Conference Paper
15th Annual Joint Conference of the Retirement Research Consortium
August 1-2, 2013
Washington, D.C.

Health and Wealth In a Lifecycle Model

John Karl Scholz and Ananth Seshadri*

Department of Economics

University of Wisconsin-Madison and NBER and University of Wisconsin-Madison

July 2013

Abstract

We develop a model of health investments and consumption over the life cycle where health affects longevity, provides flow utility, and health and consumption can be complements or substitutes. We solve each household's dynamic optimization problem using data from the Health and Retirement Study from 1992 through 2008 and Social Security earnings histories. Our model matches well household out-of-pocket medical expenses, self-reported health status, and wealth. The model also does a nice job matching the evolution of health status in old age and changes in wealth between 1998 and 2008. We illustrate the importance of endogenizing health investment by examining the effects on mortality and wealth of eliminating our stylized representation of Medicare. Medicare has meaningful effects on mortality, particularly at the bottom of the income distribution.

*The research reported herein was pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium (RRC). The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government or the RRC. We are also grateful to the NIH for generous financial support through grant R01AG032043. We thank seminar participants at the Chicago Fed, Stanford, NBER, UCSB and the Institute for Fiscal Studies for their very helpful comments and Mike Anderson, Junjie Guo and Hsueh-Hsiang Li for fine research assistance. Two anonymous referees provided excellent comments.

1 Introduction

Health and consumption decisions are interlinked, yet the ways that consumption and health interact are hard to untangle. Health changes, such as disability or illness, affect labor market decisions and hence income and consumption possibilities. But causality also operates in the other direction, where consumption decisions such as smoking or exercise affect health. There are also unobserved differences between people in their ability to produce and maintain health and human capital, leading to correlations between health and lifetime income and wealth. This paper examines links between health, consumption and wealth.

There are many possible ways to examine these links. Our analysis starts from ideas dating back at least to Grossman (1972), who argued that health is the cumulative result of investment and choices (along with randomness) that begin *in utero*. We model household utility as being a function of consumption and health, where individuals make optimizing decisions over consumption and the production of health. In our model, health affects not just utility but also longevity. Surprisingly, given the centrality of health to economic decision-making and well-being, numerical models of lifecycle consumption choices generally treat health in a highly stylized fashion. Authors commonly do not model health as being an argument of utility and do not allow health to affect longevity (see, for example, Hubbard, Skinner, and Zeldes, 1995; Engen, Gale, Uccello, 1999; Palumbo, 1999; and Scholz, Seshadri, Khitatrakun, 2006). Instead medical expense shocks that proxy for health shocks affect the lifetime budget constraint. Households in these papers respond to exogenous medical expense shocks by decreasing consumption, saving for precautionary reasons.

In this paper we formulate a lifecycle model that we solve household-by-household, where health investments (including time-use decisions) affect longevity and health affects utility. By modeling investments in health, longevity becomes an endogenous outcome, which allows us to study the effects of changes in safety net policy, for example, on mortality as well as wealth. Our model also captures the effects of poor health on sick time and hence on earnings and retirement.

In the lifecycle consumption papers noted above, households will respond to cuts in safety net programs by increasing precautionary saving. In our model households might maintain consumption at the cost of activities that degrade health and consequently affect longevity. In practice, these health-reducing activities might include working an additional job (and foregoing sleep); foregoing exercise; or eating high-calorie, inexpensive fast food rather than healthier home-cooked meals. Over the long run, the consequences of these decisions can be large. In a world without health-related social insurance, young forward-looking households may recognize the futility of accumulating wealth to offset expected late-in-life health shocks and simply enjoy a higher standard of living for a shorter expected lifetime. Depending on lifetime earnings or the economic environment, other households may sharply increase precautionary saving in a world without health-related social insurance. Our model provides quantitative insight about these responses.

We, of course, are not the first to examine the links between health, consumption, and wealth. Clear discussions are given in Smith (2005) and Case and Deaton (2005) and many other places. De Nardi, French and Jones (2010) and Palumbo (1999) are more closely related to our work. In their models, the only response that households have to the realization of medical expense shocks is to alter consumption. Death occurs through the application of life tables with random longevity draws.[1] They document

[1] In section 9 of De Nardi, French and Jones (2010) they write down and estimate key structural parameters of a model where consumption and medical expenditures are arguments of utility, and where health status and age affect the size of medical-needs shocks. Their model is estimated on a sample of single individuals age 70 and over. They find that endogenizing medical expense shocks has little effect on their findings that medical expenses are a major saving motive and that social insurance affects the saving of the income-rich and the income-poor.

Two other related papers model intertemporal consumption decisions and include health in the utility function. Fonseca, Michaud, Galama, and Kapteyn (2009) write down a model similar to ours and that late-in-life health shocks, including nursing home expenses, and social insurance play a substantial role in old age wealth decumulation.

We build on the past lifecycle consumption and health literature in at least three ways. First, our specification of utility is different. Most prior papers that add health or medical expenditures to utility assume it is separable from consumption in preferences. Two important exceptions are Murphy and Topel (2006), who use a utility function that features consumption-health complementarity to value improvements in health, and Yogo (2009) who models health and portfolio choices of the elderly in an economy that features complementarity between consumption and health. Health is the object of interest in our approach and we model health production. We allow consumption and health to be complements or substitutes in preferences. In practice, we find consumption and health are complements and complementarity is quantitatively important to understanding the evolution of health and wealth as individuals age. In particular, consumption will optimally decline in old age, tracking the inevitable deterioration of health, which implies consumption will be shifted to earlier periods in the life cycle relative to models that ignore health-consumption interactions.

Second, most papers do not examine health investments and consumption decisions of households younger than 65. Health capital, however, may be well-formed by prior decisions and expenditures by the time an individual reaches 65. We model health production from the start of working life.[2] Forward-looking households will respond to income shocks, health shocks, or to changes in institutions by altering their health investments and consumption during their working lives.

Third, the literature including Palumbo (1999) and De Nardi, French and Jones

---

solve the decision problem for 1,500 representative households. Consumption and health are separable in utility in their model and the focus of their work is on explaining the causes behind the increases in health spending and life expectancy between 1965-2005. Yogo (2009) solves a model similar to ours for retired, single women over 65 to examine portfolio choice and annuitization in retirement.

[2] Health is undoubtedly influenced by shocks and decisions made *in utero* and in childhood. We do not have data on these experiences, however, so lack of data and computational demands lead us to start our analysis at the beginning of working life.

(2010) has shown that anticipated and realized medical expenses are an important determinant of wealth decumulation patterns in old age. The focus of our work differs. We develop a model of consumption and longevity to study how health and income shocks affect consumption plans, and how health and income shocks affect investments in health capital over the lifecycle. If death occurs when health capital falls below a given threshold, households may respond to policy or exogenous shocks by reducing or increasing consumption and hence altering longevity relative to a world where health is not an argument in preferences. Studying the trade-off between consumption and health investments on health, longevity, and wealth offers new insights into household behavior.

After calibrating our model to match key moments for the typical household, we find the model is able to match the cross-sectional variation in medical expenses, longevity, the stock of health and consumption in the Health and Retirement Study. We also match changes in wealth, health spending, and health status between 1998 and 2008. In addition, we match patterns of medical spending, health stocks as well as longevity earlier in the lifecycle. Our analysis reveals that the degree of complementarity between consumption and health capital plays a critical role in explaining various features of the data. Finally we conduct policy experiments to highlight various features of the model as well as to bring out the differences relative to a model with exogenous medical expenditures. We find that Medicare has meaningful effects on life tables, especially at the bottom of the income distribution.

2 Descriptive Facts

We use Health and Retirement Study (HRS) data from 1992 through 2008. The Health and Retirement Study (HRS) is sponsored by the National Institute of Aging and conducted by the University of Michigan with supplemental support from the Social Security

Administration. The HRS is a national panel study with a sample (in 1992) of 12,652 persons in 7,702 households. It oversamples blacks, Hispanics, and residents of Florida. The sample is nationally representative of the American population 50 years old and above. The baseline 1992 study consisted of in-home, face-to-face interviews of the 1931–41 birth cohort and their spouses, if they were married. Follow up interviews have continued every two years through 2010. As the HRS has matured, new cohorts have been added. Our sample includes households from the AHEAD cohort, born before 1924; Children of Depression Age (CODA) cohort, born between 1924 and 1930; the original HRS cohort, born between 1931 and 1941; the War Baby cohort, born between 1942 and 1947; and the Early Boomer cohort, born between 1948 and 1953. The sample is a representative, randomly stratified sample of U.S. households born before 1953.[3]

We start with 30,548 (19,058 unique households) individuals in the RAND HRS Version J (RAND, June 2010). We keep 11,494 households in which either the head or the surviving spouse responded in 2008. We drop 287 households with insufficient earnings to estimate the household fixed effect in the earnings model. Next, we drop 35 households in which both the household head and their spouse are the same gender. This leaves with 11,172 households in our sample.

In addition to a wide range of health information, the HRS has excellent measures of household financial well-being. To measure household net worth we use respondents' reports of the value of primary and secondary residences, other real estate, vehicles, businesses, as well as a wide range of personal savings instruments (IRA, Keogh accounts, stocks, mutual funds, investment trusts, checking accounts, saving accounts, money market accounts, CDs, government savings bonds, Treasury bills, bonds, bond funds, and "other savings"). Household financial liabilities are subtracted from the sum of household wealth and include the value of all mortgages, land contracts, and "other debt."

[3] Comprehensive information on the HRS is available at http://hrsonline.isr.umich.edu/.

Observed household medical expenses are reported as both out-of-pocket and total expenses. In this paper we use the out-of-pocket measure, which includes the costs respondents pay for hospitals, nursing homes, doctors, dentists, outpatient surgery, prescription drugs, home health care, and care in special facilities. While we have reasonable confidence in reported out-of-pocket medical expenses in the HRS, total expenses are considerably more difficult for a household to report accurately in an interview survey. Because of this we use moments for total medical expenses by age, cross-classified by insurance status, drawn from the 2008 Household Component of the Medical Expenditure Panel Survey (MEPS) to calibrate parameters that map out-of-pocket expenses into total expenses.[4] The MEPS is administered by the Agency for Healthcare Research and Quality and contains detailed information on health care expenditures.

Our model must be capable of matching several descriptive facts about health and wealth. The first fact is perhaps obvious, but self-reported health declines with age. The HRS asks respondents about their self-reported health status, where respondents can respond on a 5-point scale (excellent, very good, good, fair, or poor). Figure 1 plots the average responses for two groups of responses in 2008 by cohort. The modal response for the four oldest cohorts is good while it is very good for the youngest cohort, the Early Boomers. The percentage of respondents reporting excellent or very good health declines monotonically with age across cohorts. The percentage of respondents reporting poor or fair health rises with age across cohorts. Recognizing this biological fact, health depreciates in our model of health production.

The second fact highlighted is that exercise is positively correlated with lifetime income as shown in Figure 2. This relationship is potentially important in a model of health production as there is abundant evidence that exercise, smoking, and diet influence health and hence longevity.[5] Nevertheless, the computational demands that

[4] Information is available at http://www.meps.ahrq.gov/data_stats/mepsnet/mepsnethc08.shtml

[5] See, for example, Paffenbarger et al. (1993), Willette (1994), Mokdad et al. (2004), and Warburton et al. (2006).

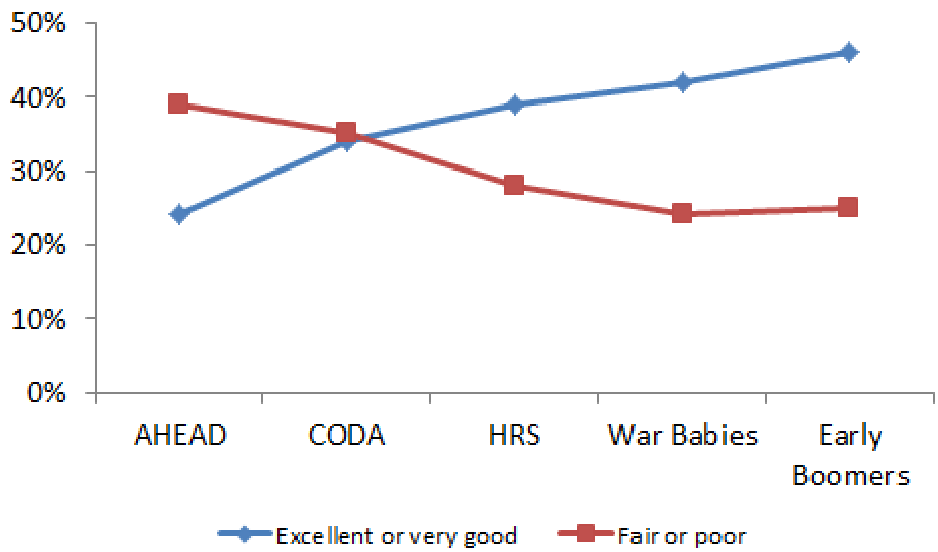

Figure 1: Self-Reported Health and Age (Cohort), 2008 HRS

arise in solving our dynamic programming model household-by-household with endogenous consumption and health production requires parsimonious modelling. Given this requirement, we assume health can be improved by investments of money and by investments of time. Specifically, time investments in health production reduce leisure. Both working and retired households face a combined time and financial budget constraint, which we describe in greater detail below. In this way, we capture the essential trade-off between non-health related consumption and health investment.

The third fact is "the gradient:" health is positively related to socioeconomic status, whether measured by lifetime income, net worth, or related measures. As Figure 3 makes clear, the positive relationship between self-reported health and net worth is strongly present in the HRS. The Figure is similar when households are sorted by lifetime income quintile as opposed to net worth quintile. Illuminating economic decisions over the lifecycle that result in the joint distribution of health and wealth, household-by-household, is a central challenge for this paper.

