemployed) is calculated using estimated regression equations in which hours worked depend on a set of individual characteristics (listed below) and an autoregressive error term.	See above.
Unemployment	The probability that an individual experiences any unemployment spells during a year depends on age, sex, race, education and marital status.	See above.

Table 3, (cont.)

Model Sector	DYNASIM	PRISM
Labor Force Activity (cont.)		
Duration of Unemployment	The fraction of labor force hours in which an individual is unemployed, given that some unemployment occurs, is calculated using an estimated regression equation which depends on a set of individual characteristics (listed below) and an autoregressive term. The individual characteristics which enter the participation, hours in labor force, and unemployment duration equations are: age, sex, race, marital status, disability status, education, region, school enrollment status, the presence of children, spouse's earnings, a set of age-education interaction variables, expected wage, and previous transfer income.	See above.
Wage Rates	Individual's wage rate is calculated using estimated regression equations in which the wage rate depends on age, sex, race, marital status, disability status, region, school enrollment, a set of age-education interaction variables, and an autoregressive error term.	Wage rates determined by applying growth rate assumptions which vary by age, sex, and whether employed. Wage rates are adjusted to be consistent with forecasts produced by the Macroeconomic-Demographic model (see below).

Table 3, (cont.)

Model Sector	DYNASIM	PRISM
Labor Force Activity (cont.)		
Employment	Probability that an individual changes jobs in a given year depends on age, job tenure, and industry of employment. If an individual changes jobs, the new industry is determined by applying probabilities which depend on sex and the previous industry of employment.	Probability of changing jobs depends on age, full-time vs. part-time status, and job tenure. Probability of attachment to a given industry after the job change depends on the previous industry, age, and sex.
Pension and Social Security Benefit Accumulation		
a. Private Pension Benefit Accumulation		
(i) Pension Coverage	Probability of being covered by a private pension plan is applied to job changers and new workers, and depends on industry, real earnings and sex. Probabilities assumed constant over time.	Probability that a job changer is covered depends on industry, hourly wage, and age. Probabilities assumed to rise over time.

Table 3, (cont.)

Model Sector	DYNASIM	PRISM
Pension and Social Security Benefit Accumulation (cont.)		
(ii) Plan Assignment and Benefit Calculation	Participation probabilities depend on the number of hours worked, job tenure, age and sex. Workers selected to participate are randomly assigned a plan type using probabilities which vary by industry. Benefit formulas and constants used in them are selected using probabilities which depend on industry. Benefit amounts are calculated using the formula selected and stochastically applying limitations, eligibility types (early vs. normal vs. special retirement), vesting rules, and election of joint survivor's option.	Workers selected to participate in plans from the ICF pension plan data base. Probabilities and plans depend on industry of employment. Actual plan provisions used to calculate benefits.
b. Social Security Retirement, Disability and Survivor's Benefit		
(i) Retirement Benefits	Retirement benefits calculated according to the provisions of the law prior to the 1983-amendments, using the individual's history of coverage and earnings.	Retirement benefits calculated according to the statutory provisions pre-1983-amendments, using the individual's history of coverage and earnings.
(ii) Disability Benefits	Individuals are eligible if simulated to be "disabled" (PSID self-report indicator) and have zero earnings in two consecutive years. Benefits calculated according to statute.	Disability claimants simulated directly. Benefits calculated according to statute.

Table 3, (cont.)

Model Sector	DYNASIM	PRISM
(iii) Survivors Benefits	Wife's benefits, survivor's benefits, and a woman's own retirement benefit calculated using the earnings and benefit history of the head of household.	Wife's benefit, survivor's benefit, and a woman's own retirement benefit calculated using the earnings and benefit history of the head of household.
Supplemental Security Income	Federal benefits calculated using the disability eligibility proxy described above, age, and earnings. The asset test is ignored. A fraction of recipients in a region receive state benefits which are set at the average state benefit for that region.	Federal benefits calculated using age and earnings, ignoring disability. The asset test is simulated stochastically. State benefits are simulated for the 13 states which apply a higher income floor.
Individual Retirement Acounts (IRA's)	Workers randomly selected to establish IRA's when they begin a new job, and they contribute the allowable maximum each year. At retirement the account ceases to earn interest and is paid out in equal annual payments.	Workers randomly selected to establish IRA's each year, and then contribute at rates (70 percent to 90 percent) which vary with earnings. When they begin receiving Social Security or private pension benefits, workers use the IRA to purchase annuities.

Table 3, (cont.)

Model Sector	DYNASIM	PRISM
Retirement and Benefit Acceptance	Estimated regression equations are applied each year after age 58 to determine if an individual leaves a job, and if so whether he or she retires. Equations depend on age, disability status, marital status, private pension benefit eligibilities, Social Security pension benefit eligibility, wage rate, earnings, the level of Social Security and of private pension wealth, and the loss in each of these from retiring one year later.	Workers eligible for Social Security retirement benefits are randomly selected to begin receiving benefits, using probabilities which depend on age and sex. If so, they automatically accept any private pension benefits for which they qualify. If not, they are randomly selected to retire anyway, using probabilities which depend on age, sex and earnings. Individuals who are vested but not yet eligible for a private plan are assumed to stay on the job until the normal retirement age for that plan.
Linkages to Macroeconomic Time Series	Adjustment factors, set by the user, enter the determination of labor force participation, and wage rate equations, by age-sex category. Nominal interest rate and price inflation rate are given exogenously. Birth rates, mortality rates, and marriage rates are aligned to Alternative II-B of the Social Security Administration.	Employment rates and average hourly wage rates are 22 age-sex categories are constrained to replicate forecasts of the "Macroeconomic-Demographic Model." Nominal interest rate and price inflation rate are given exogenously. Birth rates, mortality rates, and marriage rates are aligned to Alternative II-B of the Social Security Administration.

3. DEMOGRAPHICS

Both DYNASIM and PRISM simulate death, marriage, divorce, childbearing, and disability as the first step in processing the model. Both begin with a sample population in a given year and apply the various demographic modules to each individual in each year to determine the probability of a given event for that individual in that year. Random numbers are then drawn to determine whether the events actually occur. The entire sample is processed for a given year before moving on to the next year.

In both models the probabilities of various demographic events are functions of socioeconomic characteristics of the individuals. These are limited, of course, to characteristics observed for the sample, and those characteristics the model builders choose to carry through the simulation. As the description below will indicate, the two models differ in this regard. In general, the probabilities in DYNASIM's demographic modules depend on a richer set of individual characteristics than do the probabilities in PRISM's demographic modules. However, certain features of PRISM's structure in these modules are designed to more realistically simulate events important to retirement benefit levels--the simulation of disability status is an example (see section 3d.). DYNASIM contains two modules--education and location--which PRISM does not contain. Given PRISM's model structure, they are not needed, since location is never used as an explanatory variable in the determination of pension benefits, and the educational history of their sample population is known (at least to age 25) at the beginning of the simulation period.

3a. Mortality

DYNASIM

The current DYNASIM mortality module is an extension of the original model of Steven Caldwell (Orcutt, et al., 1976). The Caldwell model consists of estimated relationships between mortality rates and a number of individual characteristics. Base mortality rates are assumed for 26 age, race and sex groups.[4] These are assumed to decline exponentially over time from current levels to approach given asymptotes. These group mortality rates are then adjusted multiplicatively based on an individual's age, marital status, and years of education.[5] An additional multiplicative factor is applied to women aged 45 to 64 which varies by the number of children ever born to that woman.

PRISM

The PRISM mortality rates are derived from Social Security Administration data and forecasts. Different mortality rates are applied to disabled and nondisabled individuals. For the nondisabled, these rates vary by one-year age, and sex. For disabled workers, rates vary by sex, age (one-year age cohorts), and by the number of years disabled. These mortality rates are adjusted each year to equate the percentage change in the rates with the Social Security Administration Alternative II forecast of mortality rates by five-year age cohorts, sex (1990, 2000, 2020, and 2040). The same percentage changes in mortality are assumed to apply to disabled and nondisabled.[6]

COMPARISON

The DYNASIM mortality model goes further towards using information available about an individual's characteristics in determining the applicable mortality rate. The white/nonwhite adjustments, the education adjustments, and the marital status adjustments are all lacking in the PRISM mortality module. This could bias retirement benefit estimates, since these factors are likely to be correlated with economic outcomes; for example, PRISM could be underestimating the mortality of low-income workers.

Although disaggregation of mortality rates to one-year age cohorts in the PRISM module adds some detail, it is not likely to significantly affect the pension benefit estimates. As discussed below, the model output is compiled for a given population as individuals reach retirement age, rather than being compiled as "snapshots" of the population at given years. Thus, the age distribution of mortality within five year cohorts is unlikely to greatly affect retirement age outcomes.

The relation of disability and mortality modelled in PRISM is also important in estimating retirement behavior. Since health status is known to affect propensities to retire, it is important that the health characteristics of the surviving population be accurately simulated. The disaggregation by disability in PRISM is potentially even more significant, since at every age mortality rates are higher for the disabled than for the non-disabled by about an order of magnitude. This distinction is likely to be an important factor in developing accurate forecasts of disabled-worker benefit experience.

In summary, DYNASIM allows race, education, and marital status to affect mortality and PRISM does not, while PRISM allows disability to affect mortality and DYNASIM does not. Omitting any of these factors could alter the projected distribution of retirement income.

3b. Childbearing

DYNASIM

The DYNASIM module first simulates a woman's desire to bear a child. The probability of a married woman desiring another child varies with the number of children already born. Unmarried women are assumed not to desire a child. If a woman desires a child the probability that she will have a child is an exponential function of MAX[0, A-24], where A is the woman's age. If a woman does not desire a child the probability that she will have a child is the same exponential function (but with different parameter values) multiplied by a "contraceptive efficiency" adjustment factor which varies by race and education.[7] If a birth occurs to a woman, the arrival of twins or triplets is determined by probabilities which vary by race. The sex of the child is determined with probabilities which also vary by race. The various probabilities and parameter values are chosen to reproduce data in the U.S. Public Health Service's Vital Statistics for 1969.[8]

PRISM

The PRISM childbearing module relies on recent CPS data. Fertility rates are compiled by age group, marital status, employment status, and number of children.[9] These rates are adjusted each year to conform to

SSA Alternative II birth rate forecast assumptions. The PRISM module never delivers twins or triplets.

COMPARISON

The DYNASIM approach allows a woman's race and education to affect childbearing, while PRISM does not. PRISM, however, allows employment status to affect childbearing while DYNASIM does not. Neither model directly simulates the spacing of childbirths, but spacing patterns will emerge as the by-product of the intertemporal interdependence of simulated childbearing. The simulated spacing of childbirth will affect the labor force behavior of women of childbearing age, particularly job tenure patterns. The documentation of model specification does not provide a reason to believe that either model is biased, either relative to each other or relative to actual childbirth spacing.

The difference between the two models in the role assigned to race and education seems most likely to make an economic difference; PRISM may be underestimating births to low-income poor women, and overestimating them for women with more education. The affect of employment status on childbearing (and vice versa) is potentially very important, but this relationship might be captured in the PRISM model by accounting for the affect of childbirth on employment experience (see the discussion on the labor force activity models, particularly "hours worked," in section 4).

Between the two models, the role assigned to employment status is likely to be among the most significant differences. Labor force participation by women is well known to be related to childbearing. Since assessing this factor requires examination of how childbearing affects labor force outcomes, discussion of this area will be postponed to section 4.

The DYNASIM fertility rates apparently are not constrained to meet any given forecast, while PRISM's are constrained to meet SSA fertility assumptions. This is potentially very important as well. Suppose the models were identical in capturing the effects of childbearing on a mother's labor force behavior, but one model simulated more births than the other. Then the model which simulated more births would simulate lower labor force participation for women, and thus lower pension coverage and pension benefits.

3c. Marriage, Divorce, and Household Formation

DYNASIM

The DYNASIM marriage module has three parts. The first is a model of the probability of a first marriage for never married individuals between ages 18 to 29. These probabilities are calculated using separate estimated regression equations for 12 age-race-sex groups.[10] The equations were estimated using 1981 CPS data (see Orcutt, et al., 1976, Chapter 4). The variables which determine first marriage probabilities in these equations include "3 regional, 5 education, 6 hours worked, and 5 hours worked categories." The constant terms in the equations are adjusted over time so that the equations replicate the most recent Vital Statistics data.

The second part of the marriage module applies to individuals who are not age 18 to 29 or who have never married. Marriage rates for these individuals depend only on sex, race, age, and previous marital status.[11]

Applying these marriage rates and the first marriage equations to unmarried individuals each year produces a list of males and females selected to marry in that year. In the third part of the module, these lists are segregated by sex and are ordered randomly. The first male and female are selected from the list and it is randomly determined whether or not they marry each other. The probability of their union is calculated as a declining exponential function of the sum of the squared differences in their ages and in their education. This probability is applied, and if they are not wed, the next female on the list is selected and evaluated. A male who strikes out five times is matched with the female with whom his probability of marriage was highest. Each male on the list is matched in this fashion, and leftovers are returned to the single population.

Divorce is simulated in a three stage process. First, a nonlinear time trend determines the probability that the marriages contracted in a given year will ultimately be dissolved through divorce. The second stage distributes these divorces over time. This is done by applying estimates of the proportion of divorces in any given marriage cohort which occur to marriages of given durations. The third stage distributes the divorces which occur during any year (to all marriage cohorts) among the surviving marriages. This is done by applying an equation which gives the effect of various household characteristics on divorce probabilities.[12]

A person who is between ages 14 and 34 and who is not the head of a household or married may leave home and become an independent family. The probability of doing so in any given year depends on age, race and sex.[13]

PRISM

The PRISM marriage module randomly selects individuals to become married using probabilities which depend on sex, age, and marital status.[14] For males who have been selected to marry during the year, a probability distribution is used to determine the age group of his wife. The probability distribution depends only on the male's age group. If no female is available in the age group for a given male, the next lower and higher age groups are used. Leftovers are listed to marry next year.

Divorce is determined by applying to each married couple the 1977 central divorce rates by age group of husband and age group of wife.[15]

COMPARISON

The DYNASIM marriage module is rich in socioeconomic detail, designed to capture the effects of many household characteristics on marital status. To the extent that these particular effects are important--clearly an empirical issue that will vary with the issue being analyzed--DYNASIM has an advantage. The relations between marital status and labor force behavior are likely to be relevant in estimating retirement benefits, and some of these relations are modelled in the first marriage probability equations, and in the matchmaking and the divorce algorithms. In addition, the divorce algorithm is designed to produce marriage durations near historical patterns; divorce is unrelated to duration of marriage in PRISM. This may also have a bearing on retirement benefit levels, since for example, marriage duration will affect qualification for survivor's and social security benefits.

In its favor, PRISM ensures that the relative ages of spouses bears a relation to historical data. The tables for PRISM show that males tend to marry women younger than them on average, whereas DYNASIM treats age differences symetrically (i.e., it is equally likely for the male to be 5 years older as for the female to be 5 years older). Furthermore, the DYNASIM matchmaking procedure is completely *ad hoc*; there is no guarantee that it replicates any observed pattern concerning the relative ages of spouses. This difference in the two models has the potential for affecting certain types of retirement level benefits, such as survivors benefits and social security, since if women are, on average, younger than their spouse they are more likely to be widowed than if they are, on average, the same age. The more likely women are to be widowed before age 65, the higher their retirement income is likely to be. Thus, DYNASIM's treatment of age differences could bias downward their estimates of female retirement income.

3d. Disability

DYNASIM

DYNASIM simulates disability by determining in each year whether an individual enters or exits from disabled status. A person is defined as disabled if they report a work-impairing disability on the PSID. The probability of becoming disabled is calculated using a logit probability function of a vector of dummy variables indicating an individual's race, age, sex, and marital status.[16] The probability of exiting from disabled status is based on a logit function that depends on education, as well as the factors listed above.[17] These equations were estimated on four years of PSID data.

