NBER WORKING PAPER SERIES

THE ROLE OF HEALTH IN RETIREMENT

Alan L. Gustman
Thomas L. Steinmeier

Working Paper 19902
http://www.nber.org/papers/w19902

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2014

This paper was supported by a grant from the National Institute on Aging, 5R01AG30854-5, to the National Bureau of Economic Research. All opinions and views expressed herein are those of the authors and do not necessarily reflect the views or opinions of the National Bureau of Economic Research or of the National Institute on Aging. We would like to thank our colleagues Erzo Luttmer, Steven Venti and participants in meetings held by NIA and SSA for their helpful comments.

The Role of Health in Retirement
Alan L. Gustman and Thomas L. Steinmeier
NBER Working Paper No. 19902
February 2014
JEL No. D31,D91,E21,H55,I1,I3,J14,J18,J26,J3,J32

ABSTRACT

This paper constructs and estimates a dynamic model of the evolution of health for those over the age of 50 and then embeds that model of health dynamics in a structural, econometric model of retirement and saving.

The health model traces the effects of smoking, obesity, alcohol consumption, depression and other proclivities on medical conditions, including hypertension, diabetes, cancer, lung disease, heart problems, stroke, psychiatric problems and arthritis. These in turn influence an overall index of health status based on self-reported health, work limitations and ADLs, which is used to classify the population into good, fair, poor or terrible health.

Compared to a situation where the entire population is in good health, the current health status of the population reduces the retirement age of the entire population by an average of about one year. While poor health or terrible health have a great impact on the disutility of work and thus on retirement, fair health as opposed to good health has a relatively minor effect. Smoking depresses full-time work effort by up to 3.5 percentage points by those in the early sixties, reducing the average retirement age by four to five months. Effects of trends in health care and health policies on retirement are also analyzed.

Including detailed measurement of health dynamics in a retirement model improves understanding of the effects of health on retirement. It does not, however, influence estimates of the marginal effects of economic incentives on retirement.

Alan L. Gustman
Department of Economics
Dartmouth College
Hanover, NH 03755-3514
and NBER
ALAN.L.GUSTMAN@DARTMOUTH.EDU

Thomas L. Steinmeier
Department of Economics
Texas Tech University
Lubbock, TX 79409
thomas.steinmeier@ttu.edu

An online appendix is available at:
http://www.nber.org/data-appendix/w19902

This paper constructs and estimates a model of the evolution of health over time and embeds the dynamic model of health in a structural econometric model of retirement and saving. The model is estimated for a sample of 51 to 61 year old married men from the original 1992 cohort of the Health and Retirement Study (HRS).

The health model has three parts and is considerably more detailed than the relatively simple one-dimensional health variable that is typically used in retirement models. It begins with some basic characteristics or proclivities of individuals that affect health, such as smoking, drinking, obesity, and cognition. These proclivities, which are fixed over the retirement period, influence the course of medical conditions, such as whether the individual has ever had diabetes, cancer, and so on. Proclivities and medical conditions in turn feed into a health index reflecting self-reported health status, mobility limitations, pain, and the ability to do activities of daily living (ADL's).

The individual starts each period with an initial set of medical conditions and an initial health index. Depending on the basic characteristics and proclivities, medical conditions, and the initial index of health limitations, the individual has a certain probability of surviving the period. Conditional upon survival, additional medical conditions may appear with probabilities again determined by the basic characteristics and proclivities, the previous medical conditions, and the previous index of health limitations. These changes in medical conditions and health limitations then become the initial conditions for the next period.

By repeating this process for multiple periods, the health sub-model can track the evolution of health and survival over time. It is possible to project future mortality, conditions, and health status for individuals starting with different initial conditions and health status at a given age. It is also possible to see how changing the probabilities of various health transitions among medical conditions or among categories of health status resulting from advances in medical treatments would play out in terms of the evolution of mortality, other medical conditions, and health status.

The retirement model contains variables for labor supply and assets, with asset accumulation endogenously determined by decisions regarding consumption vs. saving. The basic retirement sub-model incorporates developments from a recent version of our stochastic dynamic model of retirement (Gustman and Steinmeier, 2013), allowing reversals from states of lesser to greater work. Individuals are assumed to maximize an expected utility function whose arguments are consumption and leisure. The principal decision variables are how much to work and how much to save rather than consume.[1] The budget set includes potential income from work, which may depend on the amount of work done, receipts from pensions and Social Security, which may depend on previous work decisions, and income from assets.

The health sub-model feeds into the retirement sub-model. Health status enters the retirement model by affecting the value of leisure (or disutility of work), which in turn may affect work choices, including the decision as to whether to accept part-time employment. Health expenditures enter through the budget constraint.

The model is designed to incorporate the major sources of uncertainty facing older workers, including their health status over time, uncertain returns on accumulated assets, and unanticipated layoffs with resulting penalty to future wages. For those who are uninsured and below Medicare age, a fourth element is the possibility of large medical expenses, which depend on current health status. Fifth, a long nursing home stay preceding death may exhaust the resources that the surviving spouse needs to sustain living standards.

[1] While expected health and health expenditures influence retirement and saving, we do not model health expenditures as a form of investment in human capital (Grossman, 1972). Smith (1999) suggests that linkages from wealth to health are more likely to operate for children and youths, while the health to income and wealth linkage operates for the population of older persons around retirement age whose behavior we analyze. For a survey of the relevant literature, see Currie and Madrian (1999).

It takes dozens of equations to describe the evolution of the various elements of the health model. Many additional equations are required to describe the constraints governing the retirement and saving model, including the constraints governing health related expenditures and nursing home expenditures, basic constraints related to earnings over the life cycle, asset accumulation, survival and other sources of income. Once these equations are estimated, we use them to estimate the dynamic health model and the life cycle model of retirement and saving. In an effort to produce a readable version of the paper, and to focus on major elements of the model and conclusions, we present some representative equations for each type of constraint. We then confine the remaining equations describing the majority of the constraints to four appendices. These appendices are posted on line. Similarly, some tables reporting the results of simulations of policy alternatives are discussed in the paper, but appear in the appendices.

Section II specifies the health model. Estimates of the various equations are in Section III. The retirement model is specified in Section IV, with estimates of the model presented in Section V. Various simulations with the model are in Section VI. Section VII concludes.

II. The Health Index and Its Components

Our health index incorporates two approaches to measuring the health of the retirement age population. One focuses on functional limitations, such as the inability to climb a flight of stairs or raise your arms above your head. The other is centered on various medical conditions, such as cancer, diabetes, or heart problems. Both Soldo, Mitchell, et al. (2007) and Meijer, Kapteyn, et al. (2013) use functional limitations, but not conditions, in their construction of health indices. Poterba, Venti, and Wise (2013) use both, but the principal components weightings clearly show that the limitations have a much larger weight than conditions.

A. Health index from self-reported health, physical limitations and ADLs.

The health index consists of five dimensions of health: 1. Self-reported health, divided into five categories ranging from excellent to poor. 2. Lower body mobility, an index ranging from 1 to 7 reflecting the number of questions to which the respondent answers yes. All seven have the form of "Because of a health problem do you have any difficulty with …?" The seven activities are: (a) walking several blocks, (b) walking one block, (c) sitting for about two hours, (d) getting up from a chair after sitting for long periods, (e) climbing several flights of stairs without resting, (f) climbing one flight of stairs without resting, and (g) stooping, kneeling, or crouching. 3. Upper body agility. This variable ranges from zero to four, reflecting the number of yes answers to questions about: (a) reaching or extending your arms above shoulder level, (b) pulling or pushing large objects like a living room chair, (c) lifting or carrying weights over 10 pounds, like a heavy bag of groceries, and (d) picking up a dime from a table. 4. Pain has four values ranging from no pain to severe pain. 5. ADL's. This variable counts the number of the following five ADLs: (a) dressing, (b) walking across a room, (c) bathing or showering, (d) eating, such as cutting up your food, and (e) using the toilet, including getting up and down.

For all of the variables in the health model, the data set includes all the cohorts in the first nine waves of the Health and Retirement Study. Typically, there were around 147,000 observations with few missing observations. About 8.7 percent of the observations registered the best category in all five dimensions of the health index, while less than 0.09 percent were in the worst category in all five dimensions. *Note that the health index does not include measures of medical conditions, which are included separately.*

Health is a state variable. To make the model tractable, we summarize the five measures of health with a single health index. Following Soldo, Mitchell et al. (2007) we use item response theory (IRT) to build the health index. IRT can consistently include variables that are neither binary nor have a natural numeric scale, but are categorical (e.g., no pain, mild pain, moderate pain, or severe pain) in a way that principal components cannot. Item response theory begins with an underlying and unobserved health variable θ, which is presumed to be normally distributed in the population. For each of the observed health measures, the probability of falling into one of the response groups is given by an ordered probit, influenced by the value of the underlying health measure.

Since it is computationally infeasible to treat health as a continuous variable, we divide it into four regions. We may think of these four regions as being good health, fair health, poor health, and terrible health. About 35 percent of the sample is in good health and another 30 percent is in fair health. 20 percent is in poor health and 9 percent is in terrible health.

B. Medical conditions.

The HRS consistently asks about eight medical conditions. In each case, the questions are of the form "Has a doctor ever told you that you have ...." The eight conditions considered are: hypertension (high blood pressure), diabetes (high blood sugar), cancer, lung disease (chronic bronchitis or emphysema), heart problems (heart attack, coronary heart disease, angina, or congestive heart failure), stroke, psychiatric problems (including emotional or nervous problems), and arthritis (including rheumatism). Arthritis and hypertension are by far the most common conditions, with percentages between 50 and 60 percent. Heart conditions affect around a quarter of the observations, and diabetes and psychological problems affect almost 20 percent. Cancer and lung problems occur in around 12 percent of the observations, and the least common is stroke, at around 7 percent. Given the nature of the question, it is assumed that once an individual has answered positively in one survey, the answer to that question in all subsequent surveys will be positive.

C. Personal characteristics and proclivities.

Five personal characteristics affect medical conditions and the health index. These are smoking, alcohol abuse, obesity, depression, and cognition. About 17 percent of the observations are positive for smoking. The respondent is considered obese if the BMI value is over 30. About 26 percent of the observations for obesity are positive. For reasons described in Appendix 1, we consider the 19 percent of respondents with a score of 4 or more on the standard depression scale used by the HRS, the so-called CES-D (Center for Epidemiological Studies) scale, to have an indication of depression. The cognition measure has four components related to recall, knowledge of current events or functions of items such as scissors, ability to count backward from 20 to 10, and by 7's from 100, and word definitions. As with each of the elements of the health index and medical conditions, details on the distribution of personal characteristics and proclivities and correlations are in Appendix 1.