The fourth fact that our model must accommodate is that there is a strong relation-

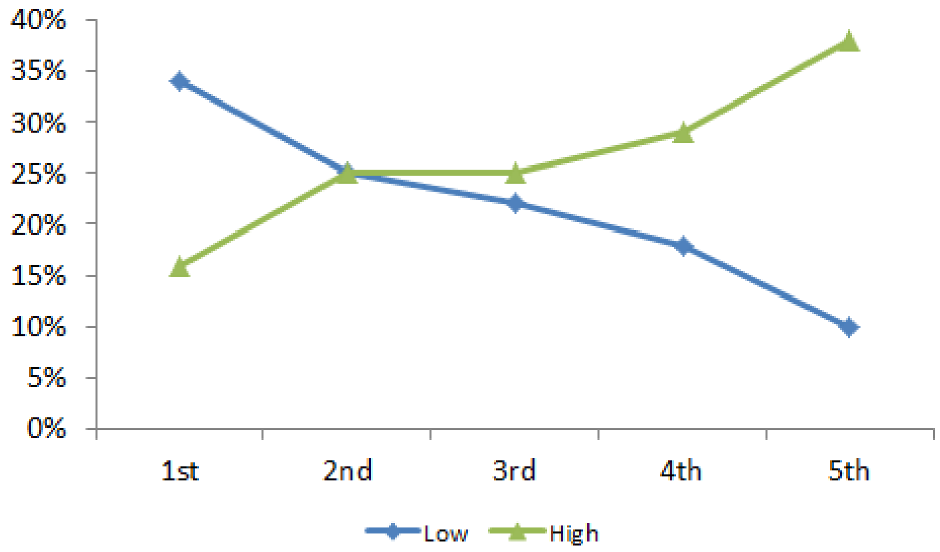

Figure 2: Physical Activity and Lifetime Income Quintile, 2008 HRS

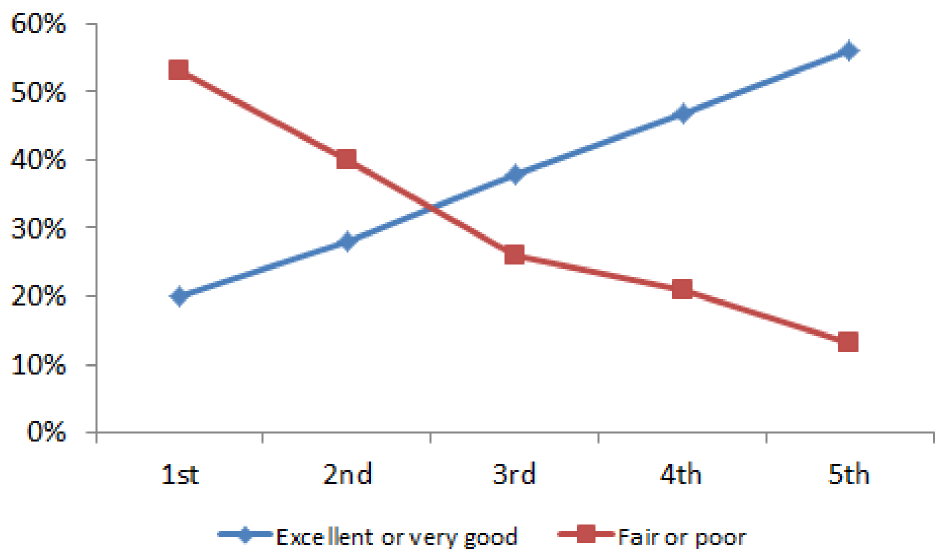

Figure 3: Self-Reported Health and Household Net Worth Quintiles at Age 65, 2008

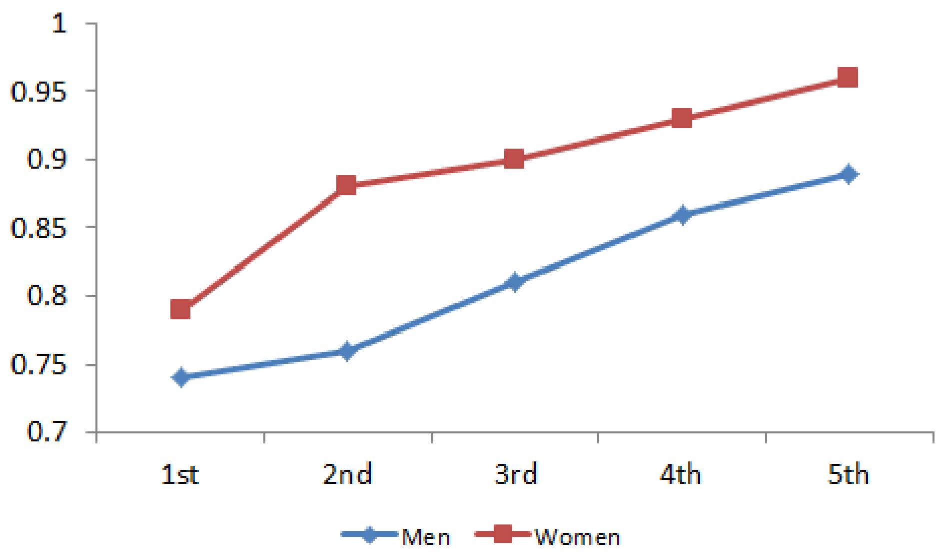

Figure 4: Ten-Year Survival Probabilities for Men and Women by Lifetime Income

ship between lifetime income and survival in the HRS. To show this, we restrict the sample to birth years that, in principle, would allow someone to reach a specific age by the last year of our HRS sample, 2008. So, for example, when we look at patterns of survival to age 70, we restrict the sample to those born before 1938. We also drop all sample members who were over 60 years old in the year they entered the HRS sample. The ten-year survival probabilities to age 70 shown in Figure 4 increase monotonically with lifetime income, from 74 percent for men in the lowest lifetime income quintile to 89 percent for men in the highest. The gradient for women goes from 79 percent in the lowest lifetime income quintile to 96 percent in the highest.

There are many likely explanations for the positive relationship between lifetime income and survival. We write down and solve a model that captures several of these explanations. Households in our model have different draws on annual earnings and hence different lifetime incomes. They differ in the timing of exogenous marriage and fertility. Given differences in incomes and demographic characteristics, their consumption and health investments choices will respond to health shocks (that vary by age), earnings shocks (which are also affected by health), and government programs in different ways. Moreover, we allow consumption and health to be gross complements or gross substitutes in utility. The work that follows, therefore, illuminates the channels through which health, consumption, and wealth are related.

3 Model Economy

In this section, we present the basic elements of our model and then proceed to describe the dynamic programming problem more formally. Even though our HRS sample begins when individuals are older than 50, we use restricted access earnings data for HRS households that typically starts when household heads are between the ages of 18 and 25. Denote the age at which we begin to observe earnings of a household by S. We start the decision problem for a household at age S and assume, at this starting age, all households have zero assets and all household members, husbands and wives, have an identical stock of health.

Demographics: With the exception of life expectancy which we model, other demographic variables are treated as exogenous and deterministic by each household. The number of children a household has varies over the life-cycle and this affects consumption needs during the period of time they are attached to the household. This varies across households and is provided in the HRS. Households are either single or married throughout their lives - we do not model marriage and divorce. The marital status is set as of the first HRS wave. We do, however, model transitions from married to single status upon the death of a spouse that occurred after the first HRS wave. For the cohorts we study, divorce rates and remarriage were not as common as they are now. Retirement is also treated as exogenous and the age of retirement is taken directly from data.[6] One advantage with our data is that we have detailed information on all these

[6] We performed a sensitivity analysis by restricting our attention to 'always married' and 'always

demographic characteristics as of the first HRS wave. For each household, we specify in a deterministic fashion the exact ages in which children arrive and leave, whether they are married or single and the age of retirement.

Stochastic shocks: There are three sources of uncertainty in our model. First, there are health shocks, $\varepsilon_{j,g}$, that are assumed to be i.i.d across individuals and drawn from the distribution $\Xi_{j,g}(\varepsilon_{j,g})$, where j denotes age and g stands for gender that can either be female (2) or male (1). These shocks vary by gender and by age and adversely affects the stock of health. The variation by gender is essential to match the differential mortality rates of women relative to men. The increasing likelihood of these shocks as households age is critical to obtaining declining health status with aging. Second, we model *household* earnings as an AR(1) process where the i.i.d shocks vary by number of earners (n_e), marital status (k), education (edu) and birth cohort of the household head. Specifically, the distribution of earnings at age $j+1$, e_{j+1}, conditional on earnings at age j, e_j, is given by $\Omega_{j,edu}^{f,n_e,k}(e_{j+1}|e_j)$ where f denotes household-specific variation that incorporates variation in the intercept term and k stands for marital status that can either be married (2) or single (1). Birth cohort is implicitly indexed in this distribution through f. Third, the probability of surviving into the following period depends on the stock of health. Healthier households are more likely to survive into the next period, but there is a chance that any individual can die at a given age. The probability of surviving into the next period is given by the function $\Psi(h)$ where h denotes health stock. This function satisfies two properties. First, as h goes to ∞, $\Psi(h)$ converges to 1. Second, $\Psi(h) = 0$ for $h \leq 0$. This ensures that as soon as h goes to zero, the individual dies. The health stock affects utility and also affects the probability of surviving into the next period. Our formulation captures the notion that healthier people are less likely to die.

Preferences and Choices: A household maximizes expected lifetime utility by choos-

single' households and the fit between model and data for the stock of health, medical expenses as well as net worth were all slightly better.

ing consumption, health investments, and leisure. Incorporation of health capital into an otherwise standard consumption-savings problem involves two additional choice and one additional state variable. It is nevertheless a significant complication. In addition to affecting longevity, we assume that households derive direct satisfaction from health. Lifetime utility for a single household at any age j in retirement, $V^1_{j,g}$, is given by

$$V^1_{j,g} = \underset{c_j, i_j, l_j, m^{oop}_j}{Max} \{ n_j U(c_j/n_j, l_j, h_j) + \beta E_j[\Psi(h_{j+1}) V^1_{j+1,g}] \}.$$

This household maximizes expected lifetime utility by choices of consumption, c_j, time investments in health, i_j, leisure, l_j, and out-of-pocket medical expenses, m^{oop}_j that affects the stock of health in the next period through a production function for health. The first argument inside the parenthesis denotes momentary utility during that age while the second term stands for the (expected) continuation value. The expectation operator E_j denotes the expectation over future health shocks. g is the gender of the head of the household, β is the annual discount factor, h_j is the individuals' stock of health, n_j is a household equivalence scale and is a function of the number of adults, A_j, and children, K_j, in the household, so $n_j = g(A_j, K_j)$.

We assume that health affects the time endowment of the husband and the wife through Grossman's formulation of sick time: households experience some loss in their time endowment, $s(h_j)$, which is inversely related to their health status h_j. Upon retirement, an individual splits his or her time endowment of $1 - s(h_j)$ in each period between leisure l_j and activities that augment health investments i_j. Before retirement, we assume that an individual spends an indivisible amount of time $\omega(h_j, g, k, j, a_j, e_j)$ working each period (this function is given exogenously) and spends the rest of his time endowment $1 - s(h_j) - \omega(h_j, g, k, j, a_j, e_j)$ on either leisure l_j or on health investments i_j.[7] In this formulation, whether people with lower stocks of health have less time to

[7]In our model, we assume that poor health adversely affects the time that an individual spends in the labor market. While poor health may well affect investments in human capital and consequently

spend on investments in health and leisure depends on the strength of two effects. On the one hand, lower health stocks are associated with more sick time. On the other hand, labor supply is increasing in the stock of health. Which effect dominates depends on the relative strength of the two effects. The decision problem during the working phase is very similar to the decision problem specified above with one notable difference. There is an additional source of uncertainty - uncertain future earnings for ages prior to retirement. A complete description of the dynamic programming problem for the working and the retirement phase is given below.

The married household's decision problem at age j in retirement involves taking into account the choices of two decision makers. The value function is given by

$$V_j^2 = \underset{c_j, i_{h,j}, l_{h,j}, m_{h,j}^{oop}, i_{w,j}, l_{w,j}, m_{w,j}^{oop}}{Max} \left\{ \begin{array}{c} n_j[\mu U(c_j/n_j, l_{h,j}, h_{h,j}) + (1-\mu) U(c_j/n_j, l_{w,j}, h_{w,j})] \\ +\beta E_j \left[\begin{array}{c} \Psi(h_{h,j+1})\Psi(h_{w,j+1}) V_{j+1}^2 \\ +[1-\Psi(h_{h,j+1})]\Psi(h_{w,j+1}) V_{j+1,2}^1 \\ +\Psi(h_{h,j+1})[1-\Psi(h_{w,j+1})] V_{j+1,1}^1 \end{array} \right] \end{array} \right\}$$

The first two terms inside the parenthesis above stand for momentary utility for the couple where $\mu \in [0,1]$ is the weight on the husband's utility in household utility, $h_{h,j}$ is the husband's stock of health, $l_{h,j}$ is leisure of the husband, $h_{w,j}$ and $l_{w,j}$ are corresponding health stock and leisure of the wife, and $m_{h,j}^{oop}$ and $m_{w,j}^{oop}$ are out of pocket medical expenses for the husband and wife, which affect their stocks of health in the

the wage rate that an individual faces, we observe data on earnings and we do not observe either hours worked or the wage rate. Consequently, data limitations prevent us from disentangling the impact of bad health on the wage from the effect of bad health on hours worked. Furthermore, the analysis in French (2005) suggests that health status has a much larger impact on labor supply and labor force participation than on the wage rate. To be sure, if we did have data on hours worked, we would be in a position to introduce a labor supply dimension to our model. The lack of information on hours worked leads us to approximate a labor supply function for individuals in different states using data from the PSID. Besides health (h), gender (g), marital status (k), age (j), wealth (a) and earings (e), labor supply also depends on whether an individual is a union member, which is assumed to be exogenous and taken from data. Details are in the Appendix.

next period through health production specified below. The expectation operator E_j now denotes the expectation over future health shocks facing both the husband and the wife. The three products of functions Ψ and V inside the expectation operator give the continuation values of the household when both the husband and the wife live to next period, only the wife lives to next period and only the husband lives to next period respectively. Setting μ to be 0 (1) will give us the corresponding lifetime utility for a household headed by a single female (male). This representation of preferences captures the notion that while consumption is a public good within the family, leisure and health are largely the result of individual choices. Understanding the complex decisions made by members of a given family requires us to recognize that they are independent actors - something that our collective model does.