PRISM

In PRISM, an individual is defined as disabled if they apply for and are classified as totally disabled by the Social Security Administration's Disability Insurance Program. The probability of becoming disabled in a given year depends on the age (one-year cohorts) and sex of the individual. The probability of exiting from disability status depends on age, sex and the number of years disabled. These probabilities were estimated from Social Security Administration data on disability benefit claimants.

COMPARISON

The inclusion of race and marital status as determinants of disability are a potentially important advantage for DYNASIM. These two characteristics are clearly related to economic status, and capturing their effect would improve the realism of the depiction of the economic status of disabled workers. On the other hand, PRISM's termination probabilities depend explicitly on the number of years disabled. This allows more realistic portrayal of the effect of the duration of disability on exit from that status. PRISM's mortality rates also depend on disability status and the number of years disabled, and this should improve PRISM's forecasts of the number of disabled retired workers.

The most important difference between the two models is their definitions of disability status. The DYNASIM equations were estimated on PSID data and thus were forced to rely on a self-reported definition of disability distinct from Social Security disability benefit eligibility. In the section of DYNASIM where it is determined whether or not a given individual qualifies for disability benefit, a cumbersome ad hoc procedure checking the individual's labor force participation is employed in order to simulate Social Security disability status.

For the purpose of estimating retirement income, estimating disability benefits is the most important use of the information on the individual's disability status. It therefore seems much preferable to simulate Social Security disability status directly, as in PRISM.

On the other hand, disability status has been found to be an important determinant of retirement propensities for older workers. Both models allow their respective disability indicators to influence retirement decisions. Because DYNASIM's indicator is broader than PRISM's, DYNASIM is allowing some people to become disabled that PRISM does not. DYNASIM gives these individuals different labor force behavior than nondisabled individuals, while PRISM does not. This implies that PRISM might be simulating these individuals' retirement behavior inaccurately. On these grounds, the broader disability definition would be preferred.

3e. Education

DYNASIM

DYNASIM simulates the schooling of individuals year-by-year by applying various transition probabilities which depend on the individual's age, race, sex, and on the parents' education. Children aged five, six or seven are eligible to begin school. The probability of enrolling depends only on age, and all chidren enroll before age eight. All children remain enrolled until after seventh grade. Then, students may either repeat seventh grade, enter eighth grade (and graduate the next year) or drop out. The probabilities depend on the race of the student, the number of years of schooling completed by the student, and the number of years of schooling by the parents. Students held back face the same probabilities the next year. The probability of entry into high school depends on the race and sex of the student and the education of the parents. A student completing a year of high school either enrolls for another year (to complete another grade) or drops out. For students enrolled for four or more years may graduate, drop out, or enroll for another year. The drop out rates depend on the student's race, age, sex, the number of grades completed, and on the parents' education. The probability that a high school graduate enrolls in college depends on the student's age, sex and the parent's education (data by race apparently were not available). One third of high school graduates who continue onto college delay their entry by one year. The probability of continuing in college for another year depends on race, sex, years of college attended, and parents' education. Some students require five years to graduate. The probability of entering graduate school depends on race, sex, and parents' education. All graduate students attend for two years and then graduate.

PRISM

Because PRISM employs a baseline sample population of 25 to 64 year olds, simulating their behavior until retirement age, there is no need to simulate educational experience before age 25. PRISM implicitly assumes that no additional education occurs after age 25.

COMPARISON

For the projections on which the two models can be compared, only educational attainment after age 25 matters. PRISM assumes none, while some schooling takes place for individuals over age 25 in DYNASIM, although the probability is very small.[18] One would have to prefer DYNASIM's explicit modelling here although the effect of the differences--via labor force participation while in school--is unlikely to have a significant effect on estimates of retirement income.

3f. Location

DYNASIM

DYNASIM simulates the mobility of its sample population in three stages. First, households are selected to move using probabilities which depend on the duration of the marriage (for married couples), sex of the head of the household, and age.[19] Second, some of the households that have been selected to move are selected to move across county lines using probabilities that depend on the age, education, and sex of the head of the household.[20] Third, each of these intercounty migrants has a destination region selected, depending on their current region and their race. The regions are: Northeast, North Central, South and West. The documentation does not define each region. The region in which a household resides affects the labor force participation equation and the unemployment duration equation, and it enters these equations as a dummy variable which is one if the household is in the South and is zero otherwise. Region of residence also enters calculation of Supplemental Security Income (see section 5c.).

PRISM

PRISM does not model migration. The state in which an individual resides is recorded in the baseline data. It is assumed that individuals remain in the state in which they initially reside for the entire simulation. This information is used only in simulating the state supplements to Supplemental Security Income (see section 5c.).

COMPARISON

DYNASIM allows location to influence labor force behavior, while PRISM does not. DYNASIM allows net migration, while PRISM does not. This omission is important to the degree that regional location affects otherwise independent determinants of retirement income, such as various aspects of labor force behavior. For example, labor force participation and unemployment duration may be correlated across individuals due to the association of each with location. Omitting location as a determinant of each could lead to simulation results in which participation and unemployment duration are not as correlated as in the underlying population. This could bias estimates of retirement income, since each enters nonlinearly into the determination of pension and Social Security benefits. The direction of this bias, however, is difficult to assess.

Location also affects the state supplements to Supplement Security Income in each model. As our discussion of the PRISM procedure notes (see below), 13 states account for 88 percent of all state supplemental benefits. All but four states--Hawaii, Nevada, California, and Washington--are in the Northeast or North Central regions.[21] These are states among which one would expect to see significant net migration over the next few decades. DYNASIM can be expected to capture some inter-regional flows and will for this reason will have an advantage over PRISM in estimation of Supplemental Security Income. If, for example, net emigration is expected for those 13 states, PRISM will overestimate state supplements to SSI. On the other hand, DYNASIM does not actually simulate state of residence, only region. DYNASM's procedure for assigning state supplemental benefits may introduce a separate bias (see section 5c. for further discussion).

4. LABOR FORCE ACTIVITY

4a. Participation, Hours Worked, and Unemployment

The procedures employed for simulating labor force activity are crucial for forecasting patterns of pension coverage, the distribution of pension benefits, and the aggregate public and private costs of pensions. Coverage of private pensions often depends on industry and occupation. Benefit levels depend heavily on the years of work and on the individual's wage rate. And, the actual benefits paid depend on these factors plus decisions on when to retire, which are themselves dependent on labor force activity. In this section, we describe and compare the procedures used by the two models in this area. Table 4 presents a simple scheme describing the labor force modelling procedure in DYNASIM and PRISM.

DYNASIM

DYNASIM simulates an individual's labor force activity in any given year using four sets of equations estimated on data from the Michigan Panel Study of Income Dynamics. The first equation determines the probability that an individual is or is not in the labor force. A second equation determines the number of hours spent in the labor force. A third equation determines the probability that an individual has a spell of unemployment during the year. A fourth equation determines what fraction of the number of hours spent in the labor force are spent unemployed. Adjustment factors enter each equation, which factors insure that the aggregate labor force outcomes are constrained to replicate exogenously specified aggregate labor force time series (see section 7).

Table 4

Schematic Overview of Simulation of Labor Force Participation, Hours Worked, and Unemployment in DYNASIM and PRISM

DYNASIM

Does Individual Participate?	→	If So, For How Many Hours?	→	Did Individual Experience Any Unemployment?	→	How Long Was Unemployment Spell?

Each equation depends on individual characteristics and serially correlated error terms.

PRISM

Determines the individual's category and then applies the hours worked transition matrix for that category.

Individuals first categorized by:	Individuals then categorized by:	Individuals further categorized, if resulting group had sample size of at least 30:
*Women experiencing "special circumstances"	*Divorced *Widowed *Experiencing childbirth *Child present *Young	*Hours "normally" worked
*Individuals receiving retirement income	*Type of retirement income: SS, pension, and age	*Sex *Marital status
*All others	*Age categories	*Sex *Marital Status *Hours "normally" worked *Education

Labor force participation is determined by separate estimated probit equations for 16 age, sex and race groups.[22] Let S(i,t) be an indicator for a given age-sex-race group such that individual i is in the labor force at time t if $S(i,t) \geq 0$, and is not a participant if $S(i,t) < 0$. S(i,t) depends on a vector of characteristics, X(i,t), an error term, E(i,t), and an adjustment factor, A(t):

$$S(i,t) = X(i,t) * \beta + E(i,t) + A(t).$$

The individual characteristics included in X in various group's equations are: education, region, disability status, marital status, age, whether enrolled in school, presence of a child under six in the family, number of children in the family, spouse's earnings, and a set of age-education interaction variables. Not all variables appear in all equations.[23]

The error term has two components. One is a "lifetime error" which is selected randomly for each individual at age 16, and which remains with the individual throughout the simulation. The other error component is a "transitory error" which is assumed to follow a first-order autoregressive process--the transitory error at time t, v(t), is given by $v(t) = r \cdot v(t-1) + w(t)$, where w(t) is selected randomly each period, and r is the estimated serial correlation coefficient. The serial correlation coefficient varies over age, sex, and race groups, but is the same within each group. The stochastic term in the autoregressive process, w(t), is drawn from a mean zero normal distribution with the appropriate (estimated) variance. The adjustment factor, A(t), can be altered by the model user so that aggregate labor force participation replicates a given time series forecast. Note that a separate adjustment factor is available for each age-race-sex group.

In the module calculating the number of hours in the labor force for an individual there are separate estimated equations for 14 age-sex-race groups.[24] The number of hours in the labor force for an individual i at time t, H(i,t), is given by:

$$H(i,t) = [X(i,t) * \beta + U(i,t)] \exp[n(i,t) - \sigma_n^2].$$

The vector of characteristics, X(i,t), includes: age category, transfer income in the previous year, expected relative wage, disability status, marital status, presence of children under six, income of other household members in the previous year, and the number of children present.[25] Not all variables appear in each equation; the number of children present and the income of other household members appear only in the female equations.

The error U(i,t) is assumed to follow a first order autoregressive process:

$$U(i,t) = r(i) \cdot U(i,t-1) + Y(i,t).$$

The serial correlation coefficient, r(i), varies over age-sex-race groups. The stochastic term in the process Y(i,t) is drawn from a normal distribution with a mean of zero and the appropriate variance. The n(i,t) error term also enters the wage equation and is designed to capture the negative correlation between hours worked and wage rates. This term is also drawn from a normal distribution with a mean of zero. New entrants are assigned a number of hours worked equal to α times the number of hours predicted by that individual's equation, where α is drawn from a uniform probability distribution over the interval (0,1).

DYNASIM determines unemployment for an individual in two steps. It first determines whether or not the individual experiences any unemployment during the year. Then, for those individuals who do experience some unemployment, the model determines what proportion of hours spent in the labor force are spent unemployed.

Whether an individual experiences any employment in a year is determined by a table of probabilities which depend on age, sex, race, education, and marital status.[26] If an individual is unemployed during the year, an estimated equation determines what fraction of the numbers of hours in the labor force are spent unemployed. There are separate equations for persons age less than 21, persons age 65 and over, and by race and sex for persons age 21 to 64. Various of these equations depends on education, age, region, marital status, disability status, presence of a child under six, and race and sex.[27] An error term is also included in the equations. The error is assumed to be a first order autoregressive process exactly analogous to the autoregressive errors described above. The serial correlation coefficients are different in each estimated equation. The stochastic term is drawn randomly each period.

PRISM

In PRISM, the number of hours worked for an individual is simulated using a set of transition probability matrices. Individuals are classified into one of five annual hours worked categories: zero; one to

500; 501 to 1000; 1001 to 1500; 1501 and greater. The transition matrices specify the probability that an individual in a given hours worked category this year, t, is in some other given hours worked category next year, t+1. Individuals are all assumed to work the median number of hours for the hours worked category to which they are assigned in a given year. Note that the decision to participate in the labor force and the determination of whether or not an individual is unemployed and if so for how long are both subsumed in these transition probabilities.

ICF has estimated transition matrices for 113 nonoverlapping and exhaustive socioeconomic groups. The data set was pooled from 1975 through 1980 Current Population Surveys, and from SSA work history records matched to the March 1976 and March 1978 CPS's. These sources together contain a sample of 140,000 adults for whom hours worked information is available for two consecutive calendar years. The transition probabilities were estimated by subdividing this sample according to certain socioeconomic classifications. Separate subdivisions were created for categories of "women experiencing special circumstances," such as childbirth or divorce, and for categories of "individuals receiving retirement income." Some of these were further subdivided by age, sex, marital status and "hours normally worked" (defined below).[28]

An individual was defined as "normally" works full-time if they were covered by social security for four quarters in each of the second and third years prior to the survey (t-1 and t-2). Anyone who earned between 1 and 3 quarters of coverage in one or both of t-1 and t-2 were classified as "normally" works part-time. Any individual who earned zero hours of coverage in both t-1 and t-2 is classified as "normally" does not work. Workers who were employed in the federal government or in non-covered state and local government were classified on the basis of the number of hours worked during the survey year, t. During model simulation, individuals are reclassified each year on the basis of their hours worked in the two previous years, t-1 and t-2.

COMPARISON

The simulation of individual entry to and exit from the labor force is clearly very important for modeling the accrual of retirement benefits. Qualification for both private pensions and social security and their benefit levels are closely linked to the individual's employment record. The more frequently an individual is simulated to depart the labor force, the smaller that individual's retirement income is likely to be. In addition, the number of hours which an individual worked will help determine earnings, on which many pension benefit formulas are based.

Each of these models has strengths relative to the other in modelling labor force behavior. As we will suggest below, the breadth of individual characteristics which help determine labor force outcomes in DYNASIM is a strong advantage, both in capturing the interaction of these characteristics with labor force outcomes and in projecting retirement income. However, certain features of PRISM give that model important advantages in replicating intertemporal dependencies in labor force behavior, which in turn have important effects on projections of retirement income. In both cases, unfortunately, the net direction of bias is difficult to assess.

There is a strong rationale for a detailed approach like that of DYNASIM--essentially an argument that omitted variable bias is a potentially important problem. The DYNASIM approach to modelling labor force behavior uses sets of estimated equations which utilize very detailed information about an individual's characteristics. PRISM, by contrast, is relatively parsimonious, depending on only a few classifying traits. There are important correlations between these characteristics and labor force behavior. The functional forms of the equations estimated for the DYNASIM modules conform generally to what much previous research has indicated are appropriate for capturing the relationship between these characteristics and the expected values of hours worked, wage rate, and labor force participation. In contrast, PRISM attempts to estimate separate transition probability matrices for subgroups categorized by some of these characteristics. Using this approach, degrees of freedom are quickly exhausted and the detail of explanatory factors is sharply constrained. For example, in PRISM only some groups are allowed to have transition matrices which vary by the individual's education, and even for these the variable is just a simple binary one (some college versus no college). In DYNASIM, education enters the regression equations as dummy variables indicating the individual's presence in one of five classifications of number of years of schooling; in addition, these dummy variables are multiplied by age to capture interaction between the effects of age and education.

The structure of DYNASIM's approach can thus be expected to reliably depict more interactions between demographic characteristics and labor force activity than that of PRISM. To the extent that the association of these demographic characteristics with labor force behavior is responsible for associations between demographic characteristics and pension plan coverage, participation and benefit levels, DYNASIM can be expected to be more reliable in depicting the associations between demographic characteristics and retirement income. For example, individuals with less education are more likely to be unemployed and are likely to earn lower wages. If retirement benefits are positively related to lifetime earnings, DYNASIM would capture the association between low schooling and low retirement income. Notice that this example suggests an additional reason DYNASIM's approach could be expected to produce more reliable estimates of retirement income; wage rates and hours worked both influence an individual's retirement income and both are related to individual characteristics and thus will be correlated. While this correlation is captured in DYNASIM, they are treated as independent in PRISM. To the extent that different factors (e.g., hours worked, industry of employment) which independently affect retirement benefits are related to common characteristics of the individuals (e.g., education, race), DYNASIM, by modelling the effect of the individual characteristics on these factors, will more accurately depict the correlation between these otherwise independent factors.

