

Retirement and Consumption in a Life Cycle Model

David M. Blau

University of North Carolina at Chapel Hill

August 2004

David Blau
Department of Economics
Gardner Hall
University of North Carolina
Chapel Hill, NC 27599-3305
david_blau@unc.edu
919-966-3962, fax: 919-966-4986

I am grateful to the National Institute on Aging and INSEE for financial support and to CREST-INSEE and DEEE-INSEE for hospitality during the work on this paper. I thank Eric French and Jon Skinner for very detailed and helpful comments. I am also grateful to Rob Alessie, John Bailey Jones, Pierre-Olivier Gourinchas, Steven Haider, Winnfried Konniger, Francis Kramarz, Anne Laferrere, Guy Laroque, Jacques Mairesse, Gerard van den Berg, and seminar participants at the University of Basel, CREST-INSEE, IZA, the Finnish Doctoral Program in Economics, and the Tinbergen Institute for helpful comments. I thank Alan Gustman and Tom Steinmeier for helpful advice on pension data. None of the above are in any way responsible for the content.

Retirement and Consumption in a Life Cycle Model

Abstract

Consumption expenditure declines sharply at the time of retirement for many households. Some analysts argue that this is inconsistent with the behavior implied by forward-looking life-cycle models. However, most studies of saving for retirement take the timing of retirement as given, and analyze consumption behavior conditional on retirement. This study builds a life cycle model in which both consumption and employment are choices. The model incorporates key constraints facing older households, including Social Security retirement and disability programs, employer pensions, stochastic earnings and asset returns processes, layoff risk, job offer risk, and health and mortality risk. The model is solved and calibrated to the behavior of a cohort of U.S. men, using data from the Health and Retirement Study. The main finding is that there is a drop in consumption at retirement in simulations of the model. Three features of the model are critical to this result: (a) retirement is a discrete event, (b) retirement is a choice, and (c) uncertainty. Eliminating the retirement decision or eliminating uncertainty from the model results in no drop in consumption at retirement. Adding more hours-of-work choices to the model makes the retirement choice less lumpy, and results in a smaller drop in consumption at retirement. The magnitude of the drop in consumption at retirement is closely related to the amount of wealth at retirement. A drop in consumption at retirement is avoidable for most households, but appears to be part of the optimal life cycle plan in many cases.

JEL: J260, H550

1. Introduction

Consumption expenditure declines substantially following retirement for many U.S. households (Bernheim, Skinner, and Weinberg, 2001; Hurd and Rohwedder, 2003). Some analysts argue that this is inconsistent with the behavior implied by forward-looking life-cycle models in which consumption smoothing is desirable. Alternative “behavioral” models of saving, in which consumers have limited ability or willingness to plan for the future or to carry out their plans, imply that a laissez-faire approach to saving policy is likely to be inefficient (Ameriks et al., 2003; Bernheim 1994; Laibson et al., 1998; Poterba, Venti, and Wise, 1996; Thaler and Benartzi, 2004). Others argue that the evidence is consistent with adequate retirement savings, and that policies intended to encourage greater savings are likely to be ineffective and costly (Lazear, 1994; Engen, Gale, and Uccello, 1999; Engen, Gale, and Scholz, 1996). This is clearly a major public policy issue in view of the rapid aging of the population.

Most studies of saving for retirement take the timing of retirement as given, and analyze saving behavior conditional on retirement. However, just as there is wide dispersion in wealth at retirement, there is also wide dispersion in the age of retirement in the U.S. If consumers exercise choice over both the timing of retirement and saving for retirement, then conditioning on one behavior in the analysis of the other possibly closely-related behavior could produce misleading inferences. More generally, factors that affect saving behavior may also affect employment decisions. Such factors include preferences, the probability distribution of future income, realizations of the income process, and government programs that affect income and consumption.

Conversely, studies of the timing of retirement have focused mainly on health, Social

Security, and pension incentives as explanations for empirical regularities and variation. Most empirical studies either ignore wealth as a determinant of labor force exit behavior, or treat accumulated wealth as given in the retirement decision. French and Jones (2004) and Van der Klaauw and Wolpin (2003) analyze saving and retirement behavior jointly, but do not discuss the implications of their models for consumption following retirement. French (2003) analyzes consumption and labor supply jointly and finds that the only way to explain the drop in consumption at retirement is with a utility function that is non-separable in consumption and leisure within periods.

This study builds a life cycle model in which both consumption and employment are choices. The model incorporates key constraints facing older households: Social Security retirement and disability programs, employer pensions, a consumption floor provided by government welfare programs, stochastic processes for earnings and asset returns, layoff risk, job offer risk, health risk, and mortality risk. The model is calibrated to data from the Health and Retirement Study (HRS) and solved, and its implications for consumption behavior are analyzed. The issue examined here is whether the model can generate simulated consumption and retirement behavior that resemble the general patterns of behavior we see in the data, in particular the drop in consumption at retirement, *without* resorting to non-separable preferences.

The main finding of the study is that the model generates a drop in consumption at retirement, despite the assumption of separable preferences. The critical features of the model that generate this behavior are (a) retirement is a “lumpy” event, (b) retirement is a choice, and (c) uncertainty. A model with these features generates a drop in consumption at retirement for a wide range of parameter configurations regardless of the other features of the model. A model in

which retirement is not a choice or there is no uncertainty does not generate such behavior, regardless of other features of the model. And making retirement less lumpy by adding more hours-of-work choices to the model reduces the magnitude of the drop in consumption at retirement.

The explanation for why a drop in consumption at retirement is consistent with a forward-looking life cycle model with separable preferences, despite the consumption-smoothing motive built into such a model, is as follows. Income drops substantially at the time of retirement for the typical U.S. household, as Social Security and pension benefits replace only a fraction of earnings. In an environment of uncertainty, the age of retirement is *a priori* unknown. Individuals may attach much higher probability to retiring at some future ages than at others, but shocks to health, earnings, preferences, etc. can cause changes in retirement plans. When retirement actually does occur, an important source of uncertainty is resolved: namely, the age at which income drops substantially. Consumption drops at this age because retirement is a negative shock to lifetime resources. The discreteness of the retirement decision is also a critical part of the explanation: if hours of work could be continuously adjusted, then consumption would not change as abruptly at retirement. When consumption is the only continuous choice variable, it is optimal to change consumption abruptly at retirement. The finding that increasing the number of discrete hours-of-work options in the model reduces the magnitude of the drop in consumption at retirement is consistent with this intuition. The findings reported here imply that reasoning about saving for retirement based on unrealistically simple life cycle models can produce misleading inferences. The discrete nature of retirement and the fact that retirement is a choice are critical features of a model of saving for retirement.

Several recent analyses of retirement consumption behavior recognize that retirement may be endogenous to consumption (Banks, Blundell, and Tanner, 1998; Bernheim, Skinner, and Weinberg, 2001; and others). These studies use Instrumental Variables(IV) approaches to estimation, treating retirement as endogenous in a consumption equation. Even after instrumenting retirement, using a variety of identification assumptions, they find a drop in consumption at retirement. Bernheim et al. (2001) take this as evidence that the drop in consumption at retirement is inconsistent with forward-looking life cycle behavior. However, these IV approaches use a linear model for consumption with a low order polynomial in age. I replicate this approach on simulated data from the life cycle model, and find the same result as in these studies: there is still a drop in consumption at retirement even after instrumenting, albeit smaller than without instrumenting. Yet the data used here are generated from a life cycle model, so by construction they are consistent with such a model. Hence I conclude that the IV approach based on a linear consumption equation is not successful at controlling for the endogeneity of retirement, most likely because the consumption equation implied by a life cycle model is highly nonlinear in age.

The main determinant of the magnitude of the drop in consumption at retirement in the model is the amount of wealth held at the time of retirement. All agents start with the same wealth endowment in simulations of the model, but they accumulate different amounts of wealth because they experience different sequences of realizations of the stochastic processes. Low wealth accumulation leads to a larger drop in consumption at retirement, other things equal. Simulated agents who accumulate the median amount of wealth observed at retirement in the HRS data experience a fairly small drop in consumption at retirement. Nevertheless, the typical

U.S. household experiences a fairly large drop in consumption at retirement. The most plausible explanation for this discrepancy is that households do not consider home equity as part of saving for retirement consumption purposes. The majority of wealth is held as home equity, and when this portion of wealth is excluded from the analysis, the drop in consumption at retirement predicted by the model is much closer to the observed decrease.

The next section of the paper describes the drop in consumption at retirement and discusses how previous studies attempt to explain it. The following sections describe the model and solution approach; data, calibrated parameters, and model fit; and the results of simulations of behavior based on the model. The final section offers conclusions.

2. The Drop in Consumption at Retirement

The drop in consumption expenditure at retirement has been documented in the U.S. with data from the Panel Study of Income Dynamics (PSID), the Retirement History Survey (RHS), the Health and Retirement Study (HRS), and other data sets. However, these surveys have significant limitations for analyzing consumption expenditure. Most cover only a few consumption items, typically food, utilities, and housing. They are general purpose surveys, and consumption expenditure is not measured in the same careful way as in consumption surveys. In the subsequent calibration analysis, I use the HRS because it has the most complete and up-to-date information on the older U.S. population. In this section, I document consumption expenditure at the time of retirement using the Consumer Expenditure Survey (CEX), which contains the highest quality consumption data available in the U.S. The CEX is a short rotating panel; households are interviewed five times at three month intervals and then dropped from the

survey. The first interview collects baseline information. Consumption expenditure data are recorded at the remaining four interviews for each of the three calendar months preceding the month of the interview. Combining the 12 months of data for a household that responds to all surveys yields a measure of annual consumption expenditure. Data for households that responded to some but not all of the last four surveys are scaled up to provide an estimate of annual expenditure.

In the calibration analysis, I use a relatively homogeneous sample in order to avoid mixing life cycle patterns with other sources of heterogeneity. Specifically, the HRS subsample that I use consists of married, white, non-Hispanic men with exactly 12 years of education, born from 1921-1950. Here I use a subsample of the CEX drawn from the same population, and I also present some results for a broader population. I use two measures of consumption expenditure: total and non-durable. And I use two measures of employment status: employed, and employed year round full time. The latter corresponds to the definition of employment used in the calibration analysis. I report results derived from CEX surveys covering the period 1980-2000. The analysis in this section is descriptive, so there is no effort to account for joint determination of employment and consumption. I report results based on both cross section and cohort analysis.