The Production of Health: A challenge when modelling health is that there is at best mixed evidence that marginal expenditures on medical care in the U.S. buy greater health, and hence longevity.[8] This phenomenon is sometimes referred to as "flat of the curve" medicine. It is noteworthy just how hard scholars need to look to find evidence that expenditures on medical care have a discernible, positive effect on health and particularly mortality outcomes. Card, Dobkin, and Maestas (2008), for example, is one of a small number of studies that find expenditures are positively correlated with survival. Their work is based on a very large sample of people admitted to emergency rooms in California: they find the positive effects of spending apply to a small subset of the conditions that lead people to show up in emergency rooms. Doyle (2010) shows that men who have heart attacks when vacationing in Florida have higher survival probabilities if they end up being served by high- rather than low-expenditure hospitals.

[8] See, for example, the Dartmouth Health Atlas (http://dartmouthatlas.org/), which documents little relationship between regional variation in health spending and health outcomes. Finkelstein and McKnight (2008) find little effect of Medicare on mortality when the program was initiated. Chay, Kim and Swaminathan (2010) challenge this assessment.

In addition to evidence that health investments enhance health, Oster et al (2012) show that those diagnosed with a terminal disease are less likely to quit risky behaviors. Specifically, they study the effect of Huntington Disease on health investments. Individuals who learn they carry the HD mutation through genetic testing or symptom onset are much less likely to quit smoking than comparable individuals without this information. Those with earlier symptom onset are less likely to have ever undergone cancer screening (conditional on age). Other studies suggest that *marginal* medical expenditures have little discernible effect on health.

In addition to the evidence above, it is clear that some expenditures improve health. Antibiotics can effectively cure strep throat. Treatment can help people survive cancer. A good orthopedist can help people recover fully from broken bones. Given this, we assume that household members possess a health stock and investments improve health. The accumulation process of the stock of health for a household member is given by

$$h_{j+1} = F_t(m_j, i_j) + (1-\delta)h_j - \varepsilon_{j,g},\ j \in \{S, ...\}, g \in \{1, 2\}$$

The stock of health at the next age, h_{j+1}, is determined by the production of health, given by $F_t(m_j, i_j)$ which depends on calendar time t because we allow productivity of health technology to change over time. Health capital is produced using time, i_j, which could be exercise or other health-producing activities, and medical expenditures. Total medical expenditures, m_j, are a function $M^{ins}(\cdot)$ of out of pocket medical expenses, m_j^{oop}, where the function $M^{ins}(\cdot)$ is determined by health insurance status (ins) and will be specified later when we discuss calibration. In the above equation, δ stands for the depreciation rate of health. Introducing age-dependent shocks to health, $\varepsilon_{j,g}$ is both realistic and necessary if we are interested in matching biological processes and the data. They vary by gender. In typical lifecycle models, medical expenditures have only financial consequences. Here medical expenditures have financial consequences and affect health capital which, in turn, affects utility and longevity. The modeling approach mimics the modeling of human capital – additions to human capital can be either consumption or investment as in Becker (1964), Mincer (1974) and the subsequent, vast human capital literature.

Budget Constraints: Consumption, health investments and leisure are chosen to maximize expected utility subject to the constraints.

$$y_j = e_j + ra_j + T_t(e_j, a_j, j, n_j), j \in \{S, ..., R\}$$

$$y_j = SS\left(\sum_{j=S}^{R} e_j\right) + DB(e_R) + ra_j + T_{R,t}\left(e_R, \sum_{j=S}^{R} e_j, a_j, j, n_j\right), j \in \{R+1, ...\}$$

$$c_j + a_{j+1} + m_{h,j}^{oop} + m_{w,j}^{oop} = y_j + a_j - \tau_t(e_j + ra_j), j \in \{S, ..., R\}$$

$$c_j + a_{j+1} + m_{h,j}^{oop} + m_{w,j}^{oop} = y_j + a_j - \tau_t\left(SS\left(\sum_{j=S}^{R} e_j\right) + DB(e_R) + ra_j\right), j \in \{R+1, ...\}$$

In these expressions y is household income, and e is household earnings, a is household assets, r is the interest rate, T is a transfer function that depends on earnings, assets, age and the number of adult equivalents in the household. The husband and wife in a household are assumed to enter the labor market simultaneously at age S of the head and retire simultaneously at age R of the head. Social security (SS) is a function of lifetime earnings, defined benefit pensions (DB) are a function of earnings in the last year of life, τ is a payroll and income tax function, and the transfer function for retirees (T_R) is a function of the last earnings observation before retirement (which approximates DB pensions), aggregate earnings over the lifetime (which approximates social security income), assets, age, and family structure. Transfer functions (T and T_R) and tax function (τ) are year-specific and thus indexed by calender time t.

Timing: The number of children is exogenous in the model as is health insurance status. Household members are assumed to have perfect foresight on the entire paths of both fertility and health insurance, social security rules (SS), the defined benefit pensions function (DB), the time varying transfer functions (T and T_R) and time varying tax function τ. If the household is not retired, the household realizes its earnings shock at the beginning of each period and then makes decisions on consumption, health investments and leisure. The health shock is realized at the end of each period after the decisions have been made.

3.1 Working Household's Dynamic Programming Problem

A working single household between ages S and R obtains income from labor earnings and assets. The dynamic programming problem at age $j < R$ for a working single household is given by

$$W^{f,n_j,n_e,1}_{t,edu,g}(e_j, E_{j-1}, a_j, j, h_j) =$$

$$\max_{c_j, i_j, m_j^{oop}} \left\{ \begin{array}{c} n_j U(c_j/n_j, 1 - s(h_j) - \omega(h_j, g, 1, j, a_j, e_j) - i_j, h_j) + \\ \beta \int_{\varepsilon_{j,g}} \int_{e_{j+1}} C^{f,work}_{t+1,g}(j+1) d\Omega^{f,n_e,1}_{j,edu}(e_{j+1}|e_j) d\Xi_{j,g}(\varepsilon_{j,g}) \end{array} \right\}$$

subject to

$$y_j = e_j + ra_j + T_t(e_j, a_j, j, n_j)$$

$$c_j + a_{j+1} + m_j^{oop} = y_j + a_j - \tau_t(e_j + ra_j)$$

$$h_{j+1} = F_t(M^{ins}(m_j^{oop}), i_j) + (1-\delta)h_j - \varepsilon_{j,g}$$

$$E_j = E_{j-1} + e_j$$

where $C^{f,work}_{t+1,g}(j+1) = \Psi(h_{j+1})W^{f,n_{j+1},n_e,1}_{t+1,edu,g}(e_{j+1}, E_j, a_{j+1}, j+1, h_{j+1})$. In the above equation, $W^{f,n_j,n_e,1}_{t,edu,g}(e_j, E_{j-1}, a_j, j, h_j)$ denotes the expected present discounted value of lifetime utility for household f at age j in year t. E_{j-1} stands for cumulative earnings up to the current age while $C^{f,work}_{t+1,g}(j+1)$ gives the continuation value of the household. We integrate over health and non-health-related earnings shocks. The other variables are defined above.

The dynamic programming problem for a single household at age R, the last working period, is almost the same as the above dynamic programming problem at age $j < R$. The only difference is that, at age R, the continuation value, $W^{f,n_{j+1},n_e,1}_{t+1,edu,g}(e_{j+1}, E_j, a_{j+1}, j+1, h_{j+1})$, should be replaced by $V^{n_{R+1},1}_{t+1,g}(e_R, E_R, a_{R+1}, R+1, h_{R+1})$, the value function for single retirees introduced below, because the household will be retired in the next period.

Similarly, the dynamic programming problem at age $j < R$ for a working, married household is given by

$$W^{f,n_j,n_e,2}_{t,edu}(e_j, E_{j-1}, a_j, j, h_{h,j}, h_{w,j}) =$$

$$\max_{c_j, i_{h,j}, i_{w,j}, m^{oop}_{h,j}, m^{oop}_{w,j}} \left\{ \begin{array}{c} n_j \left[\begin{array}{c} \mu U(c_j/n_j, 1 - s(h_{h,j}) - \omega_{h,j} - i_{h,j}, h_{h,j}) + \\ (1-\mu) U(c_j/n_j, 1 - s(h_{w,j}) - \omega_{w,j} - i_{w,j}, h_{w,j}) \end{array} \right] + \\ \beta \int_{\varepsilon_{j,1}} \int_{\varepsilon_{j,2}} \int_{e_{j+1}} C^{f,work}_{t+1}(j+1) d\Omega^{f,n_e,2}_{j,edu}(e_{j+1}|e_j) d\Xi_{j,2}(\varepsilon_{j,2}) d\Xi_{j,1}(\varepsilon_{j,1}) \end{array} \right\}$$

subject to

$$y_j = e_j + ra_j + T_t(e_j, a_j, j, n_j)$$

$$c_j + a_{j+1} + m^{oop}_{h,j} + m^{oop}_{w,j} = y_j + a_j - \tau_t(e_j + ra_j)$$

$$h_{h,j+1} = F_t(M^{ins}(m^{oop}_{h,j}), i_{h,j}) + (1-\delta)h_{h,j} - \varepsilon_{j,1}$$

$$h_{w,j+1} = F_t(M^{ins}(m^{oop}_{w,j}), i_{w,j}) + (1-\delta)h_{w,j} - \varepsilon_{j,2}$$

$$E_j = E_{j-1} + e_j$$

where

$$C_{t+1}^{f,work}(j+1) =$$

$$\begin{aligned}
&\Psi(h_{h,j+1})\Psi(h_{w,j+1})W_{t+1,edu}^{f,n_{j+1},n_e,2}(e_{j+1}, E_j, a_{j+1}, j+1, h_{h,j+1}, h_{w,j+1}) \\
&+[1-\Psi(h_{h,j+1})]\Psi(h_{w,j+1})W_{t+1,edu,2}^{f,n_{j+1},n_e,1}(e_{j+1}, E_j, a_{j+1}, j+1, h_{w,j+1}) \\
&+\Psi(h_{h,j+1})[1-\Psi(h_{w,j+1})]W_{t+1,edu,1}^{f,n_{j+1},n_e,1}(e_{j+1}, E_j, a_{j+1}, j+1, h_{h,j+1})
\end{aligned}$$

$W_{t,edu}^{f,n_j,n_e,2}(e_j, E_{j-1}, a_j, j, h_{h,j}, h_{w,j})$ denotes the expected present discounted value of lifetime utility for a married household f at age j in year t. $\omega_{h,j} = \omega(h_{h,j}, g = 1, k = 2, j, a_j, e_j)$ and $\omega_{w,j} = \omega(h_{w,j}, g = 2, k = 2, j, a_j, e_j)$ are labor supply of the husband and wife respectively. The three product terms in $C_{t+1}^{f,work}(j+1)$ give the continuation values of the household when both the husband and the wife live to next period, only the wife lives to next period and only the husband lives to next period respectively. We integrate over health shocks facing both the husband and the wife and non-health-related earnings shocks facing the household. The other variables are defined above.

3.2 Retired Household's Dynamic Programming Problem

A retired single household between ages $R+1$ and death obtains income from social security, defined-benefit pensions, and assets. The dynamic programming problem at age j for a retired single household is given by

$$V_{t,g}^{n_j,1}(e_R, E_R, a_j, j, h_j) =$$

$$\max_{c_j, i_j, m_j^{oop}} \left\{ \begin{array}{c} n_j U(c_j/n_j, 1 - s(h_j) - i_j, h_j) + \\ \beta \int_{\varepsilon_{j,g}} \Psi(h_{j+1}) V_{t+1,g}^{n_{j+1},1}(e_R, E_R, a_{j+1}, j+1, h_{j+1}) d\Xi_{j,g}(\varepsilon_{,.g}) \end{array} \right\}$$

subject to

$$y_j = SS(E_R) + DB(e_R) + ra_j + T_{R,t}(e_R, E_R, a_j, j, n_j)$$

$$c_j + a_{j+1} + m_j^{oop} = y_j + a_j - \tau_t(SS(E_R) + DB(e_R) + ra_j)$$

$$h_{j+1} = F_t(M^{ins}(m_j^{oop}), i_j) + (1 - \delta)h_j - \varepsilon_{j,g}$$

In the above equation the value function, $V_{t,g}^{f,n_j,1}(e_R, E_R, a_j, j, h_j)$, denotes the expected present discounted value of maximized utility from age j until the date of death for this single household. Total earnings up to the current period are denoted by E_R while the last earnings draw at the age of retirement is e_R. Note that these values do not change once the household is retired. Relative to the working phase, household indicator f, number of earners n_e and education of the head edu do not appear in the value function during retirement because these variables only affect earnings.

Similarly, the dynamic programming problem at age j for a retired married household is given by

$$V_t^{n_j,2}(e_R, E_R, a_j, j, h_{h,j}, h_{w,j}) =$$

$$\max_{c_j, i_{h,j}, i_{w,j}, m_{h,j}^{oop}, m_{w,j}^{oop}} \left\{ \begin{array}{c} n_j \left[\begin{array}{c} \mu U(c_j/n_j, 1 - s(h_{h,j}) - i_{h,j}, h_{h,j}) + \\ (1-\mu) U(c_j/n_j, 1 - s(h_{w,j}) - i_{w,j}, h_{w,j}) \end{array} \right] \\ + \beta \int_{\varepsilon_{j,1}} \int_{\varepsilon_{j,2}} C_{t+1}^{retired}(j+1) d\Xi_{j,2}(\varepsilon_{j,2}) d\Xi_{j,1}(\varepsilon_{j,1}) \end{array} \right\}$$

subject to

$$y_j = SS(E_R) + DB(e_R) + ra_j + T_{R,t}(e_R, E_R, a_j, j, n_j)$$

$$c_j + a_{j+1} + m_{h,j}^{oop} + m_{w,j}^{oop} = y_j + a_j - \tau_t(SS(E_R) + DB(e_R) + ra_j)$$

$$h_{h,j+1} = F_t(M^{ins}(m_{h,j}^{oop}), i_{h,j}) + (1-\delta)h_{h,j} - \varepsilon_{j,1}$$

$$h_{w,j+1} = F_t(M^{ins}(m_{w,j}^{oop}), i_{w,j}) + (1-\delta)h_{w,j} - \varepsilon_{j,2}$$

where

$$C_{t+1}^{retired}(j+1) =$$

$$\begin{aligned}
&\Psi(h_{h,j+1})\Psi(h_{w,j+1})V_{t+1}^{n_{j+1},2}(e_R, E_R, a_{j+1}, j+1, h_{h,j+1}, h_{w,j+1}) \\
&+[1-\Psi(h_{h,j+1})]\Psi(h_{w,j+1})V_{t+1,2}^{n_{j+1},1}(e_R, E_R, a_{j+1}, j+1, h_{w,j+1}) \\
&+\Psi(h_{h,j+1})[1-\Psi(h_{w,j+1})]V_{t+1,1}^{n_{j+1},1}(e_R, E_R, a_{j+1}, j+1, h_{h,j+1})
\end{aligned}$$

In the above equation the value function, $V_t^{n_j,2}(e_R, E_R, a_j, j, h_{h,j}, h_{w,j})$ denotes the expected present discounted value of maximized utility from age j until the date of death for this married household. The three product terms in $C_{t+1}^{retired}(j+1)$ give the continuation values of the household when both the husband and the wife live to next period, only the wife lives to next period and only the husband lives to next period respectively. We integrate over the distribution of health shocks facing the husband and the wife in the married couple.