However, for the purpose of projecting retirement benefit levels, the functional form of the PRISM specification explicitly recognizes serial correlation and, hence, is likely to be superior to that of DYNASIM. As mentioned above, job tenure is an important determinant of vesting and benefit levels. If one model predicts longer average job tenure than the other, that model will predict ceteris paribus, larger retirement income.

Depicting job tenure essentially involves depicting the intertemporal persistance in an individual's labor force activity. In PRISM, individuals who are currently working full-time will be assigned a number of hours worked next year using probabilities different from those used to assign a number of hours worked to an indivdiual working zero hours this year. In addition, for many groups, current full-time workers are assigned a number of hours worked next year differently depending on their work experience in the previous two years. Keep in mind that, in PRISM, workers in the more than 1500 hours worked category are assigned the median number of hours worked in that category (for 1500 or more hours worked this is almost always 2000 hours). In the sample transition probability tables displayed in the documentation, the probability of working more than 1500 hours next year, given the individual is working more than 1500 hours this year is usually quite large--generally over three-fourths. Thus, PRISM's simulations are very likely to produce workers with a string of 2000 hours worked.

DYNASIM also captures some serial interdependence but less satisfactorily in some respects. DYNASIM's equations predict a number of hours worked for an individual, and then add a serially correlated error term with normally distributed disturbance. As an example, suppose that in a population over which one of the DYNASIM equations is estimated someone with a certain set of characteristics works full-time but is unemployed for a full year during every fourth year. If the estimated equation gives an unbiased prediction of the expected value of the number of hours worked for individuals with these characteristics, it will predict 1500 hours when the error term is set to zero. The error term will be chosen in each simulation year by the rule U(i,t) = r · U(i,t-1) + e(i,t), where r is the estimated serial correlation and the e(i,t)'s are independent draws from a normal distribution with a zero mean and the appropriate variance. Neglecting the wage rate equation error term, the number of hours worked will be 1500 + u(i,t). If the characteristics do not change from one year to the next, the number of hours worked next year, H(i,t+1), will be H(i,t) + e(i,t). In this simplified example, the specification is unlikely to simulate these individuals working 2000 hours very often. This illustrates that a specification designed to unbiasedly predict the expected value of, say, hours worked per year conditional on a set of observed characteristics may be inappropriate for simulating, say, the number of consecutive years of full-time work.

PRISM takes a different approach to capturing the serial dependence of labor force activity. PRISM simulates hours worked next year separately for individuals in different categories of hours worked this year. This, by itself, is a Markov model approach which assumes that all of the relation between hours worked this year and the indivdiual's labor force history can be summarized by the relation with hours worked last year. This is extended in PRISM by classifying much of the sample by three categories of hours "normally" worked. Another PRISM feature which affects the time pattern of labor force activity is that individuals are assigned a number of hours worked equal to the median number of hours worked in the category to which they are assigned. Thus workers simulated to work full-time in PRISM will all work 2000 hours; similarly for "half-time" workers. This structure is more likely to reproduce certain time patterns in hours worked than the DYNASIM structure. Individuals are likely to stay at the same number of hours worked for a number of consecutive years, an occurance that is quite unlikely with the first order autoregressive specification in DYNASIM.

The case is less clear with respect to labor force participation. PRISM ignores the distinction between the unemployed and those not in the labor force, lumping them both in "zero hours worked." In DYNASIM, zero hours worked can result from either nonparticipation or through an unemployment spell of 100 percent of hours in the labor force. Recall that the former involves an indicator function with an autoregressive error term. An unemployment spell of 100 percent involves simulation of the occurance of unemployment and prediction of a duration of 100 percent of participation. The occurance of an unemployment spell involves the (serially independent) application of a table of probabilities. The duration of the unemployment spell involves an estimated regression equation with a first order autoregressive error term. So some serial dependence can arise in the occurance of zero hours worked for an individual in DYNASIM through the unemployment duration equations. A serially correlated error added to an indicator function will make participation less likely this year if an individual did not participate last year, and thus could, in principle, simulate the duration (over calendar years of nonemployment) as realistically as the transition matrix approach of PRISM. Unlike the case of hours worked, it is unclear a priori which structure will produce the more "realistic" simulation of the inter-temporal pattern of nonemployment.

The specification of the number of hours worked in PRISM has its own particular weaknesses. As discussed above, the transition probabilities are invariant with respect to many individual characteristics which research has shown to be significantly associated with patterns of labor force behavior. Although the PRISM specification is designed specifically to capture serial dependence in labor force activity, part of the serial dependence actually observed in labor markets arises because individuals with different characteristics exhibit different labor force behavior over their entire lifetime. In PRISM, serial dependence fades away over time since hours worked depends on, at most, the past three years of work history. DYNASIM could capture these lifetime differences more accurately than PRISM if they are systematically related to individual characteristics such as education. In addition, the DYNASIM equations include an "individual effect;" an error term randomly assigned for the individual's entire lifetime. Thus, there are certain aspects of the intertemporal pattern of labor force behavior for which PRISM is not necessarily more suitably specified than is DYNASIM. An additional disadvantage of the PRISM specification is that it artificially compresses the number of hours worked within each hours-worked category. This implies some understatement of the variability of income and employment.

Nonetheless, our overall assessment of the two models is that the specification of PRISM is better suited to simulating the intertemporal pattern of the number of hours worked, but that DYNASIM is superior in its modelling of the influence of numerous individual characteristics on labor force outcomes, characteristics which are likely to have important effects on retirement income via independent effects on participation, hours worked and unemployment. Unfortunately the two modelling approaches are effectively exclusive--neither model could "add on" the features of the other. In a sense, each modelling approach uses its available degrees of freedom differently, and thus it is to be expected that they yield final estimates with different relative strengths. Again, we would emphasize that without the ability to secure comparable simulations from each of the models, no overall assessment of the biases implicit in each--with respect to each other and to actual labor market performance--is possible.

4b. Wage Rates

DYNASIM

The wage rate module in DYNASIM consists of separate estimated equations for 16 age-sex-race groups.[29] The wage rate for an individual, at time t, W(i,t), is given by:

$$W(i,t) = \exp[X(i,t) * \beta + n(i,t)].$$

The characteristics specified in X(i,t) include education, region (South vs. nonsouth), disability status, marital status, age, whether or not enrolled in school, and a set of age-education interaction variables (see footnote 23). Not all appear in each equation.

The error term, n(i,t), consists of a permanent "lifetime error," plus a transitory error. The same term appears in the hours worked equation, in order to capture the correlation between an individual's hours worked and wage rate. The transitory error is assumed to follow a first order autoregressive process. The serial correlation coefficients vary over age-sex-race groups, but are the same within groups. The wage rate for an individual which is used in the model is the wage rate predicted by the equation, multiplied by a wage adjustment factor for the individual's age-sex-race group. This adjustment factor allows the outcomes to be aligned with aggregate time series (see section 7).

PRISM

Wage rates for a given individual in PRISM are calculated using that individual's wage rate in the initial sample year, and applying wage rate growth assumptions each year. These wage rate growth assumptions were estimated from data in a match of the May 1977 and May 1978 CPS. They vary by sex, age, and whether or not the individual changes jobs during the year.[30] Unemployed individuals are imputed a wage rate growth of 80 percent of an employed persons' wage rate growth; this determines their wage rate when and if they again become employed.

In the module which links the microdata with macroeconomic time series, wage rate forecasts (and simultaneously, employment rate forecasts) are calculated for 22 age-sex groups (see section 7 below). The individual wage rates predicted using the assumptions described above are then adjusted in order to replicate these group average wage rates.

COMPARISON

Wage rates are important in forecasting retirement income insofar as many plans link benefits to earnings. Within given age-sex groups it seems clear that DYNASIM's approach using estimated regression equations on individual characteristics will produce a larger variance in individual wage rates both across individulas and over time for given individuals than will PRISM's.[31] The only variation in wage rates within age-sex groups in PRISM is the variation in observed wage rates in the initial sample year and the variation due to unemployment. Once again, to the extent that wage rates are associated with individual characteristics (e.g., education or rate) which show independent links to other determinants of retirement benefit levels, (e.g., frequency and duration of unemployment spells), DYNASIM will capture correlation between retirement income determinants that are neglected by PRISM.

4c. Industry of Employment

DYNASIM

DYNASIM's simulation of job change and industry of employment is performed in the Jobs model, after a complete demographic and labor force behavior record has been created for each individual over his or her lifetime. In any given year, each individual who is in the labor force in that year and the previous year is eligible for a job change. The probability of a job change depends on age, job tenure, and the previous industry of employment.[32] These probabilities were estimated using January 1973 CPS data.

Every individual in the labor force--whether employed or not--is assigned one of ten industries (see Table 5). Labor force entrants are assigned to industries using probabilities which depend on sex and education.[33] For individuals already in the labor force who are selected to change jobs during the year, a (possibly new) industry is assigned using table of "inter-industry transition probabilities." There are two of these matrices--one for each sex. A row in one of these tables specifies the probability that an individual employed in that

Table 5

Industry Classification Used in DYNASIM and PRISM

DYNASIM	PRISM
Agriculture	Agriculture
Construction and Mining	Construction
	Mining
Manufacturing	Manufacturing
Transportation	Transportation
Utilities and Communication	*
Trade	Trade
Finance, Insurance, and Real Estate	Finance, Insurance, and Real Estate
Services	Services
	Self-Employed
State and Local Governments	State and Local Governments
Federal Government	Federal Goverment

*Because PRISM does not define the industry classifications used, it cannot be determined from the documentation how the Utilities and Communications industries are classified.

row's industry last year will move to any given industry this year. The row entries sum to one. Many job changers, however, do not change industry, so the diagonal entries are typically the largest in the row (the exceptions are the federal government and the state and local government sectors).

PRISM

PRISM assumes that any individual who is assigned zero hours worked in a year terminated the job they held the previous year, if they worked in that year. Workers assigned zero hours worked in the previous year and nonzero hours worked this year are entrants. Workers assigned non-zero hours in both years are eligible to change employers. Job change probabilities for these workers depend on age, full-time vs. part-time status, and job tenure.[34]

Individuals who are selected to change employers during a year are assigned a (possibly new) industry using inter-industry transition probability tables. Four separate tables were constructed by age (under 35 and over 35) and sex. These tables were estimated using Janaury 1978 CPS data. Inspection reveals that the diagonal elements are much larger than the corresponding elements in the DYNASIM inter-industry transition tables. New entrants are assigned industries using probabilities which depend on age, education, sex, and for individuals age 25 and over, part-time vs. full-time status.[35]

COMPARISON

The remarks in the previous section concerning the relation between labor force behavior and retirement income are equally applicable here.

Retirement benefits are closely linked to job tenure with a given employer, so realistic simulation of change of employer is important. Industry change is relevant here also, since the types of plans offered and the average coverage rates vary greatly from industry to industry.

The structure of DYNASIM and PRISM are very similar for job change and industry of employment, but there are minor differences in the variables included in certain relationships. PRISM breaks down inter-industry probabilities into two age categories for each sex, and breaks down assignment of entrants by age and full-time vs. part-time status while DYNASIM does not. PRISM disaggregates the self-employed as a separate industry while DYNASIM does not, but DYNASIM disaggregates the communication and utilities industries while PRISM does not. DYNASIM allows the current industry of employment to affect the probability of job change.

Probably the most crucial difference between the two models is in the construction of the inter-industry transition probability tables. The tables are significantly different. For example, in DYNASIM, the probablity of a male job changer staying in the manufacturing industry is 0.527, while in PRISM it is 0.7732 for men under age 35 and 0.7354 for men over age 35. The probability of a female job changer staying in the federal government is 0.200 in DYNASIM and is 0.7505 and 0.8187 for females under and over age 35, respectively, in PRISM. In general the diagonal elements are much larger in PRISM than in DYNASIM. The DYNASIM inter-industry transition probabilities are based on work by James Schulz,[36] and were estimated from a January 1973 CPS data set describing 3331 workers who changed jobs between January 1972 and January 1973, and were employed at both dates. The PRISM transition probabilities were derived from the January 1978 CPS data, describing workers who changed jobs between January 1977 and January 1978, and were employed at both dates. Besides using different data sets, nothing in the documentation of either model indicates any particular reason for the substantial disparity in the resulting transition probabilities; both apparently set the probabilities equal to observed transition rates.

Because of the difference in the size of the diagonal elements, PRISM can be expected to simulate job changers to stay in their industry much more often than DYNASIM. Both models control aggregate employment by industry to match given time series so that the differences in inter-industry job transition probabilities will not be responsible for differences in the predicted distribution of employment across industries. (The adjustment is made to the number of new entrants in an industry.) However, these probabilities clearly can have an effect on patterns of job tenure by industry, and on a given individual's job history. This difference will directly affect retirement income predictions because in certain industries, such as construction and mining, multi-employer plans are prevalent, but the group of employers participating in the plan are generally in the same industry. Both models assume that multi-employer plans are only provided within given industries. Thus if one model "under-predicts" job changers staying in these industries, it will underpredict enrollment in multi-employer plans. For this reason, the difference in the inter-industry transition matrices implies that PRISM will simulate more retirement income from multi-employer plans than DYNASIM will. However, without scrutiny of independent data on inter-industry employment transitions it is impossible to determine which model is more realistic in this regard.

5. PENSION AND SOCIAL SECURITY BENEFIT ACCUMULATION

5a. Private Pension Benefit Accumulation

DYNASIM and PRISM differ substantially in their simulations of the accumulation of private pension benefits. Both of the models utilize an industry-by-industry method of projecting coverage rates (the percentage of an industry's employees covered by a pension plan), although their coverage rate assumptions are somewhat different. However, both the simulation of: 1) plan participation (the employee's decision on whether or not to participate in an available plan), and 2) the specification of plan characteristics, are very different in the two models. DYNASIM follows a probabilistic approach, randomly assigning benefit formulas and coefficients to individuals using probabilities derived from data in the May, 1979 CPS, in which separate questions on private pension participation and plans were included. In effect, data from the CPS reflect the allocation of participating individuals to various types of plans--given the characteristics of individuals covered by pension plans, the probability that any individual with given characteristics will be participating in any given plan can be calculated. PRISM, on the other hand, draws on data from an ICF employer survey describing the characteristics of actual pension plans. In effect, PRISM assigns individuals to specific plans and then employs the characteristics of the plan to determine eligibiity and benefit levels for those individuals.

In principle, the two approaches could yield comparable results. However, because of the enormous difference in both the basic data on which they rely and on the simulation logic employed such similarity is unlikely. Given the absence of documented real world results against which to compare simulation results, and the lack of funding required to obtain simulation runs of the two models which would enable comparison of the effects of the data and logic, we will rely on a description of the two procedures and present a judgmental evaluation of the impact of the differences on the final estimate of pension benefit accumulation.

5a. (i) Pension Coverage

DYNASIM

In DYNASIM, complete lifetime demographic histories and labor force participation, hours worked and wage rates for each family member are first simulated in the Family and Earnings History Model. Then, in the Jobs and Benefit History (JBH) Model, job change and industry of employment are simulated. Using this job mobility and industry information, workers are assigned a pension plan (no plan is an option). Workers who change jobs or who enter the labor force in a given year are first assigned to covered or non-covered status. The probability of being covered varies by industry (see Table 5 for a listing of industries), real earnings, and sex.[37] As noted above, these probabilities are derived from the May, 1979, CPS. Workers maintain their assigned covered or noncovered status until they change jobs. Moreover, industry-specific coverage probabilities do not change over the forecast horizon.

PRISM

PRISM is structured similarly in simulating coverage rates. Job changers or new job holders are probabilistically assigned to covered or noncovered status. The probabilities--also derived from the May, 1979, CPS--depend upon industry, hourly wage, and age.[38] And again, workers maintain their coverage status as long as there is no job change.

