



The Interaction between Consumption and Health in Retirement

John Karl Scholz and Ananth Seshadri



The Interaction between Consumption and Health in Retirement

John Karl Scholz

University of Wisconsin-Madison

Ananth Seshadri

University of Wisconsin-Madison

September 2016

Michigan Retirement Research Center

University of Michigan

P.O. Box 1248

Ann Arbor, MI 48104

www.mrrc.isr.umich.edu

(734) 615-0422

Acknowledgements

The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium through the University of Michigan Retirement Research Center Award RRC08098401. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA or any agency of the federal government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States government or any agency thereof.

Regents of the University of Michigan

Michael J. Behm, Grand Blanc; Mark J. Bernstein, Ann Arbor; Laurence B. Deitch, Bloomfield Hills; Shauna Ryder Diggs, Grosse Pointe; Denise Ilitch, Bingham Farms; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park; Katherine E. White, Ann Arbor; Mark S. Schlissel, *ex officio*

The Interaction between Consumption and Health in Retirement

Abstract

We study the interaction between consumption and health in retirement. Our main contribution is the estimation of a consumption Euler equation taking health into consideration. The Euler equation is derived from a model of consumption in retirement with three important building blocks of health: health shocks, health as an investment and health as a provider of utility. We estimate the Euler equation using data from the Health and Retirement Study (HRS) and Consumption and Activities Mail Survey (CAMS). The estimates suggest that health is an important determinant of utility. We use the estimated model to study the empirical significance of the three building blocks of health. We find that health shocks play an important role in slowing down the decline of consumption with age in retirement. We also find that including health into the utility function could help explain the heterogeneous consumption-age profiles related to health. Finally, we find that health investments, such as physical exercise, have a significant effect on the evolutions of both health and consumption in retirement.

Citation

Scholz, John Karl, and Ananth Seshadri. 2016. "The Interaction Between Consumption and Health in Retirement." Ann Arbor, MI. University of Michigan Retirement Research Center (MRRC) Working Paper, WP 2016-344.
<http://www.mrrc.isr.umich.edu/publications/papers/pdf/wp344.pdf>

Introduction

In this paper, we seek to further our understanding of the process through which consumption and health interact and how this affects the well-being of the elderly. In particular, we integrate health capital into a life-cycle model of consumption to better understand the interrelationship between health, health investments, mortality risk, and wealth in retirement.

Incorporating health into a life-cycle model of wealth accumulation is, in our view, very important. Existing models of consumption behavior for working households largely ignore health production, at least for working-age households. But one potentially important response to adverse income or health shocks or to governmental policy changes is for individuals to adjust health investment in addition to consumption. Changes in health investment can alter individual's health and have longer-run implications for mortality. Our innovation is not just to incorporate health as an investment that increases an individual's health stock but also to incorporate health in the utility function. This allows health to affect the marginal utility of consumption and the allocation of consumption over the life cycle.

Health and consumption decisions are intertwined, 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. Surprisingly, given the centrality of health to economic decision-making and well-being, numerical models of life-cycle 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 the life-cycle consumption papers noted above, households will respond to cuts in safety net programs by increasing precautionary saving. In the real world, 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. 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 are not the first to examine the links between health, consumption, and wealth. Clear discussions are given in Smith (2005), Case and Deaton (2005), and many other places. We build on the important recent papers by De Nardi, French, and Jones (2010), and Fonseca, Michaud, Galama, and Kapteyn (2009). Our paper is also related to Finkelstein et. al. (2013). Since we have a fully specified optimizing model with data on consumption, health status, and lifetime earnings, we can get around the issue of having to measure utility.

Our prior work shows that these links between health, wealth, and mortality are potentially important. For example, we examine the effects of altering the Medicare program — the health care social insurance program for those 65 and older. For example, if Medicare were

unanticipatedly removed, we show that there would be only negligible effects on mortality rates in the first 10 years following its removal. This is due to the fact that health status at age 65 is largely determined by decisions households made at younger ages.

Our main contribution in this paper is the estimation of a consumption Euler equation taking health into consideration. The Euler equation is derived from a model of consumption in retirement with three important building blocks of health: health shocks, health as an investment, and health as a provider of utility. In our model, a retiree maximizes lifetime utility by making decisions on consumption and health investments in each period. The utility flow in each period depends on both consumption and health capital. Health investments could affect both the probability of surviving to the next period and future health capital conditional on survival. There are health shocks in each period such that both survival and future health capital conditional on survival are stochastic.

We estimate the Euler equation of our model using data from both the core Health and Retirement Study (HRS) and the supplemental Consumption and Activities Mail Survey (CAMS). Our estimation takes two steps. In the first step, we use all respondents in the HRS to estimate inputs into the structural model including the survival probabilities as a function of current age, health status, and health investments. In the second step, we use respondents at age 65 or older with valid information on consumption and health to estimate the structural parameters of the model, including the coefficient of relative risk aversion and the share of consumption in the consumption-health composite good that produces utility. We estimate the model separately for each gender. Our estimates suggest that health is an important determinant of utility.

We use the estimated model to study the empirical significance of the three health building blocks. We find that health shocks play an important role in slowing down the decline of consumption with age in retirement. Without health shocks, retirees will run down their wealth at a much higher speed. We also find that including health into the utility function provides interesting interactions between health and consumption, and could help explain the heterogeneous consumption-age profiles related to health. Finally, we find that health investments, such as physical exercise, have a significant effect on the evolutions of retirement health and consumption. This finding suggests that public programs like Medicare and Medicaid play an important role in shaping the consumption and health behavior of the retirees, as well as the retirement and saving decisions of working people.

In terms of policy, the central insight that arises from our work is that long-run adjustments to changes in the institutional environment can be substantial. As an example, consider the removal of Medicare or Medicaid. Relative to a standard life-cycle model of consumption without endogenous health production, the consumption responses in our model 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 social insurance changes on consumption relative to standard life-cycle models.