4 Model Parameterization and Calibration

In this section we specify functional forms and parameter values that we use to solve the model. We start by specifying functional forms for utility and health production. We then set some parameter values based on information from the literature or from reduced form estimates from the HRS. We identify the other parameters by fitting the predictions of the model for the typical household to data on wealth accumulation, medical expenses and survival probabilities. Once we have these parameter values, we then solve the model household-by-household and examine predictions for each household in our sample.

Preferences: We assume that momentary utility for a household member has a constant relative risk-averse form. We further assume the sub-utility function over consumption-leisure composite and health has a constant elasticity of substitution. Hence the period utility takes the form

$$U(c/n,h,l) = \frac{\{\lambda[(c/n)^{\eta}l^{1-\eta}]^{\rho} + (1-\lambda)h^{\rho}\}^{\frac{1-\sigma}{\rho}}}{1-\sigma} + B.$$

Following Hall and Jones (2007), B is a large enough constant to guarantee that utility is positive. The elasticity of substitution between the consumption-leisure composite and health is $1/(1-\rho)$. The discount factor (β) is set at 0.97, the value used in Hubbard, Skinner, and Zeldes (1995); and Engen, Gale, and Uccello (1999). We also set $\eta = 0.36$ from Cooley and Prescott (1995). Finally, we set σ, the coefficient of relative risk aversion equal to 3, a value commonly used in many studies including Hubbard, Skinner, and Zeldes (1995). We analyze the sensitivity of our results to β, σ and η. We calibrate B, ρ and λ.

Equivalence Scale: This is obtained from Citro and Michael (1995) and takes the form

$$n = g(A,K) = (A + 0.7K)^{0.7}$$

where A indicates the number of adults and K indicates the number of children in the household.

Rate of Return: We assume an annualized real rate of return, r, of 4 percent. This assumption is consistent with McGrattan and Prescott (2003), who find that the real rate of return for both equity and debt in the United States over the last 100 years, after accounting for taxes on dividends and diversification costs, is about 4 percent.

Taxes: We use a specification for the tax function taken from Gouveia and Strauss (1994). The specification for effective taxes for household f in year t with income y (in thousands) is:

$$\tau_{f,t} = b_t[y_{f,t} - (y_{f,t}^{-\rho_{\tau,t}} + s_t)^{-\frac{1}{\rho_{\tau,t}}}]$$

where b, ρ_τ, s are year-specific parameters to be estimated. To obtain the tax function parameters for our sample window, we assembled data from 1951 to 2007 using the Statistics of Income volumes available electronically through the Boston Public Library. For each year, the SOI data gives the mean tax liability for a range of income classes (AGI). These data were used to fit a tax function in each year: 1951 to 2007. The criteria for the fit was to minimize the sum of squared errors in the average effective tax rate:

$$\sum_f (\frac{\text{estimated tax}_f \text{ - observed tax}_f}{\text{taxable income}_f})^2.$$

Earnings and Earnings Expectations: Earnings data come from three sources: Social Security Administration Summary Earnings files, SSA earnings detail files (W2 information), and HRS self-reports. In the process of assembling the earnings data priority is given to each of these sources in the order listed. Earnings data in the Summary Earnings files is subject to top-coding. Before imputing the top coded earnings observations we first check to see if W2 earnings records exist; these data are available for most respondents starting in 1978. If W2 data is not available, HRS self-reports of earnings are used (if available).

The remaining top-coded earnings observations are split into two windows, 1951-1977 and 1978-2007. In the first period no top-coded earnings are recovered from W2 or HRS data. A censored regression model is estimated to predict the top-coded earnings in each year using the following covariates: gender, education, birth year, race, census region, marital status, average percentile in the earnings distribution over the past 5 years (if available), average percentile in the earnings distribution over the next 5 years (if available), number of children in the household, total years reported working, and average real household net worth over the HRS study years (1992,1994,...2008). The covariates used are taken from the first wave the respondent appears in the HRS.

In the second window, 1978 – 2007 many top-coded earnings observations are recovered using W2 data. An earnings model with the same covariates is estimated on the high-income observations that were recovered using W2 and HRS data. The parameters of these estimates are used to predict earnings for the high-income observations that remain top-coded. Starting in 1992 a new covariate, labor force status, is added and the covariates used for prediction are taken from the nearest HRS interview. Missing earnings are filled in when possible using HRS responses. Missing earnings in years following the respondents' last year of work or retirement year are set to zero. Missing earnings are set to zero for respondents who report never having worked. Missing earnings for respondents younger than age 17 are also set to zero. The remaining missing earnings are imputed via an earnings model using most of the variables listed above. The difference is that instead of using the spot in the earnings distribution, the respondent's average real earnings in the past/next five years are used when available.

Earnings expectations are a central influence on life-cycle consumption and health accumulation decisions, both directly and through their effects on expected pension and social security benefits.[9] We aggregate individual earnings histories into household earnings histories, putting earnings in constant dollars using the CPI-U. The household

[9] Due to data and computational limitations, we assume that earnings expectations are independent of health status. Credibly relaxing this assumption would require data on wage rates, hours, and health prior to when households enter the HRS.

model of log earnings (and earnings expectations) is

$$\log e_j = \alpha^f + \beta_1 j + \beta_2 j^2 + u_j$$

$$u_j = \rho_e u_{j-1} + \epsilon_j$$

where, as mentioned above, e_j is the observed earnings of the household f at age j in 2008 dollars, α^f is a household specific constant, u_j is an AR(1) error term of the earnings equation, and ϵ_j is a zero-mean i.i.d., normally distributed error term. The estimated parameters are α^f, β_1, β_2, ρ_e and σ_ϵ.

We divide households into six groups according to education, marital status and the number of earners in the household, resulting in six sets of household-group-specific parameters, which we then estimate separately for each of the five HRS cohorts (resulting in 30 sets of parameters).[10] Estimates of the persistence parameter, ρ_e, across groups range from 0.69 to 0.82.

Transfer Programs: We model public income transfer programs using the specification in Hubbard, Skinner and Zeldes (1995). Specifically, the transfer that a household receives while working is given by

$$T = \max\{0, \underline{c} - [e + (1+r)a]\}$$

whereas the transfer that the household receives upon retiring is

$$T_R = \max\{0, \underline{c} - [SS(E_R) + DB(e_R) + (1+r)a]\}$$

[10] The groups are (1) married, head without a college degree, one earner; (2) married, head without a college degree, two earners; (3) married, head with a college degree, one earner; (4) married, head with a college degree, two earners; (5) single without a college degree; and (6) single with a college degree. We estimate the parameters separately for the AHEAD, CODA, HRS, War Babies, and Early Boomer cohorts. A respondent is an earner if his or her lifetime earnings are positive and contribute at least 20 percent of the lifetime earnings of the household.

This transfer function guarantees a pre-tax income of $\underline{c}$ and implies that earnings, retirement income, and assets reduce public benefits dollar for dollar. To set $\underline{c}$ for each year we use information from Moffitt (2002) for 1960, 1964, 1968 to 1998 and extend the series using data from The Urban Institute, Mathematica Policy Research Inc., Center for Medicare and Medicaid Services, and the UKCPR National Welfare Data.[11] These data are at the state level so we take a weighted average according to state population in each year. Benefits have trended down since 1974 when the consumption floor for a single parent, two-child family peaked at \$14,767 (in year 2008 dollars). In 2007 the same family would have received transfers worth \$11,308.

Defined benefit pensions: Pension expectations and benefits come from an empirical defined-benefit pension function estimated with HRS data. The function includes indicator variables for having a defined benefit plan and belonging to a union, and variables for years in the pension by the retirement date, household earnings in the last year of work and the fraction of household earnings earned by the male and the fraction earned by the female.

Health Shocks: We assume health shocks follow a log normal distribution with mean $\mu^{\varepsilon}_{j,g}$ and variance σ^2_{ε}. Notice that we allow the mean to vary by gender and age. In practice, we discretize the support of log health shock, which is the real line, into five grid points and call the one that gives the worst health outcome the bad shock. These five grid points are fixed and do not vary over age or across gender. The probability of getting a bad health shock, however, varies both over age and across gender because of $\mu^{\varepsilon}_{j,g}$.

Health production: We assume that the production of health is given by $F_t(m,i) = A_t\left(m^{\chi}i^{1-\chi}\right)^{\xi}$, where total medical expenses are a function of out-of-pocket expenses, $m = M^{ins}(m^{oop})$ and health is also produced with time, i. We assume A_t grows at 2 percent per year reflecting aggregate improvements in productivity of health technology.

[11] See http://www.ukcpr.org/AvailableData.aspx

Total medical expenditures are related to out-of-pocket expenditures by a linear function that depends on insurance status. For the uninsured ($ins = 0$) this function takes the form, $m = \begin{cases} m^{oop} + \underline{m}, & \text{bad shock} \\ m^{oop}, & \text{no bad shock} \end{cases}$. In the absence of a bad health shock, health care expenditures come directly out of the uninsured household's pocket. In the event that the uninsured household suffers a bad health shock, a baseline level of care, $\underline{m}$, is provided via charity care.

For an insured household, total medical expenses are paid partially out of pocket and partially through insurance, $m = \underbrace{D + \zeta(m - D)}_{\text{OOP}} + \underbrace{(1 - \zeta)(m - D)}_{\text{insurance}}$. There are two parts of out-of-pocket expenses, the deductible D and a fraction $\zeta \in [0, 1]$ of the balance, $(m - D)$, that remains after the deductible has been paid.

We use the Medical Expenditure Panel Survey (MEPS) to calibrate the parameters of the medical expense model for six different insurance categories. Households in which the head is younger than 65 may be: uninsured, insured with public insurance only, or insured with any sort of private insurance. Three more categories capture older households: Medicare only, Medicare with supplemental public insurance (but no private), or Medicare and any private insurance.

To calibrate the value of charity care for the uninsured, we draw from Doyle (2005) who suggests the previous estimates "center around forty percent less care for the uninsured."[12] The average total medical spending for the insured (under age 65) in the event of a health shock in the 2008 MEPS data was $\overline{m}_i = \$3,768$. Average out-of-pocket spending for the uninsured was $\overline{m}_u^{oop} = \$861$. Using the relationship that $0.6\overline{m}_i = \overline{m}_u = \overline{m}_u^{oop} + \underline{m}$ we recover the average value of charity care in the event of an adverse health shock, $\underline{m} = \$1,400$.

To calibrate the "generosity parameter," ζ, for each of the insurance types, we use

[12] See, for example, Currie and Gruber (1997), Currie and Thomas (1995), Haas and Goldman (1994), Long, Marquis, and Rodgers (1997), and Tilford *et al.* (1999) who provide information on medical care use for the insured and uninsured.

estimates of the average deductible, average total medical spending and average out-of-pocket spending. The spending model implies that $m^{oop} = D + \zeta(m - D)$ which can be rewritten to solve for $\zeta = \frac{m^{oop}-D}{m-D}$ for each insurance types. The resulting values are $\zeta = 0.039$ for households under 65 with any private insurance; $\zeta = 0.063$ for households under 65 with only public insurance; $\zeta = 0.159$ for households over 65 with Medicare only; $\zeta = 0.145$ for households over 65 with Medicare and some private insurance; and $\zeta = 0.042$ for households over 65 with Medicare and supplemental public insurance.

Survival Probability: The survival function is given by the cumulative distribution function $\Psi(h) = 1 - \exp(-\psi_1 h^{\psi_2})$.

Working Time: As mentioned previously, working time $\omega(h, g, k, j, a, e)$ depends on health, gender, marital status, age, assets, earnings and union status and this is calibrated from PSID.

Sick Time: We assume that the amount of sick time is given by $s(h) = h^{-\alpha}$.

Initial conditions: The age S at which a household enters the labor market is taken to be the age of the household head when we first observe the household in our data, and thus could vary across households. Initial assets are set to be zero for all households. The initial stock of health is assumed to be the same for all husbands and wives. Other individual level heterogeneities include education, gender, marital status and age of retirement and household level heterogeneities include health insurance status and number of children. As mentioned earlier, these are taken to be what they were when a household first enters the HRS.

4.1 Calibration

While several parameters are set based on estimates from the literature or by estimating reduced form empirical models from the HRS, additional critical parameters still need to be specified. We use information on asset holdings, life tables and medical expenses for the typical household in the HRS to pin down these parameters. The 19 parameters we calibrate are λ, the utility weight on consumption relative to health; ρ, which determines the elasticity of substitution between consumption and health; μ, the weight on the husband in the household utility function; B, the constant in utility to guarantee that it is positive; ψ_1, the coefficient on health in the survival function; ψ_2, the curvature of the survival function with respect to health; ξ, the curvature of the health production function; χ, the share parameter of monetary input in health production; δ, the annual depreciation rate of health; α, the elasticity of sick time with respect to health status; σ_{ε}, the standard deviation of the i.i.d health shock, $\mu^{\varepsilon}_{65,1}, \mu^{\varepsilon}_{75,1}, \mu^{\varepsilon}_{85,1}, \mu^{\varepsilon}_{85+,1}$, the mean of the health shock for men less than 65 years of age, between 65 and 75, between 75 and 85 and above 85 respectively; and $\mu^{\varepsilon}_{65,2}, \mu^{\varepsilon}_{75,2}, \mu^{\varepsilon}_{85,2}, \mu^{\varepsilon}_{85+,2},$, the corresponding values for women.