However, in PRISM these coverage rates are permitted to change over time. The model user specifies industry-specific rates of growth in coverage. These, in turn, are applied to 1979 coverage rates by industry and age group to determine the expected future path of coverage rates by industry. In each year of the forecast (and for each industry) the expected coverage rate is compared with the actual coverage simulated for the year before and an adjustment factor is calculated. This adjustment factor is then used to alter the coverage rates applied to both job entrants and job changers (but not the rates applied to individuals staying on the job held last year). The coverage growth rate assumptions in the PRISM simulations vary from zero to 0.87 percent per year.

COMPARISON

In comparing the treatment of pension coverage in DYNASIM and PRISM, the most prominent difference is the feature which allows for annual growth in coverage rates in PRISM. If the user specifies any growth above zero, PRISM will, ceteris paribus, project higher private pension benefit levels than DYNASIM. Moreover, with any specified growth in coverage rates, the divergence in pension benefits between the two simulations will be larger for cohorts retiring later than for those retiring earlier. Moreover, the difference between the two models will be larger for women than men. First, since women experience more job turnover and job entrance than men, the application of the coverage growth rate adjustment to new workers will cause their coverage growth to be larger. Second, coverage growth rates are high in the services and trade sectors, which tend to employ more women than men.

How reasonable are the alternative projections for coverage rate growth? This seems to be a matter over which analysts differ. Alicia Munnell notes that "the widespread introduction of new pension plans between 1940 and 1960 led to a large increase in the percentage of workers covered by such plans.[39] Aggregate coverage rose from 14.6 to 40.8 percent of private nonagricultural wage and salary workers. After 1960 the rate 'increased' very slowly, rising to only 48.1 percent in 1980." This "reflected the growth of employment in firms that already had pension plans, (and) liberalization of qualification for pension coverage ..." (p. 53). She goes on to say

> Because of the influence of industry structure on pension coverage, the percentage of the work force covered by pension plans is not expected to increase significantly in the future. Industries with traditionally high pension coverage, such as manufacturing, are expected to employ a declining share of workers while employment in industries with low pension coverage, such as retail trade and services, is projected to increase. Moreover, small businesses, which employ the bulk of noncovered workers, are unlikely to adopt pension plans (p. 200).

Other analysts have projected significant coverage rate growth in the aggregate. Sylvester Schieber and Patricia George argue that the slowdown in the growth of coverage rates which occured after 1960 was due to certain temporary influences.[40] They point out that the growth in the numbers of workers covered was very similar in the three decades between

1950 and 1979; 8.9 million between 1950 and 1960; 7.4 million between 1960 and 1970; and 9.1 million between 1970 and 1979. They attribute the slowdown in the growth of the coverage rate in the late 1960s and the 1970s to the rapid growth of the labor force associated with the aging of "baby boom" cohorts and the entrance of women into the workforce. They argue that these developments reduced growth in the coverage rate because young workers have below average coverage rates and because female workers are more likely to be employed part-time or in the trade and services industries, both traditionally low coverage categories. They go on to argue that since these trends are expected to be reversed or slowed, the slowdown in coverage rate growth is temporary and further significant increases in the coverage rate can be expected.

Thus various analysts have different views of the future path of private pension coverage rates. Because an evaluation of the creation and growth of private pension plans is beyond the scope of the present report, we cannot assess the accuracy of either model's coverage rate assumptions.

However, as indicated above, we can identify the direction of the difference in retirement income projections caused by the different coverage rate assumptions. PRISM assumes higher future coverage rates and this clearly implies a higher forecast of retirement income than DYNASIM; ceteris paribus. The magnitude of the difference implied by the different coverage rate assumptions could be gauged if results from PRISM assuming no growth in coverage rates were available.

5a. (ii) Plan Assignment and Benefit Calculation

DYNASIM

In DYNASIM, private pension plan type and pension benefits are assigned as follows: First, a covered worker is selected to either participate or not participate in a private pension plan. Second, for those who participate, a particular plan type is assigned. Third, the type of retirement benefit eligibility--normal, early or special retirement--is determined. Fourth, workers not eligible for one of these retirement benefit options are assigned a vesting status. Fifth, within each plan type a particular benefit formula is assigned. Finally, the calculated benefits are adjusted for maximum years of service limitations, minimum benefit limitations, and election of a joint survivor's benefit option.

The probability of a covered worker participating in a pension plan depends upon the number of hours worked, tenure on the job, age, and sex.[41] These probabilities are derived from the May, 1979 CPS and are applied to job entrants and to any covered worker who was not previously selected to participate. All workers selected to begin participation are randomly assigned to one of four plan types--single employer defined benefit, single employer defined contribution, multiple employer defined benefit or multiple employer defined contribution.[42] The probabilities used for assignment to plan type vary by industry of employment, and were derived by ICF from estimates in their survey of the distribution of participants in 1975.

The assignment of a specific benefit formula and the calculation of benefits can occur in two situations. One is when a worker is simulated to experience a change in jobs. In this case benefits are calculated for the job that the individual just left. The other situation in which retirement benefits are calculated is if a worker "contemplates" a retirement decision. This occurs during each year past age 58 and is handled by a separate "retirement module," described more fully in section 6 below. This module requires calculations of the retirement benefits received if the worker retired this year and the benefits received if the worker retired in the subsequent year. This is used as input to the decision of whether to stay in the labor force for one more year. Once a particular plan type and benefit formula is assigned to a worker it remains with the worker as long as they retain that job.

The calculation of benefits begins by determining under which type of eligibility a worker qualifies for benefits. The types of eligibility included in the model are normal retirement, early retirement, and special early retirement (the "30 and out" rule in the automobile industry). First, the worker is selected to be either eligible or not eligible for <u>normal</u> retirement benefits. The probabilities for this selection depend on industry group, years of service, and age.[43] If an individual is not selected for normal retirement eligibility, he/she can be selected to be eligible for <u>early</u> retirement. These probabilities depend on industry, years of service, and age.[44] If a worker is not eligible for normal or early retirement, he or she may be eligible for "<u>special early retirement</u>." A worker is selected for this "<u>30 and out</u>" option if he/she was employed in the automobile industry in the initial survey year (1973) and has had 29 years of continuous service. The automobile industry is not separately disaggregated for simulation calculations, so workers entering the labor force after 1973 are not eligible for special retirement.

For those workers who are not selected to be eligible for benefits under any of these retirement classifications (normal, early or special), vesting status is determined. Workers who have been in an industry for 10 years are simulated to be fully vested or not on the basis of probabilities which depend on their industry. Other workers are simulated to be vested in one of two phased vesting arrangements: "the graded 15-year service rule" (25 percent after five years, with five percent additional vesting for each of the next ten years), or the "rule of 45" (50 percent vesting when age and service total 45 years, with a minimum of five years of service). The procedures for assigning workers between the two arrangements are not specified in the model documentation.

If a worker changing jobs or contemplating retirement is simulated to be eligible for a pension benefit, a particular pension formula is assigned. If the worker qualifies for a defined benefit, a particular defined benefit formula is selected on the basis of probabilities associated with the pension structure of the industry in which the individual is employed, and whether the worker is participating in a single- or multiple-employer plan. The formulas are listed in Table 6.

If the worker is eligible for a defined contribution plan, a prototypical defined contribution plan benefit is calculated. The prototypical plan assumes that the employer's contribution is 7 percent of the employee's salary, and that the accumulated contributions earn 7 percent

Table 6

DYNASIM Defined Benefit Plan Formulas

Formula	Specification
	Single-Employer Plans
1	$b_1 = d * n$
2	$b_2 = a * x * n$
3	$b_3 = a_1 * x * n - a_2 * s$
4	$b_4 = a_1 * \text{MIN}\ [x, c] + a_2 * \text{MIN}\ [x - c, o]$
5	$b_5 = \text{MAX}\ [b_1, b_2]$
6	$b_6 = \text{MAX}\ [b_1, b_3]$
7	$b_7 = \text{MAX}\ [b_1, b_4]$
	Multiple-Employer Plans
1	$b_1 = d * n$
2	$b_2 = a * x * n$
3	$b_3 = \text{MAX}\ [b_1, b_2]$

b_i is the annual benefit amount under formula i.

n is the number of years of service.

x is the average of the highest five of the last ten years' earnings.

d is a dollar amount coefficient (selected randomly).

a, a_1, a_2 are percentage coefficients (selected randomly).

s is the worker's social security benefit.

c is a cutoff for the split percentage formula, b_4 (selected randomly).

Source: The probability distributions for each of the randomly selected parameters were derived from the 1974 BLS Defined Benefit Plan Survey.

nominal interest every year.[45] The amount of the annual benefit payment is determined by the amount of accumulated contributions and specific age-sex-life expectancy assumptions.

Once a particular formula has been assigned to a worker, the parameters of the formulas are randomly assigned on the basis of probabilities which depend on the industry in which the worker is employed, and on which formula is selected. The dollar valued constants in the benefit formulas (d and c) are brought forward to the year for which the calculation is performed in two ways. One method indexes them using the overall consumer price index. The other uses an exogenous user-supplied index. Benefits are calculated using both methods, and both results are saved. All further adjustments are made to both results. Once a particular formula is assigned it remains with that individual as long as he/she holds that job.

Before the final benefit calculation is made, workers are simulated to be subject to a maximum number of years of service limitation on the basis of probabilities which depend on the individual's industry, and on whether a single- or multiple-employer plan applies.[46] If a worker is selected to fall in a limitations category, that limitation governs how many years of service are used in benefit calculations from that job. The benefit is then calculated.

If a normal or an early retirement benfit is calculated, a benefit minimum is assigned on the basis of probabilities which depend upon the industry of the worker. The minimum can be either a flat amount, or an amount conditional on the number of years of service. These probabilities were obtained from the 1974 BLS Survey of Defined Benefit Plans and the 1975 Banker's Trust Survey of Pension Plans.

If an early retirement benefit is calculated, the worker may be subject to a reduced benefit which depends on the difference between the individual's age and 65, multiplied by a randomly selected reduction coefficient. The distributions of these coefficients were derived from the 1974 BLS Defined Benefit Plan Survey.

The calculated benefit may then be adjusted for post-retirement joint survivor's option. Under this option, the spouse continues to receive a benefit if the original benefit recipient dies after retirement. The model assumes that the spouse's benefit is 50 percent of the original benefit. Election of this option may require reduction of the initial annual benefit.[47]

The preretirement survivor's option is assigned in the simulation to 2 percent of eligible workers, and a corresponding reduction is made to their current benefits. This option provides a reduced annual benefit to the spouse even if the worker dies before retirement. Again, no empirical basis exists for the relevant magnitudes.

PRISM

PRISM makes extensive use of the ICF Pension Plan Data Base in simulating pension benefits to individual workers.[48] Job entrants who are selected to be covered by a pension plan (see section 5a(i)) are selected for particular plan sponsors from probabilities implicit in the data base. The sponsor selected for the worker depends on industry and hourly/salaried status. The worker's pension benefits are then determined by the actual provisions of the particular plan which that sponsor provides. In general, a worker terminates coverage under the assigned plan simultaneous with termination from the job. The exception is that individuals covered by a multi-employer plan who change jobs but stay in the same industry are assumed to remain covered by the same plan.[49]

The actual provisions of each of the plans assigned to the worker, along with the simulated history of hours worked, earnings and coverage for the worker, are used to calculate retirement benefits. Different procedures are generally used for the four basic types of defined benefit plans and the defined contribution plans.[50] PRISM assumes that the dollar amounts used in computations of benefits in "unit" benefit plans are indexed to total average wage rates. Other defined benefit plans use some measure of earnings and thus are "self-adjusting" for wage trends. This adjustment only applies before the retirement is accepted; once a worker retires with a pension calculated using a unit benefit formula the pension is set, and is no longer indexed.[51] If the individual retires early, PRISM reduces the benefit as provided in that individual's pension plan. If the plan provides that benefits be offset for social security benefits or that temporary benefits are available between the age of retirement and the age of social security eligibility, such offsets are taken account of in benefit calculations. Disability benefit provisions of the plan(s) are invoked if the individual is simulated to be disabled.

The choice of the post-retirement survivors benefit option, mandated by ERISA, is simulated for married individuals on the basis of probabilities which depend upon sex and the amount of the benefit originally assigned. For benefits less than \$3000, 30 percent of married males and 25 percent of married females select the option; for benefits over \$3000, 70 percent of married males and 50 percent of married females select the option.[52] The model assumes that all individuals who select the post-retirement survivor's benefit receive a "50 percent joint and survivors annunity." Under this plan, the surviving spouse received 50 percent of the initial benefit, and the initial benefit is accordingly reduced to equate the actuarial present value of electing and not electing the option.

PRISM assumes that 75 percent of all eligible married individuals elect a pre-retirement survivor's benefit option in plans which provide such an option. The particular plan's formulas are used to calculate the surviving spouse's benefit if the worker dies before retirement.

PRISM models defined contribution plans by assuming that employers contribute an amount equal to eight percent of the employee's annual salary; information necessary for tying contributions to individual firm profitability is unavailable. Individuals in such defined contribution plans are assumed to accept payment when they terminate the job with which the plan is associated. In certain circumstances accumulated defined contribution plan benefits are "rolled over" into an IRA account: if the value of the accumulated benefits is over $1750 in 1980 dollars <u>and</u> the worker is age thirty or over. Otherwise the individual is assumed to spend the money on current consumption, and it does not show up as retirement income.

The rate at which eligible individuals participate in savings and thrift plans is assumed to vary with the individual's wage rate and the employer "match rate."[53] The probabilities associated with these categories are from the ICF Pension Plan Data base. Individuals are assumed to contribute the amount required to receive the maximum matching contribution by the sponsoring employer.

All individuals who receive a defined contribution plan benefit are assumed to convert the accumulated contributions to an annuity providing annual payment until the individual dies. The annuity is assumed to begin at the earlier of the age at which social security benefits are accepted or the age at which defined benefit plan benefits are accepted, but not before age 55.[54]

COMPARISON

From this discussion, it is clear that DYNASIM and PRISM differ substantially in their modelling of pension benefits. PRISM relies heavily on compiliations of actual plan provisions from their survey of plan sponsors. DYNASIM is designed to reproduce probabilistically the distribution and characteristics of plans in the May 1979 CPS data base. There are, consequently, many reasons why the two models are likely to produce quite different projections of retirement benefit levels. Here, we will attempt to identify the primary characteristics of each which will lead to different estimates between them, and to appraise the relative strengths and weaknesses of each.

The primary differences between the two model's treatment of pension benefits are of three sorts. First, PRISM assigns the provisions of actual plans to participating individuals while DYNASIM independently assigns various plan characteristics to each individual. Thus, PRISM is designed to capture more detail concerning the interrelation of various plan characteristics. Second, PRISM is identifiably more generous in certain arbitrary assumptions than DYNASIM. Third, DYNASIM omits a number of pension plans which PRISM includes.

In DYNASIM, various provision characteristics which are randomly assigned are, in fact, likely to be correlated. While DYNASIM may accurately reproduce the aggregate incidence of various characteristics, the model may underestimate or overestimate benefit levels by not modelling the interrelationship of pension plan features. For example, workers who have ten-years of job tenure and who have not been selected by the simulation to be eligible for normal, early, or special retirement are randomly assigned to a vesting status. The probability of full vesting depends only on industry. If, in reality, liberal vesting is associated with stingy benefits, the DYNASIM method will probably overestimate average benefit levels for this group. In contrast, PRISM assigns a worker to a certain plan and uses all of that plan's characteristics to determine vesting, eligibility, benefit formulas, etc., and in principle should capture any existing correlations among these factors. The direction of bias due to this difference in approaches is very difficult to assess from the model documentations since it depends on the interaction of these correlations with the various nonlinearities in the determination of retirement income.