We believe our work has important potential benefits to Social Security policy-making. Our model recognized the varied ways that individuals build health capital. In particular, an extensive literature documents that there appears to be very low returns to marginal investments in health, at least given health capital levels observed in the United States. Nevertheless, there is

a strong positive relationship between affluence, whether measured by lifetime income or wealth, and health, whether measured by health status, functional limitations, or mortality. Our model allows people to invest in health through direct expenditures on medical inputs and through time spent on health-producing activities, such as exercise or preparing nutritious meals. Although we only consider physical exercise in the empirical application of this paper, an extended version of the model would yield predictions on the patterns of medical expenditures (in dollars) over the life cycle, time used in health-producing activities over the life cycle, life-cycle patterns of health, and predictions about mortality.

The rest of this paper proceeds as follows. In the first section below, we introduce the data and provide some facts that motivate our study. We present our model in the second section and estimate it in the third section. In the fourth section, we use the estimated model to study the channels through which health interacts with consumption. We conclude the paper in Section 5.

Data and Motivation

We use data from the core Health and Retirement Study (HRS) and its supplement, the Consumption and Activities Mail Survey (CAMS). The HRS is sponsored by the National Institute of Aging and conducted by the University of Michigan with supplemental support from the Social Security Administration. It is a national panel study with a sample (in 1992) of 12,652 persons in 7,702 households. 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. As the HRS has matured, new cohorts have been added. We use the RAND HRS Version O, which is a representative, randomly stratified sample of U.S. households born before 1959.

The CAMS is an ongoing supplement to the HRS. One of its primary objectives is to measure total household spending during the previous 12 months. In September 2001, the first CAMS survey was mailed to 5,000 households selected at random from households that participated in the HRS 2000 core survey. Follow up surveys have been fielded biennially in odd-numbered years. We use the RAND CAMS Version D, Release 2. The measure of total household consumption is available for six waves. We use CPI-U to convert the consumption measures in different years into constant dollars.

We link the two datasets to study the interaction between consumption and health in retirement.

Consumption, Health and Health Investments

As a motivation for the importance of including health in modeling consumption of the elderly, we combine wave nine of the HRS and CAMS and report in Table 1 mean household consumption by lifetime earnings¹ and self-reported health status. The mean consumption of households in the third earnings quintile who reported to be in good health is normalized to be one.

¹ Lifetime earnings are constructed using restricted-access social security earnings data, as well as self-reported earnings in the HRS. See Scholz, Seshadri, and Khitatrakun (2006) for details.

Table 1. Mean Household Consumption by Earnings and Health

Health	Lifetime Earnings				
	First	Second	Third	Fourth	Fifth
Excellent	0.93	1.02	1.15	1.29	1.68
Very Good	0.75	0.91	1.2	1.23	1.55
Good	0.69	0.88	1	1.13	1.38
Fair	0.64	0.73	0.87	1.04	1.29
Poor	0.61	0.71	0.79	0.92	1.24

Consumption rises with income for each health status and consumption rises with health status for each income quintile. For example, for respondents in the third quintile of lifetime earnings, those in excellent health on average spend 15% more on consumption than those in good health, while the consumption of those in poor health is only about 79% of the consumption of those in good health. The fact that consumption co-moves with health status in the cross-section is consistent with consumption-health complementarity in preferences. We allow for this possibility in the model presented in the next section. We will also look at changes across the retired years by age to identify the relationship between consumption and health in the utility function.

If health and consumption are related, a model of health is necessary in order to understand the evolution of consumption in retirement. Starting from Grossman (1972), health has been modeled as a capital good where people can make investments. The HRS provides extensive measures of health investments. These measures can be categorized into three groups. The first group includes measures of medical care utilization, such as hospital stays, nursing home stays, doctor visits, home care, and whether the respondent reported regular use of prescription drugs, outpatient surgery, dental visits, or use of special facilities or services during

the reference period. The second group includes measures of medical expenditures. The last group includes health behaviors like participation in physical activity/exercise, drinking, smoking, and preventive behaviors like a blood test for cholesterol, a flu shot, monthly self-checks for breast lumps, a mammogram, a pap smear, and a check for prostate cancer.

To provide suggestive evidence of the effects of health investments, we use the variables measuring participation in vigorous physical activities to study the effect of exercise on mortality. The first six waves of the HRS record whether a respondent participated in rigorous physical activities three or more times each week. Starting from the seventh wave, this variable is now replaced by another one indicating whether a respondent participated in rigorous physical activities for two or more times each week. We combine these two variables to create a dummy that equals to one if a respondent participated in rigorous physical activities for either three or more times (first six waves) or two or more times (subsequent waves). We will refer to this dummy variable as exercise.

We calculate the effect of exercise on mortality as follows. For respondents at a given age with a given health status, we calculate the fractions of these individuals surviving to the next survey according to their exercise status. We calculate the survival probabilities separately for each gender and report them in Figures 1 and 2. Because the survey is conducted every two years, this calculation gives us the two-year survival probabilities. We use the LOWESS command in STATA to smooth the raw survival probabilities over age. To save space, we only report the results for the two extreme self-reported health statuses.