The upper panel of Table 1 shows that consumption expenditure is 22 log points higher among employed men compared to non-employed men at ages 50-80 in the relatively homogenous sample. This regression controls for single year-of-age dummies and single year-of-birth dummies. Adding controls for family structure, region and urban location, and spouse characteristics reduces the difference to 18 log points, and controlling for the spouse's employment status reduces it to 16 log points. Comparing men employed year-round full-time to

other men yields differences that are 6-8 log points smaller. The associations between employment and non-durable consumption are very similar to those for total consumption. The results were also very similar when (a) the analysis was limited to the typical age range of retirement of 55-65; (b) 10-year birth cohorts were analyzed separately; and (c) occupational expenses were eliminated from the consumption measure. The results in the lower part of Table 1 for a broader population from the same birth cohorts yield differences that are substantially larger, even controlling for the respondent's characteristics. In rows 2-3, the differences are 15-26 log points, compared to 10-20 log points in the upper panel. Overall, these findings are similar to those reported in other U.S. studies of the retirement-consumption puzzle, and confirm that there are large differences in consumption expenditure among older households as a function of the employment status of the male household head. This does not demonstrate that consumption falls at retirement, however, because the analysis is based on a time series of cross sections in which individuals cannot be followed as they move from employment to non-employment.

In order to examine whether consumption falls at retirement, I follow Banks, Blundell, and Tanner (1998) by studying how average consumption expenditure of birth cohorts changes over time as their employment status changes. To do this, I use eight birth cohorts of men, in four-year intervals from 1921-24 through 1949-52. Averaging log consumption and employment within each birth cohort in each year and taking first differences yields 67 observations in which the cell size is at least 100 for the eight cohorts between the ages of 51 and 72. The regression results reported in the upper panel of Table 2 control for birth cohort dummies and a linear age term. The results show a strong association between the year-to-year change in employment

status and the change in log consumption, with effects ranging from .57 to .71. Results are also reported for the more homogenous sample in the lower panel, although the underlying cell sizes are much smaller. The effects for this group are similar in magnitude. The magnitude of the drop in consumption at retirement implied by the cohort analysis in Table 2 is consistent with the magnitude implied by the cross-section estimates in Table 1. The most common age of retirement in the HRS data is 62, and as shown below the year-round full-time employment rate falls by 20 percentage points from 61 to 62. Multiplying the coefficient estimate of .74 for the narrow cohort in Table 2 by .2 yields an implied drop in consumption at retirement at age 62 of .148, which is close to the estimates of .100 to .141 for the corresponding column in Table 1.

How can the drop in consumption at retirement be explained in economic terms? A large literature on household saving behavior specifies life cycle models in which the main motive for saving is to finance consumption during retirement (see Browning and Lusardi, 1996, for a review). In these models, risk-averse consumers prefer consumption trajectories that keep the marginal utility of consumption relatively constant over the life cycle. A large drop in annual income at the time of retirement is typical in the U.S., so smoothing the marginal utility of consumption requires saving while employed and dis-saving during retirement. Bernheim, Skinner, and Weinberg (2001) argue that strong tastes for work-related consumption and leisure substitutes could rationalize a drop in consumption at retirement and low wealth at retirement, in anticipation of lower consumption after retirement. Their empirical results using data from the PSID indicate that total consumption expenditure falls discontinuously at retirement, but expenditure on work-related items and leisure substitutes does not fall. They note that a consumer who retires earlier than planned due to an unanticipated event would have lower

wealth at retirement and a bigger drop in consumption at retirement than a consumer who retired at the planned date. They re-estimate their consumption model treating the age of retirement as endogenous, using age dummies and education as identifying instruments. They still find a strong negative correlation between retirement wealth and the magnitude of the drop in consumption at the time of “predicted” retirement. Bernheim, Skinner, and Weinberg conclude that their evidence is inconsistent with the implications of forward-looking life cycle models, but potentially consistent with behavioral models of saving.

Banks, Blundell, and Tanner (1998) find that the bulk of the fall in consumption at retirement in the U.K. can be accounted for by anticipated changes in demographics and labor market status - i.e., non-separability of consumption and leisure. But an important proportion is unexplained. They argue that the only way to fully reconcile the fall in consumption around the time of retirement with the life cycle model is with the arrival of unexpected adverse information at the time of retirement. For example, consumers might systematically overestimate their retirement income, and become aware of their actual retirement income only upon retiring.

Several recent papers provide further analysis of the retirement-consumption puzzle. Hurd and Rohwedder (2003) report that the majority of HRS households that are not yet retired anticipate that consumption will drop following retirement. The average decline in consumption expenditure forecast by these households is quite close to the average actual decline reported by households that have already retired. They investigate substitution of home production for market goods as a possible explanation for the drop. Their results indicate that hours per week spent on home activities that are possible substitutes for purchased goods (cleaning, washing, shopping, preparing meals, etc.) are significantly higher for retired men compared to men of the

same age who are not retired. Aguiar and Hurst (2004) find that food *expenditure* decreases at retirement but food *consumption* does not. Using data from diaries recording food eaten during a 24 hour period, they show that retired households consume on average almost exactly the same number of calories per day as non-retired households, other things equal, while food expenditure for retired households is about 17 percent lower than for non-retired households. They also investigate home production as an explanation, and find that retired households spend more time on home production of food than do non-retired households. Haider and Stephens (2004) find a nine percent drop in food consumption at retirement in the U.S. Retirement History Survey (RHS) data from 1969-79. Using subjective retirement expectations as an instrument for actual retirement does not change the magnitude of the estimate. See Lundberg, Startz, and Stillman (2003) and Miniaci, Monfardini, and Weber (2003) for additional evidence.

The model developed in this paper deliberately abstracts from most of the issues discussed in these papers. In this model, consumers are *ex ante* identical. They have different *ex post* outcomes because they experience different realizations of the stochastic processes for earnings, layoffs, asset returns, and so forth. In this respect, the model resembles those of Engen, Gale, and Uccello (1999) and Scholz, Seshadri, and Khitatrakun (2004). These papers find that most American households are saving at least as much as the “optimal” amount implied by their calibrated life cycle models, *given* the actual earnings histories and Social Security and pension entitlements of the households in their samples. The key difference between their approach and mine is that consumers have no control over the timing of retirement in their models, while in my model employment is a decision variable. Furthermore, in the model developed here, consumption and leisure are additively separable in the period-specific utility function. Thus I do

not allow for either non-separability of consumption and leisure or home production. The question posed here is whether a life cycle model can generate retirement-consumption behavior like that observed in data, *without* resorting to heterogeneity in preferences, expectations, and constraints, and without assuming non-separability of consumption and leisure.

3. The Model

A. Model Specification

The agent in the model is assumed to be married, but behavior of the spouse is not modeled. Time is discrete and is indexed by t . All shocks dated period t are realized at the beginning of the period, before period t choices are made. There is a finite horizon, T , and death is certain at $T+1$. $t = 1$ corresponds to age 51, so age in period t is $a_t = 50 + t$. A period is one year.

Health and mortality. The probability of health changing from h to i at the beginning of period t is π_{hit} , $h = 0, 1$; $i = 0, 1, 2$, where 0=good, 1=bad, and 2=dead. Health transition rates are specified as functions of age, but do not depend on other state variables.

Employment and earnings. The employment choice variable is j_t , which takes on the value 1 if the agent chooses to be employed in period t and 0 if he chooses non-employment. There is a layoff probability of λ each period. If the agent was employed in period $t-1$ and not laid off at the beginning of t , then his choice set includes employment. If he was not employed in $t-1$ or was laid off at the beginning of t , then he receives a job offer at the beginning of t with probability ϕ_t . With probability $1 - \phi_t$, there is no job offer and the only option available is non-employment. The log hourly wage offer in period t is $\ln w_t = \beta + \eta_t$, where $\eta_t \sim iid N(0, \sigma_\eta^2)$.

Wages are constant with respect to age and experience because this is what the data show at older ages. Annual earnings are $W_t = 2000w_t$. There are no differences among jobs *except* that the job held entering period 1 provides a pension, described below. In any period in which the agent remains employed at the pension job and has not yet reached the mandatory pension-entitlement age, there is a probability χ that he receives an offer from another employer that does not provide a pension. If he receives such an offer, then in addition to choosing between employment and non-employment, he also chooses whether to remain on the pension job or take the new offer. Leaving the pension job is irreversible, although re-entry to employment at another job is possible. The job choice variable is k_t , where $k_t = 0$ if the individual chooses to remain on the pension job, and $k_t = 1$ if he leaves the pension job and either accepts an offer from a different employer or leaves employment. Once he has left the pension job, there are no differences among jobs and job choice is no longer modeled. The log wage offer received from the alternative job in periods in which alternative jobs are relevant is $\ln w_t^a = \beta + \xi_t$, where $\xi_t \sim iid N(0, \sigma_\xi^2)$.

Social Security. The Social Security retirement (Old-Age and Survivors Insurance, or OASI) benefit is determined by the formula $B_t = B(AIME_t, fe, j_t W_t, a_t)$, where $AIME_t$ is average indexed monthly earnings, and fe is the age at which the individual first becomes entitled to OASI. Earnings (jW) matter because there is an earnings test at some ages. Age (a_t) matters because there is a minimum age of eligibility, and because the earnings test is age-specific. Details of the formula are described in the Appendix. Initial entitlement is assumed to take place in the first period on or after the minimum age of eligibility (62) in which the individual chooses non-employment. Thus entitlement is not a choice variable. Initial entitlement is a one-time

irreversible choice, although the benefit can vary over time as a result of the earnings test and subsequent adjustment for benefits lost as a result of the earnings test. Entitlement is assumed to take place no later than age 70, since there is no further gain to postponing entitlement past age 70 and no earnings test

The Social Security Disability Insurance (SSDI) benefit for an individual who is enrolled in SSDI is determined by the formula $B_t = D(AIME_p, t)$, for $a_t < \min\{fe, anr\}$, where anr is the age of normal retirement (here, taken to be 65). Eligibility for SSDI benefits is assumed to be determined as follows. An agent can become entitled to SSDI benefits only after one year spent in non-employment (the waiting period, which in reality is 5 months) and in bad health. Application is assumed to be costless, so everyone who is in bad health and chooses non-employment at age t (and has not yet entitled to the OASI benefit or reached the age of normal retirement) applies. At the beginning of $t+1$, an individual who applied for SSDI in period t is accepted to the program with probability ψ . An individual cannot be on SSDI if he is employed, so he is automatically terminated in any period in which he chooses to be employed. Otherwise, he remains on the program until he reaches the age of normal retirement, since being enrolled is assumed to be costless conditional on non-employment. The SSDI benefit formula is similar to the retirement benefit formula except that there is no reduction for early entitlement. An individual who is enrolled in SSDI upon reaching the age of normal retirement has his SSDI benefit converted to a retirement (OASI) benefit, following Social Security rules. This specification is an approximation to the complicated application, appeal, and eligibility process that characterizes the actual SSDI program.