Because DYNASIM attempts to capture the diversity of pension plan characteristics with a limited number of randomized samplings, the possible variation in each characteristic is necessarily constrained to some unknown extent. For example, as described above, the vesting formula for workers who do not yet qualify for retirement is very simple; a worker is either fully vested after ten years, or not vested at all. Such a simplification does not reflect the actual diversity in vesting provisions. Or, again, while seven formulas are used to calculate the benefits from defined benefit plans, these are a significant simplication of the wide diversity of plans in use. It is very difficult to assess the direction of bias due to these features from merely examining the documentation.

The procedures used for calculating "benefit reductions" in simulating early retirement and the survivors option in DYNASIM are also streamlined. A number of maximum service categories are defined and randomly applied to workers. Early retirement causes the benefit to be reduced by the product of the number of years remaining until age 65 and a random coefficient. The resulting amount may have little relationship to actuarial values. The benefit reduction applied to workers electing the joint survivor's option depends upon the worker's and the spouse's age category, rather than being based on an explicit actuarial calculation, as in PRISM. These simplications do not necessarily imply large inaccuracies--either vis-a-vis reality or vis-a-vis another modelling approach. However, they are a likely source of bias. Again, it is impossible to assess the likely direction or magnitude of the bias without reconstructing the model itself.

Another important area of difference in the benefit calculations performed in the two models involves certain numerical assumptions made. One of these is the method of indexing defined benefit formulas. In both of the models some defined benefit formulas apply fixed or varying percentages to an average of recent earnings. These formulas adjust benefit levels automatically for inflation, since the earnings to which they are applied increase over time as individual wage rates increase. Other formulas apply a coefficient to the number of years of service the worker has accumulated. This coefficient gives the dollar amount of benefit per year of service. Benefits calculated using this type of formula do not adjust for inflation unless the coefficients themselves are adjusted for inflation. This category of defined benefit plans is not insignificant in the two models; in DYNASIM the percentage of single-employer defined benefit plan participants assigned to a formula which uses a dollar-valued coefficient is 98 percent in construction, 73 percent in manufacturing, 66 percent in mining, and 56 percent in transportation. All multiple-employer defined benefit plans involve dollar-valued constants. Single-employer defined benefit plans account for between 20 percent (finance) and 67 percent (manufacturing) of all pension plan participants in DYNASIM.

PRISM indexes all of the dollar-value coefficients for the defined benefit plans in their pension plan data base using the economy-wide average wage rate. DYNASIM, in contrast, indexes these coefficients using an overall price index. Since both models assume positive productivity growth, the dollar-valued coefficients in PRISM will grow faster than those in DYNASIM. (The inflation and productivity assumptions of each model are described in section 7 below.) In effect, DYNASIM assumes that the coefficients are constant in real terms, while PRISM implicitly assumes that the real value of the coefficient is proportional to average real wages. The difference between the two models in the resulting coefficients--approximated by the overall rate of productivity growth in the economy--will be large and will grow over time. For example, for a productivity growth rate of 2 percent per year, the coefficients in PRISM will be 24 percent higher after 10 years, 49 percent higher after 20 years, and 81 percent higher after 30 years. For a productivity growth rate of 3 percent per year these figures are 34 percent, 81 percent, and 143 percent, respectively. Thus, this is a major and quantitatively significant methodological difference between the two models.

Another significant difference between assumptions of the two models has an identifiable but less significant effect on the estimates of pension levels. For defined contribution plans, PRISM assumes that employer contributions equal eight percent of the employee's salary. In DYNASIM this rate is seven percent. This implies that retirement income from defined contribution plans is 14 percent higher in PRISM than in DYNASIM, ceteris paribus.

A similar difference between the two models involves election of the post-retirement joint survivors option. DYNASIM assumes that 75 percent of eligible married males and 25 percent of eligible married females elect this option. In PRISM, different percentages are applied depending on the size of the benefit; for benefits less than $3000, 30 percent of males and 25 percent of females elect; for larger benefits, 70 percent of males and 50 percent of females elect the joint survivors option. Thus, in DYNASIM more males but fewer females elect the option. Recall that selecting this option usually lowers the current annual payment to the vested spouse while still alive. For surviving married couples DYNASIM thus understates the current benefits males receive and overstates the current benefits females receive, relative to PRISM. The opposite is true for surviving widows and widowers. This potential bias is somewhat different from others noted here, in that it affects only reported annual pensions being received, not the actuarial value of benefits.

There is a third potentially significant methodological difference between the two models. A number of pension benefit plans modelled by PRISM are omitted by DYNASIM. One set is public sector pension plans. DYNASIM does not assign pension plans to federal, state or local government employees; PRISM assigns the Civil Service plan to federal workers and assigns survey-based state and local government plans to state and local government workers.

A more general problem in DYNASIM is the omission of supplemental pension plans. Many plan sponsors in PRISM offer each covered and eligible employee a regular plan and a supplemental plan. The supplemental plan is usually a voluntary savings or thrift plan in which employers match a certain percentage of employees' contributions. PRISM assumes that roughly one-half of those eligible for a savings plan participate, and that those who participate contribute enough to receive the maximum matching contribution from the employer. Data on savings plan participation and contributions are scarce; hence, rather arbitrary participation assumptions are difficult to defend.

While DYNASIM makes similar assumptions for those employees selected to participate in savings plans as their primary pension plan, DYNASIM does not provide supplemental plans to be chosen by individuals already assigned to primary plan. Thus DYNASIM never provides an employee with two pension plans from the same employer, while in many cases PRISM does. For those individuals who are simulated by PRISM to participate in supplemental savings plans, this difference could be quite large.

Another omitted source of pension benefits in DYNASIM are the Keogh plans for self-employed individuals. PRISM, on the other hand, is able to provide explicit estimates of this source of retirement benefits, since the self-employed are treated as a separate industry.

In summary, the modelling of private pension benefits is identifiably more generous in PRISM than in DYNASIM--ceteris paribus, PRISM will simulate larger retirement incomes from private pensions than DYNASIM. This occurs for three primary reasons. First, as was detailed in the previous section, PRISM assumes increases in pension plan coverage rates in many industries, whereas DYNASIM assumes that these rates remain constant. This is essentially an issue of comparative prognostication about which various analysts have disagreed. Second, certain key assumptions are more generous in PRISM: dollar-valued constants in defined benefit formulas are indexed to wage rates in PRISM and prices in DYNASIM; and employer contributions to defined contribution plans are larger in PRISM. The former is likely to be more significant. Third, PRISM includes models of many pension plans which DYNASIM omits: federal, state, and local government worker pensions, self-employed person's pensions, and employer-sponsored supplemental savings plans.

Only in this third area--the omission of certain plans--does the direction of bias between the two models correspond to a bias in one model (DYNASIM) relative to "reality." In the other areas, the two models differ in an identifiable direction, but which more closely approximates reality is very difficult to assess.

5b. Social Security Retirement, Disability, and Survivor Benefits

While DYNASIM and PRISM are similar in their modelling of social security benefits, there are important differences. Each begins the simulation with the actual earnings and coverage histories of the individuals in its sample up to the initial year of simulation. Each model simulates work and earnings stories for each individual through their (potential) retirement age. At retirement age, then, each sample individual has a complete actual plus simulated earnings and coverage history. Both models then calculate and assign Social Security retirement benefits, spouse's benefits (including survivor's benefits) and disability benefits. DYNASIM also calculates children's benefits. Since PRISM only estimates retirement benefit levels, the omission of children's benefits does not affect their results.

DYNASIM

DYNASIM creates an entire work history for a family before retirement income is simulated. The social security model starts by calculating retirement benefits for heads of households who have reached the minimum retirement age (62) and have the required number of quarters of covered earnings. First, the individual's record of quarters worked and covered earnings is checked to determine insured status.[55] Then, the individual's potential primary insurance amount (PIA) is calculated. The pre-1979 formula and the wage-indexed formula established in the 1977 amendments are both employed along with the grandfathering provisions that apply. The wage-index time series used in the latter calculation must be provided by the user. The third step in calculating retirement benefits is to assess the adjustments to the PIA required if an individual is over or under the normal retirement age (currently 65). Finally, the earnings test is applied by determining if an individual's simulated current earnings disqualify them for retirement benefits.

In practice, individuals are eligible for disability benefits if they are permanently and totally disabled, and after a five-month waiting period after applying. The disability indicator simulated by DYNASIM only shows whether the individual has a condition which "prevents or limits" work (see section 3d. above), and only on a calendar year basis. In DYNASIM individuals are eligible for disability benefits if the disability indicator occurs in two consecutive years and the individual is simulated to have zero earnings in both years. This is designed to capture only those disabilities which prevent employment and to approximate the waiting period requirement. The calculation of the actual amount of disability benefits follows that of social security retirement benefits.

The accrual of spouse benefits to husband's is "relatively rare," and in DYNASIM, only women can receive spouse benefits. Spouse benefits are calculated only after the benefit history of the household head is simulated. A married woman can receive a wife's benefit, a survivor's benefit, or her own retirement benefit.[56]

PRISM

PRISM has a substantially more streamlined approach to simulating social security benefits. The retirement benefit calculation is as specified in the 1977 amendments, except that the grandfathering provision for individuals whose first year of retirement is 1979 through 1983 is not included. This will not affect cohorts retiring in 1984 or beyond. The calculation of PIA, the adjustment for early or late retirement, and application of the earnings test all follow Social Security provisions, in much the same way as does DYNASIM. The formulas are indexed over time as provided in the 1977 amendments, again with a user provided deflator.

An individual receives Social Security disability benefits if they are simulated to be disabled according to PRISM's disability definition, and if they meet the applicable criteria for the number of quarters covered. Recall that in PRISM an individual is defined as disabled if they qualify as totally disabled by Social Security Administration standards and receive benefits (see section 3d.). Thus the disability indicator simulated by PRISM corresponds directly to qualification for and acceptance of Social Security disability benefits. Note that this indicator is actually the conjunction of participation and qualification, rather an indicator of physical condition only.

Benefits for dependent spouses are calculated taking into account the spouse's PIA, adjusting for early or late retirement, and applying the earnings test. Survivors benefits require calculating the benefits for the deceased spouse as well as the benefits for which the survivor qualifies on his or her own. The benefit received by the survivor depends on the difference in the two, adjusted for early or late retirement, and is subject to the earnings test.

COMPARISON

Since both models set out to replicate the same well-defined social security benefit formulas, and have available the same information on each individual, there are few important differences between them. One significant difference is the procedure for determining eligibility for Social Security disability benefits. PRISM has a clear advantage in that the occurrence of disability status which qualifies for Social Security disability benefits is modelled directly. DYNASIM, on the other hand, relies on self-reported work impairing disability status, and then models the relation between this indicator and actual acceptance of Social Security disability benefits. This relation is complicated, and involves covered earnings criteria, differences in the definition of disability, the waiting period stipulation, and the individual's decision to apply for benefits. DYNASIM adopts an admittedly imperfect procedure for translating their disability indicator into acceptance of benefits, a procedure which the documentation claims will err on the side of underestimating the number of recipients. This is an identifiable source of bias, both relative to PRISM and relative to an unbiased forecast.

The reason for this difference deserves comment. The procedure used in each model was probably determined by the data sets chosen for use in estimation of the models. PRISM's disability rates are from a Social Security Administration Actuarial Study published in 1980, reporting data for 1974 through 1978. The rates used in simulation only vary with age, sex, and the number of years already disabled. DYNASIM estimates the equations predicting disability status on PSID data for 1969 through 1972. These equations depend on age, race, sex, and marital status. As we discuss elsewhere, DYNASIM relies heavily on PSID data, while PRISM make relatively little use of this data. This difference, in turn, reflects the modelling strategies embodied in each model. DYNASIM is designed to capture a broad range of socioeconomic interactions and takes advantage of the detail available in the PSID. PRISM, on the other hand, is more narrowly focused on retirement income issues, and is more parsimoniously specified with respect to certain demographic and socio-economic individual characteristics.

5c. Supplemental Security Income (SSI)

SSI benefits are available to individuals who are 65 or older or disabled and whose income and asset fall below certain specified levels. Benefits are equal to the amount by which "countable" income falls short of a national income floor. In addition, some states provide supplemental SSI benefits, either by applying a higher income floor or through different eligibility requirements.

DYNASIM

DYNASIM first certifies eligibility for all individuals who are simulated to be "disabled" (see section 3d.),[57] or who are age 65 or over and who meet income tests. This is the same disability criteria used in the Social Security module. In addition to the disability test, DYNASIM requires the individual to have annual earnings below $1680 (in 1978 dollars). The unreliability of asset information leaves the model unable to apply the asset tests. The benefit provided is the difference between the national income floor (including any state supplement) and the individual's "countable" income.[58]

DYNASIM does not simulate state of residence but does simulate region of residence (North, South, Central, West; see section 3f.). In order to estimate state supplemental benefits, individuals are randomly determined to be eligible for state supplemental benefits, using probabilities which vary by region. Those who qualify are assigned the average state supplement for that region. These average benefit supplements are indexed to the CPI.

Participation of eligible individuals in the SSI program is simulated using participation probabilities which depend on the annual benefit levels for which the individual would qualify. These probabilities are .25 for those with simulated benefits of less than $500, .50 for those with benefits between $500 and $1500, and .75 for those with benefits over $1500 (in 1974 dollars). This pattern of participation rates corresponds with those estimated in the research literature.

PRISM

PRISM determines eligibility for SSI benefits on the basis of age, income, and assets, ignoring disability. Rather than simulating individual asset accummulation, probabilities are applied to determine whether the asset test is satisfied or not. These probabilities are derived from March 1980 CPS data on individual income from various types of assets. The average rate of return on these assets for 1979 is then applied to construct estimates of asset holdings. These holdings are, in turn, used to estimate the proportion of individuals who pass the SSI asset eligibility test by the age of the family head, family income, sex and marital status.[59]

PRISM uses individual earnings, and individual pension and Social Security income, to calculate the federal benefit using the statutory formula. State supplemental payments are simulated for the 13 states which provide supplements by using a higher income floor than the federal benefits. The appropriate state maximum benefit is applied instead of the federal maximum for individuals residing in these states.[60] Individuals are assumed to reside for the entire simulation in the state in which they resided in 1979.

Participation rates in PRISM vary by size of the monthly benefit and by whether it is an individual or a married couple filing.[61] These rates are between 0.37 and 0.46, except for individuals qualifying for monthly benefits of $220 or more, in which case a rate of 0.74 is applied.

COMPARISON

There are numerous differences between DYNASIM and PRISM in modelling SSI benefits, but the net effect of these is impossible to assess from the documentation.

DYNASIM allows some individuals to be eligible for SSI benefits on the basis of the self-report disability indicator described above, while PRISM ignores the disability component of SSI. DYNASIM is to be preferred here since it simulates at least some SSI disability recipients. PRISM, however, applies a stochastic simulation of the asset test, while DYNASIM ignores assets. Although asset accumulation is quite difficult to model in a microsimulation framework, some attempt to capture this means of disqualification is desirable. In both of these areas DYNASIM will yield higher benefit estimates. In the case of disability benefits, DYNASIM's higher benefits are more realistic; in the case of the assets test, PRISM is more realistic.

Because PRISM models state supplements separately for each state, it seems more likely to capture the distribution of state supplemental benefits across individuals than DYNASIM, which assigns selected individuals the average state supplement for the region. PRISM omits a small fraction of the state supplements--those not implemented via an income floor above the national floor--but these only amount to 3 percent of all SSI benefits. On the other hand, DYNASIM explicitly models interregional migration, while PRISM assumes people stay within their 1979 state of residence (see section 3f.). This may be responsible for some bias since state supplements vary significantly by region. The direction of bias is uncertain a priori, however, since states in regions of expected net immigration as well as states in regions of expected net emigration offer state supplemental programs.[62]

PRISM disaggregates participation rates by marital status, and benefit size, while in DYNASIM participation rates vary only with benefit size. The disaggregation by marital status seems worthwhile, since participation rates vary substantially over marital status.[63] This could significantly affect the final estimates of the relative retirement incomes of married versus unmarried couples in the two models.