Figure 1. Survival Probabilities by Age, Health, and Activity Status: Men


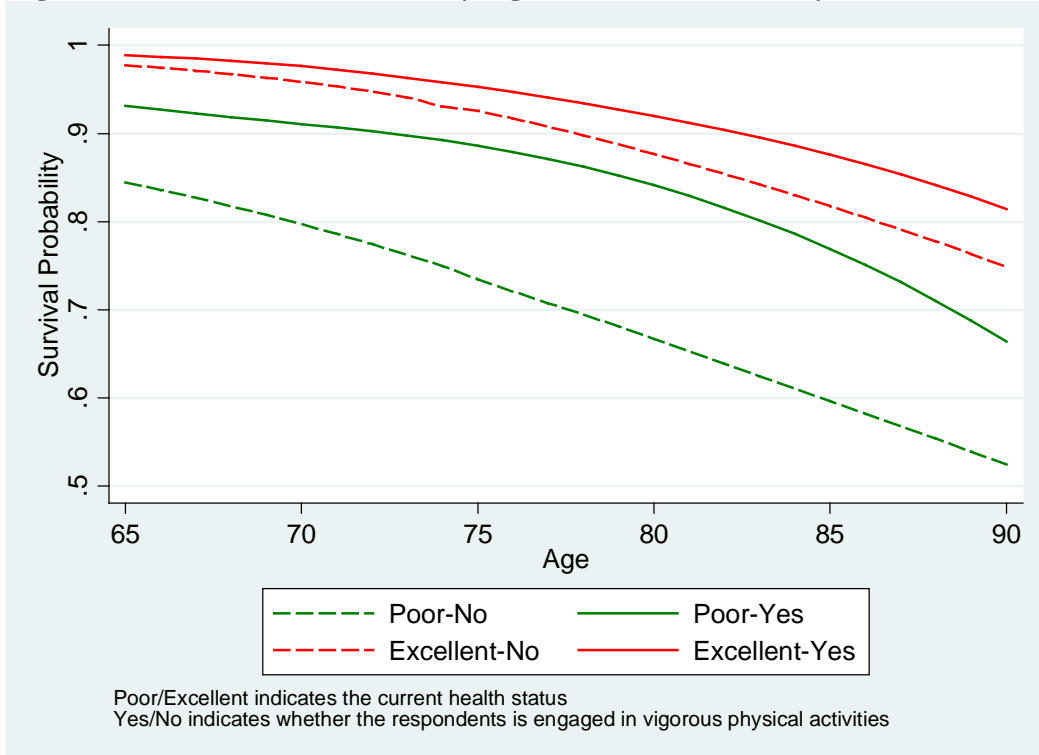


Figure 2. Survival Probabilities by Age, Health, and Activity Status: Women



Not surprisingly, the survival probabilities are higher for respondents in the better health status at all ages. Conditional on self-reported current health status, those who participate in rigorous physical activities are more likely to survive than those who don't. For example, among women in poor health at age 65, around 84% of those who don't exercise can survive for another two years, and the number increases to about 93% for those who exercise. The difference is about 9%. This difference seems to increase with age. For example, at age 80, the difference is about 15%. The facts that those who exercise are more likely to survive than those who don't and this advantage increases with age hold for men and women and for both health statuses. It also appears that exercise's effect on survival decreases with health: Those in poor health could benefit more from exercise than those in good or excellent health.

To summarize, Figure 1 suggests that health investments, such as physical exercise, have a positive effect on future health and the effect is larger for individuals in poor health. We allow for this possibility in the model presented below.

Model

We model the consumption decision of single (divorced, never married, or widowed) retirees taking into account the effect of health. Retirees face different paths of income ($y = \{y_t\}$) and prices ($p = \{p_t\}$) of health investments over the remaining years of their lives.² Income (y) is determined by lifetime earnings before retirement, and the prices (p) of health investments is determined by factors such as health insurance. Both y and p are taken exogenously.

At each age t , a retiree starts with a stock of financial wealth a_t and health status h_t , and chooses consumption (c_t) and health investments (i_t) to maximize the present value of lifetime

² Both y and p , as well as other variables introduced below such as consumption, health, health investments and wealth, are individual-specific. We suppress the individual subscripts for simplicity.

utility. Let $u(c_t, h_t)$ be the flow utility, and $\varphi(h_t, i_t, t)$ be the probability of surviving to the next age: The individual problem can be described using Bellman equation as follows

$$V(a_t, h_t; t, y, p) = \max_{c_t, i_t} u(c_t, h_t) + \beta \varphi(h_t, i_t, t) E_t V(a_{t+1}, h_{t+1}; t + 1, y, p)$$

subject to the following budget constraint

$$c_t + p_t i_t + \frac{a_{t+1}}{1 + r} = a_t + y_t$$

and the health transition function

$$h_{t+1} = f(h_t, i_t, \epsilon_{t+1}; t)$$

where β is the discount factor, r is the annual interest rate, and E_t is the expectation operator over the health shock ϵ_{t+1} realized at the beginning of age $t + 1$. To keep the model simple, we ignore the bequest motives and normalize the utility in the case of death to be zero. In a different model, De Nardi, French, and Jones (2010) cannot reject the joint hypotheses that (1) there is no bequest motive and (2) health does not enter the utility function directly.

The model allows health to affect consumption through three channels. First, health enters the utility function directly. If the marginal utility of consumption is positively related to health, other things equal, individuals will prefer to consume more when healthy and less when ill. As health deteriorates over age, we expect consumption to decline. The opposite will be true if the marginal utility of consumption depends negatively on health.

Secondly, following the seminal work of Grossman (1972), we allow individuals to affect the evolution of health, and consequently the optimal path of consumption through health investments. We allow health investments to affect both the survival probability and the health status conditional on survival. With a high survival probability and a better health conditional on survival, individuals who invest more in health should experience a faster consumption growth (or a slower consumption decline) when the marginal utility of consumption depends positively

on health. In this case, individuals could reduce future consumption by lowering current investments in health. Finally, we allow health shocks to affect future health and consumption.

Euler Equation and Moment Conditions

Combine the first order condition for c_t and the envelop condition for a_{t+1} , we obtain the following consumption Euler equation

$$u_c(c_t, h_t) = \beta(1 + r)\varphi(h_t, i_t, t)E_t u_c(c_{t+1}, h_{t+1})$$

where $u_c(c_t, h_t) = \frac{\partial u(c_t, h_t)}{\partial c_t}$ is the marginal utility of consumption. Let

$$\varepsilon_{t+1} = \beta(1 + r)\varphi(h_t, i_t, t) \frac{u_c(c_{t+1}, h_{t+1})}{u_c(c_t, h_t)} - 1$$

be the expectation error satisfying $E_t \varepsilon_{t+1} = 0$, and let Z_t be a vector of variables realized before or at the beginning of age t , the Euler equation can be estimated using the Generalized Method of Moments (GMM) with the following moment conditions

$$E_t(Z_t \varepsilon_{t+1}) = 0$$

In the rest of this paper, we use these moment conditions to estimate the model parameters and use the estimated model to study the interaction between consumption and health in retirement.