If the individual has not yet reached the age of eligibility for early retirement and is not

on SSDI, then average earnings ($AIME$) evolve according to the rule

$$AIME_t = (AIME_{t-1} * (t^* - 1) + j_t \min\{W_t, W_m\}) / t^*$$

where W_m is the maximum amount of taxable Social Security earnings, $t^* = 25 + t$, and the agent is assumed to have worked for 25 years prior to the first period of the model. If he is between the early and normal retirement ages, then $AIME$ follows this rule or the rule $AIME_t = AIME_{t-1}$ depending on which yields the higher value. Beginning at the age of normal retirement, $AIME_t = AIME_{t-1}$.

Employer-Provided Pension. The job held at $t=1$ provides a defined benefit (DB) pension plan with a benefit formula that is representative of such plans in the U.S. DB pension plans typically strongly penalize exit before the early retirement age, while providing relatively little incentive to remain after the early retirement age, and strong incentives to exit by the age of normal retirement. The pension benefit P_t is typically based on average pension-job earnings during the last few years on the job. However, in order to make solution of the model feasible, I assume that the benefit is determined by the Social Security $AIME$ at the time of pension entitlement, as in French (2003). The age of early retirement is assumed to be 55, and the age of normal retirement 65. These are the most common ages for early and normal retirement in such plans. If the agent remains on the pension job at age 65, I assume that he nevertheless becomes entitled to his pension since there is no further increase in the benefit from further delay, and he can receive the benefit while remaining employed at the pension job. See the Appendix for further details.

Assets and returns. The law of motion for assets is $A_{t+1} = A_t^* R_{t+1}$, where A_t^* is assets at the end of t , A_{t+1} is assets at the beginning of $t+1$, $R_{t+1} = 1 + r_{t+1}$, and r_{t+1} is the rate of return on

assets. R_{t+1} , the return factor, is determined by $R_{t+1} = \bar{R} \exp\{\theta_{t+1}\}$, where $\theta_t \sim iid N(0, \sigma_\theta^2)$. Thus, the model allows for uncertainty in the rate of return, with a stationary process. Returns are defined to include capital gains, so $R_t < 1$ corresponds to a capital loss. I assume there is a liquidity constraint, $\bar{A} > 0$, and a consumption floor, $\bar{C} > 0$. The liquidity constraint prevents agents from borrowing against uncertain future income. The consumption floor is a simple way of modeling government welfare programs such as Supplemental Security Income (SSI), Food Stamps, and Medicaid that allow individuals with no other sources of income to survive (Hubbard, Skinner, and Zeldes, 1995). If cash on hand (income plus assets) for given employment and job choices is less than \bar{C} , then the government provides the individual with consumption of \bar{C} and confiscates all cash on hand.

Other income and medical expenditure. To account for income other than earnings and Social Security and pension benefits, I specify a simple exogenous process for the log of other income given by $\ln O_t = o_0 + o_1 t + \zeta_t$, where $\zeta_t \sim iid N(0, \sigma_\zeta^2)$. In the data, most other income is from the spouse, and here this income is assumed to be taken as given by the agent. The log of out-of-pocket medical expenditure is also modeled as an exogenous stochastic process that can be affected by health status, $\ln M_t = m_0 + m_1 t + m_2 h_t + \nu_t$, where $\nu_t \sim iid N(0, \sigma_\nu^2)$. There is no utility from medical expenditure, and health insurance is not explicitly modeled.

Income, taxes, and consumption. Non-asset income net of out-of-pocket medical expenditure is $I_t = j_t W_t + B_t + P_t + O_t - M_t - \tau_t$, where τ_t includes federal income and payroll taxes, calculated using the rules in effect for 1992. Cash on hand at the beginning of period t net of out-of-pocket medical expenditure and taxes is $A_t + I_t$, and assets carried forward to the next period are $A_t^* = A_t + I_t - C_t$, where C_t is non-medical consumption expenditure.

Utility function. Utility in period t for employment choice j , job choice k , consumption choice C , and health status h is given by

$$U_t = C_t^{1-\alpha}/(1-\alpha) + (\gamma_{h1} + \gamma_{h2}t)j_t + \gamma_{h3}(1-j_{t-1})j_t + \varepsilon_{hjkt}$$

where α is the coefficient of relative risk aversion, the γ_{hi} are health-specific determinants of the disutility of employment, and ε_{hjkt} is an *iid* (across periods) shock to utility: $\varepsilon_{hjkt} \sim N(0, \sigma_\varepsilon^2)$. Note that leisure and consumption are additively separable within periods. The coefficient γ_{h3} allows for a psychic cost of returning to work, which is a simple way of preventing excessive “churning” in employment behavior.

The model includes a bequest motive, because the agent is assumed to be a married man. The typical married man can expect his wife to outlive him, and a bequest motive allows him to place value on his wife’s consumption after his death. The utility received in period t in the event of death in $t+1$ is specified as $\kappa_t \ln(\max\{A_t^*, \bar{A}_t\})$, where $\kappa_t = \kappa_0 - \kappa_1 t$, and $\kappa_1 > 0$. The strength of the bequest motive declines with age because the longer one lives, the more likely it is that one’s spouse has already died, in which case the bequest motive would presumably become weaker. This approach avoids incorporating the wife’s presence as an explicit state variable.

The agent’s goal is to choose j_t , k_t , and C_t , $t=1, \dots, T$, to maximize the expected present discounted value of lifetime utility, with discount factor δ . The model is solved by backward recursion starting at period T , using a variant of the monte carlo simulation and regression approximation approach developed by Keane and Wolpin (1994). The solution method is sketched in the Appendix, and a more detailed description is available on request.

Several limitations of the model should be noted. First, there are no permanent shocks to earnings and asset returns. Second, there is no asset allocation decision. Third, the employment

decision is binary. Fourth, there is no defined contribution pension (DC) plan. Finally, the decision to claim Social Security retirement benefits is not modeled directly. Below, I discuss the consequences of relaxing some of these restrictions.

4. Data and Calibration

Data from the U.S. Health and Retirement Study (HRS) are used to characterize labor force behavior, income, health, and mortality. The model is calibrated to generate employment behavior that resembles the average by age of key variables in the HRS. The HRS began in 1992 with baseline in-person interviews with a sample of households containing individuals aged 51-61 and their spouses. Telephone interviews have been conducted every two years since 1994. New cohorts have been added periodically. I use data from the RAND HRS files (www.rand.org/labor/aging).

The model is calibrated to data for white non-Hispanic men born between 1921 and 1950 who graduated from high school but did not attend college, and are married. I chose a relatively homogeneous group because the model does not allow for heterogeneity.¹ The sample size is 1,970, with 8,025 person-year observations from the first six waves of the HRS (1992-2002). Employment is defined here as working full-time (35+ hours per week) year round (36+ weeks per year). The average hourly earnings of individuals who were employed by this definition is the wage measure (excluding cases with hourly earnings less than \$2 or greater than \$500).

¹Employment patterns for ten-year birth cohorts (1921-30, 1931-40, 1941-50) within this sample are very similar, but employment patterns for different education groups within ten-year birth cohorts are quite different. Hence, I use a narrowly defined education group to preserve homogeneity, and a relatively wide birth cohort definition in order increase sample size.

Health is defined as bad if the answer to the question on self-reported health status is “fair” or “poor” and as good if the answer is “good,” “very good,” or “excellent.” Information on the date of death is available from relatives and the National Death Index. Assets are measured by the total net worth of the household. The ages at which the respondent first claimed his OASI and pension benefits are recorded, and an indicator for whether the respondent is receiving SSDI benefits is available as well. The amounts of each of these benefits are also recorded. All dollar amounts are inflated by the Consumer Price Index to 2003 dollars.

The model was calibrated as follows. Logit health transition equations were estimated as a function of age to characterize the health transition and mortality risk parameters. The wage data were used to estimate the parameters of the log wage distribution. The layoff, other income, and medical expenditure process parameters were estimated from HRS data. The mean rate of return on assets was estimated from aggregate data covering the 1992-2003 period, with weights on the different components of assets determined by observed shares in the HRS sample. The pension parameters were chosen to be representative of defined benefit plans in the U.S. See the Appendix for details on asset returns and pensions.

Alternative arbitrary combinations of the coefficient of relative risk aversion (α), time preference (δ), the job offer rate (ϕ_t) and the variance of the asset return process (σ_θ^2) were specified. The utility function parameters (γ 's and σ_ε 's), bequest parameters ($\kappa_t = \kappa_0 + \kappa_1 t$), and SSDI acceptance rate (ψ) were then chosen so as to match the observed age profiles of employment, employment transitions, OASI entitlement, SSDI enrollment, and assets. This was done visually, so this process should be treated as nothing more than an “existence” proof; there is no presumption that the parameters used here are the only ones consistent with the patterns

observed in the data or that they maximize the value of some objective criterion function. Table 3 shows the parameter values for a baseline case with a risk aversion coefficient of 2, a rate of time preference of 3 percent, variance of log asset returns of .0016, and an initially high but steeply declining job offer rate profile.²

The model solution was used to simulate life cycle trajectories for 10,000 artificial agents. Each agent begins at age 51 with given values of assets (\$100,000) and *AIME* (\$32,600), chosen to match initial median assets and mean *AIME* in the data. Agents are randomly assigned an initial health status, employment status, and SSDI status to match the proportions observed in the data. They are randomly assigned a draw from the distribution of the disturbances, and they then make optimal employment, job, and consumption choices for $t=1$ (age=51). The realizations of the stochastic processes together with the choices they make determine the values of the state variables at the end of the period, which are passed forward to the next period. The process repeats until period T or death.