In summary, the treatment of SSI benefits seems more generous in DYNASIM than in PRISM. DYNASIM includes SSI disability benefits, although imperfectly, and should be preferred to PRISM in this regard. However, PRISM models the assets test (also imperfectly), models state supplements by state rather than region, and disaggregates participation rates by marital status. These should be counted as advantages for PRISM.

5d. Individual Retirement Accounts

ERISA (1974) allowed individuals who were not active in any qualified retirement plan or government retirement plan to establish tax deductible retirement savings accounts known as individual retirement accounts

(IRA). Contributions were limited to the smaller of $1500 or 15 percent of an individuals' earnings. In 1981, ERTA altered some of these provisions. All working persons up to age 70 1/2 were eligible to establish IRAs, and the allowable annual contribution was set at the smaller of $200 or 100 percent of an individual's earnings. In addition, nonworking spouses can contribute $250 per year.

DYNASIM

In DYNASIM, an IRA is treated as another pension plan. Workers are randomly selected to establish an IRA when they enter or leave the labor force or when they change jobs. During the years 1975 through 1981, probabilities, which depend upon income and sex and are derived from the May 1979 CPS, are applied to workers not covered by an employer plan; from 1982 on they are applied to all workers. Individuals selected to establish IRA's are assumed to contribute the allowable maximum, including the allowed contribution of $250 by a nonworking spouse. Contributions accumulate at an interest rate which is user supplied. When a worker changes jobs, the account is treated as a vested benefit, and accumulates interest until distributed. The interest rate used is not specified in the available documentation.

Distribution is allowed as early as age 59 1/2 by law. DYNASIM treates IRA distributions as a potential retirement benefit when simulating the retirement decision. Once an indiviudal is selected to retire, the account is paid out in annual amounts equal to the total value of the account divided by the individual's life expectancy. The account ceases to earn interest once the individual retires. Any funds remaining in the account at death are paid out similarly to the surviving spouse. Life expectancy figures vary by age, race, and sex.

PRISM

The simulation of IRA's is somewhat more streamlined in PRISM. During each year of the simulation workers who have not already done so are randomly selected to establish IRA's on the basis of probabilities which are derived from IRS and May 1979 CPS data and which depend upon age and earnings.[64] Once a worker has established an IRA, they are assumed to contribute to the account at annual rates that vary with earnings and whether or not a nonworking spouse is also contributing. These rates vary between 70 and 95 percent of the maximum allowed annual contributions, and were estimated using IRS data. Individuals are assumed to receive distributions from the accounts as annuities beginning at the earlier of the age they first accept Social Security benefits or the age they first receive a defined benefit pension (but not before age 60).

COMPARISON

It is difficult to compare the rates at which individuals establish IRA's in DYNASIM and PRISM, since this decision is simulated at different points in each model. Once established, DYNASIM simulates more generous contributions to the accounts. On the other hand, the formula for paying out the account, as described in the documentation of DYNASIM, is not equivalent to an annuity--interest earnings after retirement are neglected, and the account value is divided by life expectancy rather than equated with a payment stream of uncertain duration. For both of these reasons, DYNASIM would tend to understate retirement income from IRA's, relative to PRISM. However, because account contributions in DYNASIM are uniformly larger, the net effect on future retirement income from IRAs is unclear. And because there is so little evidence on the establishment of IRAs under the new regulations, it is unclear which model is more realistic.

6. RETIREMENT AND BENEFIT ACCEPTANCE

DYNASIM and PRISM are very different in their approach to simulating the decision to depart permanently from the labor force and the decision to begin accepting retirement benefits. DYNASIM relies on two sets of estimated equations to determine when an individual over the age of 58 decides to leave a job and then, contingent on that decision, to leave the labor force. This then determines whether the individual begins drawing a private pension or Social Security retirement benefits. The individual is confronted with this decision annually after age 58. PRISM simulates directly the decision to accept pension or Social Security benefits and this determines labor force participation. Distinctions are made in PRISM between the benefit acceptance rates of workers of different vesting and eligibility status, whereas in DYNASIM these factors enter along with others in the estimated equations.

DYNASIM

The DYNASIM model simulates a complete labor force history for each individual, prior to simulating his/her benefit history (see section 2). After these simulations are performed the individual's record is reexamined in a retirement module at age 58 and each year thereafter. This module uses equations estimated from the "Retirement History Survey"[65] to determine the year in which individuals choose to retire. One set of equations specifies the probability of leaving a job in a given year. A second set specifies the probability of leaving the labor force, i.e., retiring, given that an individual has left a job. Together these equations determine the retirement behavior of individuals age 58 and older. These equations are applied each year to every individual aged 58 and over who is in the labor force. The retirement module essentially resimulates labor force participation for older individuals. Before the module is invoked each individual already has a complete employment history with an implicit retirement date; the retirement module alters this employment record. The employment simulation of the Jobs model (see section 4a.) is based on equations estimated on PSID data. Thus the retirement module is added on to allow the employment record of individuals age 58 and over to reflect the behavior uncovered in a more recent data set.

In the survey data on which the equations are estimated, respondents were interviewed at two year intervals. The estimated equations give the probability of job change and labor force exit over two year intervals. These are converted to one-year rates by assuming that the probability of remaining on a job or in the labor force is constant during each of the two years. Therefore, for the job change probability:

$$\text{P(job exit in 1 year)} = 1 - \text{P(job stay for 1 year)}$$

$$= 1 - \sqrt{\text{P(job stay for 2 years)}}$$

$$= 1 - \sqrt{1 - \text{P(job exit in 2 years)}}.$$

The probability of leaving a job within two years is estimated as a "logit" function with the following set of independent variables:

1. disability status (as described above),
2. marital status (married, not married),
3. "full employer pension benefit status" ("yes" and "otherwise"),
4. "reduced employer benefit status" ("yes" and "otherwise"),

5. "Social Security and full employer benefit status" ("both" and "otherwise"),

6. "Social Security and reduced employer benefit status" ("both" and "otherwise"),

7. Social Security eligibility status (yes or no),

8. wage rate,

9. earnings,

10. loss in Social Security wealth if retirement is delayed one year,

11. loss in employer pension wealth if retirement is delayed one year,

12. Social Security wealth, and

13. employer pension wealth.

Separate equations are estimated for the age groups 58 to 60, 61, to 64, and 65 to 67. The equations for the probability of a job changer accepting a new job (as opposed to retiring) are estimated over the same age categories and depend upon the same variables.

The Social Security wealth and employer pension wealth variables are calculated as the expected value of the discounted sum of the annual expected pension payments. The loss of wealth variables are the differences in these sums calculated for successive retirement years.

If these two equations jointly determine that a particular individual has retired, the individual's earnings are zero in subsequent years. This triggers the acceptance of Social Security benefits via the formula for calculating retirement benefits; since the individual has withdrawn from the labor force, the earnings test will clearly be passed.

The break in service implied by retirement triggers the simulation of an employer pension benefit for the individual. The timing of the acceptance of retirement benefits is thus modelled indirectly by first modelling labor force participation (and, implicitly, retirement) and then modelling retirement benefits with the retirement decision given.

PRISM

In the PRISM model, the retirement decision is essentially equivalent to the decision to accept pension benefits--either Social Security benefits or employer pension benefits. First, the model determines whether an individual would be eligible for Social Security benefits if retired by applying the eligibility criteria appropriate for that cohort to that individual's earnings history. Then, individuals who are eligible, but who are not yet receiving benefits, are randomly selected to accept benefits on the basis of probabilities which depend on age and sex. These probabilities are designed to replicate actual acceptance patterns and are obtained from Social Security Administration data.

Exceptions to this procedure are applied to the disabled, surviving spouses, and the unemployed. Disabled individuals are assumed to accept benefits in the first year in which they become eligible. Surviving spouses are assumed to accept the survivor's benefits as soon as the spouse dies and they reach age 62. An eligible individual who is unemployed and age 65 and older, or who is unemployed and receiving a private pension, is assumed to accept Social Security benefits during that year.

Individuals who are vested and eligible for retirement benefits under private pension defined benefit plans are selected, as soon as benefits are available, to possibly accept benefits and retire. If these individuals have already been selected to receive Social Security benefits, they are automatically retired and the simulated employer pension benefit is assigned to them. If they have not yet been selected to receive Social Security benefits, retirement probabilities are applied which depend on age and sex.[66] Separate probabilities are applied to private, state and local, and federal workers. These probabilities are obtained from the May 1979 CPS, and are designed to reproduce actual private pension acceptances.

If an individual is vested but not yet eligible for retirement, they are assumed to begin receiving benefits at that plan's normal retirement age or at age 65, which ever occurs first. This selection mechanism for vested, eligible, defined benefit plan workers overrides the usual job change simulation, described in section 4a., implying that vested, eligible defined benefit plan workers do not leave their job until they accept the pension benefit.

In simulating the number of hours worked by an individual in a given year, PRISM treats separately individuals who have been selected to receive retirement income (see section 4a.). In effect, PRISM first determines whether an individual starts receiving retirement income--private pension or Social Security benefits--and then determines whether or not they work. Thus in PRISM, the retirement decision is embedded in the decision to accept retirement benefits.

COMPARISON

DYNASIM and PRISM differ substantially in their approach to modelling the benefit acceptance and retirement decisions. DYNASIM models retirement explicitly and allows this decision to affect benefit acceptance through its benefit computation and assignment routines. PRISM models the retirement decision implicitly in its treatment of the benefit acceptance decision. Since the focus of our comparison of these two models is the projection of retirement benefit levels--that is, the income of a cohort at some future point in time--it seems appropriate to view the modelling of retirement itself as an instrument in securing this income estimate. Projections of future patterns of retirement may be interesting in their own right, and DYNASIM's retirement equations seem to have been designed with this type of purpose in mind. However, in the context of estimating future pension benefits, the projection of the timing and patterns of retirement benefit acceptance, as opposed to retirement, is the crucial issue. The comments which follow will focus on this aspect of the modelling of the retirement decision.

Perhaps the major difference between the two models in this area is the direction in which causations are modelled. In PRISM, individuals are selected to accept Social Security benefits or private retirement benefits. This then affects their labor force behavior in that year and in subsequent years. In addition, individuals who are vested and eligible in private defined benefit plans, but who have not yet accepted a benefit, have the usual labor force transition model overridden by a special set of retirement probabilities. So in PRISM benefit acceptance determines labor force activity.

In contrast, DYNASIM allows mutual interaction between labor force activity and benefit acceptance. The estimated retirement equations determine the labor force activity of individuals over age 58, and this labor force activity (e.g., breaks in service, job change, etc.) is viewed as influencing the decision as to whether pension benefits are accepted or not. (Note that DYNASIM only allows acceptance of benefits at a break-in-service.) The individual's pension vesting status--full versus partial vesting for private pensions and "vested" or not for Social Security--enters into the regression equation to influence the retirement decision. In addition, the effect on potential pension benefits, both private and Social Security, of delaying retirement one year enters through the pension and Social Security wealth-loss variables in the retirement equation.

As noted above, interrelations between two given variables can, in principle, be modelled equally well by assigning variable x randomly and capturing the effect of x on y, or vice versa. Therefore, we have no a priori reason to prefer one method of capturing the relation between x and y to the other in the present context. In this sector of the models, it is very difficult to uncover a potential relative bias in one model or the other, given the documentation which is available. Determining the direction and extent of differences would require extensive detail from simulation runs, which detail is not available.

7. LINKAGES TO MACROECONOMIC TIME SERIES

Both DYNASIM and PRISM attempt to link the simulation of individual behavior to time series aggregates. Some of the linkages employed consist of specifying exogenously the values of individual economic variables, such as the nominal interest rate. Other, more basic linkages involve imposing constraints on the aggregate of the outcomes of simulated individual behavior. This section will systematically describe these linkages, although several aspects of them have been mentioned above.

Each of the models imposes an important set of macroeconomic linkages by imposing constraints on the characteristics of the aggregate labor force over time. However, DYNASIM and PRISM take very different approaches in setting these constraints.

Labor Force Behavior in DYNASIM and PRISM

As we indicated earlier (section 4.), DYNASIM simulates labor force behavior via four sets of estimated equations, in each case dependent upon sets of individual characteristics. The first equation determines whether or not an individual participates in the labor force in a given year. The second equation determines the number of hours the individual participates in the labor force. A third equation determines the fraction of those hours which the individual spends unemployed. The second and third equations jointly determine the number of hours worked by the person. A fourth equation determines the individual's wage rate. Given the results of these simulations, DYNASIM calculates the participation rate, average number of hours worked, and average wage rate for each age-sex category, by year.

The model user can alter the time series of these results by setting adjustment factors which enter the individual equations. For each age-sex category there are separate adjustment variables for the participation, hours worked, and wage rate equations. In the hours worked and the wage rate equations the adjustment enters as a scale factor increasing or reducing the result for every individual. The labor force participation equation calculates a number which indicates participation if it is greater than zero and nonparticipation if not. The adjustment to the participation equation enters this equation additively. Hence, changing the value of this adjustment changes the participation status of only some individuals, which in turn alters the participation rate for age-sex category. Notice that the user does not supply the actual target values for these time series, only adjustments to the unconstrained predictions. Presumably the adjustment factors chosen would be designed to constrain results to some exogenously forecasted time series.

In constrast, PRISM explicitly models aggregate labor force participation rates, unemployment rates and average hourly wage rates for 22 separate age-sex groups over the forecast horizon. These projections are produced outside of the micro-data simulation using a "Macroeconomic-Demographic Model" of the U.S. economy.[67] This macroeconomic forecasting model is a modification of the Hudson-Jorgenson growth model which allows substitution among the 22 age-sex categories. The model produces a projection of labor force participation rates and unemployment rates over time for each of the categories. These projections are combined to form employment rates for each category, i.e., the ratio of employment to population in each category. In turn, these employment rates are used as targets and as adjustment factors for PRISM's simulation of labor force behavior. The process is as follows: Using transition probability matrices, the PRISM simulation predicts the number of hours worked by each individual in a given year. Individuals who are simulated to work more than zero hours during the year are counted as employed, and the results are aggregated by age-sex category to obtain a predicted category-specific employment rate. This rate is then compared with the target employment rate for that category for that year, and the discrepancy between the two is used to adjust the probability of working zero hours.[68]

The Macroeconomic-Demographic Model also projects average hourly wages for each age-sex group over time. These category-specific real values are converted to expected nominal values using an assumed inflation rate equal to that incorporated into the alternative II-B assumptions from the 1982 Annual Report of Trustees of the Federal Old-Age and Survivors and Disability Insurance Trust Funds (Washington, D.C.: Social Security Administration). These expected category-specific nominal wage rates are compared with the simulated average wage rates for each category, obtained by applying assumed percentages changes in individual nominal wage rates each year, depending on age, sex, and whether an individual changes jobs or not, and weighting these rates by the number of hours worked by the individual prior to aggregation. (See section 4b.) The ratio between the two category-specific rates is used to rescale all of the individual wages rates in that category. This procedure preserves the relative wage gains implied by the PRISM's individual wage change assumptions, while insuring that the aggregate projections conform to those derived from the macroeconomic model.

Nominal Interest Rate in DYNASIM and PRISM

A second linkage between the models and macroeconomic time series is the nominal interest rate. In both models, this linkage takes the form of an exogenously specified rate, which in both models affects the earnings received on IRAs and on accumulated defined contribution pension plans. In the DYNASIM simulations this rate is 7 percent per year throughout the simulation; the interest rate used in PRISM not specified in the documentation. The conformity of these assumptions to true current and future real interest rates is open to question. For example, the alternative II-B inflation rate assumptions range from 7.9 percent in 1983 to 4.5 percent in 1990 to 4.0 percent in the years after 1990. This implies that the real rate of return on IRA and defined contribution pension plans is very small or negative for most of the first decade. Given the incorporation of these expected near-term inflation rates in the models, the interest rates forecasts would appear to be too low (yielding underestimates of the real benefits from IRAs and defined contribution plans).