Estimation

To estimate the Euler equation, we specify the flow utility as

$$u(c_t, h_t) = \frac{(c_t^\alpha h_t^{1-\alpha})^{1-\gamma}}{1-\gamma}$$

where γ is the coefficient of relative risk aversion and α is the share of consumption in the consumption-health composite good that affects utility. This specification leads to the following expectation error

$$\varepsilon_{t+1} = \beta(1+r)\varphi(h_t, i_t, t) \left(\frac{c_{t+1}}{c_t}\right)^{\alpha(1-\gamma)-1} \left(\frac{h_{t+1}}{h_t}\right)^{(1-\alpha)(1-\gamma)} - 1$$

which can be used in the moment conditions to estimate the model parameters.

Our estimation proceeds in two steps. In the first step, we estimate the survival function $\varphi(h_t, i_t, t)$ using all respondents in the HRS. For simplicity, we consider here only one type of health investments i_t : whether a respondent participated in rigorous physical activities or not. The construction of this variable and the calculation of survival probabilities $\varphi(h_t, i_t, t)$ were described earlier in the data section.

In the second step, we set the discount factor β and the interest rate r to typical values in the literature and estimate the two parameters (α and γ) in the utility function using GMM. We vary the values of β and r to check the robustness of our estimates. In practice, because what matters is the product $\beta(1+r)$, we vary this product instead of individual values of β and r .

We estimate the moment conditions implied from the Euler equation using single (divorced, never married, and widowed) retirees 65 and older. For a retiree at age t , the vector of instruments Z_t includes age t , current health status h_t and the consumption growth between last period and current period $\frac{c_t}{c_{t-1}}$. These instruments should be valid for two reasons. First, both age t and health status h_t should be correlated with the change in health $\frac{h_{t+1}}{h_t}$, while past consumption growth $\frac{c_t}{c_{t-1}}$ should be correlated with future consumption growth $\frac{c_{t+1}}{c_t}$. Second, because all three variables are known to the agent at time t , they should be uncorrelated with the expectation error ε_{t+1} .

It should be noted that, because the HRS is fielded every two years, both $\frac{c_t}{c_{t-1}}$ and $\frac{c_{t+1}}{c_t}$ represent consumption growth in two years, and $\frac{h_{t+1}}{h_t}$ represents the two-year health change. We should also adjust the discount factor β and the interest rate r accordingly.

Estimation Results

Table 2 reports the estimation results. There are three panels. The upper panel reports the estimates using both men and women. The middle and bottom panel reports the results for men and women, respectively. There are five columns, each with a different value of $\beta(1 + r)$.

Table 2. Estimated Model Parameters

Parameter	$\beta(1+r)$				
	1	0.98	0.99	1.01	1.02
Both Men and Women					
α	0.394*** (0.048)	0.358*** (0.056)	0.376*** (0.052)	0.411*** (0.045)	0.428*** (0.042)
γ	0.219*** (0.080)	0.284*** (0.089)	0.251*** (0.084)	0.187** (0.076)	0.156** (0.072)
Men					
α	0.505*** (0.117)	0.484*** (0.141)	0.494*** (0.128)	0.515*** (0.107)	0.527*** (0.098)
γ	0.184 (0.126)	0.285** (0.141)	0.235* (0.133)	0.134 (0.119)	0.084 (0.113)
Women					
α	0.476*** (0.058)	0.453*** (0.071)	0.464*** (0.064)	0.488*** (0.052)	0.501*** (0.047)
γ	0.262*** (0.083)	0.343*** (0.094)	0.303*** (0.088)	0.220*** (0.078)	0.178** (0.073)

Note: Standard errors are in the parentheses. *** p<0.01, ** p<0.05, * p<0.1.

As a baseline, we start with $\beta(1+r) = 1$. Using both men and women, we obtain an estimate of $\alpha = 0.394$ with a standard deviation of 0.048. Note that our specification of the utility function implies that the larger α is, the less important health is in the utility function. In the extreme, when $\alpha = 1$, health no longer has a direct effect on utility and consumption.

Because our estimate of α is far below 1, it suggests that health plays an important role in utility and it potentially has a large direct effect on consumption. We examine this effect through simulation in the next section.

Turning to the estimate of γ , we obtain an estimate of 0.219 with a standard error of 0.08. This estimate is much lower than typical estimates from Euler equations during the working periods where the estimates are usually larger than 1. It is, however, not a surprise when death is considered. In particular, as we normalize the utility at death to be 0, we need $\gamma < 1$ to guarantee a positive utility while alive.

When the model is estimated separately by gender, we obtain larger estimates of α . However, the estimates for both men and women are still well below 1. The estimated γ is around 0.2 for both men and women, although the estimate for men is not precise due to small sample sizes.³

Moving across columns, we find that the estimate of α increases while the estimate of γ decreases with $\beta(1+r)$. However, both estimates are relatively stable, with α between 0.4 and 0.5 and γ between 0.2 and 0.3.

Discussion

We use the estimated model to study the interaction between health and consumption in this section. Recall that our model allows health to affect consumption through three channels: health shocks, health in utility, and health investments. We discuss each channel in turn.

³ Because women live longer than men on average, the sample of retired women is much larger than the sample of retired men. Specifically, our final sample includes 181 men and 791 women. These numbers are small relative to the total number of respondents in either HRS or CAMS because we use only single (divorced, never-married and widowed) retirees 65 or older with at least three observations of consumption (C_{t-1}, C_t, C_{t+1}) and two observations of health (h_t, h_{t+1}).

The simulations presented in this section use $\alpha = 0.476$ and $\gamma = 0.262$ estimated from women when $\beta(1 + r) = 1$. Using other estimates leads to qualitatively similar results.

Health Shocks

Starting with health shocks. We make inferences about the effect of health shocks on consumption by comparing the actual consumption-age profile with the hypothetical profile generated from our model assuming there is no health shock ϵ_{t+1} .