Figures 1-12 illustrate the fit of the model to the age profiles of key variables in the HRS data. Figure 1 shows that the model fits the employment profile quite well, capturing the acceleration in the rate of non-employment around age 62 and the much slower rate of change after age 65. These changes in the slope of the age-employment relationship are mostly due to Social Security and pension incentives in the model, since health status changes smoothly with age and the disutility of employment is constant with respect to age. In Figure 2, the model reproduces the general pattern of the two-year entry rate to employment: high at younger ages

²The job offer rate when non-employed is about 0.80 at age 51, 0.50 at age 60, 0.20 at age 70, and 0.04 at age 80.

and then declining. The model over predicts the early spike in entry. Figure 3 shows that the model is able to mimic the two year transition rate from employment to non-employment reasonably well, although the size of the age-62 spike is under predicted.

Figure 4 shows that the model over predicts the rate of entitlement to OASI benefits at all ages, but captures the general pattern of entitlement occurring predominantly at ages 62-65. In the model, anyone who is not employed at age 62 is assumed to entitle, which is clearly not the case in the data. The OASI benefit is over predicted by about \$2,000 (Figure 5). Figure 6 shows that the model predicts the level and age pattern of receipt of SSDI benefits reasonably well.³ The SSDI benefit is over predicted by about \$2,000 (Figure 7), similar to the OASI benefit.

Figure 8 shows the cumulative distribution of the age of pension entitlement in the data and simulation. In the simulation, entitlement can occur beginning at age 55 and must occur by age 65, by assumption. The data show that few people entitle before age 55 and the entitlement rate increases gradually to about 55 percent by age 65. The simulation shows 30 percent entitlement at age 55 and almost linear growth to 100 percent at age 65. The average pension benefit conditional on receipt shown in Figure 9 declines with age in the data and grows with age in the simulation. It is clearly difficult for a model with a single pension plan to accurately capture patterns in the data that are generated by heterogeneous pensions. In the data, 80 percent of the sample report ever receiving a pension benefit *or* ever being covered by a pension plan from which they expect to receive a benefit. And about 80 percent of this group reports being

³In the model, SSDI receipt is zero by assumption beginning at age 65, since SSDI is converted to OASI at age 65 as required by law. In the data, SSDI and Supplemental Security Income (SSI) receipt are included in the same variable, so the 5 percent rate of receipt at age 65 in the data is presumably SSI.

covered by a DB plan, either alone or in combination with a DC plan. Thus it is clearly important to include a DB pension plan in the model. The divergence between the simulations and the observed pension data is due to a combination of heterogeneity in pension coverage, pension type, and pension plan characteristics in the data, mortality (dying before entitlement), and reporting error. Given the importance of pensions in this sample, I keep the pension in the model despite the inability of the model to accurately mimic pension patterns in the data. However, all of the main qualitative results reported below concerning the drop in consumption at retirement are robust to excluding the pension from the model.

Figure 10 shows that median assets in the HRS data rise from \$100,000 at 51 to around \$250,000 at 65 and are roughly constant or slightly declining thereafter. The model mimics this pattern closely, thanks to the bequest motive. In the absence of a bequest motive the simulated asset trajectory is very different, as discussed below. The sharp increase in median assets observed in the data is partly a result of a high saving rate in the 50s. In the CEX, income exceeds consumption by an average of about \$13,000 between ages 51 and 60. The HRS asset data also show considerable appreciation of home values and other assets.

Mean annual total household income is under predicted by about \$10,000 at younger ages, but is well-captured beginning in the early 60s (Figure 11). The pattern of income declining with age in the data is mimicked by the model. The under prediction of mean income at younger ages is mainly due to under predicting asset income, as shown in Figure 12. The data show mean asset income of \$5-15,000, while the simulation generates \$2-5,000. Mean asset income in the data is heavily influenced by the extreme skewness of the asset distribution.

The model fits the age patterns of mean hourly wage rates, annual earnings, other

income, health transition rates, health, and mean medical expenditure very well (not surprisingly, since these variables were directly calibrated to the HRS data), but in order to conserve space these are not shown.

Figure 13 compares the average consumption trajectory generated by the model with the trajectory derived from the CEX for the same cohort of men. The model predicts a gently rising profile, while the data show a rather steeply declining profile. In the simulation, the rate of time preference (3 percent) exceeds the mean rate of return (2 percent), and mortality risk makes the effective rate of time preference even larger. This would tend to result in a declining consumption profile, but the bequest motive decreases with age, which favors a rising consumption profile. It is not clear why consumption declines with age in the data, but one possibility is cohort effects. The CEX is a cross-section survey and the profile shown in the figure mixes true life cycle patterns with effects resulting from differences across cohort in lifetime income. Banks, Blundell, and Tanner (1998) find a declining consumption-age profile in similar data from the U.K.

5. Implications of the Model for the Retirement-Consumption Puzzle

Figure 14 shows mean consumption and net (of medical expenditure) after-tax income trajectories in the simulations, by employment status. Agents save on average while employed and until age 63 dis-save while not employed. Non-employment at younger ages is often temporary and not due to permanent exit from the labor force, so dis-saving while non-employed at younger ages is not surprising. Beginning at age 63, consumption tracks net income very closely among the non-employed. There is a gap of about \$7,000 between consumption when

employed and not employed at age 51. The gap falls gradually to about \$3,000 at age 62 and remains constant thereafter. This is a manifestation of a drop in consumption at retirement.

To explore this further, Figure 15 displays the mean change in the logarithm of consumption from one period to the next as a function of the change in employment status between the adjacent periods. Changes in consumption associated with remaining employed or non-employed are very close to zero on average. There is an increase in consumption associated with labor force entry. The average change in consumption associated with exit from the labor force is negative at all ages. At the typical ages of retirement in the late 50s to mid 60s, the average drop in consumption associated with exit from employment is 3-7 log points. In order to focus exclusively on the change in consumption associated with retirement, as opposed to temporary withdrawal from employment, Figure 16 plots consumption profiles for selected employment sequences in which there is a clear date of retirement. In all of the cases shown, several consecutive periods of employment are followed by several consecutive periods of non-employment. In each case, there is a clear drop in consumption at retirement, ranging from about \$4,000 (10 percent) for retirement at 55 to about \$2,000 (5 percent) for retirement at 65. Consumption continues its gradual pre-retirement increase following the drop.

Why does consumption drop abruptly at retirement in this model? We can rule out a number of explanations suggested in the literature, since they are not part of the model. These include work-related expenses, non-separability of leisure and consumption, home production, changes in household composition, and lack of self control. Two changes in model specification can eliminate the drop in consumption at retirement: removing employment as a choice variable and eliminating all sources of uncertainty. If consumers are required to work until a specified age

and then forced to retire, consumption does not drop abruptly at retirement. This is illustrated in Figure 17, which shows the consumption path generated by a specification that is identical to the one described above *except* that consumers are required to work in every period through age 61 and then retire at age 62. There is no drop in consumption at retirement in this case; in fact, there is a slight increase. Figure 17 also shows the consumption profile for the case of perfect certainty. In this specification, retirement occurred at age 62, and there is a small (\$1,000) drop in consumption at age 62, but it is similar to small consumption declines at several other ages and seems to be part of an ongoing trend of declining consumption rather than a distinct phenomenon associated with retirement.⁴

These findings suggest that the combination of uncertainty and choice of retirement date is critical for generating a drop in consumption at retirement. Retiring earlier provides more lifetime leisure, and depending on their preferences consumers may be willing to sacrifice a substantial amount of consumption in exchange for additional years of leisure. In a world of certainty, one could plan for this precisely, and the entire lifetime consumption path would be lower for agents who retired earlier. There would be no need for an abrupt drop in consumption. In an uncertain world in which the date of retirement is not known in advance, it can be optimal to reduce consumption abruptly upon retiring as a result of the negative shock to lifetime

⁴Another way to eliminate or substantially reduce the magnitude of the drop in consumption at retirement is to make retirement easily reversible. Setting the utility cost of moving from non-employment to employment (γ_{h3}) to zero, and setting the job offer rate when not employed (ϕ_t) to one, results in an average change in log consumption at the time of exit from employment of only -.012, compared to -.053 in the baseline model. This is sensible: if employment decisions are reversible at low cost, then shocks that result in labor force exit do not have persistent consequences for lifetime resources. However, this specification fits the employment transition data poorly: it generates a much higher rate of movement from non-employment to employment in the simulations than is observed in the data.

resources that results from retirement.

This discussion suggests that “unexpected” or “unplanned” retirement should be associated with a larger drop in consumption than “expected” or “planned” retirement. This idea has been discussed in the literature, but without reference to a specific model the concept of “unexpected” retirement is difficult to precisely define and measure. In the model developed here, there is a natural definition and a straightforward way to measure this concept. Let V_t represent the value function and \mathbf{S}_t the vector of state variables. Define $E_t \max(work) = E_t \max_{C_{t+1}, k_{t+1}} V_{t+1}(\mathbf{S}_{t+1}, j_{t+1}=1 | j_t, C_t, k_t, \mathbf{S}_t)$ as the expected value of working in period $t+1$ given period t choices and states, and $E_t \max(nowork) = E_t \max_{C_{t+1}, k_{t+1}} V_{t+1}(\mathbf{S}_{t+1}, j_{t+1}=0 | j_t, C_t, k_t, \mathbf{S}_t)$ as the expected value of not working. Note that in both cases the optimal choices of C_{t+1} and k_{t+1} are assumed to be made, given the employment decision. Define $EG_t = E_t \max(work) - E_t \max(nowork)$ as the expected gain in lifetime utility from working in $t+1$ relative to not working. A natural interpretation is that the larger is the value of EG_t , the more likely it is that the agent “plans” to work in $t+1$, and the more “unexpected” it is if he actually does not work in $t+1$ given the realizations of the $t+1$ shocks. Thus a higher value of EG_t should be associated with a bigger drop in consumption at retirement.

I used the simulated data to estimate a regression explaining the change in log consumption at the time of labor force exit as a function of age, assets, health change, experience (consecutive periods of employment), lagged consumption (a useful summary statistic for past decisions), and EG_t . To provide a flexible characterization of the effects of assets and expectations, the regression uses a series of dummy variables for intervals of assets and EG_t . Cells in the tails of the asset distribution with small numbers of observations were merged, and

similarly for the EG_t distribution. Cases in which employment was not an option as a result of a layoff and lack of a job offer are excluded. Figure 18 shows the estimated relationship between EG_t and the drop in consumption at retirement, holding the other variables constant at their means. The graph shows that the drop in consumption at retirement is larger when retirement is more unexpected, but also indicates that the effect of expectations is very small. The difference in the consumption decline at retirement between the least and most surprising retirements by this metric is only one log point. Thus it does not appear that “unexpected retirement” is a major part of the quantitative explanation for the drop in retirement at consumption. This is consistent with the finding of Hurd and Rohwedder (2003) that the majority of older households anticipate that consumption will decline at retirement. It is also consistent with the results in French (2003), who finds that in a model in which hours of work is a continuous choice variable, non-separability of leisure and consumption is a much more important explanation for the drop in consumption at retirement than are shocks that result in unexpected retirement.