Demographic Sectors in DYNASIM and PRISM

A final linkage of the models to macroeconomic time series occurs within the demographic sectors. Both models align their birth rate, mortality rate, and marriage rate projections to the Alternative II-B projections of the Trustee's report. The one exception to this demographic linkage is in the disability rate assumptions. As noted in Section 5b., DYNASIM predicts the occurrence of disability status which is distinct from and broader than eligibility for Social Security disability benefits. Eligibility for disability benefits is then modelled using information from the individual's earnings history. There is no indication that these are aligned with the current disability rates. PRISM, on the other hand, simulates eligibility for Social Security disability benefits by applying actual Social Security disability rates, which are assumed to remain constant throughout the forecast period.

COMPARISON

The linkages between each of the models and various time series aggregates is constrained by the structures of each of the models. Both models simulate wage rates, and can control average wage rates by category. DYNASIM produces separate estimates of labor force participation and unemployment and each can be separately adjusted. PRISM combines these two into a single outcome, employment, and thus can only monitor and control employment rates. But since employment itself is a more important determinant of pension benefits than whether or not an individual is looking for work, this distinction between the two models is probably not critical.

Both models use nominal interest rate and inflation rate assumptions which together imply a low but steadily rising real interest rate. A perhaps more realistic approach would be to set the nominal rate at a constant amount above the inflation rate, with the constant determined by a "modified golden rule"--the productivity growth rate plus the labor force growth rate. This would imply a real rate of return on IRAs and defined contribution plans that is constant over the forecast horizon.

In summary, there is probably little reason to prefer either model's method of constraining the aggregate results of the microsimulations.

8. CONCLUSION

In sections 2. through 7., we have described both the DYNASIM and the PRISM models in detail and have compared and critiqued each model's components. Our purpose is to assess how the large number of identifiable differences in the two models affect estimates of retirement income for future retirees. Table 7 summarizes the primary differences between the two models and describes for each the probable effect on estimates of future retirement income.

While numerous structural and methodological differences exist between the models, three factors appear to be primarily responsible for the discrepancy between the two models in projections of retirement income from private pensions. First, PRISM assumes that private pension coverage rates rise significantly in certain industries over time. DYNASIM, on the other hand, assumes that coverage rates remain constant over time. These assumptions are arbitrary and, in the absence of additional information, there is little reason to prefer one set over the other. The fragmentary evidence which does exist suggests a moderate growth in industry coverage rates over time.

Second, DYNASIM omits supplemental savings plans and Keogh plans for the self-employed. The former are usually lucrative "matching" plans and the latter can account for a substantial portion of retirement income for the self-employed. Since there is little data on savings plan participation, PRISM's assumptions concerning contributions to these plans are essentially arbitrary. Nevertheless, omission of these plans in the DYNASIM models is a distinct deficiency.

Table 7

Summary of Differences Between DYNASIM and PRISM

Model Sector	Difference	Direction of Effect	Size of Effect	Assessment
Demographics				
Mortality	D allows race, education and marital status to influence mortality rates; P does not.	?	Probably very small.	D approach preferred; potentially important in forecasting the income distri- of retired persons but is probably not important for retirement income.
	P uses separate mortality rates for the disabled; D does not.	?	Some effect on disability benefits, recipiency rates; effect on other retirement income probably negligible.	P approach preferred; for forecasts of disability benefits this may be important.
Childbearing	D allows race and education to affect childbearing; P does not.	?	Probably very small.	D approach preferred; potentially important for income distribution of retirees.
	P constrains fertility rates to SSA Alternative II-A assumptions; D does not constrain fertility rates.	?	?	P approach preferred; potentially very important in determining female work histories.
Marriage and Divorce	D allows many individual characteristics to affect marriage and divorce probabilities; P does not.	?	Probably small.	D approach preferred; potentially important for income distribution of retirees.
	P more realistically depicts relative ages of spouses.	?	Probably small.	P approach preferred; potentially important for survivors benefits.

Table 7 (cont.)

Model Sector	Difference	Direction of Effect	Size of Effect	Assessment
Demographics (cont.)				
Disability	D disability definition broader.	Implies lower retirement income estimate for D.	Probably small.	D approach preferred; potentially important for work history simulation.
	D allows race and marital status to influence disability probabilities.	?	Probably small.	D approach preferred; potentially important for income distribution of retirees.
	P allows termination probabilities to depend on the number of years disabled; D does not.	?	Probably very small.	P approach preferred; implies a more realistic intertemporal pattern of disability.
Education	D allows schooling after age 25; P does not.	Implies slightly lower retirement income estimate for D.	Very small.	D approach is more general.
Location	D allows interregional migration; P does not.	Only affects SSI.	Very small.	Allowing migration preferred.
	P carries state of residence but assumes it is fixed.	Only affects SSI.	Very small.	State of residence important for SSI; information worthwhile but cumbersome to incorporate into a migration model.
Labor Force Activity				
Participation, Hours, and Unemployment	D allows many individual characteristics to affect labor force activity; P allows fewer and restricts range of effects.	?	?	D approach has important advantages in capturing covariations in otherwise independent determinants of retirement income.
	P explicitly designed to capture intertemporal pattern of employment.	?	?	P approach has important advantages in simulating job tenure although variance in hours worked is compressed.

Table 7 (cont.)

Model Sector	Difference	Direction of Effect	Size of Effect	Assessment
Labor Force Activity (cont.)				
Wage Rates	D allows many individual effects on wage rates which P omits.	?	?	D approach strongly preferred since P approach sharply constrains variance in wage rates.
Industry of Employment	Interindustry transition matrices are numerically very different, especially in the diagonal elements.	Probably implies higher retirement income estimate for P.	Might be large.	Without independent data we cannot assess the two sets of assumptions.
	P allows age, and full-time/part-time status to affect probabilities; D does not.	?	Probably very small.	P approach preferred.
Pension and Social Security Benefit Accumulation				
Private Pension Benefit Accumulation:				
Pension Coverage	P assumes growth in some industry's coverage rates; D assumes no growth.	Implies higher retirement income estimate for P.	Probably large.	Analysts disagree over prospects for future growth in coverage rates.
Plan Assignment and Benefit Calculation	P assigns actual plans to individuals while D randomly assigns plan characteristics and formulas.	?	?	P uses valuable information not used in D. P strongly preferred.
	D omits supplemental savings plans and Keogh plans which P includes.	Implies higher retirement income estimate for P.	Probably large.	P strongly preferred.
	P indexes some nominal benefit formulas to nominal wages rather than the price level as in D.	Implies higher retirement income estimate for P.	Approximately 50 percent higher after 20 years for those benefit formulas. Those categories are large.	Neither assumption strongly preferred to the other, *a priori*.

Table 7 (cont.)

Model Sector	Difference	Direction of Effect	Size of Effect	Assessment
Pension and Social Security Benefit Accumulation (cont.) Private Pension Benefit Accumulation: Social Security Retirement, Disability and Survivors Benefits	P models SS Disability claimant directly, while D requires ad hoc adjustments to its indicator.	?	?	P strongly preferred.
Supplemental Security Income	P ignores disability claimants of SSI, D includes them.	Implies higher disability benefit income estimate for D.	Probably small.	D preferred.
	P applies simulation of assets test; D does not.	Implies higher retirement income estimate for D.	Probably small.	P preferred.
	P models state supplements using state of residence, while D only carries region of residence.	?	Probably small.	P preferred.
	P disaggregates SSI participation by marital status; D does not.	?	Very small.	P preferred.
Individual Retirement Accounts	Contributions to IRAs are assumed to be more generous in D.	Implies higher retirement income estimate for D.	D retirement income estimates approximately 20 percent higher.	Without a longer historical record under the latest IRA legislation we have no strong reason to prefer either set of assumptions.
	Distributions are annuities in P, but are not exactly annuities in D.	Implies higher retirement income estimate for P.	Probably small.	P preferred.

Table 7 (cont.)

Model Sector	Difference	Direction of Effect	Size of Effect	Assessment
Retirement and Benefit Acceptance	P constrains workers near date of vesting to stay on the job.	Implies higher retirement income estimate for P.	Probably small.	Uncertain: neither preferred *a priori*.
	D allows many individual characteristics to affect retirement decisions, which P omits.	?	Probably very small.	D preferred; might be important for income distribution of retirees.
Linkages to Macroeconomic Time Series	Nominal interest rate and inflation rate assumptions in both models imply low and steadily rising real interest rates.	No effect relative to each other.		Nominal interest rate assumptions should be derived from a real rate assumption and the assumed inflation rate.

Third, PRISM indexes the nominal-valued constants in private defined benefit plan formulas to nominal wages, while DYNASIM indexes these to more rapidly growing price levels. These formulas apply to a substantial portion of private pensions, and the difference implied is large. Because of the more rapid growth of prices than wages, the values for the estimated outlays in defined benefit plans in PRISM would be approximately 50 percent higher after 20 years than those in DYNASIM, given the same initial values. However, without historical time series on these plan parameters we have no way of assessing which assumptions will lead to results which will be closer to actual performance. Of these three factors, only one--the omission of supplemental plans--leads to a judgement that one model (PRISM) is clearly preferable to the other.

In addition to these three factors, there are additional important differences between the models which could contribute indirectly to some of the observed discrepancy in forecasts of retirement income. These differences are chiefly in the modules simulating labor force activity. They would have an effect on retirement income projections through their impact, for example, on average job tenure, or on the variability over time in the number of hours worked by an individual, or in his or her wage rates. While PRISM has been designed to capture these intertemporal aspects of labor force activity, DYNASIM employs a more traditional regression equation approach that is ill-suited to simulating intertemporal tenure patterns. For this reason, it seems likely that DYNASIM exaggerates the year-to-year variation in an individual's labor force outcomes. On the other hand, certain features of PRISM--the compression of hours worked within given hours--worked categories and the wage rate algorithm itself--artificially reduce the year-to-year variation in an individual's labor force outcomes. These are discussed above in Section 5.

Another feature of the labor market component of both models is a set of matrices giving the probability that a job-changer moves from one industry to another. These matrices are very different in each model, leading both to be suspect. For the purpose of projecting retirement income, more work needs to be done developing models of labor market outcomes which are unbiased with respect to the determinants of pensions benefits--i.e., job tenure, variability of hours worked and wage rates, and industry of employment.

Certain general observations concerning the modelling strategies behind each model bear emphasizing. DYNASIM was designed to capture the interactions of many socioeconomic factors influencing individual economic outcomes, and thus contains a very rich set of detail on individual characteristics--for example, race, education, regional location, and disability. PRISM captures some of these influences on labor market outcomes but only in a very streamlined way, necessarily constraining their importance. It is not that information was discarded in the construction of PRISM, but rather that the choice was made to devote scarce degrees of freedom to other uses, such as the intertemporal pattern of hours worked. The decision by PRISM's constructors to attempt to capture these intertemporal factors has important benefits in accurately projecting retirement income, as we have stressed above. But costs of this choice--restricting the influence of certain individual characteristics--are easy to underestimate. As we have argued above in sections 3 and 4, individual characteristics are important determinants of outcomes, such as frequency of unemployment spells and wage rates, which have independent effects on an individual's future expected retirement income. In addition, these characteristics have important interactions with the distribution of income and employment opportunities, and thus will have important effects on the distribution of income among retirees. Since income distribution is a critical issue in pension policy, an appropriate modelling strategy requires the accurate depiction of these interactions.

Overall, we found both models to be impressive and highly innovative pieces of model-building research. Their impressiveness, however, is simultaneously their weakness. The enormous complexity which they embody makes them, effectively, black boxes. The input into them can be seen, understood, and judged. The projections which they yield can be understood and are illuminating. Yet how the assumptions and other inputs came to yield the printed-out projections cannot be seen, understood, or judged. The interaction of the complex relationships, transition matrices, time-triggered status changes, random drawings from unknown pools, and constraints to insure comparability is so complicated that little intuition or "feel" is possible for why the resulting projections are what they are. The evidence which would lead a reviewer to believe the predictions of one model more than that of another is slim, indeed.

Three approaches are ultimately possible to determine if the structure of one model is superior to the other, in the sense that the projections of the model will be closer to real world results, holding initial values, assumptions and data constant. First, each model could be used for backcasting--run over a past time period whose results are known, with identical initial conditions and starting data. The question then is: Which model best tracks reality? In fact, neither model has been subjected to this test. Second, each model could be run several times with various assumed relationships, transition matrices, and time-triggers varied systematically. Such a procedure would be a form of sensitivity test, the results of which would give the model user additional information about which characteristics of each model are having a large impact on the reported projections. Neither model has been subjected to this test either, and the resources available to this project were insufficient to permit even a limited set of sensitivity comparisons to be made. The third approach, and the one we have followed, is to attempt to understand the various procedures and implicit assumptions present in each of the components of each model, to appraise these procedures and assumptions in terms of their conformance to ideal standards and accepted evidence, and on this basis to attempt to understand which factors are the primary determinants of divergent projections and which areas of the model require supplementation or modification. This approach is, obviously, the least attractive, but it is better than having no evidence--or only informed suspicion--regarding which factors are the primary contributors to divergent projections, or which model characteristics seem to stack up least well relative to some ideal.

With respect to evaluating the predictive accuracy of large scale micro-econometric models such as DYNASIM and PRISM, the following comments by Kenneth Arrow, made at a conference designed to evaluate such models, are illuminating:[69]

The development of microanalytic simulation models raises methodological issues, which may not be totally new but certainly appear in new versions. The first one I want to mention is the problem of estimation. One key methodological hope for microanalytic models is that parameters can be estimated "directly," in some sense. But strictly speaking all parameters are parameters of statistical relations, and the general principles of statistical inference should apply. No matter how detailed the observations in a particular sector are, there is still a statistical inference of some kind to arrive at the estimate of a parameter. It is very important, in particular, to know that parameters are estimated with uncertainty. Unfortunately, as far as I can see, in all uses of models for policy purposes there is no confidence or error band.

In some models there is only one observation. In those cases it is impossible to get any estimate of variance at all. There is no way of deriving a confidence band from a single measurement. But we know that no model is correct. It couldn't be; there are so many factors omitted, even under ideal circumstances. Estimates based on a single observation, therefore, must also be unreliable, even though we cannot obtain measures of their unreliability. What is needed is replication, repeated observations within a time series or a cross-section context (although the latter has other difficulties). So it has to be understood that even direct observation should be tested by repeated observations, at several points in time or for several individuals.

Suppose a model is not doing very well at prediction or plausibility. It may become difficult to see what is wrong. It is possible to spot a particular relation that isn't fitting very well. But, more often, what is needed for improved prediction is not just putting one new variable in and taking one away but rethinking the model more comprehensively. If what is needed is a change in concepts, for example, definitional changes will be needed everywhere in the model. In effect, scrapping a large intellectual investment is called for, and this gets harder and harder the larger and more complex the model.

Whether or not these remarks are fair, they lead to the questions of validation. How do we validate these relationships?

One possibility is to use the model for forecasting for a different period. Forecasting or hindcasting is a way of validating a whole system. One would like, really, to be able to validate individual relations as well because, if the whole system doesn't work well, it is necessary to know where the repair job is needed. This problem requires a methodological discussion which I have not seen yet.

Notes

[1]This original model is described in detail in Guy Orcutt, Steven Caldwell, and Richard Wertheimer, et al., Policy Exploration Through Microanalytic Simulation, The Urban Institute, Washington, D.C., 1976.

[2]Part of the difference in recipiency rates is due to the difference in the individuals included in each group. The DYNASIM results pertain to both married and unmarried individuals, while the PRISM results pertain only to unmarried individuals. Married males have a higher recipiency rate than unmarried males and unmarried females have a higher recipiency rate than married females, chiefly because of differences in work histories.

[3]The Macroeconomic-Demographic model is described in Volume I of the ICF Final Report to the National Institute on Aging and The President's Commission on Pension Policy, September, 1981.