III. Estimates of the Health Model.

In this section, we present an overview of how the transition equations for the health model are obtained. Included equations include estimates of mortality transitions, transitions into the various medical conditions, and transitions among the four health categories. These equations are combined to form a dynamic model of health, where the results for one year become the explanatory variables for the next year.

We begin with a mortality equation where health status and conditions at time $t-1$ determines mortality between $t-1$ and t. The mortality hazard is given by the functional form:

$$m = 0.01\, e^{\beta X}$$

where m is the mortality hazard and X is a list of variables influencing mortality. The X variables include a constant, age – 65, $(\text{age} - 65)^2$, and interactions between these three terms and each of the health status, medical conditions, and personal characteristics variables.

Most of the health status, medical conditions, and personal characteristics have at least some significant effect on mortality. About half of the medical conditions have effects that vary with age and/or age^2. The signs of the coefficients are generally as expected. The more serious medical conditions and health problems, as well as smoking, all generate an elevated probability of mortality. The main factor going in an unexpected direction is obesity, but this effect may be overwhelmed if obese individuals are in poor or terrible health.

Using individual hazards, such as each of the eight diseases in various combinations, would create hundreds of state variables. Accordingly, we use the individual mortality hazards to group the conditions to reflect their relative effect on mortality. Based on the hazards reported in Appendix 2, we form three groups of conditions with regard to mortality. These are:

Minimal Effect on Mortality	Hypertension Psychiatric Problems Arthritis
Moderate Effect on Mortality	Diabetes Lung Disease Heart Problems Stroke
Large Effect on Mortality	Cancer

Individuals are assigned to the groups based on the most severe medical condition they have. Individuals with no medical conditions are assigned to the first group, with minimal effect on mortality. Table 1 reports the results of the mortality hazard.

Because of attrition in the HRS sample, the mortality equations overstate the survival probabilities. Specifically, when an individual who was interviewed in one survey is not able to be interviewed in the next survey, an effort is made to determine why. Some individuals simply refuse but are known to be alive, while in other cases the interviewers are told that the individual has died. For individuals who simply disappear, the HRS uses the National Death Index to ascertain whether the cause of non-interview is death. However, it is known that this procedure probably does not capture all the non-interviewees who have died.

Table 1
Mortality Estimates with Grouped Conditions

	Main Effect	Interaction Terms × 0.1 × (age – 65)	× 0.001 × (age – 65)^2
Base	-0.608 ***	0.951 ***	0.166
Mortality Group			
Moderate	0.693 ***	-0.395 ***	0.933 ***
High	1.314 ***	-0.736 ***	1.376 ***
Personal Characteristics			
Smoking	0.454 ***	-0.020	-0.150
Drinking	0.011	0.028	-0.484
Obesity	-0.435 ***	0.037	-0.229
Depression	0.157 ***	-0.021	-0.488 *
Cognition – Impaired	0.276 ***	0.013	-0.101
Cognition – Excellent	-0.161 **	0.040	-0.758
Health Status			
Fair	0.534 ***	-0.197**	0.353
Poor	1.109 ***	-0.073	-0.399
Terrible	1.771 ***	0.017	-0.810

Significance levels: * 90%, ** 95%, and *** 99%

Given that it is likely that the HRS probably underestimates actual mortality, when the relationship between mortality and health status, medical conditions, and personal characteristics is used in the dynamic simulations, the overall level of mortality by age is then adjusted to match the aggregate mortality rate as reported from other sources. This procedure should preserve the differences in relative mortality among the various groups while giving overall results that match the actual mortality experience.

B. Transitions Among Health States.

The transitions among health states are taken to depend on the previous health state, previous medical conditions, new medical conditions, and personal characteristics. Given the previous health state, the probability of the current health state, denoted from 1 to 4, is taken to be an ordered probit, where the value of the health states are good, fair, poor, and terrible. The four ordered probit equations (one for each of the health states in the previous year) are estimated jointly. The results of the estimation are presented in Appendix Table 2.3. Each panel of the table corresponds to a particular health state in the previous year.

To use these estimates in the structural health and retirement model, we group the medical conditions according to the degree to which they make transitions to a worse health state more likely. This is done by comparing the average negative shift of the βX term over the four categories of health in the previous year for both existing conditions and new conditions for the same three ages that were used in the mortality analysis.

Hypertension appears to have the least effect on health transitions, perhaps because existing medications can control this disease fairly well, while stroke has the most effect on health for those who survive it. Among the remaining conditions, diabetes, cancer, and heart problems have relatively mild effects on the health transitions while lung disease, psychiatric problems and arthritis have more noticeable effects. This yields the following four groupings:

Little or No Effect on Health Transitions:	Hypertension
Mild Effect on Health Transitions:	Diabetes Cancer Heart Problems
Moderate Effect on Health Transitions:	Lung Disease Psychiatric Problems Arthritis
Large Effect on Health Transitions:	Stroke

Individuals with multiple medical conditions are assigned to the group with the largest effect of any of the individual conditions. Those with no medical conditions are assigned to the first group, with little or no effect on health transitions. Detailed estimates of the health transitions based on these four groups of conditions can be found in Appendix Table 2.4. The coefficients are generally fairly similar to the corresponding coefficients estimated with the full set of conditions.

C. Transitions Among Medical Condition Groups.

The previous analysis of mortality has grouped the medical conditions into three groups related to their effect on mortality, and the previous analysis of health transitions has grouped the medical conditions into four groups on the basis of their effects on health transitions. We may view the related groupings in a two-way chart as follows:

		Effect on Mortality		
		Little or None	Moderate	Large
Effect On Health Transitions	Little or None	Hypertension		
	Mild		Diabetes Heart Problems	Cancer
	Moderate	Arthritis Psychiatric Problems	Lung Disease	
	Large		Stroke	

Since an individual can have more than one condition, and since the effect on mortality and health transitions is presumed to be at the greatest level of any of the individual conditions, some of the cells of this table may be occupied by combinations of the conditions. For instance, someone with diabetes and arthritis would be considered to be in the cell with a moderate effect on mortality (due to the diabetes) and a moderate effect on health transitions (due to the arthritis). Cells that are not occupied by individual conditions may nevertheless be occupied by combinations of the conditions. For example, someone with cancer and arthritis would be in the cell with a large effect on mortality and a moderate effect on health transitions. Taking these combinations into account, eight of the twelve potential cells in this table are occupied. They may be labeled as follows:

		Effect on Mortality		
		Little or None	Moderate	Large
Effect On Health Transitions	Little or None	C1	o	o
	Mild	o	C2	C3
	Moderate	C4	C5	C6
	Large	o	C7	C8

Cells marked with a "o" do not correspond to any combination of conditions. By focusing on these eight combinations of conditions, rather than all 256 permutations of the eight conditions, we are able to provide reasonable variation in the ability of the conditions to affect mortality and the health transition probabilities, while limiting the number of states to make the structural health and retirement computationally feasible.

Transitions occur when the respondent reports a new condition that has either a greater effect on mortality or a greater effect on health transitions than his previous combination of conditions.[2] All of the combinations except C6 and C8 can be the result of single condition. For instance, an individual who previously had no conditions but now has lung disease would transition from C1 to C5. A transition to C5 might also result if an individual who previously reported diabetes now reports arthritis. This would be a transition from C2 to C5. A transition to C6 can occur if an individual who previously reported cancer now also reports lung disease (a transition from C3 to C6) or an individual who previously reported arthritis now also reports cancer (a transition from C4 to C6).

[2] Transitions among these eight combinations are constrained by the fact that the questions for the conditions are of the form "Has a doctor ever told you that you have …." Once this question is answered in the affirmative, it cannot logically be answered negatively in future years. This means that transitions up or to the left in the table above are logically not possible.

Instances where individuals report two new conditions are quite rare, so we abstract from them and assume that, for instance, a transition directly from C1 to C8 does not occur because it would require both a new report of cancer and a new report of stroke. It is possible, however, for a transition from C1 to C3 to occur in one year due to cancer and from C3 to C8 to occur in a later year due to a stroke.

Transitions to the various condition combinations are taken to depend on previous conditions, health status, and the personal characteristics. Each transition equation is taken to be the transition to a particular cell, and the sample for that equation is the group of individuals who can reach that cell with one new condition. For instance, the equation for transitions to cell C3 has as its sample the individuals who were in either cell C1 or C2 the previous period. Individuals who were in cells C4, C5, C6, C7, and C8 already have conditions which have a greater effect on health transitions than C3. Transition equations into the groups can be found in Appendix Table 2.5.

D. Dynamic Simulation of the Health Model.

The mortality, health transition, and condition transition equations may be combined to form a dynamic model of health, where the results for one year become the explanatory variables for the next year. The dynamic framework we use may be sketched as follows:

(1). The health status and condition combination at time t-1 determines mortality between t-1 and t. (2). For those who survive, the health status and condition combination at time t-1 determines the condition transition probabilities to the condition combination at time t. (3). The health status and condition combination at time t-1, along with the condition combination at time t, determine the health transition probabilities to the health status at time t. This allows, for instance, a new instance of heart problems at time t to have a different impact on health transitions than heart problems which already existed at t–1. (4). The new health status and condition combination at time t, as determined in steps (2) and (3), becomes the basis for the future changes at times t+1 and times future to t+1.

To ensure that the dynamic simulation of the health sub-model tracks the population statistics, we again adjust for underestimation in the HRS of aggregate mortality and other factors by adjusting the constants in the various estimated transition equations up or down. This allows the dynamic simulations to track the observed frequencies fairly well while maintaining the differential effects of the conditions and health state on the various transition frequencies.[3]

Table 2 compares the probabilities of being in the various health states. The major trends in health are tracked quite well, and the simulation succeeds in reflecting the peak of fair health, which rises in the early age ranges as people fall from good health to fair health and then falls in the later age ranges as people fall from fair health to poor or terrible health.

Table 2
Frequencies of Health States

Age Range	50-59	60-69	70-79	80-89	90-99
Health State			Simulated		
Good	54.6%	42.5%	30.3%	18.9%	11.0%
Fair	26.9	32.5	35.1	33.2	27.1
Poor	12.7	19.1	25.8	33.2	37.4
Terrible	5.9	5.9	8.7	14.7	24.5
			Observed		
Good	54.5	43.3	29.5	18.6	11.9
Fair	26.8	32.1	37.3	33.3	25.1
Poor	12.8	18.0	24.6	32.9	39.0
Terrible	5.9	6.6	8.6	15.2	24.1

[3] These adjustments are described in Appendix 2, Section D. After adjustment, the absolute difference between simulated and actual survival rates never exceeds 0.8 percentage points, and the average arithmetic difference is on the order of 0.1 percentage points. The average absolute deviation between the simulation and the observed frequencies of being in various condition groups is around one percentage point.