Figure 19 shows the estimated relationship between assets at the time of retirement and the drop in consumption at retirement, using the regression evaluated at the means of other variables. The predicted drop in log consumption at the time of labor force exit is close to linear in assets, with a drop of about -.2 for assets of \$100,000 or less, -.1 for assets of \$200,000, zero for \$300,000 in assets, and an increase of +.1 for \$400,000 or more in assets. As shown above in Figure 10, median assets are about \$210,000 at age 62, a typical age of retirement, so a drop in log consumption of about -.07 at retirement would be typical in the model according to Figure 19. Agents in the model could avoid a drop in consumption, since they are endowed with \$100,000 in assets in the first period and accumulate assets steadily. Apparently, they gain

greater expected lifetime utility by cutting consumption at the time of retirement, rather than building up additional assets in order smooth consumption.

The recent literature on the retirement-consumption puzzle recognizes that the timing of retirement may be endogenous to consumption. For example, Bernheim, Skinner, and Weinberg (2001) argue that unforeseen events that result in early retirement will also cause a downward revision in expected future income and lower than expected wealth at retirement. This could lead to a drop in consumption at retirement that is consistent with the life cycle framework. To eliminate this source of endogeneity, they replace retirement in their consumption equation with predicted retirement, using as instruments education, family size, gender of the household head, and marital status, all interacted with single-year-of-age dummies. By replacing actual retirement with predicted retirement based on fixed or easily anticipated factors, they argue that the Instrumental Variables (IV) estimate of the effect of retirement on consumption will not be contaminated by the effects of unexpected retirement. A drop in consumption at retirement estimated by this approach, they argue, is inconsistent with the life cycle framework because consumers can easily smooth consumption around “anticipated” retirement. Banks, Blundell, and Tanner (1998), and Haider and Stephens (2004) make similar arguments but use different instruments - lagged retirement, consumption growth, and income growth in the former case, and subjective retirement expectations in the latter.

As noted above, there is a sense in which all retirements are unexpected in the presence of uncertainty, although it does not appear that unexpected retirement is a major source of the drop in consumption in the results presented here. However, it is worth asking whether the IV approaches used in the literature to deal with “unexpected retirement” do in fact accomplish the

desired goal of eliminating the variation in consumption associated with unexpected retirement. I can test this by replicating these approaches using simulated data from the calibrated life cycle model. By construction, *any* drop in consumption in the simulated data is consistent with the life cycle model, since the data are generated by such a model. Thus, if the IV approaches used in previous studies are successful, they should completely eliminate the drop in consumption at retirement in the simulated data.

I attempted to replicate the empirical approaches of Bernheim, Skinner, and Weinberg (2001) and Banks, Blundell, and Tanner (1998). I cannot replicate the approach of Haider and Stephens (2004) because it is not clear how to construct a measure of subjective retirement expectations with simulated data. Details of the replication are provided in the Appendix. Bernheim et al. report in their Table 7 an estimate of $-.685$ (t-statistic 4.3) for the effect of the predicted probability of retirement on the logarithm of consumption for households in the lowest wealth and income quartiles of their sample. My corresponding estimate is $-.038$ (std. err. $.012$) using predicted retirement, compared to $-.052$ (std. err. $.001$) using actual retirement instead of predicted retirement. The IV estimate using simulated data is closer to zero than the OLS estimate, but it is significantly different from zero. Banks et al. report an IV estimate of the effect on log consumption of being out of the labor force of $-.258$ (standard error $.067$) (page 779). My corresponding estimate is $-.013$ ($.001$), compared to an OLS estimate of $-.067$ ($.0004$). The IV estimate is quite small in this case, but is significantly different from zero. The failure of the IV approaches to fully eliminate the drop in consumption in simulated data suggests that the approaches may not be valid. Both IV approaches rely on the assumption that the consumption equation is linear in age (and age squared, for Bernheim et al.). But the model used here (which

is quite standard as far as consumption is concerned) implies a nonlinear effect of age on consumption that may not be well captured by a low order polynomial. Thus, excluding age dummies from the consumption equation may be invalid. On the other hand, the IV approaches do succeed in reducing the estimated magnitude of the drop in consumption at retirement.

6. Robustness

An extensive analysis of the robustness of the drop in consumption at retirement to alternative parameter values and model specifications is summarized here. Results are presented for two samples: all simulated cases of transition from employment to non-employment, and simulated individuals who worked for at least three consecutive periods, left employment, and remained non-employed for at least three consecutive periods as of age 67. Cases in the latter sample have a well-defined date of retirement, although the specific criterion is somewhat arbitrary. In each case, the employment preference and bequest parameters and the SSDI acceptance rate (γ 's, σ_e , κ 's, ψ) were calibrated so that the simulated employment and asset data matched the HRS data by age as closely as possible. The baseline specification is the one shown in Table 3, and the results of the robustness analysis are shown in Table 4.

The first row of Table 4 shows that the baseline specification generates a drop in consumption of -.053 for all labor force exits and -.039 for clear cut cases of retirement. As noted above, the difference in consumption between men employed year round full time and other men in the CEX data is 10-14 log points. To explain this why the model generates a smaller consumption drop than we see in the data, I hypothesize that households do not consider home equity as a source of wealth for consumption smoothing purposes. The average share of

home equity in net worth is 57 percent in the HRS sample in 1992, so non-home wealth is much smaller than total wealth. I re-calibrated the model to fit the trajectory of median non-home wealth rather than total net worth. Figure 20 shows that median non-home wealth rises to a peak of about \$130,000 in both the simulation and the data. In this specification, the average drop in consumption at labor force exit is 9.0 log points (see Table 4). This is closer to the consumption difference by employment status observed in the CEX, and indicates that the hypothesis that households do not consider home equity as a source of wealth for consumption smoothing purposes is one way to reconcile the size of the consumption drop in the model with the data.

I noted above that the lumpiness of the employment decision in the model is one of the probable causes of the drop in consumption at retirement. To test this, I extended the model to allow both full time and part time employment. This makes the employment choice less lumpy and provides agents with another mechanism for adjusting to shocks. I calibrated this version of the model to the HRS data on full time and part time employment. The drop in consumption at retirement in this case is $-.027$, compared to $-.039$ in the baseline case. This confirms the intuition that lumpiness of the employment choice is a factor in the drop in consumption at retirement.

Table 4 shows that the magnitude of the drop in consumption at retirement is robust to alternative values of the risk aversion parameter except for the case of risk neutrality. Risk-neutral agents are indifferent to consumption fluctuations, so a large drop in consumption at retirement in this case is not a surprise. The results are also quite robust to alternative values of the rate of time preference. Increasing the variance of asset returns or earnings causes a bigger drop in consumption at retirement. This is consistent with the finding reported above that

eliminating all sources of uncertainty eliminates the drop in consumption at retirement. Making retirement irreversible by setting the job offer probability to zero for the non-employed increases the drop in consumption at retirement from $-.039$ to $-.068$ in the second column. A shock that leads to retirement clearly has more persistent consequences when the retirement decision is irreversible. Setting the job offer probability to one has little impact. Setting the consumption floor to \$1,000 instead of \$5,000 also has little impact.⁵

7. Conclusions

The results in this paper show that a drop in consumption at retirement is not a puzzle for a life cycle model. In a life cycle model with a reasonable amount of uncertainty, the date of retirement is inherently uncertain. Retirement is a discrete event and is not easily reversible, so when it does occur it often represents a shock to lifetime resources. Thus it should not be surprising that consumption behavior is discontinuous at retirement. Agents could save more in anticipation of this shock, but they are already saving a lot in the 50s and early 60s. Additional saving may be more costly in lifetime utility terms than cutting consumption at the time of retirement. A life cycle model provides agents with the opportunity to smooth consumption, but does not guarantee that it will always be optimal to do so. Reasoning about saving for retirement based on models in which retirement is exogenous or perfectly foreseen can give misleading intuition.

The model proposed in this paper relies heavily on a bequest motive to account for the

⁵Another variation in model specification was to incorporate a permanent shock to earnings, as in a number of recent consumption models (Gourinchas and Parker, 2002; French, 2003). In this specification it was very difficult to obtain a stable solution, most likely because assets and earnings can grow very rapidly in a random walk model. Despite the lack of stability in the solution, the results consistently showed a drop in consumption at retirement.

buildup of assets observed in the data. In the absence of a bequest motive and given observed patterns of earnings and other income, the model predicts a sharp decline in assets and a large drop in consumption at retirement. There are other ways for the model to match the asset distribution without a bequest motive, but they rely on unrealistic parameter values such as a 9 percent real rate of return on assets or elimination of Social Security and pensions. As discussed above, a bequest motive is quite plausible for a sample of married men, given that they can expect their wives to outlive them. But there is mixed evidence in the literature on the importance of a bequest motive (see for example Hurd, 1989, and Kopczuk and Lupton, 2004), so it is somewhat troubling that the model relies so heavily on it in order to match the asset pattern observed in the data. This issue clearly warrants further research.

The model under predicts the magnitude of the observed drop in consumption at retirement unless it is assumed that agents treat housing wealth as off limits for purposes of financing consumption. This assumption is admittedly ad hoc in the context of a life cycle model. Thus, another important task for future research is to either rationalize this assumption, or find another mechanism such as non-separable preferences to allow the model to match the observed magnitude of the drop in consumption at retirement. Finally, estimation of the parameters of a life cycle model such as the one proposed here or in Van der Klaauw and Wolpin (2003) will be the most important test of whether the life cycle model provides a useful quantitative approach to understanding consumption and saving behavior at older ages.