To calculate these remaining parameters, we solve the dynamic programming problem for the 'typical' married, single male, and single female households, where 'typical' is defined as the household with average earnings and medical expenses over their lifetimes. We then use the decision rules in conjunction with observed histories of earnings and medical expenses to obtain model predictions. Notice that while we have earnings observations on an annual basis, we only have medical expenses starting in 1992. Hence we integrate out the lifetime sequence of health shocks before arriving at the model predictions for a given age. We then seek to obtain the best fit between model and data relative to the moments we seek to match for these three types of households in 1998. We emphasize that the implicit assumption employed in our strategy is that households are identical in terms of preferences and technology but face different constraints due to the evolution of shocks in the face of incomplete markets. Males differ from females in terms of the probabilities of bad health shock as they age to account for the greater longevity of women relative to men.

The moments we use to identify and pin down the parameters are:[13]

[13] Moments for net worth data and the retirement age come directly from HRS data. Mo-

Table 1A: Moments

Moments	Value
Median net worth in 1998 for married couples (husband age 63.2, wife age 60.9)	\$246,312
Median net worth in 2008 for married couples	\$281,200
Median net worth in 1998 for single males (age 64.1)	\$91,740
Median net worth in 1998 for single females (age 66.7)	\$81,708
The probability of dying between ages 50-54 for males	3.08%
The probability of dying between ages 70-74 for males	13.76%
The probability of dying between ages 80-84 for males	31.69%
The probability of dying between ages 90-94 for males	60.70%
The probability of dying between ages 50-54 for females	1.834%
The probability of dying between ages 70-74 for females	9.57%
The probability of dying between ages 80-84 for females	23.94%
The probability of dying between ages 90-94 for females	52.05%
Average annual total medical expenses for married women age 60-64	\$7,747
Average annual total medical expenses for married women age 70-74	\$12,417
Average annual total medical expenses for married women age 80+	\$17,896
Average annual total medical expenses for single women age 70-74	\$12,479
Average annual total medical expenses for married men age 70-74	\$13,255
Average annual total medical expenses for single men age 70-74	\$13,474
Sick hours relative to total work hours at age 40	0.015

The model with each calibrated parameter generates 19 non-linear equations with

ments for total medical expenses come from the National Health Expenditure Accounts (using the personal health spending totals by age for 2004), drawn from http://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/Age-and-Gender-Items/CMS1242122.html. We disaggregate the National Health Expenditure Account total for married men, married women, single men, and single women using data from the Medical Expenditure Panel Study. The mortality probabilities come from the World Health Organization, http://www.who.int/healthinfo/statistics/mortality_life_tables/en/. The data on sick hours come from the PSID.

19 unknowns. We obtained an exact match between the model predictions and the moments listed above. The resulting parameter values are given below.

Table 1B: Calibrated Parameter Values										
Parameter	λ	ρ	μ	α	B	$\mu^{\varepsilon}_{65,1}$	$\mu^{\varepsilon}_{75,1}$	$\mu^{\varepsilon}_{85,1}$	$\mu^{\varepsilon}_{85+,1}$	δ
Value	0.70	-3.6	0.43	0.17	32.1	0.30	0.42	0.65	0.86	0.034
Parameter	ψ_1	ψ_2	ξ	σ_{ε}	χ	$\mu^{\varepsilon}_{65,2}$	$\mu^{\varepsilon}_{75,2}$	$\mu^{\varepsilon}_{85,2}$	$\mu^{\varepsilon}_{85+,2}$	
Value	0.0012	1.53	0.69	1.22	0.61	0.22	0.36	0.54	0.72	

The elasticity of substitution between consumption/leisure composite and health is $\frac{1}{1-\rho} = 0.22$. Later on in the paper, we analyze the effects of changes in ρ to better understand it's effect. The change in wealth between 1998 and 2008 helps identify this parameter. Consumption and health are complements and our calibrated value is very close to the estimates in Finkelstein, Luttmer, and Notowidigdo (2013). In one of the first few papers that simulates a model with endogenous health, Yogo (2009) employs recursive preferences in a study of portfolio choices in retirement and finds that consumption and health are complements in utility. Since Yogo (2009) uses Epstein-Zin preferences, his estimates are not directly comparable to ours (we use time additive separable utility) but it is comforting to note that a substantially different approach also finds evidence in favor of complementarity. In a married household, the weight on the man's utility is 0.43, lower than the weight on the woman's utility. The rate of depreciation of health is 3.4 percent per year. The share of goods in the production of health χ is 0.61, suggesting that time and goods are both important in the production of health. Finally, note that the probability of the bad health shock increases with age since the mean (in logs) rises from 0.3 for men less than 65 to 0.86 for men above 85. For women, the same object rises from 0.22 to 0.72. The smaller probability of a bad health realization at any given age for women relative to men is instrumental in matching the higher age specific mortality rates for men.

As mentioned above, we match 19 data moments with the model to identify these

19 parameters. Clearly, altering one of the target data moments changes more than one parameter. Nevertheless, it is instructive to think about which data moments play a critical role for at least some of the more important parameters.

A lower value of ρ will lead to a higher level of assets in 1998. In addition, a lower value of ρ will have implications for asset accumulation/decumulation late in life. Predictable declines in health ought to be associated with predictable declines in consumption. Hence having asset levels in 1998 as well as 2008 helps pin down ρ.

The parameters governing the production technology for health (for males) as well as the hazard function are pinned down by the mortality probabilities as well as medical expenses. Recall that health affects utility as well as mortality. The importance of health in utility (λ) as well as the significance of health in improving longevity are both simultaneously pinned down by these moments. The probabilities of dying as people age interact with the technology for producing health to determine medical expenses. For instance if diminishing returns set in quickly, substantial medical expenses need to be expended simply to maintain the stock of health. In contrast, if the medical technology were close to linear, then additional medical expenses will have a large effect on the stock of health. Hence, all these objects (medical technology parameters, importance of health relative to consumption in utility), as well as the parameters of the hazard function, are simultaneously pinned down by the probabilities of the bad shock and medical expenses as men age.

Medical expenses for single women and probabilities of dying for men relative to women help pin down the probabilities of bad health shocks for women. In addition, mean net worth for singles relative to married couples shed light on the utility aggregator in preferences. A change in the parameter governing the importance of men relative to women in a married households (μ) will affect both the wealth of the married households as well as medical expenses for married women relative to single women.

4.2 Model Solution

With the calibrated parameters, we solve the dynamic programming problem by linear interpolation on the value function. For each household in our sample we compute optimal decision rules for assets and the stock of health from the oldest possible age (assumed to be 120) to the beginning of working life (S) for any feasible realizations of the random variables: earnings and health shocks. Recall that initial assets at age S are zero and initial health capital is normalized to the same value for all individuals at this age. These decision rules differ for each household, since each faces stochastic draws from different earnings distributions (recall they are household specific). Household-specific earnings expectations also directly influence expectations about social security and pension benefits. Other characteristics also differ across households - the age of retirement, the number of children and the ages at which these children enter and leave the household. We then use the decision rules in conjunction with the observed earnings and medical expenses to obtain the model's predictions for wealth, health, medical expenses and mortality at a given age. Since we do not have data on medical expenses before 1992, we integrate out the health shocks over this time period. To be sure, consider a working single household. Recall that the state variables are e_j, E_{j-1}, a_j, j and h_j. We start at age S with $a_j = 0$ and $h_S = \overline{h}$. The decision rule for assets is given by $a_{j+1} = A_j(e_j, E_{j-1}, a_j, j, h_j)$. We have annual observations on earnings, e_j. Knowledge of e also means we have knowledge of E. Since we do not observe health shock $\varepsilon_{j-1,g}$, which affects h_j, we integrate out the health shock and assume that $a_{j+1} = \int A_j(e_j, E_{j-1}, a_j, j, h_j) d\Xi_{j-1,g}(\varepsilon_{j-1,g})$. Beginning in 1992 (when households are around 56 years of age), we observe the medical expenses chosen by the household. From this point, we use observed medical expenses to back out the health shock. Suppose that $\widehat{m}_{j+1}$ is the observed medical expense at age $j+1$. Then $m^{oop}_{j+1}(e_{j+1}, E_j, a_{j+1}, j+1, h_{j+1}) = \widehat{m}_{j+1}$, where $h_{j+1} = F_t(M^{ins}(m^{oop}_j), i_j) + (1-\delta)h_j - \varepsilon_{j,g}$.

5 Results

As emphasized in the previous discussion, we calibrate key model parameters to the typical (married, single male and single female-headed) HRS household in 1998. The first question we address, therefore, is how the model matches household wealth, out of pocket medical expenses, and the stock of health.

5.1 Net Worth and Medical Expenses

We summarize results for household wealth and out-of-pocket medical expenses by showing median values, breaking households into lifetime income quintiles.[14] In Table 2 we present a comparison of the cross-sectional implications of the model in 1998 and in 2008. The 1998 cross-section is made up of the household heads from all birth cohorts that participated in the 1998 HRS interview (n = 9,041) and likewise for the 2008 cross-section (n = 11,172, our full HRS sample). The vast majority of the difference (2,131 households) are households in the "early boomers" cohort who were added to the HRS in 2004 and hence are not a part of the 1998 cross-section.

[14] Lifetime income is defined within four roughly equal-sized age groups: under 60, 60 to 65, 66 to 75, and over 75.

Table 2: Median Net Worth and OOP Medical Expenses, 1998 and 2008				
1998	Median Net Worth		Median OOP Medical Expenses	
Lifetime Income	Data	Model	Data	Model
Lowest Quintile	$33,588	$31,456	$413	$421
Second Quintile	60,717	53,483	646	729
Middle Quintile	97,212	93,708	904	967
Fourth Quintile	180,859	163,695	1098	1081
Highest Quintile	340,144	353,129	1,219	1,235
2008				
Lowest Quintile	$15,495	$16,394	$370	$356
Second Quintile	63,900	61,304	775	844
Middle Quintile	136,000	132,453	1,083	1,137
Fourth Quintile	238,000	248,120	1,425	1,413
Highest Quintile	443,000	432,230	1,795	1,735

There are two striking features of Table 2. First, while we calibrate the model to the average household in 1998, the model does a good job matching the wide variation in wealth across low and high lifetime income households in 1998. The correlation of actual and optimal net worth in 1998 is 0.71. Scholz, Seshadri, and Khitatrakun (2006) report a correlation between model predications and net worth in the HRS of 0.86 in 1992. There are a number of differences between our earlier work and this paper. The most important is that health affects utility and longevity, households make endogenous health investments, we model the health decisions of spouses, new cohorts have been added to the data and we now look at a more recent period, and we have new estimates of the earning process, which show somewhat more volatility in earnings than our previous estimates, among other changes. Despite these differences our earlier qualitative conclusion still holds: Most Americans appear to be preparing for retirement in a manner consistent with our life-cycle model given the current policy environment.

Predicted median out-of-pocket medical expenses also match actual expenses fairly closely. For instance, in 1998, the out of pocket medical expenses rise from $421 for the lowest lifetime income quintile to $1,235 for the highest income quintile. This tracks the data pretty closely. Richer households spend more out of pocket (despite possessing better health on average at the same age) and these investments affect both flow utility as well as longevity. The household-by-household correlation between actual out-of-pocket medical expenditures and optimal out-of-pocket medical expenditures in the model is 0.47.

The second striking feature of Table 2 is the degree to which we match the dispersion of median net worth and out-of-pocket medical expenditures by lifetime income quintile at a later date (2008). We use only one net worth moment for 2008 (the net worth of married couples): health expenses are for 2004 (due to the timing of the National Health Expenditure Accounts). Yet the behavioral model augmented with preference parameters calibrated to the average household in 1998, data on changes in household composition, and earnings realizations (for those still in the labor market) is able to closely match the 2008 distribution of median net worth and out–of-pocket health spending.

5.2 Health Status

Another feature of the HRS are questions on self reported health status, which we used in Figures 1 and 3. Households report this on a 5 point scale ranging from poor to excellent. In the model, the stock of health is a continuous variable and hence to compare with the data, we turn the continuous health variable into a discrete one. In the HRS data, 13 percent of the sample report excellent health, 28 percent report very good, 30 percent report good, 19 percent report fair and 9 percent report poor. We choose the cut-off points in the continuous distribution so that these percentages are what we observe in

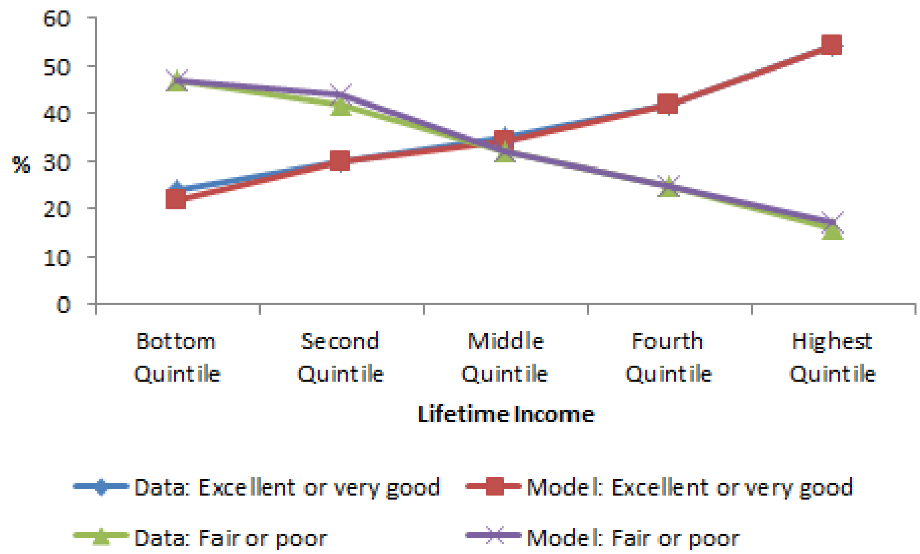

Figure 5: Self-Reported Health Status in 1998: Model vs Data

the HRS. Figures 5 and 6 depict the relationship between model and data for 1998 and 2008 in greater detail.