IV. Specification of the Retirement Model and Its Constraints.

The retirement model builds on our previous models. As in those earlier models, it incorporates uncertain asset returns and uncertain life expectancy. In addition to the direct variation of health status, sources of uncertainty that are new to this version of the model include job loss, SSDI availability, medical expenditures, and nursing home expenditures.

A. The Basic Retirement and Saving Model.

The retirement model starts out with a basic expected utility function as a function of consumption and leisure over time:

$$EU_{t_0} = u(C_{t_0}, L_{t_0}) + \sum_{t=t_0+1}^{T} \left\{ e^{-\rho(t-t_0)} \sum_{m=1}^{3} s_{m,t}\, u_m(C_{m,t}, L_{m,t}) \right\}$$

where C is consumption and L is the leisure of the husband. L can take on three values: zero (full-time work), ½ (partial retirement work), and unity (complete retirement). It would have been better to make the labor supply of the wife endogenous, but in order to accommodate the state variables related to the condition groups and health status, it is necessary at this point to take the labor supply of the wife as exogenous in order to achieve computational feasibility.[4] ρ is the time preference rate, and m indexes the three survival states for a married couple (both survive, only the husband survives, and only the wife survives). $s_{m,t}$ gives the probability that the couple is in state m at time

[4] The complications introduced by the complexity of retirement decisions and the interaction of retirement decisions by husbands and wives have led researchers to adopt a number of simplifications, including, importantly, the assumption of a unitary model. Gustman and Steinmeier (forthcoming) relaxes these restrictions. As in the current paper, the model includes three labor market states, full-time work, partial retirement and full retirement; reverse flows from states of lesser to greater work; heterogeneity in time preference; varying taste parameters for full-time and part-time work; and the possibility of changes in preferences after retirement. In addition, however, that model also includes an extended choice set created when spouses make independent retirement decisions, complicating the estimation considerably.

t. The couple is assumed intact at time t_0. Future utility is indexed according to m to allow the marginal utility of consumption to depend on the number of partners surviving in the household.

The instantaneous utility function u is given by the expression

$$u(C_t, L_t) = \frac{1}{\alpha} C_t^{\alpha} + \frac{1}{\gamma} e^{\beta X_t + \varepsilon_t} L_t^{\gamma}$$

The variables in X include a constant, a linear term in age, and dummy coefficients for fair health, poor health, and terrible health (good health is the omitted category).[5] ε is a term reflecting the relative desirability of leisure (as opposed to work). It is assumed to be constant as long as the individual is working, but after the individual leaves full-time work it may fluctuate to reflect the fact that retirement may be either more or less desirable than anticipated.

The lifetime utility function is maximized with respect to the lifetime budget constraint which is summarized by the asset evolution equation

$$A_t = (1 + r_t) A_{t-1} + w_{t,L} L_t + S_t + P_t + I_t + T_t - C_t - M_t - N_t$$

Assets are constrained to be non-negative and are assume to generate a stochastic rate of return r_t, whose distribution mimics the long-term distribution of a portfolio of roughly half stocks and half treasury bills. The wage rate for labor depends on the amount of labor supplied and whether the husband is still in his long-term, full-time job; part-time work and a break in full-time employment will generate lower wage rates. S is the total amount of Social Security benefits for the couple and depends on past work effort and, for those subject to the earnings test, on current work effort as well. S also includes SSDI payments for those on disability and below the full retirement age for Social Security. P is the amount of pension benefits received. For individuals retired from defined benefit pension jobs, the amount of the pension payments and the age they are available are calculated

[5] Although we allow the value of leisure to be affected by health status, the value of consumption is not. See Finkelstein and Luttmer (2012) for an attempt to estimate the effects of health on utility from consumption.

according to the plan rules. Individuals with defined contribution plans have their contribution amounts compounded over time at the same stochastic rate of return that applies to other assets, and the entire balance is made available on the retirement date. I is other income and includes the wages and pension receipts of the wife, which as noted above are assumed exogenous for reasons of computational tractability. M represents the stochastic medical expenses for uninsured individuals before the Medicare eligibility age. N is the amount of stochastic nursing home expenses which are incurred by a spouse who passed away at the end of the previous year. If the individual is not working and the sum of assets plus income less medical and nursing home expenses is insufficient to meet a minimum consumption standard, the household receives enough transfers T to bring consumption up to a minimal level.

The household chooses C and L each period to maximize the remaining expected utility, given the current values of the state variables. State variables include the condition group and health status of the husband, as discussed in the previous section, the mortality experience of the household (i.e., which spouses are still alive at time t), the amount of assets that the household has accumulated with prior saving. whether or not the husband is still working in the long-time full-time job (which we may call the “main” job), the accumulated balance in a defined contribution pension if the individual is still in the main job, the levels of potential Social Security and defined benefit pension benefits if the individual has left the main job, and the individual’s weight on leisure in the utility function, as given by ε.

The solution to the model follows the usual backwards induction method commonly applied to such models. At the last age, utility is calculated for each combination of the state variables; state variables that are continuous (asset levels, defined contribution balances, Social Security and defined pension benefits, and the weight on leisure) are discretized for this purpose. Since everyone is assumed to be retired at this age, the utility simply depends on consuming the available assets plus income. The calculations then shift to the next earliest age. The expected value of assets left at the end of this period is calculated, taking into account the transition probabilities into the following period. Values of utility in the following period are interpolated for the continuous variables if necessary. Then the individual decides on current period consumption and leisure, taking into account current utility and the value of the resulting state variables which arise as a result of the consumption and leisure decisions. This process is repeated each year back to the beginning of the individual's economic life. The model, in short, is an extended table showing how the individual will choose consumption and leisure at each age, given the state variables available to him at that time. The state variables, in turn, evolve according to the decisions and the stochastic processes in the model.

The model is estimated for a sample of married men from the original cohort of the Health and Retirement Study. These men were 51 to 61 in 1992, the first wave of the study. By now, almost all of them have retired, so we have a more or less complete picture of their retirement behavior on which to base the estimates.

B. Layoffs.

Although those who experience a layoff may find other jobs, they will certainly lose the wage advantage of tenure in their previous jobs and hence earnings are likely to be lower. The resulting shortfall in income relative to the anticipated amount can have implications for the retirement decision. Also likely to be affected is the individual's health insurance coverage if it was available in the previous job.

To identify involuntary layoffs, we consider men who were working full-time in one wave of the survey and inquire whether they were working in the same job in the wave two years later. If not, we look at the question as to the reason they left the job they had held two years earlier. If the reason included such things as "business closed" or "laid off/let go," the response is coded as an involuntary termination. Averaged over the first nine waves of the survey, the rate of involuntary terminations was 5.44 percent over the two year periods, or 2.72 percent per year. Note that this time frame included both the boom years of the 1990's and the relatively depressed years of the early 2000's.

C. SSDI.

If an individual becomes eligible for Social Security Disability Insurance (SSDI), he receives disability income equal to the SSDI benefit amount. But in order to keep receiving this benefit he is precluded from doing substantial work. We assume that an individual choosing to receive SSDI benefits must choose not to work at all rather than being able to work full-time or work in a partial retirement job.

Operationally in the model, the eligibility for SSDI is accomplished by forcing the value of ε, which measures the utility of leisure, to a very high value, a value high enough that the individual would never want to work. In this manner, we avoid having to introduce another state variable into the model, which would complicate the computational tractability issue, but we still are able to include the possibility of SSDI into the model in a reasonable way. This approach does, however, require that once an individual gets onto SSDI, he stays on SSDI and never returns to work. Although it is possible for individuals to leave SSDI and return to work, in practice this seems to be fairly rare, especially for those approaching retirement age, so that such an assumption does not do too much violence to reality.

The probability a person is on SSDI in the current wave is estimated with a standard probit formulation

$$\text{Prob(SSDI)} = F(\beta X)$$

where X is a vector including age, health status, and medical condition groups, and F is the cumulative normal distribution. This relationship is reported in Appendix Table 3.1. As expected, the individuals in worse health states are more likely to become eligible for SSDI, and all the health state parameters are significant. The coefficient attached to the terrible health status is particularly large, which is not surprising considering that individuals in this state usually have multiple ADL's. The coefficients for the condition groups are in the expected direction but are generally smaller than the health status coefficients, with only some of them reaching statistical significance. The exception is the last condition group, which combines cancer and stroke. Individuals in this group are quite likely to be eligible for SSDI, especially since many of them will be in the worst health status group. Somewhat surprisingly, conditional on health status and condition group, age does not appear to affect the probability of becoming eligible for SSDI.[6]

D. Medical Expenditures and Health Insurance.

A third major uncertainty concerns medical costs for those who are uninsured. For the retirement model, there are really two issues relating to medical expenditures. First, is the individual insured, either by employer provided insurance or by some other means?[7] Secondly, if the answer to the first question is negative, what is the distribution of annual medical costs? These costs will surely depend on the health state and medical conditions of the individual. Note that the issue of insurance applies to those individuals who are under the Medicare eligibility age. Individuals above this age are assumed to be covered.

---

[6] Once again, we apply a constant multiplicative factor to the transition equation so that a dynamic simulation yields approximately the same results as are observed. The results are given in Appendix Table 3.2.

[7] Regarding the availability of insurance, there are two issues. First, is the individual insured while he is working at his main job, and secondly, if he is insured while working at the main job, will he continue to be insured if he retires from that job? Individuals who are covered on their current job but may not be covered if they retire may face incentives to remain working in order to avoid being uninsured.

The treatment of health insurance in the retirement literature has continued to evolve. Early retirement studies simplified both the model (e.g., by ignoring saving), and the availability of insurance in the market.[8]

For a recent contribution see French and Jones (2011, p. 725), who estimate a structural model and conclude that medical expense uncertainty greatly increases the value of health insurance. While they conclude that "Medicare is important for understanding retirement, especially for workers whose health insurance is tied to their job.", they predict only a small effect of extending the Medicare eligibility age from 65 to 67, increasing total work over the period from age 60 to 69 by less than one month.[9]

---

[8] French and Jones (2011) argue that because Rust and Phelan (1997) assume there is no saving, they exaggerate the effect of employer provided health insurance on retirement. Rust and Phelan (1997) also assume that those not observed with health insurance in retirement were unable to purchase it, so that only those observed to have health insurance on the job were able to buy it. Since those who have already retired and chose not to purchase health insurance are incorrectly assumed not to have had the opportunity to become covered, this leads to an artificial relationship between retirement and the absence of health insurance coverage on the one hand, and coverage by employer provided health insurance and continued work on the other. Moreover, by assumption in Rust and Phelan, there is no saving, so an unexpected expenditure on health can drive consumption close to zero. This raises the implicit value of any insurance provided on the job to very high levels. The result is that a high value is imputed to insurance even for a person who could have purchased health insurance on the private market, but chose not to do so. Following Rust and Phelan (1997), van der Klaauw and Wolpin (2008) do not allow for private purchase of health insurance if it is not available on the current job or from a previous employer, but to avoid overvaluation of health insurance, they introduce a minimum consumption value, presumably provided by transfers as in Hubbard, Skinner and Zeldes (1995). They also consider Medicare to be equivalent to employer provided health insurance, and thus seem to have no role for retiree health insurance and the purchase of a Medigap policy. Gilleskie and Blau (2006) follow the specification of Rust and Phelan (1997). They find a very small effect both of health insurance and of Medicare on retirement once they standardize for pensions and Social Security. Indeed, they find that raising the Medicare age would have little effect on retirement.