We calculate the actual consumption-age profile as follows. For respondents at age t in one wave and age $t + 2$ in the subsequent wave⁴ with valid information on consumption at both ages, c_t and c_{t+2} , we can calculate the two-year consumption growth as $\frac{c_{t+2}}{c_t}$. We assume

$\frac{c_{t+1}}{c_t} = \frac{c_{t+2}}{c_{t+1}}$ to obtain the one-year consumption growth as $\frac{c_{t+1}}{c_t} = \left(\frac{c_{t+2}}{c_t}\right)^{0.5}$. With the annual

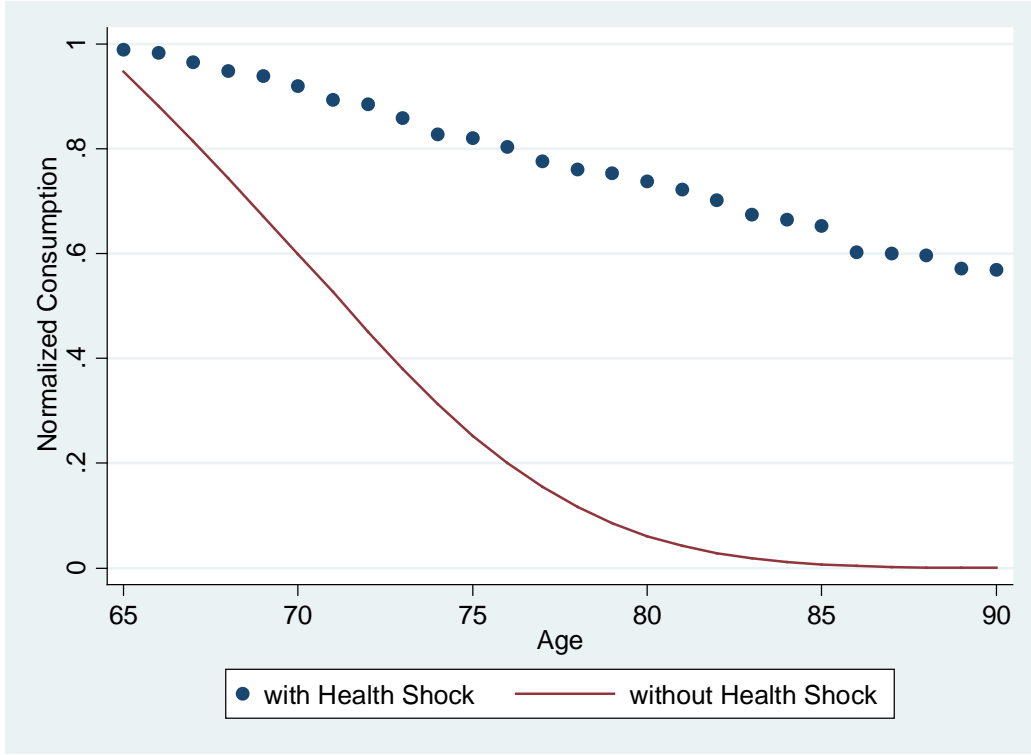
consumption growth for each respondent at age t , we calculate the average consumption growth between t and $t + 1$ as the median of $\frac{c_{t+1}}{c_t}$ among all respondents with valid information at age t .

We use median to minimize the potential influence of outliers. With the average consumption growth at each age, we obtain the consumption-age profile by normalizing the average consumption at age 65 to 1.

The dots in Figure 3 report the normalized consumption at each age. We use them as a benchmark of the consumption-age profiles with health shock. Consumption declines almost linearly with age in retirement. The average consumption of those at age 85 is about 60% of the average consumption of those at age 65. The annual rate of decline is about 2%.

⁴ Recall that both the core HRS and the CAMS are fielded every two years.

Figure 3. Consumption-Age Profiles with and without Health Shock



To obtain the simulated consumption-age profile without health shock, we proceed as follows. In the absence of health shocks, we can drop the expectation operator E_t , and the consumption Euler equation becomes

$$u_c(c_t, h_t) = \beta(1+r)\varphi(h_t, i_t, t)u_c(c_{t+1}, h_{t+1})$$

which, in combination with our empirical specification of the utility function, leads to

$$c_t^{\alpha(1-\gamma)-1}h_t^{(1-\alpha)(1-\gamma)} = \beta(1+r)\varphi(h_t, i_t, t)c_{t+1}^{\alpha(1-\gamma)-1}h_{t+1}^{(1-\alpha)(1-\gamma)}$$

Rearrange the terms of the above equation, we obtain

$$\frac{c_{t+1}}{c_t} = [\beta(1+r)\varphi(h_t, i_t, t)]^{\frac{1}{1-\alpha(1-\gamma)}} \left(\frac{h_{t+1}}{h_t}\right)^{\frac{(1-\alpha)(1-\gamma)}{1-\alpha(1-\gamma)}}$$

Given the estimated model parameters and the survival probabilities φ and health transitions $\frac{h_{t+1}}{h_t}$ in the data, we can use the above equation to calculate the implied consumption growth at each age. As before, we obtain the consumption-age profile by normalizing the

average consumption at age 65 to 1. We are essentially assuming that, conditional on survival, future health status is perfectly anticipated and there is no health shock.

The curve in Figure 3 reports the (normalized) consumption-age profile in the model without health shock. Comparing this hypothetical profile with the actual profile in data, we see a much faster decline of consumption over age in the absence of health shocks. This could be explained by the precautionary saving motive: When future health is uncertain, people will have an incentive to save more, resulting in a smaller decline in consumption. The large difference between the two profiles is consistent with De Nardi, French, and Jones (2010) who find that the risk of medical expenditures plays an important role in explaining the saving decisions of the elderly.

Health in Utility

To understand the importance of including health in the utility function, we estimate a version of the model by restricting $\alpha = 1$. In this restricted model, health no longer has a direct effect on utility. We compare the consumption profiles from this restricted model with the corresponding profiles from the unrestricted model to shed light on the importance of modeling health in the utility function.

Let the restricted utility function be

$$u(c_t) = \frac{c_t^{1-\theta}}{1-\theta}$$

which, in combination with the Euler equation, leads to the following expectation error

$$\varepsilon_{t+1} = \beta(1+r)\varphi(h_t, i_t, t) \left(\frac{c_{t+1}}{c_t} \right)^{-\theta} - 1$$

Substituting this error term into the moment conditions, we can use GMM to estimate the new coefficient of relative risk aversion, θ . The results are reported in Table 3. The estimated θ

is about 0.5 for men, and it is around 0.65 for women. As is the case for γ in the baseline model, the estimate of θ is decreasing in $\beta(1+r)$. In the following analysis, we use $\theta = 0.655$, the estimate for women when $\beta(1+r) = 1$. Results from other values of θ are qualitatively similar.