There is a very tight link between lifetime income and the self-reported health status and the model does an excellent job tracking the variation in the data. Various model features come into play here - as households age, they receive adverse shocks with greater intensity. Their ability to buffer these shocks depends largely on health investments they had made in the past (which determines their current health status) as well as their income. The pace with which health deteriorates in older ages also affects consumption (recall that consumption and health are complements) which in turn affects wealth accumulation. The fact that the model is able to match the extent to which health worsens between 1998 and 2008 adds to our confidence that the model provides a reasonable description of the evolution of health by lifetime income.

Our model makes predictions not just during the retirement phase but also throughout the working phase of the life-cycle. Unfortunately, the HRS data begin in 1992 and consequently we do not have information on the behavior of these households while

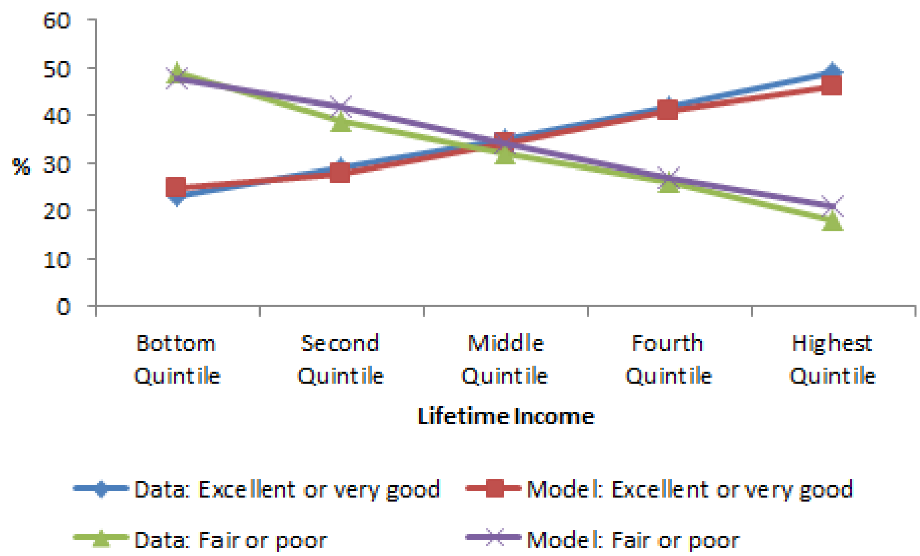

Figure 6: Self-Reported Health Status in 2008: Model vs Data

they are working. Nevertheless, it is instructive to compare model predictions with best available data.

5.2.1 Health During the Working Phase

The Panel Study of Income Dynamics (PSID) also contains information on a self-reported health status (5 point scale) much like the HRS data. Measures of the distribution of health by age and income come from the 2009 wave of the PSID. The PSID is a longitudinal panel that began in 1968. By 2009 the sample size has grown to include more than 9,000 families. The analysis in this subsection will compare model simulations with responses from 13,055 PSID respondents in 2009. The health measure gives respondents' perceived health on a Likert scale (excellent, very good, good, fair, poor). Income quintiles are defined using the household head's 2008 labor, business, and farm income.

Before proceeding with the comparison, it is useful to note that while the PSID data

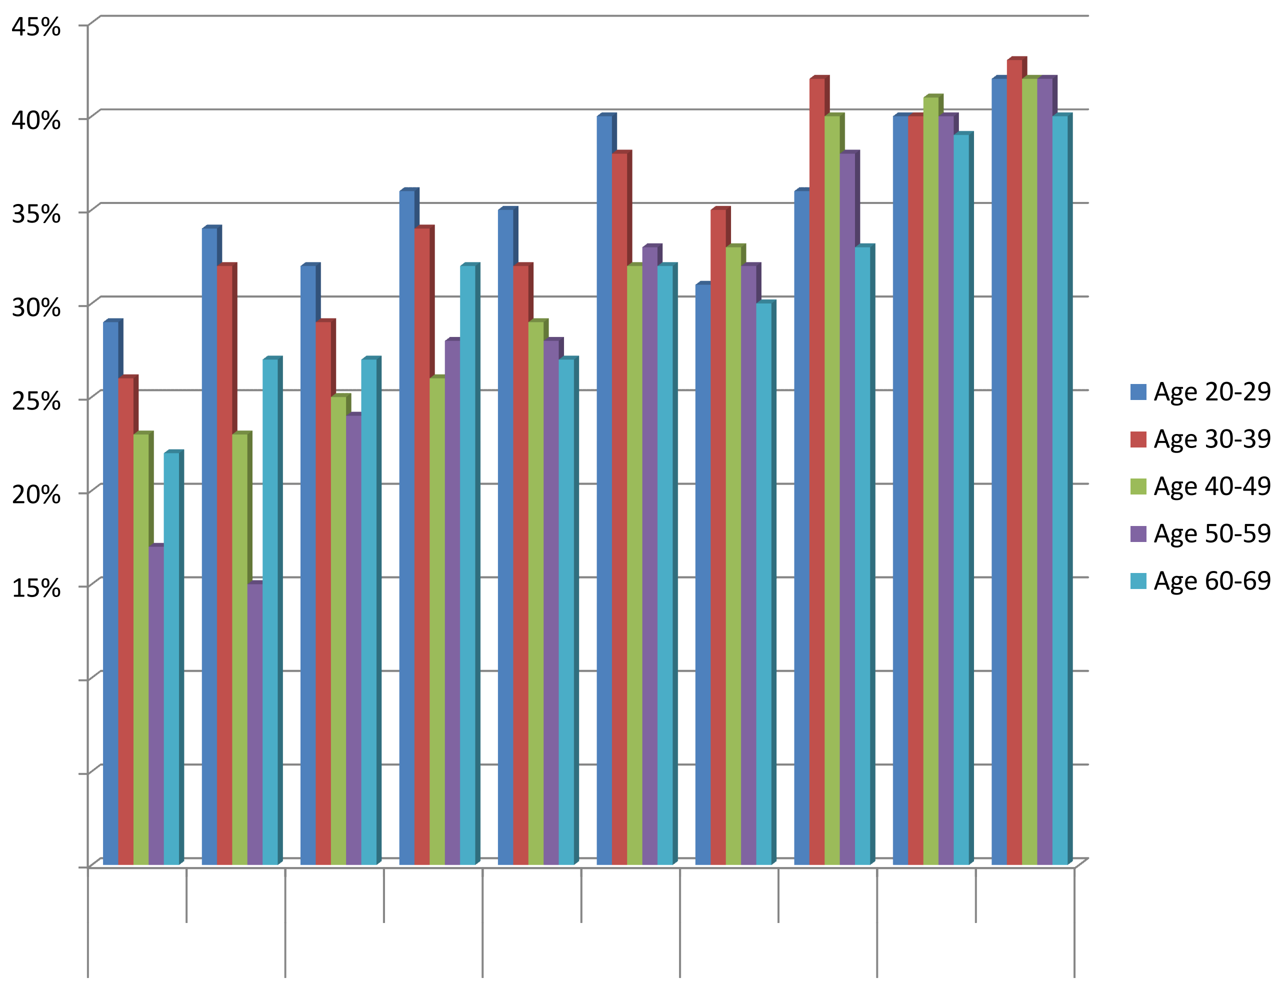

Figure 7: % Reporting 'Very Good' Health Status: Model vs PSID Data

are for the 2009 cross-section, the model simulations we present are for cohorts born much earlier. Consequently, the policies and opportunities faced by these households are quite different from the current cross-section and hence there could have been large differences in the distribution of health status by age. In the interests of space, we report the comparison between model and data for one of the self-reported health status: "very good". The results are in Figure 7.

The Figure presents a comparison between self-reports of "very good" in the PSID with it's model counterpart at various ages for 5 income quintiles (these are income levels at that age and not lifetime income). While the fit is not all that great, the model is able to track the declines in health by age as well as evolution of health with income. The fit is a bit better at higher income levels.

5.3 Consumption in Retirement

One other feature of our data set is the availability of consumption data. While consumption data are not available for the entire sample under study, hence making it impossible to compare model and data household by household, consumption data are available for a sub-sample of the population. Wave nine of the Consumption and Activities Mail Survey (CAMS) was completed by 3,587 individuals on behalf of their household. Wave nine CAMS respondents report 2009 household spending in 39 categories of nondurables and durable goods. We calculate total household spending on these categories weighted using CAMS household weights which adjust for both sample design and non-responses to both the HRS as well as the CAMS survey. The household spending is thus representative of American households aged 50 and above in 2009.

We normalize the consumption of an average household in the third quintile to 1 and report in Table 3 the consumption levels of households in the 5 different lifetime earnings quintiles for each of the self-reported health status for the cross-section of households in 2008. As can be seen from Table 3, consumption rises with income for each health status and consumption rises with health status for each income quintile. The fit between model and data is fairly good and the fact that consumption co-moves with health status in the cross-section is attributable to consumption-health complementarity in preferences.

Table 3: Mean Consumption by Health Status and Lifetime Earnings Quintile

	1st Quintile		2nd Quintile		3rd Quintile		4th Quintile		5th Quintile	
	Model	Data	Model	Data	Model	Data	Model	Data	Model	Data
Excellent	0.95	0.93	1.04	1.02	1.16	1.15	1.24	1.29	1.54	1.68
Very Good	0.79	0.75	0.93	0.91	1.12	1.20	1.20	1.23	1.44	1.55
Good	0.71	0.69	0.91	0.88	1.01	1.00	1.15	1.13	1.36	1.38
Fair	0.65	0.64	0.84	0.73	0.92	0.87	1.05	1.04	1.31	1.29
Poor	0.60	0.61	0.76	0.71	0.83	0.79	0.96	0.92	1.28	1.24

5.4 Medical Expenses During Working Years

We have information on the distribution of out-of-pocket medical expenses by age and income from the 2008 Household Component of the Medical Expenditure Panel Survey (MEPS). MEPS is nationally representative for the U.S. civilian noninstitutionalized population. The calculations shown in Table 4 under the columns labeled 'Data' were derived using the MEPSnet Query Tools and public use file HC121 with sample size 12,696 (2008 Full Year Consolidated Data File). The medical expenditure variable includes the total amount paid by the individual or their family for: medical provider visits, hospital outpatient visits, hospital emergency room visits, hospital inpatient stays, dental visits, home health care, vision aids, other medical equipment and services, and prescribed medicines. Income is a comprehensive measure of person-level income.

The same caveat that applied to the comparison between model and data for health during working years applies here as well - the data are from a cross-section while the model simulations are for the HRS cohorts. The model is able to track the rise in medical spending by age as well as the variation by income. Rather interestingly, the model's fit for the age range 60-69 is quite a bit better than for the other ages. We attribute it to the fact that the households we simulate are roughly in the 60-69 age range in 2008.

Table 4: Median Out of Pocket Medical Expenses by Age and Income Quintile										
	1st Quintile		2nd Quintile		3rd Quintile		4th Quintile		5th Quintile	
	Model	Data	Model	Data	Model	Data	Model	Data	Model	Data
Age 20-29	$131	197	145	206	156	221	198	283	202	300
Age 30-39	177	228	168	226	182	253	231	270	208	295
Age 40-49	199	268	201	336	213	315	274	333	234	354
Age 50-59	285	446	267	438	342	585	365	529	312	514
Age 60-69	652	678	635	669	647	636	701	722	671	749

5.5 Mortality

A novel feature of our economic model is that it allows us to examine the effects of policy changes on mortality. But the confidence readers have with our mortality results will depend, in part, on the ability of the model to reproduce mortality patterns in the HRS. To examine this, we take 10-year mortality probabilities in the HRS for two groups – those who are 60 years old and those who are 75 years old. Specifically, we restrict the sample to people first observed in the HRS before (or when) they reach age 60 and who, conditional on survival, would have been at least 70 in 2008. We make similar calculations for the age 75 sample. The entries in the table below under "Data" give the survival probabilities by lifetime income quintile.

The mortality calculations implied by the model require considerable calculation. For example, in the first two columns of Table 5 we take all 60 year olds. These households face many different patterns of potential health shocks ($\varepsilon_{j,g}$ paths). We integrate out over all potential sequences between the ages 60 and 70 and calculate the mass of survivors. These calculations require, of course, the optimal decision rules over the lifetime of households. We make similar calculations for households age 75. The survival rates implied by the model are given in Table 5 under the column "Model."

Table 5: Ten-Year Survival Probabilities, Model and Data				
	Age 60		Age 75	
Lifetime Income Quintile	Data	Model	Data	Model
Bottom	0.77	0.76	0.54	0.52
Second	0.83	0.81	0.54	0.53
Middle	0.86	0.84	0.52	0.55
Fourth	0.90	0.87	0.62	0.60
Highest	0.92	0.89	0.64	0.62

The model does a strikingly good job matching survival patterns in the underlying data, though we note that seven of the 19 moments that we use to calibrate the model tie down mortality probabilities by age for households with average lifetime incomes. This does not, however, imply that we would expect the model to reproduce survival patterns for high- or low-lifetime income quintile households. Both at age 60 and 75, there are substantial deviations between the survival data and predictions for households in the highest lifetime income quintiles. These are likely to be the households that are most efficient in producing health capital. At age 75 there is also a substantial deviation between data and model in the lowest lifetime income quintile. This is the pattern we expect to see as unobservable efficiency in health investment should make low-income households in the HRS who survive to age 75 healthier than the average low-income households in the model.

5.5.1 Mortality During Working Years

In our model, health shocks get increasingly likely as households age. What does our model say about mortality at younger ages? The National Office of Vital Statistics publishes life tables and these tables are available for the cohort born between 1939-41. In Table 6 we present the 10 year survival probabilities for this cohort and compare that with the model-implied survival probabilities for the HRS cohort born between 1931 and

1941 for men and women. Recall that men and women draw health shocks from different distributions. Before age 65, these shocks do not vary by age. The main age effect is the depreciation in health capital that induces different investments in health capital as individuals age and hence makes individuals more susceptible to health shocks as they age. Table 6 presents the comparison between model and data. The model implied survival probability closely tracks the life tables for the cohort born at approximately the same time adding further credibility to our modeling of health.