[9] Other investigators (e.g., Marston and Woodbury, 2006) also believe that retiree health insurance plays an important role in influencing retirement.

Unfortunately, the insurance issue is clouded by what appears to be a lot of noise in the data. Many individuals, even those whose employment state does not change, will report one insurance status in one wave and another insurance status in the next wave. Although some of these reported changes undoubtedly reflect actual changes, the sheer volume of them suggests the same kind of reporting errors that are present in other parts of the survey.

To address this issue, the following strategy is used. For individuals observed in some waves of the survey after they have left full-time work, they are considered to have retiree health insurance if they report health insurance coverage in the majority of the surveys between the time they leave full-time work and the Medicare eligibility age. It does not matter whether the insurance is from their previous employer or from some other source; the main concern is that they are covered after they are retired. For individuals who do not have retiree health insurance, there is the question of whether they were covered in full-time work. Such individuals are considered to be covered if they report coverage in the last wave in which they worked full-time.

For two cases, this yields inconclusive results. The first includes instances where the individual never is observed to retire from full-time work, either because the observations ended before retirement or because the individual worked full-time until the Medicare eligibility age. Coverage while working is not an issue here, but potential coverage while retired is. Fortunately, the survey asks questions about potential retiree coverage. The respondent is considered to have potential retiree coverage if he reported affirmatively to the retiree coverage question in the last wave he was observed to be working full-time before the Medicare eligibility age. Again, the responses tend to shift wave to wave, and it is likely that the respondent was best informed about the availability of retiree health insurance the closer he was to retirement.

The other case includes instances where in the first wave the individual was observed he already was not working full-time. If the individual has insurance while retired, we assume that he had some insurance while employed full-time. But since the questions regarding health insurance are not retrospective, it is not clear whether an individual who is uninsured after retirement had insurance when he was working full-time. We impute the insurance status for these individuals while working full-time, using variables related to the last full-time work described retrospectively.

The above procedure divides the sample into three groups: those who have no insurance, those who are insured while working full-time but are not insured afterwards, and those who always have insurance. For those who are uninsured at least part of the time, the next question is the distribution of health care costs. These were estimated for a sample of individuals who reported that they did not have any health insurance in the year. The functional form was a log-normal distribution of the type

$$\text{Prob}\ (\text{cost} < x) \ = \ F[(\ln(x) - (\mu + X\beta)) / \sigma]$$

where F is the standard cumulative normal distribution and X contains the health status categories and the condition groups. The parameter estimates are given in Appendix Table 3.3. Having poor or terrible health raises medical expenses significantly, as does having any of the medical conditions. Among the medical conditions, being in group 4 (arthritis and/or psychological problems) raises costs the least and being in group 8 (cancer and stroke) raises them the most.

Since most of the time the expenses are relatively small, the retirement model focuses on the larger expenditures which may effectively bankrupt the couple. It groups expenditures by percentile groups, with the groups divided by expenditures at the 50^{th}, 80^{th}, 90^{th}, 95^{th}, 97.5^{th}, 99^{th}, and 99.5^{th} percentiles. The probability of being in each of these eight groups defined by these breakpoints is 50%, 30%. 10%, 5%, 2.5%, 1.5%, 0.5%, and 0.5%. The average expenditure within each of these percentile ranges depends on the health status and condition group of the individual. Appendix Table 3.4 gives the average expenditure for each percentile range for the various health status and condition groups in 1992 dollars, as calculated from the log-normal distribution and the parameters estimated in Appendix Table 3.3. The table reflects the possibility of very high expenditures with relatively low probabilities. Again, bear in mind that these are estimates for uninsured individuals below the Medicare eligibility age.

E. Nursing Home Expenditures.

After retirement, the major financial uncertainties are the rate of return on assets, the possibility of needing more assets because of a long life, and the possibility that an extended nursing home stay by the first spouse to die will exhaust the assets available to support the retirement of the remaining spouse. The first two sources of uncertainty have been dealt with in previous versions of the model, but the specter of a long nursing home stay is new to this version of the model.

For HRS respondents who die between surveys, the study attempts to conduct an "exit" interview with someone who was familiar with the respondent's situation, usually a close relative. Among the questions asked are whether the individual was residing at a nursing home at the time of death, and if so, how long he had been in the nursing home. The study was careful to distinguish between where the individual was residing at the time of death and the actual location at death, which is usually in a hospital. Appendix Table 3.5 gives the percentage of HRS respondents from all cohorts and all waves who resided in nursing homes at the time of their death, broken down by health status and condition group, and Appendix Table 3.6 indicates the distribution of lengths of stays, again broken down by health status and condition group. Relatively few individuals in their 50's reside in nursing homes at the time of their death, but the number rises rapidly with age, with three-fifths of individuals over 90 years old residing in nursing homes at the time of their death. The percentage in nursing homes rises with deteriorating health status, as one might expect. The distribution of stays looks relatively stable during at younger ages. The median stay is two or three months, rising to not quite a year at the 75th percentile and about three years at the 90th percentile. As individuals age into their 80's and 90's, the distribution shifts markedly to the right, so that by the

90’s the median stay is not quite a year, the 75th percentile is a little over two years, and the 90th percentile is over four years.[10]

Equations are estimated for the probability of a final nursing home stay and for the duration of the stay. For the estimated retirement model, the distribution of durations of final nursing home stays is discretized into the seven intervals. The cost of being in each of these intervals is calculated as the average number of fitted months for those in the interval, multiplied by the 1992 cost of that number of fitted months. For example, an individual dying with a particular age, health status, and condition group might be calculated to have a 20 percent chance of residing in a nursing home at the time of death and, conditional of residing in a nursing home, as 15 percent chance that his length of stay was 6 to 12 months. The cost of such an outcome would be calculated as 8.8 months times the average cost per month of being in a nursing home.

F. Other Elements of the Budget Set.

The remaining elements of the budget set are treated in a manner similar to our previous work, so the discussion here will be minimal. Earnings from work are taken either from the Social Security administrative records, or if those are unavailable, from retrospective work history questions. Projections for future potential earnings are made using tenure and experience coefficients from wage equations estimated from the PSID. Earnings are treated as exogenous rather than stochastic to avoid computational tractability problems. This may not be such a bad assumption for older workers; the main sources of earnings uncertainty for these workers are health problems which cause a transition either to part-time work or retirement and uncertain layoffs. Both of these sources of uncertainty are captured within the model. Partial retirement earnings are used for those for whom such earnings are observed; for the remainder of the sample, partial retirement earnings are based on

[10] For further details on the distribution of length of stay and expenditures at time of death, see Appendix Tables 3.6, 3.7, 3.8 and 3.9.

a regression relating partial retirement earnings to full-time earnings plus several other explanatory variables.

Pension benefits for alternative retirement ages from the last full-time job are calculated from the information in the pension provider documents collected by the HRS. Because the pension term in the budget constraint is the pension benefit at the age it is paid, not the value of accruals at some fixed interest rate, the individual values benefits that will be paid in the future differently according to their own time preference rate. It also better reflects any liquidity constraint the individual faces. Social Security benefits are calculated from the known Social Security rules and are also included in the budget constraint on the date they are paid. These rules include the effects of any earnings test that the individual might be subject to. For both pensions and Social Security, there are state variables to keep track of the annual amount of those benefits. As previously noted, assets in the model and balances in defined contribution plans are assumed to be subject to stochastic returns from a distribution of returns reflecting a portfolio mix of roughly half short-term treasury bills and half stocks. Balances in defined contribution plans are made available when the individual retires from the job providing the pension. Assets and defined contribution pension balances have separate state variables to keep track of their changes over time.

The earnings of the wife are taken to be exogenous, as is her age of retirement. If her retirement is not observed in the survey time frame, it is taken from her expected age of retirement. Her pension benefits are likewise taken as determined by her earnings in the manner prescribed by the provisions of her pension plan. Again, the benefits are entered in the budget set on the date they are received, and not on the date they are accrued.

V. Estimates of the Retirement Model.

A. Stochastic Specification.

The first step in estimating the retirement model is to specify which parameters are to be estimated and to provide a framework for estimating them. Parameters having to do with the budget set have been discussed in the previous section. This leaves the parameters of the utility function, which include α, β, γ, ε, and ρ. α is a scalar parameter reflecting the marginal utility of consumption, and β is a vector whose values give the weight of leisure in the utility function. In this work, β includes terms for a constant, age, and three categories of the health status (fair, poor, and terrible). The omitted health status is good. The remaining parameters are considered to be individual effects, treated either as fixed effects or random effects.

γ governs the marginal utility of leisure. We take γ as a random effect so that the values of $(½)^{\gamma}$, which must logically lie between ½ and 1 in order to satisfy diminishing marginal utility, come from an exponential distribution with an exponential coefficient given by $\gamma_o + \gamma_a$ age. The age term is included to account for the fact that as individuals become older, partial retirement work appears to become more attractive relative to full-time work.

ε reflects the value of leisure relative to consumption and is also taken as an individual effect. Econometrically it is treated as a random effect with mean zero and variance σ_ε. High values of ε are associated with an increased valuation of leisure and a relatively early retirement age. Lower values are associated with a reduced valuation of leisure and a relatively late retirement age. It is assumed that the individual knows the value of ε, which stays fixed until the individual actually leaves full-time work. After retirement, the values of ε are stochastic with a year-to-year correlation of ρ_ε.

ρ, which is a central parameter in this model, is the time preference rate and is taken as a fixed effect. Given the values of the other parameters of the model, ρ is taken as that value for which the model yields an asset value equal the observed value of assets as of a specific date. The treatment of ρ as heterogeneous among the population permits the model to reflect realistically the fact that even among groups with similar lifetime incomes, there is a wide variation among the level of assets they have accumulated at any given point in their life cycles.

In the end, then, there are 10 parameters to be estimated for the model: the consumption exponent α, five parameters in the vector β (a constant, the coefficient of age, and three coefficients of health status), two values of γ determining the distribution of partial retirement preferences and its evolution as the individual ages, the variance σ_ε of the preference for leisure, and the correlation parameter ρ_ε which governs how the preference for leisure evolves after retirement. The model also estimates a separate value of ρ for each individual, and the distribution of this parameter over the population is also of interest.