Table 3. Estimated Parameters of the Restricted Model

Parameter	$\beta(1+r)$				
	1	0.98	0.99	1.01	1.02
Both Men and Women					
θ	0.643*** (0.034)	0.709*** (0.037)	0.677*** (0.036)	0.607*** (0.033)	0.569*** (0.031)
Men					
θ	0.479*** (0.067)	0.557*** (0.074)	0.518*** (0.071)	0.438*** (0.063)	0.396*** (0.060)
Women					
θ	0.655*** (0.037)	0.720*** (0.040)	0.688*** (0.039)	0.619*** (0.036)	0.581*** (0.034)

Note: Standard errors are in the parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

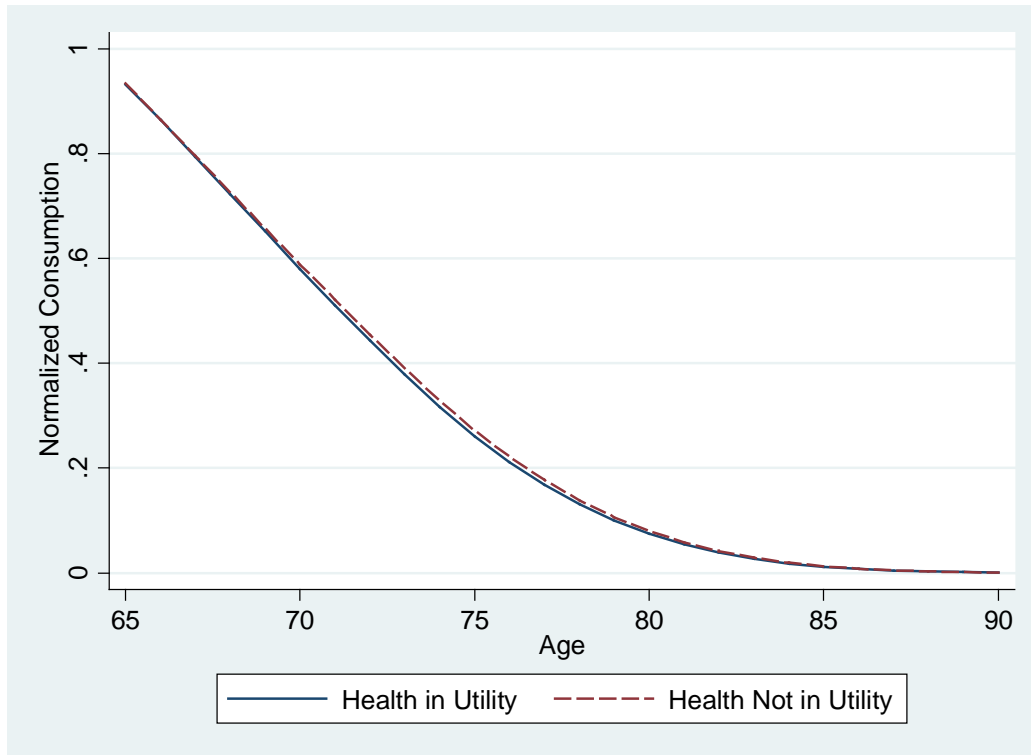
In the previous subsection we calculated the consumption-age profile from the baseline model assuming there is no health shock and plotted it in Figure 3. We can calculate the same profile from the restricted model using the following equation

$$\frac{c_{t+1}}{c_t} = [\beta(1+r)\varphi(h_t, i_t, t)]^{\frac{1}{\theta}}$$

Figure 4 plots the consumption-age profiles from both the baseline model and the restricted model. We see that the curves are almost on top of each other. Although consumption

seems to decline a little bit more slowly in the restricted model, the difference between the two curves is extremely small.

Figure 4. Consumption-Age Profiles with and without Health in Utility: All



Does it mean that including health in the utility function provides little improvement in understanding consumption in retirement? The answer is no. Although the aggregate profiles are similar between the two models, there are significant differences arising from heterogeneity. To see this, we calculate the consumption-age profiles from the two models conditional on respondents' current health status. Figure 5 plots the profiles from respondents in poor health status at each age, while Figure 6 plots similar profiles for those in excellent health.

Figure 5. Consumption-Age Profiles with and without Health in Utility: Poor Health