Appendix

Social Security Retirement (OASI)

The formula for computing the Primary Insurance Amount (*PIA*) from Average Indexed Monthly Earnings (*AIME*, measured in dollars per month) in 1992 was:

$$\begin{aligned}
 PIA &= .9 * AIME && \text{if } AIME < \$387 \\
 PIA &= .9 * 387 + .32(AIME - 387) && \text{if } \$387 \leq AIME \leq \$2,333 \\
 PIA &= .9 * 387 + .32(2,333 - 387) + .15(AIME - 2,333) && \text{if } \$2,333 < AIME
 \end{aligned}$$

The OASI monthly benefit equals: *PIA* for first entitlement at age 65; *PIA* reduced by 6.67% per year of first entitlement prior to 65 (a maximum reduction of 20% for the earliest age of entitlement, 62); *PIA* increased by 5% per year of first entitlement past 65, up to a maximum of 25% for entitlement at age 70. See Blau and Gilleskie (2004) for a description of how benefits are approximated for re-entitlement following return to employment after initial entitlement. In 1992, the Social Security Earnings Test for ages 62-64 reduced benefits by 50 cents per dollar of earnings above an annual exempt amount of \$7,440. For ages 65-70, the reduction was 33 cents per dollar of earnings above an exempt amount of \$10,200, and there was no earnings test for ages 70 and above. The maximum taxable earnings amount in 1992 was \$55,500.

Employer-Provided Pension

The formula for the pension benefit is:

$$\begin{aligned}
 P_t &= 0 && \text{if } t < t_e, && P_t &= ex\tau_p AIME_{ip} && \text{if } t_e \leq t_p \leq t < t_n \\
 P_t &= yx\tau_p AIME_{ip} && \text{if } t_p < t_e \leq t && P_t &= x\tau_n AIME_{ip} && \text{if } t_n \leq t_p \leq t,
 \end{aligned}$$

where (using age and period interchangeably) t_e is the age of early retirement in the pension plan; y is a reduction factor for exit from the pension job before the age of early retirement; x is

the fraction that converts earnings and tenure into the normal retirement benefit; τ_p is job tenure on the pension job at the date of exit from the pension job; $AIME_{tp}$ is $AIME$ on the date at which the agent entitled to the pension benefit; t_p is the age at which the individual leaves the pension job; t_n is the age of normal retirement in the pension plan; τ_n is job tenure on the pension job at the age of normal retirement; and e is a reduction factor for exit from the pension job on or after the age of early retirement and before the normal age. Here, $y = .77$, $e = .88$, $x = .03$, $t_e = 5$ (age 55), and $t_n = 15$ (age 65). I assume that tenure on the pension job at the beginning of the first period is 10 years. The probability of receiving a job offer from another employer (χ) is 0.30. Let $P(AIME_{tp}, t_p)$ denote the pension benefit formula, with the parameters implicit.

Solution Method

The value function in period t is the expected present discounted value of remaining lifetime utility from entering period t with state variables \mathbf{S}_t , including the vector of shocks $\boldsymbol{\varepsilon}_t$, and making optimal choices from t through the end of life:

$$V_t(\mathbf{S}_t) = \max_{j_t, k_t, C_t} \{U_{jkCt}(\mathbf{S}_t) + E_t V_{t+1}(\mathbf{S}_{t+1} | \bar{\mathbf{S}}_t)\},$$

where

$$E_t V_{t+1}(\mathbf{S}_{t+1} | \bar{\mathbf{S}}_{t-1}) = E_t \max_{j_{t+1}, k_{t+1}, C_{t+1}} V_{jkC_{t+1}}(\mathbf{S}_{t+1} | \bar{\mathbf{S}}_{t-1}).$$

$\bar{\mathbf{S}}_t$ is the vector of state variables at the end of period t , *after* period t choices have been made, or equivalently at the beginning of period $t+1$ *before* realization of the $t+1$ shocks:

$$\bar{\mathbf{S}}_t = (h_t, A_t^*, AIME_t, fe, un_t, t_p, j_t, SSDI_t),$$

and $V_{jkC_{t+1}}(\mathbf{S}_{t+1} | \bar{\mathbf{S}}_t)$ is the choice-specific value function in $t+1$ given $\bar{\mathbf{S}}_t$. The right hand side of the second equation above is referred to as the “E_{max}.” $U_{jkCt}(\mathbf{S}_t)$ represents current-period utility for given choices j_t, k_t, C_t and state variables. un_t is a state variable = 1 if the individual ever

returned to work following entitlement to OASI, and = 0 otherwise. It is the basis for approximating the OASI benefit for re-entitlement following return to employment after initial entitlement (Blau and Gilleskie, 2004).

The value function if alive in period T is $V_T(\mathbf{S}_T) = \max_{j, C_T} V_{j0CT}(\mathbf{S}_T)$, where

$$V_{j0CT}(\mathbf{S}_T) = C_T^{1-\alpha}/(1-\alpha) + (\gamma_{h1} + \gamma_{h2}t)j_T + \gamma_{h3}(1-j_{T-1})j_T + \kappa_T \ln(A_T^*) + \varepsilon_{hj0T}$$

(There is no pension job choice in T , so $k_T=0$). This equation holds if $I_{j0T} + A_T > \bar{C}$ and $\ln(A_T^*) > \bar{A}$, where I_{j0T} is net after-tax income conditional on employment choice j and job choice k :

$$\begin{aligned} I_{j0T} &= j_T W_T + B_T + P_T + O_T - M_T - \tau_T & P_T &= P(AIME_{jp}, t_p) \\ W_T &= 2000 * \exp\{\beta + \eta_T\} \text{ if } k=0 & W_T &= 2000 * \exp\{\beta + \xi_T\} \text{ if } k=1 \\ B_T &= B(AIME_T, fe, j_T W_T, a_T) & O_T &= \exp\{o_0 + o_1 T + \zeta_T\} \\ M_T &= \exp\{m_0 + m_1 T + m_2 h_T + v_T\} & \tau_T &= \tau(j_T W_T, B_T, P_T, O_T, M_T, A_{t-1}^*(R_t-1)). \end{aligned}$$

The state variables that enter these formulas are determined as described in the text. If $I_{j0T} + A_T < \bar{C}$, then \bar{C} replaces C_T ; and if $A_T^* < \bar{A}$, then \bar{A} replaces A_T^* . The expectation at $T-1$ of $V_T(\mathbf{S}_T)$ (the Emax) is computed by monte carlo integration over the distribution of the period T disturbances, for a random sample of points in the end-of- $T-1$ state space, $\bar{\mathbf{S}}_{T-1}$. The sample is generated by discretizing the continuous state variables (assets and $AIME$) into grids, and randomly sampling points in the discretized state space. The optimal value of consumption for each possible choice of employment in $T-1$ is determined by an iterative grid search over consumption with a successively finer grid at each iteration until a convergence criterion is satisfied. The Emax is then used as the dependent variable in a regression on the $T-1$ state variables, using the sample of $T-1$ state points for which the Emax was computed. By using a

flexible form for the regression function, the R^2 is typically greater than .990.

The regression function is used to approximate $E_{T-1}V_T(\mathbf{S}_T | \bar{\mathbf{S}}_{T-1})$ in the expression for the $T-1$ value function,

$$V_{T-1}(\mathbf{S}_{T-1}) = \max_{j_{T-1}, C_{T-1}} \{U_{jkCT-1}(\mathbf{S}_{T-1}) + E_{T-1}V_T(\mathbf{S}_T | \bar{\mathbf{S}}_{T-1})\},$$

for each trial value of consumption and employment in $T-1$, integrating over the distribution of the $T-1$ disturbances. The expectation of this expression, $E_{T-2}V_{T-1}(\mathbf{S}_{T-1} | \bar{\mathbf{S}}_{T-2})$, is computed for a random sample of state points in $\bar{\mathbf{S}}_{T-2}$, and the sample is used to estimate an approximating regression. This is repeated recursively back to period 1, with the pension job choice incorporated at relevant ages. In the solution, $T=50$, 50 draws from the error distribution were used in monte carlo integration, grid sizes of 500 and 10 were used for discretizing assets and *AIME*, respectively, and a sample size of about 1,000 state points per period was used in the approximating regression. An unpublished appendix describing the solution in more detail is available at <http://www.unc.edu/depts/econ/faculty/blauweb.html>.

Rate of return

The chart below lists the data used to measure the rate of return: the Standard and Poor 500 stock price return, the six month Certificate of Deposit (CD) return, the Moody AAA corporate bond return, and a house price index based on Fannie Mae and Freddie Mac data (www.ofheo.gov/HPI.asp). In order to construct weights for returns on the various items, I used the portfolio composition of the HRS subsample in 1992. The mean percentage shares of home equity, stocks, bonds, CDs, checking and saving accounts, IRAs, vehicles, and other assets in total net wealth were 56.8, 4.1, 0.4, 2.8, 7.5, 8.6, 10.1, and 9.5, respectively. I assumed that IRA assets were allocated equally between stocks and bonds; returns on checking and saving

accounts and vehicles were zero; and returns on other assets were equal to CD returns. The chart shows the resulting nominal return, along with the CPI and the real return. The average real return during the period shown was 2.0 percent.

<u>Year</u>	<u>Stock return</u>	<u>CD return</u>	<u>Bond return</u>	<u>House return</u>	<u>Nominal return</u>	<u>CPI</u>	<u>Real return</u>
1992	4.46	3.76	8.14	1.86	2.29	2.90	-0.61
1993	8.00	3.28	7.22	1.67	2.37	2.75	-0.38
1994	-1.13	4.96	7.97	0.79	1.35	2.67	-1.32
1995	34.11	5.98	7.59	4.50	6.53	2.54	3.99
1996	20.27	5.47	7.37	2.58	4.20	3.32	0.88
1997	31.01	5.73	7.27	4.59	6.27	1.70	4.57
1998	26.67	5.44	6.53	4.96	6.05	1.61	4.43
1999	19.53	5.46	7.05	5.33	5.68	2.68	3.00
2000	-9.19	6.59	7.62	7.63	4.74	3.39	1.35
2001	-13.04	3.66	7.08	7.54	3.98	1.55	2.43
2002	-23.37	1.81	6.49	7.72	2.96	2.38	0.58
2003	26.38	1.17	5.66	7.97	7.16	1.88	5.28

Replication analysis

Bernheim, Skinner, and Weinberg (2001) use cases with a well-defined retirement date, so in replicating their approach I also use simulated cases with a well-defined date of retirement: between ages 57 and 67 they worked at least three years in a row, left employment and remained non-employed for at least three years in a row as of age 67. I measure the ratio of (after-tax) post-retirement to pre-retirement income following the method described in Bernheim et al., and similarly for the ratio of wealth to pre-retirement income. Following their suggestion, I use the wealth ratio at age 59 as a measure of wealth “at some standardized age that is sufficiently advanced for individuals to have accumulated a significant fraction of their retirement resources, but early enough to precede retirement...” (pages 842-3). I estimate a model with the same specification as in column A of their Table 7, except that family size, marital status, female head, and household fixed effects are not relevant. The dependent variable is log consumption, and the regressors are age, age squared, and a retirement dummy (1 if retired, 0 if not) interacted with

dummies for wealth ratio quartile and income ratio quartile. All of the regressors except age and age squared are instrumented using single year of age dummies in the first stage equation. Bernheim et al. also used education interacted with age dummies as an instrument, but this is not relevant here. The model is estimated by Two Stage Least Squares, instead of the two-step probit-OLS approach used by Bernheim et al. The F-statistic for the age dummies in the first stage equation is over 20,000.