B. Solution to the Model.

The model has ten parameters as described in the previous paragraph. There are stochastic processes governing mortality, medical condition transitions, health status transitions, returns to assets and defined contribution balances, layoffs, changes in the value of leisure after retirement, health care expenditures, and end of life nursing home expenses. In addition, there are random effects whose value is known to the individual but not to the researcher, including the initial value of leisure (ε) and the relative value of partial retirement (γ).

The state variables, which specify how the decisions and stochastic events at one age affect the possibilities at later ages, include the mortality experience of the family (whether one or both spouses have survived), any medical conditions which have arisen, the health state, the amount of assets, whether the individual is still in the main job, and the value of leisure. If the individual is still in the main job and that job has a defined contribution pension, a state variable specifying the defined contribution balance comes into play. When the individual leaves the main job, any defined contribution balance is combined with the other assets, but additional state variables are added specifying the level of defined benefit pension benefits, if any, and the level of Social Security benefits. Before the individual leaves the main job, state variables for pension and Social Security benefits are unnecessary because these amounts are implied by the fact that the individual is still in the main job.

The decision variables are the level of work effort and consumption. If the individual was in the main job the previous period, the work choices are to continue in the main job or to enter partial or full retirement. The individual may also leave the main job for other full-time work, at a lower wage which reflects the loss of tenure, if the parameters of his pension plan make it advantageous to do so. If the individual has previously left the main job, the choices are to return to full-time work or to partially or fully retire. To limit the computational burden, individuals are assumed to work full-time before age 50 and to retire completely by age 70. Given the work decision, the consumption decision is also effectively the savings decision.

The solution to the model is essentially a complete enumeration of the work and savings decisions for all possible combinations of the state variables at every age. For this solution, we follow the usual backwards induction procedure applicable to stochastic dynamic models. To begin, we start at the last possible age, which for the purposes of this work is taken to be 99. For all possible combinations of the state variables, the decision which maximizes utility is found and this utility is associated with that combination of the state variables. At age 98, for every combination of state variables at that age, the decision which maximizes the sum of current period utility and expected utility of the state variables next period is calculated, allowing for the stochastic transitions from the state variables after the current period decision to the state variables one year later. This decision is associated with the combination of the state variables at age 98. The process is repeated for each successive age going backward. The final result is a complete contingency table which lists the decisions for each possible combination of the state variables at each age.

C. Estimation of the Model.

The 10 parameters of the model are estimated using the method of simulated moments (MSM). For any particular combination of values of the parameters and for each individual in the sample, the value of the time preference parameter is calculated in the manner described above. To do so, the work decisions are taken as given, so that the model reduces to a pure consumption model that can be solved. Given this value of time preference, the model is simulated numerous times. In each simulation, a draw is made for the initial values of the random effects, which are initial value of leisure (ε) and the relative value of partial retirement (γ). Assets are assumed to start out at zero, and savings and defined contribution balances are assumed to compound at the rates of return observed in the historical data. The remaining observable stochastic variables, including mortality, medical condition group, health status, layoffs, and health costs, evolve according to their observed patterns and are interpolated between waves. The decisions about saving and work effort, given the current values of the state variables, are made according to the decision tree described in the previous subsection, and the state variables evolve based on the previous decisions in combination with the observed stochastic variables. The unobserved stochastic variable, which is the change in the leisure preference (ε) after retirement, is drawn from the appropriate distribution.

The estimates use 63 moments related to retirement in various circumstances. Thirteen of the moments relate to full-time work between the ages of 54 and 66. There are a number of sets of moments that look at work at five relatively critical ages: 55, 58, 60, 62, and 65. These include moments for any work (full-time or partial retirement), full-time work by individuals in the lower third of the lifetime income distribution, full-time work by individuals in the upper third of the lifetime income distribution, full-time work by those in fair health, in poor health, and in terrible health, and any work (full-time or partial retirement) by those in fair health, in poor health, and in terrible health. Finally, there are five moments indicating whether the individual moved from a state of less work to a state of more work between successive surveys, such as returning to full-time work after partial retirement or moving from full retirement to partial retirement. For each of these moments, only observations for which the respondent was observed at the specified age and, if appropriate for the moment, at the specified level of lifetime income or the specified health status, are included.

Table 3 presents the estimated parameters from applying the method of simulated moments to the sample of HRS individuals. Most of the estimated coefficients are significantly different from zero, and all of the critical parameters of the model appear to be fairly closely estimated. From the standpoint of the effectiveness of economic incentives, the critical parameter is β_a, the coefficient of age. A high value of this parameter means that leisure preferences are shifting rapidly around retirement and leaves relatively little room for economic incentives to have an effect; a low value suggests the reverse. The estimated value is in fact fairly low, which means that things like pension and Social Security incentives may have a substantial effect. The other set of parameters which is particularly important in this work is the group of health parameters β_{hf}, β_{hp}, and β_{ht}. The progression of these parameters is as expected, and the values indicate that being in poor or terrible health is likely to have a great impact on the value of leisure, or alternatively, on the disutility of work. Being in fair health as opposed to good health, however, has a relatively minor effect of the value of leisure.

There are 63 moments and 10 parameters, yielding 53 degrees of freedom. For a χ^2 distribution with 53 degrees of freedom, the critical values are approximately 71 at the 5% significance level and 80 at the 1% significance level. The calculated value from the model is about 95, which is about 19 percent higher than the 1% significance level, indicating that the model is not reproducing at least one of the moments accurately. In examining the individual moments, the problem appears to be an overestimate of retirement, especially partial retirement, at age 60. The overestimate appears to be about 2 to 3 percentage points. Overall, however, the deviations of the moments from zero do not appear to be too severe, and the results of the base simulation in the next section appear to be reasonable.

Table 3
Estimates of the Retirement Model

		Coefficient	Absolute t-statistic
Consumption			
α	Consumption parameter	-0.18***	4.95
Leisure Preference Parameters			
β_o	Constant term in $X\beta$	-9.639***	379.01
β_a	Coefficient of age in $X\beta$	0.051***	12.65
β_{hf}	Coefficient of fair health in $X\beta$ [a]	0.67	0.88
β_{hp}	Coefficient of poor health in $X\beta$ [a]	7.01***	5.47
β_{ht}	Coefficient of terrible health in $X\beta$ [a]	14.38***	7.62
σ_ε	Variance of leisure preference random effect [a]	3.71***	5.60
ρ_ε	Correlation of post retirement leisure preference random effect	-0.16	0.11
Parameters for Distribution of Partial Retirement Preferences			
γ_o	Constant term in $X\gamma$	-4.46***	6.45
γ_a	Coefficient of age in $X\gamma$	0.02	0.14

Number of observations: 2231
q-statistic: 95.35

[a] These coefficients are expressed as multiples of the coefficient of age β_a. For example, being in poor health has the same effect on leisure preference as an additional 7.01 years of age.

VI. Simulations of the Retirement Model

This section will begin by presenting results of a base simulation, that is, a simulation with the actual environment faced by the respondents in the sample. We then ask about the overall effects of health on retirement by simulating retirement on the assumption that all respondents were in good health. We follow that with a set of simulations modifying the environment in several ways. One set of runs asks about retirement if all smokers had the same retirement transitions as nonsmokers. Another asks about the reduction in the number of individuals entering one of the conditions, such as diabetes. Others involve changes in economic incentives such as changing the Social Security early entitlement age, the distribution of health insurance coverage, or changing the probability of layoffs. In this way, we hope to shed light on the likely effects of policy changes and/or changes in the way health evolves over time. The simulations also give an indication as to the range of questions that can be investigated with this kind of model.

The model provides insight beyond previous findings of whether good, fair or poor health have a major effect on retirement. But it also imposes a cautionary note. Not all behaviors or diseases strongly affect retirement outcomes. Although certain behaviors and health related factors such as smoking exert a substantial effect on retirement, certain specific diseases, such as the incidence of diabetes, do not. Consistent with our earlier findings, health insurance, and in particular availability of retiree health insurance, has only a limited effect on retirement. Perhaps most significantly, modeling the influence of health in detail does not change our understanding of, or the estimated impact of, the effects of key economic factors on retirement outcomes.

A. Base Simulation.

Table 4 gives the simulated results for retirements, both full retirements and retirements from full-time work, by age and compares the results to the observed retirement patterns. These retirements are actually pseudo-retirements, calculated by subtracting the number retired at one age from the number retired at an age one year earlier. They are the net result of flows into the given retirement state less flows out of that retirement state.

The top half of the table gives the percentage of those still alive who are retiring at the given age. The most prominent feature of the observed data is the large spike in retirement at age 62, the Social Security early entitlement age. This spike is largest for retirements from full-time work, but there is also a substantial spike in full retirement at that age. The simulated model reproduces the spike in retirement from full-time work reasonably well. Although there is a noticeable spike in the simulated data for full retirement at age 62, it falls short of the magnitude of the spike in the observable data. The other prominent feature of the observed data is a secondary spike at age 65, which was the Social Security full retirement age for most of the sample. This spike is primarily present in retirement from full-time work; for full retirement it is more of a minor bump in the retirement rates. The simulated model does produce a bump in retirement at that age. The bump is approximately the right size for full retirements, but it falls considerably short of the spike in retirements from full-time work. In general, though, the pattern of retirements in the simulated model is relatively close to the pattern of retirements in the observed data.

The second part of Table 4 gives the associated hazard rates, which are the percentage of those still working the year before who retire at the given age. Relative to the figures in the top half of the table, the spike in the hazard rate at age 65 is much more pronounced; indeed it is marginally larger in the observed data than the spike at age 62. Although the volume of retirements at age 65 is considerably smaller than at age 62, those retirements come from a base of those still working the year before that is also considerably smaller. In the simulated data, there is also a more pronounced peak at age 65, and for the same reason, although this peak is again somewhat less than occurs in the observed data.