Table 6: Ten Year Survival Probabilities: Model (HRS cohort) vs. 1939-41 cohort

Age	Men		Women	
	Model	Data	Model	Data
20	0.98	0.97	0.99	0.98
30	0.95	0.96	0.97	0.97
40	0.91	0.92	0.93	0.94
50	0.84	0.83	0.89	0.88

Source: United States Life Tables and Actuarial Tables, 1939-1941.
http://www.cdc.gov/nchs/data/lifetables/life39-41_acturial.pdf

5.6 Complementarity

The degree of complementarity between consumption and health in the utility function plays a central role in driving our results. In Table 7 we illustrate the effect of setting ρ to zero (which would make the consumption-leisure composite and health separable in the utility function) on optimal net worth.

Table 7: The Effect of Complementarity on Optimal Wealth Accumulation, 1998		
	Median Net Worth	
	Baseline Model ($\rho = -3.6$)	$\rho = 0$
Bottom Quintile	\$31,456	\$51,463
Second Quintile	53,483	112,342
Middle Quintile	93,708	142,475
Fourth Quintile	163,695	302,253
Highest Quintile	353,129	511,391

When consumption and health are complements in the utility function, households anticipate that as health declines late in life, so will consumption. This phenomenon is absent when consumption and health are separable in preferences. Hence the model with complementarity will, all else equal, imply less asset accumulation than when ρ equals zero. Other moments are also affected. To illustrate this, we set ρ to zero and re-calibrate the model. We ignore the 2008 wealth moment and solve 18 equations in 18 unknowns. We then re-do simulations for all households in the sample. We report below the fit of the model (baseline in parenthesis)

1. The correlation between model and data for net worth in 1998: 0.49 (0.71)

2. The correlation between model and data for out-of-pocket medical expenses in 1998: 0.31 (0.47)

3. The model correctly predicts the stock of health (on a 5-point scale) for 50 (69) percent of the population in 1998

Complementarity has an important effect on the fit of the model for net worth, out-of-pocket medical expenses, and the stock of health.

Another way to illustrate the importance of consumption-health complementarity is to examine *changes* in wealth between 1998 and 2008. As individuals age, health

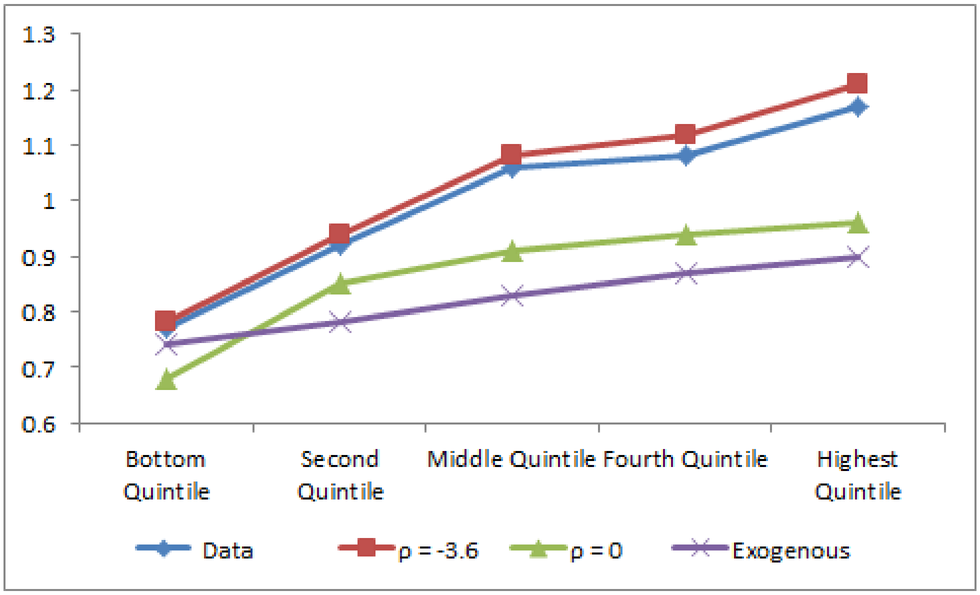

Figure 8: Median Ratio of Net Worth in 2008 to Net Worth in 1998

deteriorates. One might be worried that this deterioration in health might be associated with too steep a change in wealth when health and consumption are complements. To examine this, we compute the ratio of wealth in 2008 to wealth in 1998 for each household that is alive in 2008. We then proceed to arrange them by lifetime income quintile and report the median ratio in Figure 8.

The variation in the ratio of 2008 net worth to 1998 net worth in the data is steep across lifetime income quintiles: poor households experience a substantial decline in wealth while their richer counterparts experience a 17 percent increase in wealth. The baseline model with $\rho = -3.6$ tracks this pattern very well. For poor households, who are closer to death than a comparable rich household, it is optimal for them to run down their assets. For richer households complementarity implies that as health deteriorates, consumption declines. This results in an increase in asset accumulation, in part because richer households have longer expected lifetimes than poor households.

Figure 8 also reports what happens with separable utility, where $\rho = 0$. With separable utility the model has difficulty generating increases in wealth between 1998 and 2008 - even the rich households decumulate wealth at older ages (although at a much slower pace than poorer households). Indeed, the incorporation of a target moment for wealth in 1998 and in 2008 helps identify the degree of complementarity. Finally, to illustrate the importance of making choices that affect the stock of health, we also report the same statistics for the conventional modelling of consumption and health in a lifecycle model. Specifically we show results from a model where lifetime budgets are subject to medical expense shocks that proxy for health shocks. With exogenous health, the variation in the change (across time) in net worth between rich and poor households is smaller than in the separable utility case since medical expenses do not influence longevity.[15] We conclude that the assumed degree of complementarity is reasonable and plays an important role in explaining the behavior of mortality, health, medical expenses as well as health for older households.

5.7 The Building Blocks of the Model of Health

Before proceeding to some policy simulations, it is instructive to analyze the significance of each of the components of the model of health. We focus on three central components:

1. Health shocks (ε),
2. The impact of health on utility $(1 - \lambda)$, and
3. The impact of health on survival (ψ_1)

To examine the importance of each component in accounting for the fit between model and data, we set each of these to zero and redo our simulations. Table 8 presents

[15] Indeed, even the recent work of De Nardi *et al.* (2010) has a difficult time rationalizing the increase in asset holdings of rich singles. While their model does a very nice job explaining the behavior of the bottom quintiles, the top quintile experiences an increase in assets late in life while their model generates a decline over the same time period. We note here that these are unbalanced panels that reflect mortality bias. Recall that our sample also includes married households who are richer on average than single households. This, in large measure, explains why even our average household experiences an increase in asset holdings while in their sample, the average household does not.

the fit between model and data when each element is suppressed.

Table 8: The Contribution of the Different Elements of the Health Model			
Parameter	R^2(Health)	R^2(Med. Exp.)	R^2(Wealth)
Baseline	0.69	0.47	0.71
$\varepsilon = 0$	0.34	0.23	0.34
$1 - \lambda = 0$	0.43	0.34	0.56
$\psi_1 = 0$	0.52	0.37	0.59

When we shut down health shocks altogether, the fit worsens quite a bit. The Euclidean distance between model and data for Net Worth declines from 0.71 to 0.34. The ability to fit health stock also worsens considerably. Shutting down the contribution of health to utility as well as the effect of health on survival (when we shut down this mechanism, we replace the survival function with standard lifetables for the cohort born between 1939-41 as described under Table 6 which incorporates differential mortality by sex) also leads to substantial effects but these effects are smaller in magnitude relative to the impact of eliminating health shocks. The upshot is that each of these components adds considerable value in augmenting the standard life-cycle model to make it better suited to match the cross-sectional dispersion in health and medical expenditures.

6 Medicare and Longevity

Policy simulations are another way to gain insight into behavior and the way our model works. In Table 9 we examine the effect of removing our stylized version of Medicare, the universal social insurance program that was established in 1965 to provide health insurance to the elderly. Our modeling of this change is extreme: instead of being eligible for a Government-provided insurance program upon reaching age 65, households are uninsured.[16] There are several reasons why we focus on this policy. First, Medicare is a massive social insurance program costing $325 billion in fiscal year 2006. Second, end-of-life health shocks have been shown by several authors to have significant effects on asset accumulation. Third, Finkelstein and McKnight (2008) show in the first 10 years following the establishment of Medicare, there was no discernible effect on mortality, though Chay *et al.* (2010) challenge this conclusion. The effects of policy changes on mortality and asset accumulation in the short- and long-run are issues the model is nicely designed to address.

Suppose that Medicare were instantly eliminated in 1998 and the change was not anticipated. All assets and health capital held by households had been accumulated under the assumption that Medicare would exist. Medicare is financed by taxes on earned income and so when Medicare is eliminated, we rebate annually back to the household all future Medicare tax payments. After eliminating Medicare, we can recompute the model and examine the effects on 10-year survival probabilities. To be sure, our analysis is predicated on two considerations. First, if private markets for insurance do not exist or are not well functioning, mandatory programs can improve welfare. Second, Medicare has a redistributive component associated with it which can be welfare improving in the presence of incomplete markets (with complete markets, these programs will be distortionary in the traditional sense). Indeed, we are by no means the first to examine the welfare consequences of exogenously incomplete markets though our application is novel. We will mention three notable applications here. Aiyagari (1995) demonstrates that the optimal tax on capital in the long run is positive with incomplete markets.

[16] As noted earlier, medical care for insured households is $m = \underbrace{D + \zeta(m - D)}_{\text{OOP}} + \underbrace{(1 - \zeta)(m - D)}_{\text{insurance}}$. Total medical expenditures for the uninsured are $m = \begin{cases} m^{oop} + \underline{m}, & \text{bad shock} \\ m^{oop}, & \text{no bad shock} \end{cases}$. This policy experiment is extreme in the sense that institutions would undoubtedly evolve to provide insurance opportunities for the elderly. At the same time, there are asymmetric information problems that can inhibit the effectiveness of private insurance markets.

Varian (1980) shows that redistributive taxation can improve social insurance. And finally, Loury (1981) shows that public education improves welfare in an economy with idiosyncratic ability risk.

In Table 9 under the "No Med" column, we show the short-run effect on mortality of eliminating Medicare are sizeable. Since most of health capital and wealth is accumulated well before retirement, health status is largely fixed by age 60-65. Eliminating Medicare, therefore, has little effect on health in the years immediately following its repeal. Nevertheless, Medicare provides insurance against health shocks and, given that Medicare was eliminated unexpectedly, households did not have an opportunity to increase their savings in order to self insure.

Table 9: Short- and Long-Run Effects on Mortality of Eliminating Our Stylized Medicare Program

	Age 60			Age 75		
Lifetime Income	Baseline	No Med	LR	Baseline	No Med	LR
Bottom	0.76	0.72	0.68	0.52	0.47	0.42
Second	0.81	0.74	0.71	0.53	0.48	0.44
Middle	0.84	0.79	0.77	0.55	0.51	0.48
Fourth	0.87	0.82	0.81	0.60	0.53	0.52
Highest	0.89	0.85	0.86	0.62	0.57	0.58

There are three primary effects of removing Medicare. The first is the insurance effect, the fact that Medicare provides insurance against health shocks. The second is the investment effect, the idea that individuals are more likely to invest more in health when young if they know there is insurance available at older ages. Third, there is the income effect: removing Medicare means that Medicare taxes are rebated back to the household which results in more income available for health investments. In the short run ("No Med") experiment, only the first effect is at play.

The long-run effects of repealing our stylized Medicare program are even larger as all three effects come into play. As before, when we repeal Medicare, the Medicare tax on earnings is rebated back to the household. In the long-run analysis, households go through their entire working lives without Medicare. In the lowest and middle lifetime income quintiles there is a moderate adverse effect on survival probabilities. In the long-run, a forward-looking household with low lifetime income will recognize they have no health insurance program in retirement. They also correctly anticipate the lifecycle pattern of health shocks and the cumulative effects of health depreciation, so old-age health will be worse than health when young. Because health and consumption are complements, the life-cycle pattern of consumption mirrors the lifecycle pattern of health. Low lifetime income households will therefore invest less in health, trading off a shorter expected lifespan for greater consumption in younger ages when the marginal utility of consumption is high relative to later in life. High lifetime income households can mitigate these effects by self-insuring: they engage in precautionary saving and invest in health capital.

The effects of this experiment on wealth and out-of-pocket spending are shown in Table 10. With Medicare eliminated and many elderly people paying for all medical care out of pocket, some households engage in additional precautionary saving, self-insuring in the absence of Medicare (some still have insurance provided by Medicaid, employer-provided plans, or VA-Champus). Indeed, we see greater wealth accumulation throughout the lifetime income distribution. We also see fewer medical expenditures, especially at the lower end of the lifetime income distribution. The tables illustrate a central insight into the lifecycle model with endogenous health. Long-run adjustments to changes in the institutional environment will be made along two margins: first, households will consume less and engage in more precautionary saving. Second, private health investment will decrease. The result is that households will both consume less and die earlier in a world without Medicare. But relative to a standard lifecycle model of consumption without endogenous health production, the consumption responses will be smaller, since a portion of the response occurs through a diminution of health capital.

With less health capital, households correctly anticipate that they will die younger and hence they need to accumulate less wealth to finance consumption in retirement. Thus, the model with endogenous health mitigates the effects of changes in social insurance on consumption relative to standard lifecycle models.

Table 10: Long Run Effects of Eliminating Medicare on Net Worth and OOP Medical Expenditures

1998	Median Net Worth		Median OOP Medical Exp.	
Lifetime Income	Baseline	No Medicare-LR	Baseline	No Medicare-LR
Bottom Quintile	$31,456	$52,064	$421	$248
Second Quintile	53,483	73,452	729	535
Middle Quintile	93,708	117,405	967	689
Fourth Quintile	163,695	185,232	1081	812
Highest Quintile	353,129	364,988	1,235	1087

Endogenous versus Exogenous Health The consumption and out-of-pocket spending responses to eliminating the stylized Medicare program shown in Table 10 are fairly modest, which results in the moderate reduction in survival probabilities for low- and moderate-income households in the model. This result is quite different than what would arise from the standard modeling approach such as Scholz *et al.* (2006), where medical expenses follow an exogenous stochastic process and there is no health-consumption complementarity. In Table 11, we demonstrate the impact on net worth from removing Medicare in the model presented above as well as in a world in which medical expenses follow an AR(1) process similar to Scholz *et al.* (2006) and where health is not an argument in preferences.