Banks, Blundell, and Tanner (1998) use all observations, including cases in which non-employment may be transitory, so in replicating their approach I use all available simulated observations. Their data were not longitudinal, so they created a pseudo-panel by taking means within cells defined by birth year and time period. I use the individual simulated data since it is a panel. I replicate the model reported on page 779 of their paper. The dependent variable is $\Delta \ln C_t$, the change in log consumption from one period to the next. The regressors are age and $\Delta(\text{Head out of labor market})$, the difference in employment status across periods. Their other regressors - multiple adults, the interest rate, and mortality risk - are not relevant here. $\Delta(\text{Head out of labor market})$ is instrumented with a full set of age dummies, twice-lagged consumption growth and income growth, and two-and-three period lags of retirement.

References

- Aguiar, Mark, and Eric Hurst. "Consumption Versus Expenditure." NBER Working Paper 10307, February 2004.
- Ameriks, John, Andrew Caplin, and John Leahy. "Wealth Accumulation and the Propensity to Plan," *Quarterly Journal of Economics*, August 2003: 1007.
- Banks, James, Richard Blundell, and Sarah Tanner. "Is There a Retirement Savings Puzzle?" *American Economic Review*, 88 (4), September 1998: 769-788.
- Bernheim, B. Douglas. "Comment on 'Some Thoughts on Savings,'" in David Wise (ed.) *Studies in the Economics of Aging*, Chicago: University of Chicago Press, 1994: 171-179.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. "What Accounts for the Variation in Retirement Wealth Among U.S. Households?" *American Economic Review*, 91 (4), 2001: 832-857.
- Blau, David M. and Donna B. Gilleskie. "Health Insurance and Retirement of Married Couples," working paper, University of North Carolina at Chapel Hill, 2004.
- Browning, Martin, and Annamaria Lusardi. "Household Savings: Micro Theories and Micro Facts," *Journal of Economic Literature*, 39 (4), December 1996: 1797-1855.
- Engen, Eric M., William G. Gale, and John Karl Scholz. "The Illusory Effect of Savings Incentives on Saving," *Journal of Economic Perspectives* 10 (4) 1996: 113-38.
- Engen, Eric. M, William G. Gale, and Cori E. Uccello. "The Adequacy of Household Saving," *Brookings Papers on Economic Activity*, 1999 (2): 65-187.
- French, Eric. "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior," Working Paper, Federal Reserve Bank of Chicago, December 2003.
- French, Eric, and John Bailey Jones. "The Effect of Health Insurance and Self-Insurance on Retirement Behavior," Working Paper, SUNY-Albany, March 2004.
- Gourinchas, Pierre-Olivier, and Jonathan Parker. "Consumption Over the Life Cycle," *Econometrica* 70 (1), January 2002: 47-90.
- Haider, Steven J., and Melvin Stephens Jr. "Is There A Retirement-Consumption Puzzle? Evidence Using Subjective Retirement Expectations," NBER Working paper 10257, January 2004.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. "Precautionary Saving and Social

- Insurance,” *Journal of Political Economy* 103 (2), April 1995: 360-399.
- Hurd, Michael. “Mortality Risk and Bequests,” *Econometrica* 57 (4) July 1989: 779-813.
- Hurd, Michael, and Susan Rohwedder. “The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement,” NBER Working Paper 9586, March 2003.
- Keane, Michael and Kenneth Wolpin. “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence,” *Review of Economics and Statistics* 1994: 648-672.
- Kopczuk, Wojciech, and Joseph P. Lupton. “To Leave or Not To Leave: The Distribution of Bequest Motives,” Working Paper, Columbia University, May 2004.
- Laibson, David I., Andrea Repetto, and Jeremy Tobacman. “Self-Control and Saving for Retirement,” *Brookings Papers on Economic Activity*, 1998 (1): 91-196.
- Lazear, Edward. “Some Thoughts on Savings,” in David Wise (ed.) *Studies in the Economics of Aging*, Chicago: University of Chicago Press, 1994: 143-169.
- Lundberg, S., R. Startz, and S. Stillman. “The Retirement-Consumption Puzzle: A Marital Bargaining Approach,” *Journal of Public Economics* 87, 2003: 1199-1218.
- Miniaci, R., C. Monfardini, and G. Weber. “Is There a Retirement Consumption Puzzle in Italy?” IFS Working Paper 0314, 2003.
- Poterba, James M., Steven F. Venti, and David A. Wise. “How Retirement Savings Programs Increase Saving,” *Journal of Economic Perspectives* 10 (4) 1996: 91-112.
- Scholz, J.K., A. Seshadri, and S. Khitatrakun. “Are Americans Saving ‘Optimally’ for Retirement?” NBER working paper 10260, 2004.
- Thaler, Richard H., and Shlomo Benartzi. “Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving,” *Journal of Political Economy* 112 (1), Part 2, February 2004: S164-S187.
- Van der Klaauw, Wilbert, and Kenneth Wolpin. “Social Security, Pensions, and the Savings and Retirement Behavior of Households,” manuscript, University of North Carolina at Chapel Hill, April 2003.

Table 1: Employment Status and Consumption Expenditure

Regressor of interest:	Dependent Variable: Ln(Total Consumption Expenditure)		Dependent Variable: Ln(Non-durable Consumption)	
	Employed	Employed YRFT	Employed	Employed YRFT
Sample: Married white non-Hispanic men, educ=12, born 1921-1950, ages 50-80				
Control variables				
1. age, birth year	.222 (.032)	.141 (.032)	.231 (.028)	.130 (.029)
2. Plus core	.183 (.031)	.125 (.031)	.197 (.027)	.119 (.027)
3. Plus spouse empl. Status	.165 (.030)	.100 (.032)	.180 (.027)	.097 (.028)
Sample size	1,685			
Sample: Born 1921-1950, ages 50-80				
1. age, birth year	.485 (.012)	.345 (.011)	.472 (.011)	.330 (.010)
2. Plus core	.264 (.009)	.168 (.009)	.261 (.009)	.162 (.008)
3. Plus spouse empl. Status	.254 (.009)	.155 (.009)	.252 (.009)	.152 (.009)
Sample size	20,423			

Note: Each entry is the coefficient on the indicated employment variable from a different regression, controlling for the variables listed in the table stub. Full sets of age and birth year dummies are included. "Core" includes region, residence in a Standard Metropolitan Statistical Area, residence in an urban area, family composition, and the spouse's age and education. In the lower panel, core also includes the respondent's sex, education, race, and marital status. Standard errors are in parentheses. Source: Consumer Expenditure Survey, 1980-2000.

Table 2: Cohort Analysis of Employment Status and Consumption Expenditure

	Ln(Total Consumption Expenditure)		Ln(Non-durable Consumption)	
Regressor of interest:	Employed	Employed YRFT	Employed	Employed YRFT
Sample: Men born 1921-1952, ages 51-72				
	.605 (.057)	.714 (.063)	.572 (.056)	.684 (.061)
Sample: Married white non-Hispanic men, educ=12, born 1921-1952, ages 51-72				
	.560 (.244)	.740 (.234)	.349 (.190)	.535 (.178)

Notes: The observations are one year differences in cohort means. Sample size is 67 in the upper panel and 21 in the lower panel. The models also contain an intercept, and in the upper panel, cohort dummies and a linear age term. Standard errors are in parentheses. Source: Consumer Expenditure Survey, 1980-2000.

Table 3: Baseline Parameter Values

Parameter	Description	Value
A. Parameters calibrated to HRS data and held fixed		
\bar{R}	Mean asset return	1.02
λ	layoff rate	.03
β, σ_{η}^2	Log wage mean and variance	2.833, .3195
$o_0, o_1, \sigma_{\zeta}^2$	Log other income (intercept, age, variance)	2.697, -.024, .496
$m_0, m_1, m_2, \sigma_{\zeta}^2$	Log medical expenditure (intercept, age, bad health, variance)	-1.305, .0196, .2773, 1.1677
$\pi_{ijk}: 0 \text{ to } 1$	Health transition parameters: (intercept, slope): $\pi_{ij} = \exp\{\pi_{ij0} + \pi_{ij1}t\} / (\exp\{\pi_{i00} + \pi_{i01}t\} + \exp\{\pi_{i10} + \pi_{i11}t\} + \exp\{\pi_{i20} + \pi_{i21}t\})$, $i = 0, 1$; $j = 0, 1, 2$; the parameters for 0 to 0 and 1 to 1 are normalized to 0.	-3.3, .03
$\pi_{ijk}: 0 \text{ to } 2$		-5.8, .08
$\pi_{ijk}: 1 \text{ to } 0$		-1.4, -.025
$\pi_{ijk}: 1 \text{ to } 2$		-4.0, .07
B. Parameters determined arbitrarily and held fixed		
φ_0, φ_1	job offer rate: $\varphi_t = \exp\{\varphi_0 + \varphi_1 t\} / (1 + \exp\{\varphi_0 + \varphi_1 t\})$	1.449, -0.1575
χ	Job offer rate from other employers (if still on the pension job)	.30
σ_0^2	variance of asset return	.0016
α	Coefficient of relative risk aversion	2
δ	Rate of time discount	1/(1.03)
\bar{C}	Consumption floor	5.0
\bar{A}	Asset floor	.01
C. Parameters used to match employment, assets, and other patterns		
γ_{01}, γ_{11}	Intercept of utility of work (good health, bad health)	-.012, -.030
γ_{02}, γ_{12}	Age slope of utility of work (good health, bad health)	0, 0
γ_{03}, γ_{13}	Utility of entering employment (good health, bad health)	-.04, -.06
σ_{ϵ}^2	Variance of utility shocks	0.0002
κ_0, κ_1	bequest parameter: $\kappa_t = \kappa_0 + \kappa_1 t$	0.56, -0.0107
ψ	rate of acceptance of SSDI applications	.30

Note: All disturbance correlations are equal to zero. Monetary amounts are in thousands of dollars.