Table 5 breaks down the simulated percentages in full-time work and partial retirement jobs according to the paths taken to get to the present work status. For full-time work, an individual could have always been in full-time work (column 1), or he could have returned to full-time work after having spent some time in either partial or full retirement (column 2). As one might expect, the fraction of full-time workers who have never left full-time work starts out relatively high, but by age 67 there are almost as many full-time workers who have previously been retired at some point as

Table 4
Observed vs. Simulated Pseudo-Retirements

Age	Retirement from Full-Time Work Observed	Retirement from Full-Time Work Simulated	Full Retirement Observed	Full Retirement Simulated
	Percentage Retiring at the Indicated Age			
54	2.0%	2.8%	1.1%	1.9%
55	3.2	3.8	2.8	3.2
56	1.9	3.2	2.1	2.6
57	4.0	3.9	3.1	3.0
58	3.1	4.3	2.3	3.8
59	3.5	4.3	2.0	4.0
60	5.9	5.6	6.3	4.8
61	6.2	4.3	5.5	3.7
62	15.2	14.1	12.3	8.6
63	6.0	4.2	3.8	3.5
64	6.4	5.2	6.1	4.8
65	9.2	6.0	6.9	5.7
66	4.3	3.5	4.3	3.4
67	2.8	3.5	2.7	3.7
	Retirement Hazards at the Indicated Age			
54	2.2	3.4	1.2	2.2
55	3.6	4.8	3.0	3.8
56	2.2	4.2	2.3	3.2
57	4.8	5.4	3.5	3.8
58	3.9	6.3	2.7	5.0
59	4.6	6.7	2.4	5.5
60	8.1	9.3	7.9	7.0
61	9.3	7.9	7.5	5.8
62	25.1	28.0	18.0	14.5
63	13.2	11.6	6.8	6.9
64	16.2	16.3	11.7	10.1
65	27.9	22.4	15.0	13.4
66	18.1	16.8	11.0	9.2
67	14.4	20.2	7.7	11.0

there are full-time workers who have always worked full-time. Column 3 gives the percentage of the sample that is new to full-time work after leaving a period of full or partial retirement. On average, it appears that slightly under half of the full-time workers who have previously retired are new to full-time status at a given age.

The last three columns of the table relate to partial retirement work. The first column (column 4) gives the percentage at each age range that is engaged in partial retirement work. Comparing these figures to the number in full-time work, it appears that partial retirement work is almost as common as full-time work at the later age ranges, becoming more common from age 65 on. The next column (column 5) gives the percentage of each age range that is new to partial retirement work; it appears that around a third of partial retirement workers are new to partial retirement work that year. This implies that the average length of time that an individual spends in partial retirement work is around three years, conditional on spending any time in partial retirement work. The last column gives the percentage entering partial retirement work after full retirement. These figures suggest that not quite half of those entering partial retirement work have been fully retired previously.

Considering all ages, about 35.7 percent of the simulated sample ever return to full-time work for at least a year after having spent some time in either full or partial retirement. About 19.3 percent ever increase their work effort to partial retirement after having spent time in full retirement. And overall, around 45.8 percent of the sample ever increase work effort from the previous year, either returning to full-time work after full or partial retirement or increasing to partial retirement work after a period of full retirement.[11]

[11] The figures in this paragraph are calculated separately from those in Table 5, since individuals may return to work multiple times and/or for multiple years.

Table 5
Simulated Transitions Among Work and Retirement States

Age	Full-Time Work: Percent In Full-Time Work Never Retired	Full-Time Work: Percent in Full-Time Work After Retiring	Full-Time Work: Percent Newly Returned to Full Time Work	Partial Retirement: Percent in Part-Time Work	Partial Retirement: Percent Newly in Part-Time Work	Partial Retirement: Percent Newly Returned to Part-Time Work Previously Retired
54	73.7	6.0	2.5	5.1	2.5	0.9
55	68.7	7.1	2.9	5.7	2.7	0.7
56	64.5	8.2	3.1	6.3	2.7	0.8
57	59.9	8.9	3.4	7.1	3.1	1.1
58	54.9	9.6	3.6	7.6	3.1	1.1
59	50.1	10.1	3.8	7.9	3.0	1.1
60	44.0	10.6	4.1	8.6	3.4	1.2
61	38.8	11.5	4.5	9.2	3.6	1.4
62	29.0	7.2	3.2	14.7	9.4	1.7
63	24.1	7.9	4.1	15.4	4.8	2.6
64	19.4	7.4	3.7	15.8	4.7	2.5
65	13.7	7.0	4.1	16.2	5.0	3.0
66	10.2	7.1	4.4	16.2	4.6	3.0
67	7.4	6.4	4.0	16.1	4.4	2.9

Table 6 looks at the simulated health status and survival rates over a broader age range by 5 year intervals.[12] Not surprisingly, the percentage of individuals in good health declines monotonically as individuals age, and the percentages in poor and terrible health increases at older ages. The percentage in fair health increases at the earlier ages as individuals transition out of good health to fair health and then declines at older ages as individuals transition again out of fair health into worse health states. The column labeled "Conditional Survival" gives the percent of each age

[12] These figures are based on the estimation sample, which is limited to married men who had a substantial work career, generally defined as working the majority of years in the 10 year period before the survey began. This last selection criterion is intended to limit the sample to individuals for whom retirement was a meaningful concept.

range surviving, conditional on being alive at age 50. At the bottom of the table the life expectancy, conditional on being alive at age 60, is also calculated; for individuals in the simulation sample, that life expectancy is 21.4 years.[13]

Table 6
Simulated Percentages in Various Health States by Age Range

Age Range	Percent in Health Status Among Survivors: Good	Fair	Poor	Terrible	Conditional Survival[a]
50-54	59.1	25.2	10.8	5.0	96.3
55-59	52.2	28.7	14.4	4.6	92.3
60-64	46.2	31.4	17.2	5.2	86.6
65-69	39.8	33.5	20.4	6.3	78.5
70-74	33.1	34.9	24.1	7.9	66.1
75-79	26.4	34.8	28.2	10.5	50.6
80-84	20.3	33.2	32.2	14.4	33.7
85-89	14.6	29.9	35.6	19.8	17.7
90-94	10.3	25.8	37.2	26.8	5.7

Life Expectancy at age 60: 21.4 years

[a] Conditional on being alive at age 50

The simulated percentage of the various health groups in each work status at various ages is shown in Table 7. As might be expected, poor health has a very deleterious effect on work effort. At age 61, for instance, 58 percent of the men in good health are working full-time, while only a third of those in poor health are working to this degree, and only 10 percent of those in terrible health. And among those who are working, the percentage of those working in partial retirement jobs is much higher in the poor health categories. Again at age 61, 13 percent of those in good health who are still working are working in partial retirement jobs, but for those in poor and terrible health the corresponding figures are 25 percent and 50 percent, respectively. By the late 60's, almost all of the

[13] Additional simulations related to health are given in Appendix 4, Tables 4.1 and 4.2.

employment by those in poor or terrible health is in partial retirement jobs, while over half of the employment by those in good health is full-time.

Table 7
Simulated Work by Health Status and Age

	Percent in Work Status in Specified Health Status			
Age	Good	Fair	Poor	Terrible
	Percent Working Full-Time			
54	84.4	81.6	66.2	44.4
55	81.3	78.1	60.5	37.3
56	78.5	75.2	56.7	32.8
57	75.1	71.5	51.9	27.9
58	71.2	67.4	47.0	23.0
59	67.5	63.3	41.9	18.3
60	62.2	57.9	35.7	13.5
61	58.2	53.8	31.2	10.3
62	44.2	38.6	17.6	4.4
63	39.9	34.4	14.4	3.1
64	34.3	29.0	10.6	1.9
65	27.2	22.9	7.1	1.3
66	23.1	19.3	5.3	0.8
67	18.9	15.5	3.8	0.5

Table 7 (cont.)
Simulated Work by Health Status and Age

	Percent in Work Status in Specified Health Status			
Age	Good	Fair	Poor	Terrible
	Percent Partially Retired			
54	4.5	5.0	7.2	8.2
55	5.1	5.6	7.7	8.5
56	5.8	6.2	7.9	8.5
57	6.6	7.2	8.8	9.0
58	7.1	7.6	9.1	9.4
59	7.3	7.8	9.3	9.6
60	8.1	8.6	10.0	9.8
61	8.7	9.1	10.6	10.3
62	14.7	15.7	14.7	8.8
63	15.7	16.6	14.8	8.0
64	16.4	17.3	14.3	7.2
65	17.0	17.7	14.0	6.7
66	17.3	17.9	13.7	6.5
67	17.4	17.9	13.0	6.1
	Percent Completely Retired			
54	11.1	13.4	26.6	47.4
55	13.6	16.3	31.8	54.2
56	15.7	18.6	35.4	58.7
57	18.3	21.3	39.3	63.1
58	21.7	25.0	43.9	67.6
59	25.2	28.9	48.8	72.1
60	29.7	33.5	54.3	76.7
61	33.1	37.1	58.2	79.4
62	41.1	45.7	67.7	86.8
63	44.4	49.0	70.8	88.9
64	49.3	53.7	75.1	90.9
65	55.8	59.4	78.9	92.0
66	59.6	62.8	81.0	92.7
67	63.7	66.6	83.2	93.4

Retirement Simulations Involving Health.

To examine the effect of health status on retirement, we conduct a simulation in which everyone experiences good health with no medical conditions until age 65. The results of this simulation are presented in Table 8, which compares these results with the base simulation. There are several effects of good health on retirement, some of them offsetting. Good health increases life expectancy, which increases the need for retirement income and should induce a later retirement. However, it also reduces expectations of medical and nursing home expenses, which should have the opposite effect. Finally, good health reduces the disutility of work, which should again result in later retirement.

Table 8
Effects of Health on Work Effort
Percentage Assuming Good Health Minus Percentage with Baseline Health

Age	Percent Working Full Time: Percentage Point Increase in Work	Percent Working Full Time: Percent Increase in Work	Percent Working At All: Percent Point Increase in Work	Percent Working At All: Percent Increase in Work
54	4.6%	5.8%	4.0%	4.7%
55	5.4	7.1	4.8	5.9
56	6.1	8.4	5.5	7.0
57	6.8	9.9	6.2	8.2
58	7.6	11.8	7.1	9.8
59	8.5	14.1	7.9	11.6
60	9.2	16.8	8.5	13.4
61	9.9	19.7	9.4	15.8
62	8.7	24.0	9.8	19.3
63	8.8	27.5	10.3	21.7
64	8.6	32.1	10.6	24.9
65	7.9	38.0	10.2	27.6
66	6.9	39.9	9.4	28.1
67	5.7	41.3	8.3	27.8

The first column indicates the percentage point difference between the percent working full time in the simulation with good health and the corresponding percentage in the base simulation. The second column expresses this difference as a percentage of the full time work effort in the base simulation. The last two columns give the corresponding figures for any work, including partial retirement work. The simulations indicate that poor health substantially reduces work effort, and the magnitudes suggest that the primary cause is the reduction in the disutility of work. At its peak effect, health reduces work effort by around 9 percentage points, and the effect is much larger if the reduction is expressed as a percentage of those still working in the base simulation. Adding the figures in the third column gives the total effect of less than good health. Multiplying by baseline amounts, being in less than good health reduces the retirement age of the entire population by an average of about one year.

Simulations Involving Individual Proclivities.