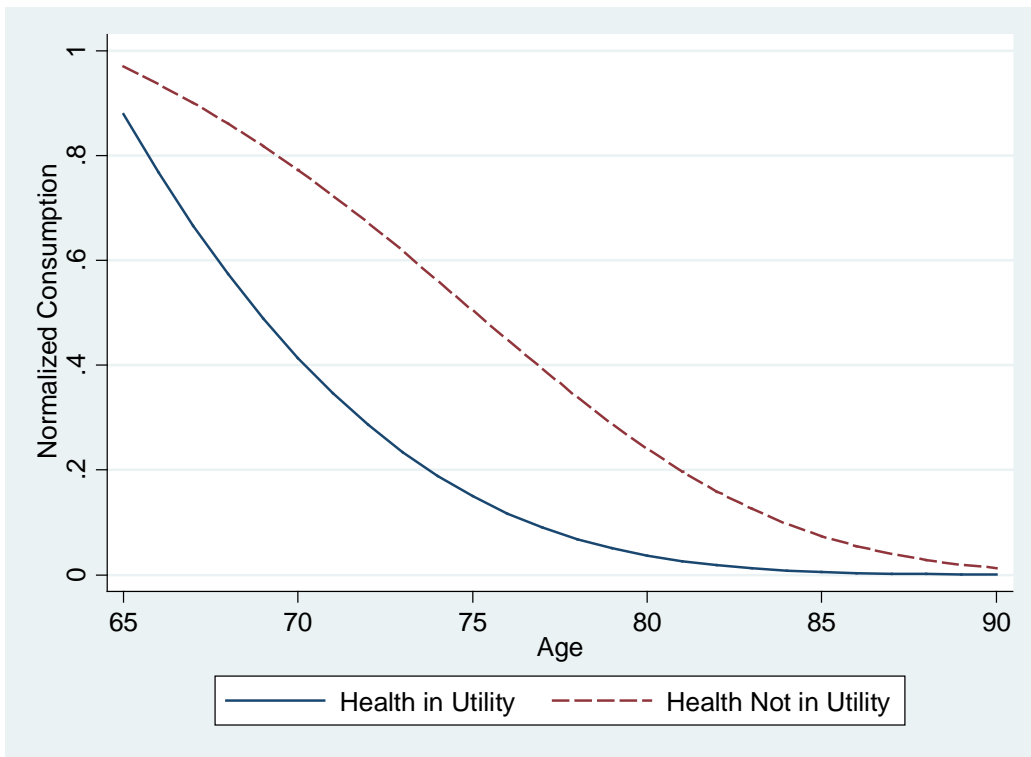


Figure 6. Consumption-Age Profiles with and without Health in Utility: Excellent Health



For respondents in poor health, the restricted model predicts a much steeper profile than the unrestricted model. The reverse is true for respondents in excellent health. The differences arise from the effect of health on the marginal utility of consumption. In the unrestricted model where health has a positive effect on the marginal utility of consumption, it's optimal for retirees to consume more while in better health. This leads to a steeper decline in consumption for those in excellent health who are expecting a health deterioration. For respondents in poor health, conditional on surviving, their health in the next period should not be much worse than what it is in the current period. A smaller decline in health thus leads to a smaller decline in consumption.

In summary, although a model without health in the utility function can generate an average consumption profile similar to data by adjusting the model parameters, a model with health in the utility function leads to richer interactions between health and consumption and has the potential to match the heterogeneous consumption-age profiles based on health.

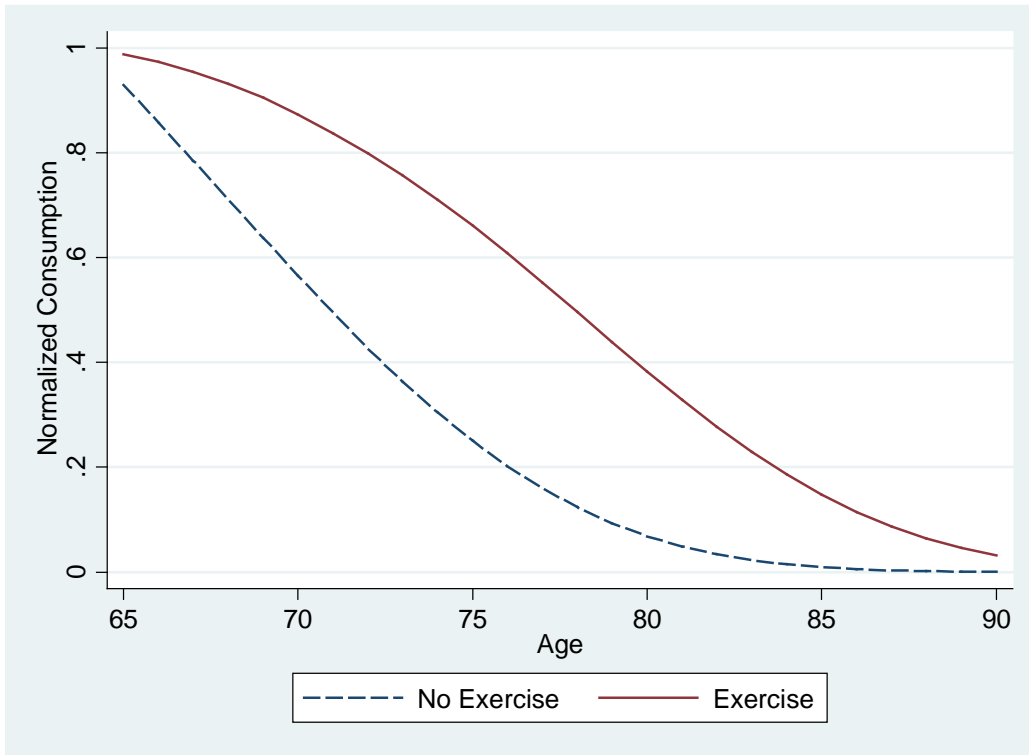
4.3 Health Investments

We turn now to the effect of health investments on consumption. Given the equation for consumption growth introduced earlier and copied below, we can calculate the consumption-age profile separately for each investment status

$$\frac{c_{t+1}}{c_t} = [\beta(1+r)\varphi(h_t, i_t, t)]^{\frac{1}{1-\alpha(1-\gamma)}} \left(\frac{h_{t+1}}{h_t}\right)^{\frac{(1-\alpha)(1-\gamma)}{1-\alpha(1-\gamma)}}$$