Table 11: Long Run Effects of Eliminating Medicare on Net Worth, Endogenous vs. Exogenous				
1998	Median Net Worth			
	Endogenous Health		Exogenous Health	
Lifetime Income	Model	No Medicare-LR	Model	No Medicare-LR
Bottom Quintile	$31,456	$52,064	$32,485	$123,697
Second Quintile	53,483	73,452	48,232	164,307
Middle Quintile	93,708	117,405	89,109	232,012
Fourth Quintile	163,695	185,232	153,075	323,433
Highest Quintile	353,129	364,988	343,607	594,457

The asset response to eliminating the stylized Medicare program is enormous in an economy in which medical expenses are exogenous. In this economy households will try to self-insure by accumulating substantial wealth. In contrast, in our model there are two channels that mitigate this response. First, when medical expenses are endogenous, when hit with a bad health shock, households can choose not to spend on medical care later in life. This option is not available in models where medical expenses are exogenous. Second, consumption-health complementarity implies optimal consumption profiles decline as health depreciates. Hence, households will optimally accumulate fewer late-in-life assets than would identical households in a model that does not recognize these complementarities. Comparing a society with social insurance and one without, we would be hard pressed to find poor households in poor economies building up large asset stocks. Hence, the fact our model implies a much smaller effect of social insurance programs on wealth accumulation than other frameworks is a desirable result in that it better accords with the evidence.

7 Sensitivity Analyses

In what follows we briefly report the effect of altering some of the exogenously set parameter values on goodness of fit. The upshot is that the fit between the model and the data is preserved with reasonable peturbations of the parameter values.

Table 12: Sensitivity Analysis

Parameter	R^2(Health)	R^2(Med. Exp.)	R^2(Wealth)
Baseline	0.69	0.47	0.71
$\beta = 0.9$	0.57	0.42	0.55
$\beta = 0.99$	0.75	0.38	0.61
$r = .01$	0.65	0.35	0.54
$r = .07$	0.56	0.46	0.74
$\sigma = 1$	0.53	0.39	0.52
$\sigma = 5$	0.62	0.48	0.59
$\eta = 0.2$	0.72	0.42	0.60
$\eta = 0.6$	0.68	0.50	0.56

When β is lower than the baseline, households discount the future more and this causes them to save less for old age. The fit worsens along all dimensions - they care less about future health and wealth. When β rises, households place more weight on the future. The results indicate that model fits better along the health dimension but the R^2 for wealth is lower. Finally, the elasticity of consumption versus leisure in utility has a fairly small impact on the goodness of fit when we vary the value between 0.2 and 0.6.

8 Conclusion

In this paper we describe a lifecycle model of consumption with endogenous investments in health. Health affects longevity as well as utility and we find that consumption and health are complementary inputs in the utility function. The model has many features:

households build health capital with investments of both time and money; insurance affects the transformation of out-of-pocket medical expenses to total medical expenses; the health status of two spouses in a marriage evolve distinctly; health affects time endowment and labor supply and earnings affect health. We solve the model household-by-household using data from the HRS. We force the model to match moments on wealth, mortality, and medical expenses for the average HRS married and single households, calibrating 19 parameters. We take these parameters as primitives for all households and vary the circumstances of the households based on observables in the HRS data such as earnings and medical expense realizations, insurance status, marital status, and demographic variables. We then ask whether this framework with the 19 parameters identified by the typical household can account for the microeconomic variation in health, wealth, mortality and retirement across the 11,172 households we analyze. We find that it can. Our study makes several contributions.

First, the model successfully accounts for the variation in medical expenses and longevity across households. In addition, the fit between the model and data on health status is excellent. We conclude that the model can rationalize a significant fraction of the variation in health across households.

Second, the degree of complementarity between consumption and health capital is critical in explaining various features of the data. Complementarity helps explain the time path of wealth decumulation late in life. While some previous authors relied on differences in discount factors and preferences to account for the steeper decline in consumption after age 50 for the less educated, we find that this decline is synchronous with the decline in health and the complementarity between consumption and health helps explain this steeper fall in consumption for poorer households relative to their richer counterparts.

Third, we find that the effect of policy in a framework in which health is endogenous is quite different from the standard framework in which medical expenses follow an exogenous stochastic process. For instance, programs such as Medicare have a much smaller effect on asset accumulation than in models without endogenous health investment. This is consistent with the available evidence.

Fourth, policy can have meaningful effects on life tables especially in the bottom few deciles of the income distribution. Further work exploring the effects of policy on life tables and longevity is an exciting area for future work.

9 Appendix

9.1 PSID Labor Supply

We use the Panel Study of Income Dynamics (PSID) to estimate the association between health status and other individual characteristics with labor supply. The long observation window of the PSID makes it well suited to studying the labor supply of individuals over the lifecycle. In order to match the household characteristics in lifecycle model, we used age, marital status, sex, self report of health, wealth, current earnings, and union status to predict annual hours worked. These data were available in a subset of the PSID observation years (1984, 1989, 1994, 1999, 2001, 2003, 2005, 2007, and 2009). The analysis sample is described in table A1 and the OLS estimation results are found in table A2. Most of the predictor variables are highly significant and all take the expected sign.

Table A1: PSID Sample

	Mean	Std. Dev.
Age	43.6	15.9
Current Earnings	$31,353	$54,437
Net Worth	$246,886	$1,375,000
	Percentage	
Female	54%	
Union Member	9%	
Health		
Excellent	22%	
Very Good	33%	
Good	29%	
Fair	12%	
Poor	4%	
Marital Status		
Married	69%	
Never Married	13%	
Widowed	6%	
Divorced	9%	
Separated	4%	
N	106,619	

Table A2: Model of Annual Labor Supply (Hours)

	Coefficient	Robust S.E.
Age	46.73***	3.15
Age^2	-0.67***	0.03
Female	-458.74***	21.41
Health		
(Excellent)		
Very Good	5.15	9.71
Good	-80.76***	15.52
Fair	-343.09***	25.62
Poor	-674.96***	32.86
Union Member	250.17***	15.28
Marital Status		
(Married)		
Never Married	-20.96	13.39
Widowed	114.95***	17.11
Divorced	165.59***	14.90
Separated	31.33	18.74
Current Earnings (1,000's)	5.26***	1.05
Net Worth (1,000's)	-0.03***	0.01
Constant	1022.92***	32.44
R-squared	0.39	
N	106,619	

Standard error adjusted for 24,819 individual clusters

* p<0.05, ** p<0.01, *** p<0.001

Dollar amounts in year 2008 dollars.

References

[1] Aiyagari, S. Rao. "Optimal Capital Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting." *Journal of Political Economy* 103(6), (December 1995), 1158-1175.

[2] Becker, G. S., 1964, *Human Capital : A Theoretical and Empirical Analysis, with Special Reference to Education.* New York: National Bureau of Economic Research.

[3] Card, David, Carlos Dobkin, Nicole Maestas, "Does Medicare Save Lives?" *Quarterly Journal of Economics,* May 2009, Vol. 124, No. 2: 597–636.

[4] Case, Ann and Angus Deaton, "Broken Down by Work and Sex: How Our Health Declines," in D. Wise (ed.) **Advances in the Economics of Aging**. University of Chicago Press, 2005.

[5] Citro, Constance F. and Robert T. Michael. 1995. **Measuring Poverty: A New Approach**. Washington, D.C.: National Academy Press.

[6] Chay, Kenneth Y., Daeho Kim, and Shailender Swaminathan, 2010, "Medicare, Hospital Utilization and Mortality: Evidence from the Program's Origins," working paper.

[7] Cooley, Thomas F. and Ed Prescott, 1995, "Economic Growth and Business Cycles," Chapter 1 in T. F. Cooley, ed., **Frontiers of Business Cycle Research**, (Princeton University Press), 1-38.

[8] Currie, Janet and Jonathan Gruber. "The Technology of Birth: Health Insurance, Medical Interventions, and Infant Health." NBER Working Paper. No. 5985. 1997.

[9] Currie, Janet and Duncan Thomas. “Medical Care for Children: Public Insurance, Private Insurance, and Racial Differences in Utilization.” *Journal of Human Resources*. 30(1). 1995: 135-62.

[10] De Nardi, Mariacristina, Eric French, and John Bailey Jones, 2010, "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy* vol. 118(1): 39-75.

[11] Doyle, Joseph J, Jr, “Health Insurance, Treatment and Outcomes: Using Auto Accidents as Health Shocks," NBER Working Paper 11099, February 2005

[12] Doyle, Joseph J, Jr, “Returns to Local-Area Health Care Spending: Using Health Shocks to Patients Far from Home,” May 2009. http://www.mit.edu/~jjdoyle/doyle_vacation_aer_may2009.pdf

[13] Engen, Eric M., William G. Gale and Cori R. Uccello. 1999. “The Adequacy of Retirement Saving.” *Brookings Papers on Economic Activity* 2: 65–165.

[14] Finkelstein, Amy, Erzo Luttmer, and Matthew Notowidigdo, forthcoming, “What Good is Wealth without Health? The Effect of Health on the Marginal Utility of Consumption,” *Journal of the European Economic Association*

[15] Finkelstein, Amy and Robin McKnight, 2008, "What Did Medicare Do? The Initial Impact of Medicare on Mortality and Out of Pocket Medical Spending," *Journal of Public Economics*, 92, 1644-1669.

[16] French, Eric, 2005. "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour," *Review of Economic Studies*, Wiley Blackwell, vol. 72(2), pages 395-427.

[17] Fonseca, Raquel, Pierre-Carl Michaud, Titus Gamama, and Arie Kapteyn, 2009, "On the Rise of Health Spending and Longevity," IZA Discussion paper #4622

[18] Gouveia, Miguel and Robert R. Strauss. 1994. "Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis." *National Tax Journal* 47(2) (June): 317–39.

[19] Gouveia, Miguel and Robert R. Strauss. 1999. "Effective Tax Functions for the U.S. Individual Income Tax: 1966–89." *Proceedings of the 92nd Annual Conference on Taxation, National Tax Association*: 155–65.

[20] Grossman, Michael, 1972, "One the Concept of Health Capital and the Demand for Health," *Journal of Political Economy*, 80, 223-255

[21] Haas, J. S. and L. Goldman. "Acutely Injured Patients with Trauma in Massachusetts: Differences in Care and Mortality, by Insurance Status." *American Journal of Public Health*. 84(10). October 1994: 1605-1608.

[22] Hall, Robert and Charles Jones. 2007. "The Value of Life and the Rise in Health Spending," *Quarterly Journal of Economic*s, February.

[23] Hubbard, Glen R., Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy*, 103(2) (April): 360–99.

[24] Long, S.H., M.S. Marquis, and J. Rodgers. "Do people Shift Their Use of Health Services over Time to Take Advantage of Insurance?" *Journal of Health Economics*. 17. 1997: 105- 115

[25] Loury, Glenn C. 1981. "Intergenerational Transfers and the Distribution of Earnings." *Econometrica* 49:843-67 (1981)

[26] McGrattan, Ellen R. and Edward C. Prescott. 2003. "Average Debt and Equity Returns: Puzzling?" *American Economic Review Papers and Proceedings*. 93(2) (May): 392–397.

[27] MEPSnet Household Component for 2008. February 2008. Agency for Healthcare Research and Quality, Rockville, MD. http://www.meps.ahrq.gov/data_stats/mepsnet/mepsnethc08.shtml

[28] Mincer, J., 1974, **Schooling, Experience and Earnings**, Columbia University Press, New York, 1974.

[29] Modigliani, Franco and R. Brumberg. 1954. "Utility Analysis and the Consumption Function: An Interpretation of Cross Section Data." In **Post-Keynesian Economics**, K.K. Kurihara (ed.), pp. 388–436. New Brunswick, NJ: Rutgers University Press.

[30] Moffitt, Robert. 2002. "Documentation for Moffitt Welfare Benefits File." Manuscript, February 22, http://www.econ.jhu.edu/People/Moffitt/DataSets.html.

[31] Mokdad, Ali H.; James S. Marks, MD, MPH; Donna F. Stroup, PhD, MSc; Julie L. Gerberding, MD, MPH. 2004. "Actual Causes of Death in the United States, 2000," *JAMA*. 2004;291(10):1238-1245.

[32] Murphy, Kevin M. and Topel, Robert H. (2006). "The Value of Health and Longevity," *Journal of Political Economy* 114 (5):871-904.

[33] Oster, Emily, Ira Shoulson, and Ray Dorsey, 2012. "Limited Life Expectancy, Human Capital and Health Investments," American Economic Review, *forthcoming.*

[34] Paffenbarger, Ralph S., Robert T. Hyde, Alvin L. Wing, I-Min Lee, Dexter L. Lung, and James B Kampert. 1993. "The Association of Changes in Physical-Activity Level and Other Lifestyle Characteristics with Mortality Among Men," *The New England Journal of Medicine*, 328: 538-45.

[35] Palumbo, Michael G. 1999. "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle." *Review of Economic Studies* 66(2): 395–421.

[36] Rosen, Sherwin. 1988. “The Value of Changes in Life Expectancy,” *Journal of Risk and Uncertainty* 1 (September): 285–304.

[37] Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. “Are Americans Saving ‘Optimally’ for Retirement?” *Journal of Political Economy*. August, 607-643.

[38] Spillman, B. C. 1992. “The Impact of Being Uninsured on Utilization of Basic Health Care Services.” *Inquiry*. 29. 457-466.

[39] Smith, James P., “Consequences and Predictors of New Health Events,” in D. Wise (ed.) **Advances in the Economics of Aging**. University of Chicago Press, 2005.

[40] Tilford, J.M., Robbins, J.M., Shema, S.J., and F.L. Farmer. “Response to Health Insurance by Previously Uninsured Rural Children.” *Health Services Research* 34(3). 1999: 761-775.

[41] Varian, Hal R. 1980. “Redistributive Taxation as Social Insurance.” *Journal of Public Economics*, 14(1): 49–68.

[42] Warburton, Darren E.R., Crystal Whitney Nicol, and Shannon S.D. Bredin. 2006. "Health Benefits of Physical Activity: The Evidence," *Canadian Medical Association Journal*, 174 (6): 801-809.

[43] Willett, Walter C. 1994. "Diet and Health: What Should We Eat?" *Science*, 264 (5158): 532-37.

[44] Yogo, Motohiro, 2009, "Portfolio Choice in Retirement: Health Risk and the Demand for Annuities, Housing, and Risky Assets," mimeo