The model can also be used to investigate the differences in health and retirement of different groups with different characteristics and proclivities. As an example of this type of investigation, we will look at the differences between smokers and non-smokers. When we simulate the effects of smoking on life expectancy, life expectancy conditional on survival at age 60 is 22.3 years for non-smokers, while it is 17.4 years for smokers. Because smokers and non-smokers may differ in other personal characteristics such as obesity, frequency of depression, cognition, and the like, we attempt to isolate the effects of smoking itself on life expectancy. When we use the smoking sample, with the observed values of drinking, obesity, depression, and cognition of this group of individuals, but impose the health transition and mortality probabilities of a non-smoking group with the same values of these personal characteristics, life expectancy at age 60 is 21.7 years. The results strongly suggest that the worse outcomes for the smoking group are due primarily to smoking itself and not to a different mix of individuals in the smoking vs. non-smoking groups. Given the health transition and mortality probabilities of the non-smokers, the smokers would have had almost as long a conditional life expectance as the non-smokers had they not smoked. Moreover, the patterns of health status by age would be fairly similar.

Table 9 looks at the employment effects of smoking. The comparison is between a simulation of the smoking group and a simulation of the same group with the same employment opportunities, but with the health transitions and mortality of non-smokers. Smoking appears to depress full-time work effort by up to 3.5 percentage points in the early 60's and to depress the incidence of any work by over 4 percentage points in the mid 60's. Again, there are offsetting effects of smoking on work effort. The reduced life expectancy reduces the need for income from work, and the worse health increases the disutility from work. On the other hand, increased medical expenses and nursing home expenses increase the need for income. On the whole, the effects reducing work effort seem to win out, with the average age of retirement being reduced by four or five months by smoking.

Table 9
Effects of Smoking on Work Effort

	Smoking Sample With			
Age	Non-Smoking Transition Probabilities	Smoking Transition Probabilities	Percentage Point Decrease in Work	Percent Decrease in Work
		Percent Working Full-Time		
54	76.0%	75.1%	0.9%	1.2%
55	73.5	72.4	1.1	1.5
56	70.4	69.1	1.3	1.9
57	66.7	65.0	1.7	2.6
58	62.0	59.8	2.2	3.7
59	59.7	57.1	2.6	4.6
60	54.9	51.7	3.2	6.2
61	51.3	47.5	3.8	8.0
62	30.8	27.3	3.5	12.8
63	27.6	24.0	3.6	15.0
64	22.4	18.8	3.6	19.1
65	18.4	15.2	3.2	21.1
66	15.4	12.5	2.9	23.2
67	12.8	10.2	2.6	25.5

Age	Smoking Sample With Non-Smoking Transition Probabilities	Smoking Sample With Smoking Transition Probabilities	Percentage Point Decrease in Work	Percent Decrease in Work
		Percent Working At All		
54	81.6	80.9	0.7	0.9
55	79.9	79.1	0.8	1.0
56	77.7	76.6	1.1	1.4
57	74.8	73.4	1.4	1.9
58	71.3	69.4	1.9	2.7
59	69.2	66.8	2.4	3.6
60	64.8	61.9	2.9	4.7
61	62.1	58.7	3.4	5.8
62	51.3	47.4	3.9	8.2
63	49.2	45.0	4.2	9.3
64	44.5	40.2	4.3	10.7
65	40.2	36.1	4.1	11.4
66	37.5	33.6	3.9	11.6
67	34.5	30.7	3.8	12.4

Effects of Changing Medical Conditions

As an example of the effect of changing a medical condition on retirement, we examined the effects of eliminating diabetes on labor market activity. Recall that diabetes was among the diseases with the least effect on health transitions in Section III. Consistent with this result, there is very little change when tracing the effects of eliminating diabetes in the health module (see Appendix Table 4.3). Not surprisingly, given the lack of much of an effect on health transitions, the elimination of diabetes has only a minor effect on work effort, increasing the average retirement age by only a week or so (see Appendix Table 4.4).

Simulations for Different Lifetime Income Groups.

The next set of simulations divides the sample into three roughly equal groups defined by their lifetime income. In Table 10, we first consider the relation between time preference rates and income group. There is a strong, positive correlation. For the highest income group, almost 60 percent have time preference rates less than 5%, while the corresponding figure for the lowest income group is only about 30 percent. At the other end of the spectrum, about 55 percent of the lowest income group has time preference rates exceeding 20%, while only about 18 percent of the higher income group have time preference rates this high. A high fraction of the individuals with imputed time preference rates above 20% have essentially no financial assets and are assumed to be consuming their entire resources every period.

Table 10
Distribution of Time Preference Rates for the Income Groups

Time Preference Rate	Income Group		
	Low	Medium	High
0-5%	29.6%	45.2%	59.4%
5-10%	9.2	15.1	16.0
10-20%	6.3	7.6	6.8
>20%	54.9	32.1	17.8

Next we consider the evolution of health status and survival for each of these income groups. The higher income group has a life expectancy over two years longer than the lower income group, and they are more likely to be in one of the better health categories and less likely to be in one of the worse health categories at every age. (See Appendix 4.5 for details.) It should be emphasized that these results are not a direct effect of income on health status and survival, since the transition equations for health status and survival do not include income as an explanatory variable. Rather, these differences arise from the effects of differing proclivities and personal characteristics such as smoking, obesity, and so forth among the income groups.

There is little difference, however, in the simulated full-time work effort by age for the three lifetime income groups. There is a slight tendency for the low income group to work full-time somewhat more before the age of 62, and somewhat less after that age. These two differences roughly offset one another, however, and the average simulated retirement age for all three groups is very nearly equal. Appendix Table 4.6 presents details.

Simulations of Retiree Health Insurance.

We first consider the effects of "job lock." Job lock occurs when a company offers health insurance to its active workers but not to retirees. If an individual has pre-existing health conditions which make it difficult to obtain insurance on the private market, or if the individual regards such insurance as being so expensive that obtaining it is infeasible, the individual has strong incentive to remain working so as to be able to obtain health insurance.

To consider the effects of job lock in the model, Table 11, first column, examines the subsample of individuals who have health insurance in their main full-time jobs, but do not have insurance after they retire from those jobs. Such individuals are identified by looking at their insurance status before and after they leave their main full-time jobs, and for individuals who have not left those jobs before their last observation, by looking at their answers to questions regarding the prospects for employer-provided health insurance after they retire. Column 1 reports percentage point differences in the percentage retired between the base simulation and simulation in which all individuals who have health insurance coverage in their main full-time jobs are assumed to be able to continue that coverage into retirement. Overall, the effect appears to be relatively modest. Retiree health insurance increases retirement by less than one percentage point at any given age, and the cumulative effect is to advance the average retirement age by less than one month.

Table 11
Effects of Retiree Health Insurance
For Sample with Insurance While Working
But No Retiree Health Insurance

Age	Percentage Point Change in Retirement from Full-Time Work	Percentage Point Change in Full Retirement
54	0.4	0.3
55	0.4	0.2
56	0.3	0.3
57	0.4	0.3
58	0.4	0.2
59	0.5	0.3
60	0.5	0.4
61	0.4	0.4
62	0.5	0.3
63	0.5	0.3
64	0.5	0.3
65	0.2	0.2
66	0.0	-0.1
67	0.0	-0.1

We also investigated the labor supply effects from the provision of insurance to the entire population. Beyond those in Table 11, who had insurance when employed but not in retirement, the additional group includes those who had no insurance when employed. We should emphasize that this simulation is strictly a supply side simulation; we did not investigate the contentious issue as to whether requiring employers to provide insurance affects employer behavior in the provision of jobs. The bottom line is that there was very little effect on the labor supply of the uninsured from providing universal insurance. Job lock is not an issue for those individuals who had no insurance even while working. As a consequence, adding health insurance both before and after retirement for these individuals has much less of an employment effect than it does for individuals who lose health insurance upon retirement.

Simulations of Effects of Layoffs.

Another issue of some importance is the role of layoffs in retirement. An older individual, especially one who has been at a firm for a number of years, who is laid off frequently has trouble finding another job which pays a comparable wage. Rather than taking a job which results in a substantial pay cut, the individual may choose to retire instead.

To implement involuntary layoffs in the model, the following strategy is used. To accommodate the possibility of a return to full-time work after retirement, the model already contains two types of full-time jobs. One of these is the main full-time job which the individual may have held for a number of years. The other is an alternative full-time job to which the individual can return after a period of retirement. The wage in the alternative job is lower, since the tenure variables in the wage equation are reset to zero following a period of retirement. In this setting, an involuntary layoff may be viewed as an end of the main full-time job, after which the individual can choose either the alternative full-time job or partial or full-retirement. In the simulations, involuntary layoffs occur stochastically with a given frequency.

Table 12 compares the base simulation to a simulation in which layoffs are eliminated after age 55. Eliminating these layoffs appears to increase full-time work, by a little over two percentage points for those in their early 60's. Given that the majority of individuals have retired from full-time work by that age range, this amounts to around a five percent increase in the number of full-time workers for that group. Looking at the results disaggregated by income group, it appears that the high income group is most affected by the possibility of layoffs. A plausible explanation would be that the wage penalty from being laid off is higher for this group than for the others.

Table 12
Effects of Eliminating Layoffs After Age 55

Age	Entire Sample	By Income Group: Low	Medium	High
	Percentage Point Increase in Full-Time Work			
54	0.0	0.0	0.0	0.0
55	0.2	0.2	0.2	0.3
56	0.6	0.5	0.6	0.8
57	1.0	0.7	1.0	1.3
58	1.3	1.0	1.2	1.5
59	1.7	1.2	1.5	2.0
60	1.8	1.5	1.6	2.2
61	2.0	1.6	1.9	2.4
62	2.2	2.0	2.1	2.5
63	2.1	1.9	2.1	2.4
64	2.0	1.8	2.1	2.2
65	1.5	1.5	1.5	1.7
66	1.2	1.4	1.2	1.3
67	1.0	1.2	0.8	1.0
	Percentage Point Increase in Any Work			
54	0.0	0.0	0.0	0.0
55	0.1	0.1	0.1	0.2
56	0.4	0.2	0.4	0.5
57	0.6	0.4	0.5	0.9
58	0.8	0.6	0.7	1.1
59	1.0	0.7	0.9	1.4
60	1.0	0.8	1.0	1.4
61	1.2	0.8	1.1	1.5
62	1.2	0.8	1.2	1.5
63	1.1	0.7	1.0	1.3
64	0.9	0.6	1.0	1.1
65	0.5	0.3	0.4	0.7
66	0.3	0.3	0.1	0.4
67	0.0	0.2	-0.1	0.1

Simulations of Increased Social Security Early Entitlement Age.