[4]Age categories are: zero, age one through four, five year cohorts through age 84, and age 85 and over. Throughout this report categorization by race refers to the division of a sample into white and nonwhite, and categorization by sex refers to the division of a sample into male and female. Only DYNASIM categorizes the samples by race.

[5]Adjustment factors vary by marital status (single, married, divorced, and widowed), age (20 to 24, ten year cohorts through age 54, five year cohorts through age 74, 75 and over), race, sex, education (zero to 4 years completed, 5 to 7, 8, 9 to 11, 12, 13 to 15, and 16 or more).

[6]The documentation does not specify the data base on which these relationships were estimated. The parameter values are said to have been taken from a September, 1979 document: Cindy McKay, "Micro Analytic Simulation System Technical Documentation," Vol. 1, The Hendrikson Corporation, Washington, D.C.

[7]Education categories are: less than 12 years, 12 years, and more than 12 years.

[8]Since birth rates can be adjusted to meet user supplied time series, since 1969. See section 7 below.

[9]Age groups are ten year cohorts. Marital status categories are married and nonmarried. Employment status categories are: full-time, part-time, and not employed. Categories of number of children are: none, one, two or three, and four or more.

[10]Age categories are: age 18 to 20, 21 to 23, and 24 to 27.

[11]Age categories are: ages 14 to 17, 18 to 19, five year cohorts through age 64, ages 65 to 74, and 75 and over. Previous marital status categories are: previously married and never married.

[12]The characteristics included in the equation are: whether children under 6 are present, whether children 6 to 17 are present, the number of children under 18, whether it is a second or thrid marriage, the age differential between the spouses, the age at first marriage, the wife's education, the number of weeks worked by each spouse, and the ratio of their wage rates. Not all characteristics are included in each equation.

[13]Age categories are: ages 14 to 17, 19 to 19, 20 to 24, 25 to 29, 30 to 34, and 35 and over.

[14]Age categories are: ages 14 to 17, 18 to 19, 20 to 24, 25 to 29, 30 to 34, 35 to 44, 45 to 64, and 65 and over. Martial status categories are: never married, divorced and widowed.

[15]Age categories are: ages 14 to 19, five year cohorts through age 54, and age 55 and over.

[16]Age categories are: less than 35 years old, ages 35 to 44, 45 to 54, 55 to 64, 65 and over. Marital status is either married or not married.

[17]Education categories are: less than 12 years, and more than 12 years.

[18]For example the child would need to enter school late, at age seven, repeat eighth grade, spend five years in high school, delay college entrance one year, spend five years in college, in order to be in college at age 25.

[19]Age categories are: ages 14 to 24, 25 to 34, 35 to 44, 45 to 64, and 65 and over.

[20]Age categories are: ages 25 to 34, 35 to 44, 45 to 64, and 65 and over. Categories of education are: zero to eight, 9 to 11, 12, 13 to 15, and 16 or more grades completed.

[21]The other states are New York, Pennsylvanis, Rhode Island, Massachusetts, Vermont, Maine, New Jersey, Michigan and Wisconsin. These are the states which provide supplemental benefits to most or all recipients, and these are the states for which PRISM simulates state supplemental benefits--see section 5c.

[22]Age categories are: ages 16 to 19, 20 to 54, 55 to 64, and 65 and over.

[23]The education categories are: less than 9 years, 9 to 11, 12 to 15, and 16 and over. The region variable is one if the individual resides in the South and zero otherwise (geographic definitions of regions--South, West, Northeast and North Central,--are absent from the documentation). Categories for the number of children in the family are: one, two, and more than two. The age-education interaction variables are constructed using two dummy variables for educational attainment; the first is one if the individual has more than 12 years of schooling and the other is one if the individual has more than 16 years of schooling. The age variables are of the form: MAX[0, Age-n], where n is different for each variable (20, 25, 30, 40, 50, 55, 60, 65). The interaction variables which enter the regressions are products of one of the educational dummy variables with one of the age variables. The spouse's earnings and children present variables only enter the female questions.

[24]Age categories are: ages 14 to 20, 21 to 64, and over 64. In addition there are separate equations for single white, single nonwhite, married white and married nonwhite females ages 21 to 64.

[25]Age categories are: ages 21 to 25, 26 to 44, 45 to 54, 55 to 60, and 61 to 64. Each category corresponds to a dummy variable which is one if the individual is in that category and zero otherwise. In addition, the variable MAX[0, Age-60] is included. Expected relative wage is calculated as the ratio of the wage predicted in the wage rate equation to the average wage for the entire population. Income of other household members only enters in the female equations. Dummy variables indicate the presence of a child under six and a total of one, two, and three or more children present. These only enter the female equations.

[26]Age-sex categories are: persons ages 14 to 20, males ages 21 to 25, males ages 26 to 64, females ages 21 to 25, females ages 26 to 45, females ages 46 to 64, and persons over 64. Education categories are: zero to 11 years, 12 to 15 years, 16 or more years. Marital status is either married or not married.

[27]The education variable is one if the individual had eight years or less and zero otherwise. The two age variables only enter the male equations and are dummies which are one if the individual is, respectively, age 31 through 40 and 41 through 50. The region variable is one if South and zero otherwise. Marital status is one if married, zero otherwise, and only enters the male equations. Presence of a child under six only enters the female equations and the equation for individuals age 21 and under. Race only enters the equation for individuals age 21 and under.

[28]Individuals were divided into three groups: (1) "women experiencing special circumstances;" (2) individuals receiving retirement income; and (3) all others. The first group was divided into: (1a) women who bear a child during the year; (1b) women with a child aged 1 to 2; (1c) women with a child aged 3 to 5; (1d) women who become divorced; and (1e) women who become widowed. Of these, groups (1b) and (1c) are subdivided by hours "normally" worked (see below for the definition of the hours "normally" worked categories). Group (2) was divided into: (2a) individuals who accept social security during the year; (2b) individuals who accept pensions during the year; and (2c) through (2g) retirement income recipients by age group (51 to 61, 62 to 64, 65 to 67, 68 to 72, and 72 and above). Each of these groups was subdivided by sex, and

groups. (2c) through (2g) were subdivided by marital status (married, never married, divorced or widowed). Group (3) was divided by age (19 to 24, 25 to 39, 40 to 54, 55 to 61, 62 to 64, and 65 and above). Each of these, except for age 65 and over, were further divided by sex. All of these subgroups in group (3) were further divided by marital status. Each of these groups were subdivided by hours "normally" worked if these subdivisions created groups whose smallest cell contained 30 or more observations. According to the documentation, "we were able to classify about 75 percent of all individuals age 19 to 61 by hours-worked, sex, marital status, and hours 'normally worked.'" These groups were further subdivided by education (some college, no college) if separate groups of at least 30 observations could be created.

[29]Age categories are: ages 16 to 19, 20 to 54, 55 to 64, and 65 and over.

[30]Age categories are: ages up to 23, 24 to 28, 29 to 34, 35 to 44, 45 to 54, and 55 and over.

[31]Strictly speaking, a sufficient condition for this to be true is that the initial DYNASIM forecast contain at least as much within-group variance as the initial year data used in PRISM, and that DYNASIM's equations add at least as much variance in subsequent years as unemployment adds in PRISM. This seems quite likely, since DYNASIM essentially forecasts the years from 1974 on.

[32]Age categories are: ages under 25, 25 to 30, 31 to 49, and 50 to 64. Job tenure categories are: zero to two years, 3 to 4 yers, 5 to 10 years, and 11 years or more. Table 5 lists the industry classifications.

[33]Education categories are: less than high school, high school graduate, and college graduate.

[34]Age categories are: under 25, 25 to 34, 35 to 44, 45 and over. Job tenure categories are: less than one year, one year, two years, 3 to 4 years, 5 to 9 years, and 10 or more years.

[35]Age categories are: 15 to 24, and 25 and over. Education categories are: less than high school, high school graduate, and "some college."

[36]James Schulz et al. (1979), cited in Jon Johnson and Sheila Zedlewski, "The DYNASIM System," Vol. II., UI working paper 1434-03, December 1982.

[37]Categories of real earnings are the $5000 intervals up to $30,000, and $30,000 and over, all in 1978 dollars.

[38]Table 5 lists the industry classifications. The hourly wage categories are: less than $4.00, $4.00 to $6.99, and $7.00 and over, all in 1979 dollars. Age categories are: less than 20, 20 to 24, ten year cohorts through age 54, and age 55 and over.

[39]Alicia H. Munnell, The Economics of Private Pensions (Washington, D.C.: The Brookings Institution, 1982).

[40]Retirement Income Opportunities in an Aging America: Coverage and Benefit Entitlement (Washington, D.C.: Employee Benefit Research Institute, 1981).

[41]Number of hours worked categories are: 1000 and more, and less than 1000. Job tenure categories are: less than one year, one to three years, and more than three years. Age categories are: less than 25, 25 to 44, 45 to 64, and 65 and over.

[42]Defined benefit plans are plans in which the employer's pension commitment is defined as the benefit the employer agrees to pay to participants. Defined contribution plans are plans in which the employer's commitment is expressed as the level of contributions the employer agrees to make on the employee's behalf. Single-employer plans are those provided by one employer. Multiple-employer plans are sponsored by a group of employers who provide benefits for all eligible employees in any participating firm. See Alicia Munnell, The Economics of Private Pensions (Washington, D.C.: The Brookings Institution, 1982), Appendix B.

[43]The years of service categories are: one to 4, 5, 6 to 9, 10, 11 to 14, 15, 16 to 19, 20 to 24, 25 to 29, 30, and over 30 years. The age categories are: less than 55, 55, 56 to 59, 60 to 61, 62, 63 to 64, 65 and over.

[44]The years of service categories are the same as in the previous footnote except that 15 to 19 and 30 years and over are each single categories. The age categories are: less than 40, 40 to 49, 50, 51 to 54, 55, 56 to 59, 60 to 61, and 62 to 64.

[45]These percentages are set arbitrarily.

[46]The service limitations caps the number of years of service which may be used in the defined benefit formula which utilize years of service. Limitations of 25, 30, 35 and 40 years are possible.

[47]9.2 percent of eligible workers are selected to receive joint survivor's option with no reduction in the initial benefit. Of the remaining workers, 75 percent of the married males and 25 percent of the married females are simulated to elect the option. For these workers the current benefit is reduced by a percentage that depends upon the worker's

age, the spouse's age and the sex of the worker. The age categories are: less than 50, 50 to 59, 60 to 64, and 65 and over. The magnitude of the assigned reduction is arbitrary and is meant to equate the actuarial present value of electing and not electing the option.

[48]The ICF Pension Plan Data Base contains the results of a large survey of private pension plan providers conducted by ICF. There are four basic types of defined benefit formulas in the data base. "Final pay plans" base benefits on an estimate of the individual's earnings in the final year or years of covered employment. "Career average salary plans" base benefits on average salary during each year of covered employment. Some of these are "integrated" with social security, others are not. "Integrated" defined benefit plans are plans which are tailored in some way to social security benefits--for example, some plans compensate the worker for the reduction in social security benefits which occurs if he/she retires before the normal retirement age. "Flat percentage of average salary" plans set the benefit as a fixed percetange of some measure of average salary. Again, some of these are integrated with social security, and others are not. "Unit benefit plans" provide a dollar amount per year of service. Note that the first three types of formulae use some measure of earnings and are thus "self-adjusting" for wage trends. The unit benefit plan is not.

[49]All federal government workers are assigned to the federal Civil Service Retirement System plan. Plan assignment for state and local government workers depends upon whether they are simulated to be covered by social security.

[50]The ICF Plan Provisions Data Base contains five basic defined contribution plan types. A "money purchase" plan sets the employer's contribution as a graduated percentage of the worker's salary. Under a "profit sharing" plan employer contributions vary with the firm's profitability. Under a "stock bonus" plan the employer contribution is in the form of firm shares. "Employee stock ownership" plans, or ESOP's, are plans in which the employer contribution is invested in the employer's own securities. A "savings" or "thrift" plan has the employer matching some proportion of the employee's contribution to the plan. The calculation of benefits under money purchase plans involves formulae based on employee earnings. The profit sharing, stock bonus, and ESOP plans involve the financial health of the firm.

[51]PRISM assumes that only the federal retirement plan provides post-retirement cost-of-living adjustments.

[52]These probabilities have been selected arbitrarily since there is little available data on election of joint survivor's option.

[53]Wage rate categories are: less than $4, $4 to $7, and more than $7. Match rate categories are: less than 50 percent, from 50 percent to 95 percent, and more than 95 percent.

[54]The nominal interest rate earned by accumulated contributions is not specified in the documentation.

[55]An individual may be either "fully, currently, or specially insured," according to the applicable Social Security rules. Only federal workers and 31 percent of state and local workers are excluded from coverage.

[56]A wife's benefit is received if the woman is a dependent of a recipient of a retirement or disability benefit, and if she either has one child in her care or has reached age 62. If she qualifies because she is over 62, the benefit is reduced by the number of months before age 65. The benefit is half of her husband's PIA. After her husband dies she is eligible for a surviving spouse's benefit if she has a child in her care, is 60 or older, or is 50 or older and totally disabled. The benefit is equal to her husband's PIA but is reduced if she is under 65. A wife may also be eligible for her own retirement or disability benefit. In this case the difference between the spouse benefit and her benefit is provided as a secondary benefit, if the former is greater than the latter. In certain cases divorced women can receive benefits as a wife.

[57]The disability indicator simulated in DYNASIM applies to individuals who have some condition which limits or prevents work, whereas an individual must be unable to engage in any substantial gainful activity for at least twelve months in order to qualify for SSDI or SSI benefits on the basis of disability.

[58]Countable income is defined as other transfer income plus asset income plus earnings minus an income deduction and minus an earnings deduction. The income deduction is $240 per year (1978 dollars); the earnings deduction is $780 per year plus 50 percent of the remaining earnings.

[59]Age of family head categories are 65 to 72; and 73 and over. Family income categories are: less than $2500, $2501 to $5000, $5001 to $7500, $7501 to $12,500, and $12,501 or more. Marital status categories are: married couples, unmarried men, never married women, divorced women, and widowed women.

[60]The remaining state supplemental benefits involve different eligibility requirements and different benefit formulae. These are said to amount to only 3 percent of all SSI benefits, and are not simulated in PRISM.

[61]Monthly benefit amount categories are: $1 to $99, $100 to $219, and $220 or more.

[62]The 13 states are: California, Hawaii, Maine, Massachusetts, Michigan, Nevada, New Jersey, New York, Pennsylvania, Rhode Island, Vermont, Washington and Wisconsin.

[63]For example, for the largest benefit size, $220 and over, the individual participation rate is 74.3 percent, while the participation rate for married couples is 42.0 percent.

[64]Age categories are: 25 to 39, 40 to 54, and 55 to 64. Earnings categories are: less than $12,500, $12,500 to $25,000, and more than $25,000.

[65]This is the only description given of this data set in the documentation.

[66]Age categories are: 50 to 54, 55 to 58, 59 to 61, and 62 and over. Earnings categories are: less than $10,000, and more than $10,000, all in 1981 dollars. Note that the decision to begin accepting Social Security benefits precedes determination of the numbers of hours worked that year, which in turn determines the actual benefit received, via the earnings test, etc.

[67]See "A Macroeconomic Demographic Model of the U.S. Retirement Income System," Volume I, Model Description; Final Report to the National Institute on Aging and the President's Commission on Pension Policy, Washington, D.C., September 1981.

[68]The underlying data for the employment rate forecasts used in the Macroeconomic-Demographic Model are based on the Bureau of Labor Statistics reported unemployment rates, and pertain to the annual average of survey week observations. These forecasts are then adjusted by constant ratios for each age-sex category to produce calendar year based employment rate projections.

[69]Kenneth Arrow, "Microdata Simulation: Current Status, Problems, Prospects," in Robert Haveman and Kevin Hollenbeck, eds., Microeconomic Simulation Models for Public Policy Analysis (New York: Academic Press, 1980).