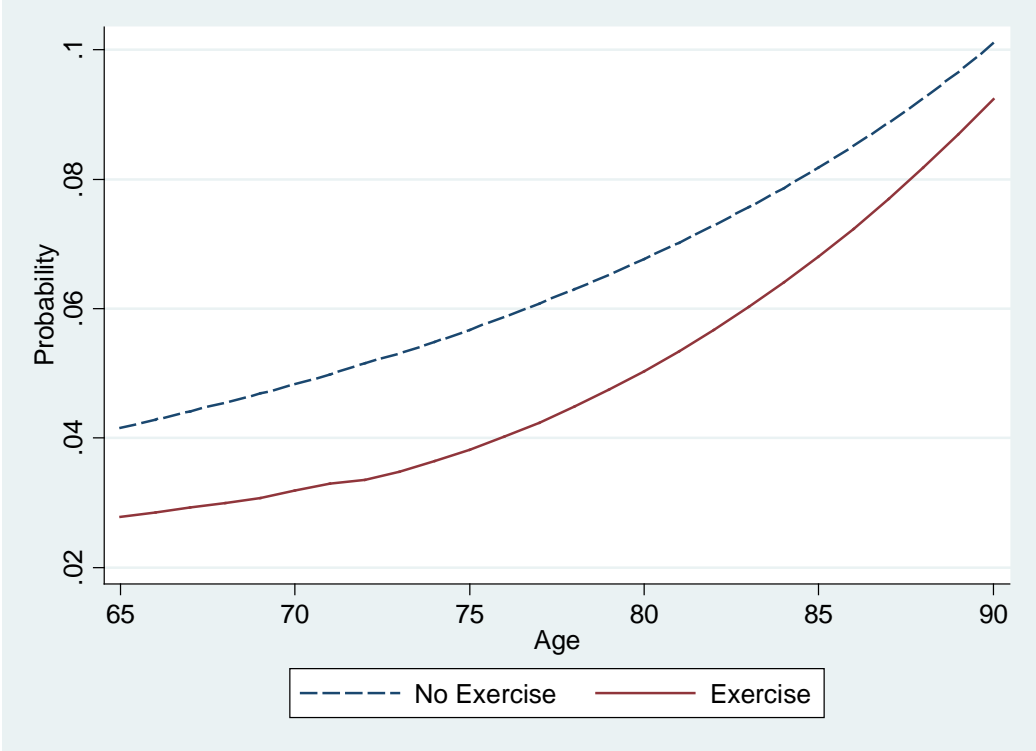
Figure 7 reports the results. Consumption declines much faster with age for those who don't exercise. There are two reasons. First, as shown earlier in this paper, exercise has a positive effect on survival. Because retirees who exercise are more likely to survive to the next period, they have a larger incentive to save, resulting in a smaller decline in consumption over time.

Figure 7. Consumption-Age Profiles by Exercise Status: Model




Secondly, conditional on survival, retirees who exercise are expected to have a better future health h_{t+1} than those who don't even if both start with the same health status in the current period h_t . For example, Figure 8 plots of probabilities of transiting from good to poor health in two consecutive surveys. Starting from the same good health status in one period, retirees who don't exercise are much more likely to end up in poor health in the next period than those who exercise. As discussed earlier, we find the marginal utility of consumption is increasing in health. Other things equal, consumption is increasing in health. A better future health for those who exercise raises their future consumption through a higher marginal utility of consumption, resulting in a slower consumption decline for them than those who don't exercise.

Figure 8. Probabilities of Transiting from Good to Poor Health



For comparison, we report in Figure 9 the actual consumption-age profiles by exercise status from data. As in Figure 7, consumption declines faster with age for those who don't exercise, although in the data it seems that majority of the differences between the two curves occur between age 65 and 70.

Figure 9. Consumption-Age Profiles by Exercise Status: Data



Although we analyze only one type of health investments in this paper, we expect other investments to have similar effects. Of particular importance is medical care utilization. With health in the utility function, medical care could affect individual consumption directly through the budget constraint and indirectly through its effect on health. While the direct effect through the budget constraint is more of an income/wealth effect that affects the consumption level in all periods and shifts the consumption-age profile up and down, the indirect effect through health, by altering the health-age profile, is likely to affect the allocation of consumption across periods and to rotate the consumption-age profile. The demand for medical care is clearly a function of its price, which depends on a series of policy instruments like Medicare and Medicaid. Through the effect on consumption and health in retirement, Medicare and Medicaid could also affect the adequacy of retirement savings and individual decisions to retire. One direction for future

research is to study the optimal design of Medicare and Medicaid by taking these effects into consideration.

Conclusion

We study the interaction between consumption and health in retirement. Our main contribution is the estimation of a consumption Euler equation taking health into consideration. The Euler equation is derived from a model of consumption in retirement with three important building blocks of health: health shocks, health as an investment, and health as a provider of utility. We estimate the Euler equation using data from the HRS and CAMS. The estimates suggest that health is an important determinant of utility.

We use the estimated model to study the empirical significance of the three building blocks of health. We find that health shocks play an important role in slowing down the decline of consumption with age in retirement. Without health shocks, retirees will run down their wealth at a much higher speed. We also find that including health into the utility function provides interesting interactions between health and consumption and could help explain the heterogeneous consumption-age profiles related to health. Finally, we find that health investments, such as physical exercise, have a significant effect on the evolutions of both health and consumption in retirement.

The findings in this paper suggest that public programs like Medicare and Medicaid play an important role in shaping the consumption and health behavior of the retirees, as well as the retirement and saving decisions of working people. One direction for future research is to further quantify the magnitude of these effects and to include these effects into the optimal design of these programs.

References

- 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
- 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.
- Doyle, Joseph J, Jr, "Health Insurance, Treatment and Outcomes: Using Auto Accidents as Health Shocks," NBER Working Paper 11099, February 2005
- Engen, Eric M., William G. Gale and Cori R. Uccello. 1999. "The Adequacy of Retirement Saving." *Brookings Papers on Economic Activity* 2: 65--165.
- Finkelstein, Amy, Erzo Luttmer, and Matthew Notowidigdo, 2013, "What Good is Wealth without Health? The Effect of Health on the Marginal Utility of Consumption," *Journal of the European Economic Association*, Volume 11, Issue Supplement s1, pages 221--258, January.
- Fonseca, Raquel, Pierre-Carl Michaud, Titus Gamama, and Arie Kapteyn, 2009, "On the Rise of Health Spending and Longevity," IZA Discussion paper #4622
- 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.
- Grossman, Michael, 1972, "One the Concept of Health Capital and the Demand for Health," *Journal of Political Economy*, 80, 223-255
- Hubbard, Glen R., Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy*, 103(2) (April): 360--99.
- Murphy, Kevin M. and Topel, Robert H. (2006). "The Value of Health and Longevity," *Journal of Political Economy* 114 (5):871-904.
- Oster, Emily, Ira Shoulson, and Ray Dorsey, 2012. "Limited Life Expectancy, Human Capital and Health Investments," *American Economic Review*, forthcoming.
- 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.
- Rosen, Sherwin. 1988. "The Value of Changes in Life Expectancy," *Journal of Risk and Uncertainty* 1 (September): 285--304.
- Scholz, John Karl, Ananth Seshadri, and Surachai Khittrakun. 2006. "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy*. August, 607-643.

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.

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