The final exercise explored for the model is one that has been done in earlier versions of the model without the health refinements, namely, an increase in the Social Security early entitlement age to 64. The results of this simulation are detailed in Table 13. The main effects of increasing the early entitlement age are to reduce retirement from full-time work substantially at ages 62 and 63. The magnitude of the reduction is between 8 and 9 percentage points. In Table 4, we see that the number retiring in the base simulation at ages 61, 62 and 63 are 4.3 percent, 14.1 percent, and 4.2 percent, respectively. Thus, the spike in retirement at age 62 is between 9 and 10 percentage points higher than in adjacent years. Since increasing the early entitlement age reduces retirement at ages 62 and 63 by approximately this magnitude, this means that such a policy change would effectively shift the current spike in retirement at age 62 to the new early entitlement age of 64. These results are very similar to those we have obtained in earlier versions of the model with the health refinements.

The simulations disaggregated by income group suggest that the effects of increasing the early entitlement age would impact those with lower incomes more so than those with higher incomes. This is a consequence of the fact that the time preference rates of the lower income groups are generally higher than those of this higher income groups, as detailed in Table 10. Those with very high time preference rates are more likely to have no assets, in which case liquidity constraints may force individuals to wait until Social Security benefits become available before they retire. Even if an individual has some assets, and hence would not be forced by liquidity constraints to wait, an individual with a high time preference rate would not perceive the future increases in benefits to be an adequate compensation to foregoing current benefits, and would find it advantageous to collect benefits as soon as possible, i.e., at the early entitlement age. If the early entitlement age is increased, such individuals may well continue working until the new early entitlement age discourages working further.

Table 13
Effects of Increasing the Social Security Early Entitlement Age to 64

Age	Entire Sample	By Income Group: Low	Medium	High
	Percentage Point Change in Retirement from Full-Time Work			
54	0.0	0.1	0.0	0.0
55	0.0	0.0	0.0	0.0
56	0.0	0.0	0.0	0.1
57	0.0	0.0	-0.1	0.1
58	0.0	0.0	-0.1	0.2
59	0.1	0.1	0.0	0.3
60	0.2	0.2	0.2	0.2
61	0.2	-0.1	0.4	0.3
62	-8.7	-20.1	-6.7	-0.8
63	-8.1	-19.1	-6.2	-0.3
64	0.6	0.8	0.6	0.2
65	0.4	0.9	0.0	0.1
66	0.5	0.9	0.3	0.2
67	0.6	0.8	0.7	0.3
	Percentage Point Change in Full Retirement			
54	0.0	0.0	0.0	0.0
55	0.0	0.0	0.0	0.0
56	0.0	0.0	0.0	0.0
57	0.1	0.0	0.0	0.1
58	0.1	0.1	0.0	0.2
59	0.2	0.1	0.1	0.2
60	0.2	0.1	0.2	0.2
61	0.2	0.1	0.3	0.3
62	-4.3	-5.3	-6.6	-1.5
63	-4.0	-5.2	-6.2	-0.9
64	1.0	1.8	1.1	0.3
65	0.6	1.6	0.4	0.1
66	0.8	1.8	0.7	0.2
67	1.2	2.1	1.2	0.4

IX. Conclusions

A fundamental question in the retirement literature has been the extent to which retirement behavior is driven by health, by economic influences, or by other factors such as preferences. A major problem is that health has not been measured comprehensively in most retirement models, with health status often based on a simple index of self-reported health. In some analyses where more detailed health information has been used, the economic influences on retirement have been simplified.

To address the question of what is the influence of health on retirement when health is comprehensively measured, we have used the very detailed health data available from the Health and Retirement Study to construct and estimate a fully integrated dynamic model of health, which we then embedded into a fully specified model of retirement and saving.

Having embedded the detailed health module into a state of the art model of retirement and saving, we are able to estimate the relation of health to retirement and saving, all in the context of a model that allows other economic and demographic factors to also affect retirement and saving outcomes through various other channels.

Incorporating a full model of health into a model of retirement and saving has enabled us to capture the major sources of uncertainty facing older workers, including (i) the possibility of job loss while working, and the accompanying possibility of loss of health insurance, (ii) the possibility of major health problems which may make it difficult if not impossible to continue working, (iii) uncertainty whether SSDI will be available if major health problems do occur, (iv) the possibility of large medical expenditures for uninsured individuals before Medicare eligibility, (v) uncertain asset returns on retirement saving which may reduce funds available to support retirement, (vi) uncertain life expectancy which will affect the period over which retirement funds will be needed, and (vii) the possibility of large nursing home expenses at the end of life which would deplete the assets available to support the remaining spouse.

In turn, we were able to isolate the effects of various health outcomes on retirement through their influence on the disutility of work, life expectancy, expenditures on current and potential future health issues, and expectations of nursing home expenditures.

Among our major findings as to the influence of health on retirement are the following:

1. When we asked about the differences in work effort and retirement between those in good health and those in poor health, as expected, we found major differences. At age 61, 58 percent of men in good health are working full time, while only a third of those in poor health, and 10 percent of those in terrible health are working full-time. This finding combines the effects of current health status, expectations of future health status, and likely behavior in future years, such as smoking and drinking.
2. To isolate the effect of current health status, we held other health related influences, such as continued future smoking, constant. Those in poor current health at age 60 retire earlier, but the effect of their health status at age 60 is not overwhelming. Compared to a situation where the entire population is in good health, the current health status of the population reduces the retirement age of the entire population by roughly one year.
3. Our model has also allowed us to examine the effects of eliminating certain diseases. When we simulated the effect of eliminating diabetes, there was only a very minor effect on retirement, however, increasing the average retirement age by only a week or so.
4. Next we analyzed the retirement effects of smoking by allowing the transitions in health and mortality to differ between smokers and nonsmokers, but holding constant employment opportunities. Overall, smoking depresses full-time work effort by up to 3.5 percentage points in the early sixties, reducing the average retirement age by four to five months.

5. We then used the estimated model to simulate the effects of health insurance on retirement. First, we looked at the effect of providing retiree health insurance to those who had health insurance on their current job, but did not have retiree health insurance. For this group, retiree health insurance increases retirement by less than one percentage point at any given age, and the cumulative effect is to advance the average retirement age by less than one month. We then examined the effects of providing universal insurance to the group of individuals who are uninsured. This group includes both individuals who are uninsured in their main full-time jobs and individuals who are insured in their main full-time jobs but uninsured after they left those jobs. Adding health insurance both before and after retirement for these individuals has much less of an employment effect than it does for individuals who lose health insurance upon retirement.

It is clear that including a detailed model of health allows us both to understand the dynamics of health among those approaching retirement age, and to determine the influence of specific aspects of health on retirement. We can trace through directly the effects of smoking, obesity, depression and other proclivities on diseases, work limitations and ADLs. This provides insights into the effects of health on retirement that were not available in previous models. These relationships enable us to quantify not only how current health outcomes affect retirement, but to estimate how changes in habits, cures of diseases, and health policies are likely to affect retirement.

The other question is whether the inclusion of health in our retirement model has changed findings as to the influence of economic factors on retirement. The portion of our model reflecting the role of economic factors is similar to the model we have used previously. There are no surprises here. We do not find evidence that integrating detailed information of health status and health dynamics into a structural retirement model changes the basic findings as to the relation of economic influences on retirement outcomes.

We are able to reproduce the key spike in retirement at the Social Security early entitlement age, the distribution of wealth, the flows into and out of various retirement states, including reverse flows from states of lesser work to states of greater work. Pensions continue to have similar effects on retirement. Involuntary layoffs have modest effects on retirement. The substantial simulated effects of raising the Social Security early entitlement age from 62 to 64 are similar to the effects found in our earlier work.

In sum, building a dynamic, stochastic model of health outcomes and embedding it in a dynamic, structural model of retirement does not change our view of how economic factors affect retirement, but it does provide a new set of tools to understand how various dimensions of health and health status affect retirement outcomes. The model provides us with a mechanism that improves our understanding of how trends in health care and health policies are likely to affect retirement outcomes.

References

Currie, Janet and Brigitte C. Madrian. “Health, Health Insurance and the Labor Market”. In Orley Ashenfelter and David Card, Editors. *Handbook of Labor Economics* 3 (Part C): 3309 to 3416. Elsevier.

French, Eric and John Bailey Jones. 2011. "The Effects of Health Insurance and Self-Insurance on Retirement Behavior." *Econometrica* 79(3): 693-732, 05.

Gilleskie, Donna B. and David M. Blau, 2006. "Health Insurance and Retirement of Married Couples."*Journal of Applied Econometrics* 21(7): 935-953.

Grossman, Michael. 1972. “On the Concept of Health Capital and the Demand for Health”. Journal of Political Economy 80(2).

Gustman, Alan L. and Thomas L. Steinmeier. 2005. “The Social Security Early Retirement Age In A Structural Model of Retirement and Wealth.” *Journal of Public Economics* 89(2-3): 441-463.

________________. 2013. “Effects of Social Security Policies on Benefit Claiming, Retirement and Saving. NBER Working Paper No. 19071, May.

________________. Forthcoming. “Integrating Retirement Models”. *Research in Labor Economics.*

Hubbard, R, Glenn, Jonathan Skinner and Stephen P. Zeldes. 1995. “Precautionary Saving and Social Insurance”. *Journal of Political Economy* 103(2): 360-399.

Kapteyn, Arie and Erik Meijer. 2013. “A Comparison of Different Measures of Health and Their Relation to Labor Force Transitions at Older Ages”. Paper presented at NBER Conference on the Economics of Aging.

Marton, James and Stephen A. Woodbury. 2006. “Retiree Health Benefits and Retirement”. W.E. Upjohn Institue staff working paper no. 06-128.

Poterba, James M., Steven F. Venti and David A. Wise. 2013. “Health, Education, and the Post-Retirement Evolution of Household Assets”. NBER Working Paper 18695.

Radloff, Lenore. “The CES-D Scale: A Self-Report Depression Scale for Research in the General Population,” *Applied Psychological Measurement*, Vol. 1, No. 3 (June 1977), pp. 385-401.

Rust, John and Christoper Phelan. 1997. “How Social Security and Medicare Affect Retirement”. *Econometrica* 65(4):781-831.

Smith, James P. 1999. “Healthy Bodies and Thick Wallets: The Dual Relation Between Health and Economic Status”. Journal of Economic Perspectives 13(2): 145-166.

Soldo, Beth J., Olivia S. Mitchell, Rania Tfaily and John F. McCabe. 2007. “Cross-Cohort Differences in Health on the Verge of Retirement”. In Soldo, Madrian and Mitchell, Editors, *Redefining Retirement: How Will Boomers Fare?* Oxford University Press, 138-158.

van der Klaauw, Wilbert & Wolpin, Kenneth I. 2008. "Social Security and the Retirement and Savings Behavior of Low-Income Households." *Journal of Econometrics* 145(1-2): 21-42.