Table 4: Robustness Analysis

Specification	Change in Ln(consumption) at exit from employment	
	All	“Retirement”
Baseline	-.053	-.039
Nonhousing assets	-.090	-.055
Full time and part time	-.050	-.027
CRRA (α) = 0	-1.017	-1.058
CRRA = 1	-.061	-.051
CRRA = 3	-.059	-.059
CRRA = 4	-.038	-.046
Time preference ($1/\delta$)= 1.00	-.058	-.052
Time preference = 1.06	-.055	-.043
Time preference = 1.09	-.057	-.045
Variance of return shock $\sigma_\theta^2 = 0$	-.046	-.034
$\sigma_\theta^2 = .0032$	-.061	-.041
Variance of wage shock $\sigma_\eta^2 = 0$	-.032	-.022
$\sigma_\eta^2 = .639$	-.065	-.046
Job offer rate $\varphi_t = 1$	-.040	-.030
$\varphi_t = 0$	-.115	-.068
Consumption floor $\bar{C} = 1$	-.059	-.047

Notes: The baseline specification is the one shown in Table 3: CRRA (α) = 2; rate of time preference ($1/\delta$) = 1.03; variance of the return shock (σ_θ^2) = .0016; variance of the log wage shock $\sigma_\eta^2 = .3195$; job offer parameters $\varphi_0 = 1.449$, $\varphi_1 = -.1575$, with $\varphi_t = \exp\{\varphi_0 + \varphi_1 t\} / (1 + \exp\{\varphi_0 + \varphi_1 t\})$, and $C = 5$. The sample in the second column contains simulated individuals who worked for at least three consecutive periods, left employment, and remained non-employed for at least three consecutive periods as of age 67. In the specification with full time and part time employment, the cases shown are for movement from full time to not employed.

Figure 1: Employment rate by age




Figure 2: Two year entry rate to employment by age

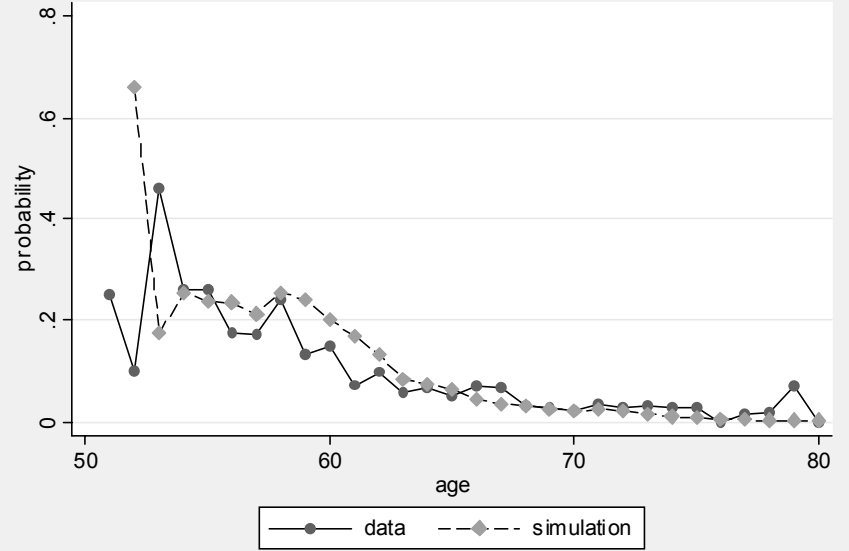


Figure 3: Two year exit rate from employment by age




Figure 4: Cumulative distribution of age of OASI entitlement




Figure 5: Mean OASI benefit by age




Figure 6: Enrollment rate in SSDI by age




Figure 7: Mean SSDI benefit by age




Figure 8: Cumulative distribution of age of pension entitlement




Figure 9: Mean pension benefit by age




Figure 10: Median assets by age




Figure 11: Mean total household income by age




Figure 12: Mean asset income by age




Figure 13: Mean consumption by age




Figure 14: Consumption and net after tax income by employment

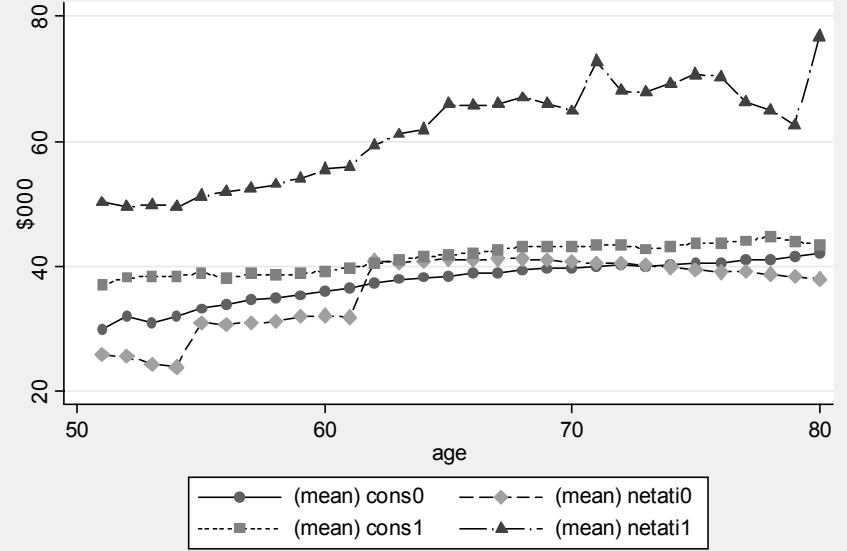


Figure 15: Change in log consumption by change in empl.

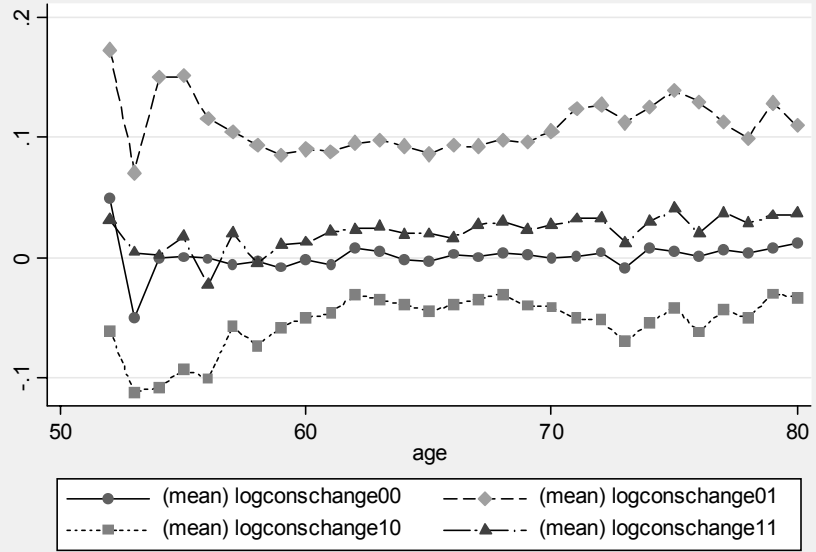


Figure 16: Consumption by employment sequence




Figure 17: Consumption paths for alternative specifications

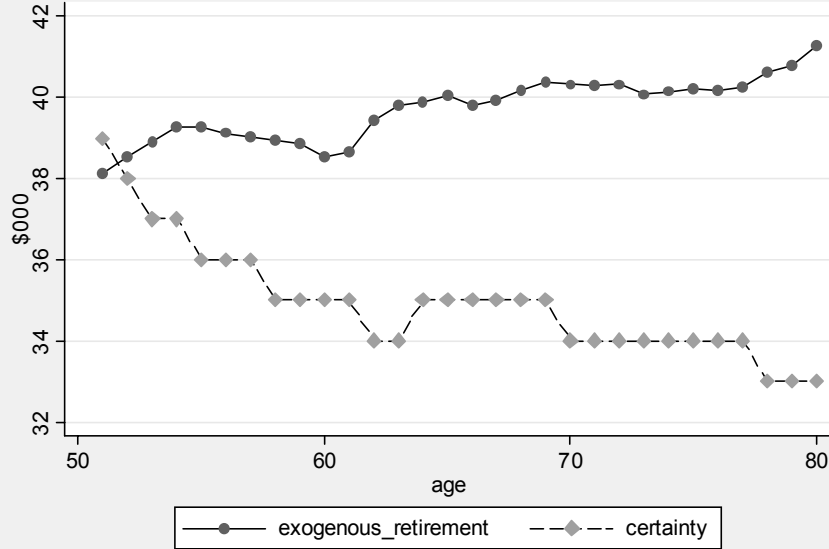


Figure 18: Change in Ln(C) at retirement by EG




Figure 19: Change in Ln(C) at retirement by assets

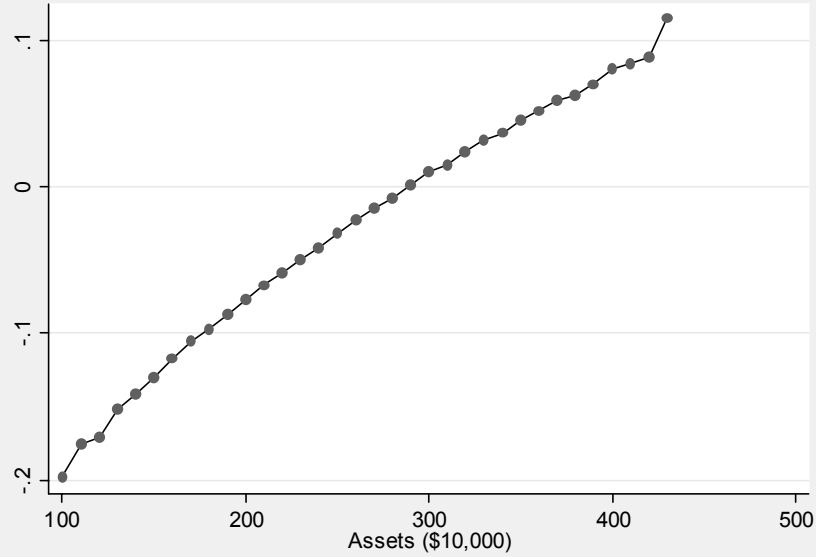


Figure 20: Median non-housing assets by age

