Preliminary
NA for Quotation

ESTIMATING LIFE-CYCLE PARAMETERS FROM THE RETIREMENT-CONSUMPTION PUZZLE

John Laitner and Dan Silverman

Department of Economics

The University of Michigan

October 18, 2004

Estimating Life–Cycle Parameters from the Retirement-Consumption Puzzle*

John Laitner and Dan Silverman
University of Michigan

Abstract

Using pseudo-panel data, we estimate the structural parameters of a life–cycle consumption model with discrete labor supply choice. A focus of our analysis is the abrupt drop in consumption upon retirement for a typical household. The literature sometimes refers to the drop, which in the U.S. Consumer Expenditure Survey we estimate to be approximately 14%, as the "retirement–consumption puzzle." Although a downward step in consumption at retirement contradicts predictions from life–cycle models with additively separable consumption and leisure, or with continuous work-hour options, a consumption jump is consistent with a setup having nonseparable preferences over consumption and leisure and requiring discrete work choices. This paper specifies a life–cycle model with these two elements, and uses the empirical magnitude of the drop in consumption at retirement to identify structural parameters — most importantly, the intertemporal elasticity of substitution. Employing our new estimates, we simulate counterfactual Social Security policy reforms.

1. Introduction

The life–cycle saving model is economists' workhorse for analyzing public policy. Papers use the framework to study a variety of social insurance programs, national debt, trade policy, and tax policy in general (e.g., Modigliani [1986], Diamond [1965], Tobin [1967], Auerbach and Kotlikoff [1987], Hubbard *et al.* [1995], Altig *et al.* [2001], and many others). The central role that the life–cycle model plays in the study of public policy generates a natural interest in estimates of its parameters and in assessments of its validity: to the extent that evaluations of policies depend on the parameters of the life-cycle model, or on the model's ability to match important features of the data, credible estimates of those parameters, and tests of the model's predictive ability, are important. This paper reconciles the life-cycle model with a now well–known feature of expenditure data, the discrete drop in average expenditure at retirement, the so–called "retirement-consumption puzzle;" the paper then employs this prominent feature of expenditure data to provide a novel and robust source of identification for parameters of a life–cycle model; and, finally, using estimated parameters identified from the drop in consumption at retirement, we evaluate the prospective impact of possible changes in the U.S. Social Security system on retirement behavior.

* We are grateful to John Bound, Gary Solon, and participants in the University of Michigan Economics of Aging Seminar for helpful comments.

A number of recent papers document a significant drop in household expenditures at retirement that some have suggested raises doubts about the validity of the life–cycle theory.[1] This decline in expenditures is sometimes viewed as a puzzle because it is inconsistent with some specifications of the model. It is inconsistent because a central prediction of the life–cycle model is that agents should smooth marginal utility across time, and, in the most basic version of model, this prediction maps into an optimal path of consumption which is itself smooth — with households saving in youth and middle age, in preparation for retirement, and dissaving in old age. When, for example, retirement is anticipated, and when consumption and leisure are additively separable, consumption should change continuously with time according to the life–cycle model.

By itself, however, a discrete drop in consumption at retirement does not necessarily challenge the validity of life–cycle analysis. There is no *a priori* reason to think that preferences over consumption and leisure should take the particularly simple, additively separable form. Moreover, a number of economists have independently argued that opportunities for market–work hours are not continuous. We show formally below that a tractable life–cycle specification with intratemporal utility that is nonseparable in consumption and leisure, and with work options that are discrete, forecasts a jump in consumption at retirement. Our view that consumption–puzzle outcomes are consistent with life–cycle planning finds support from Hurd and Rohwedder [2003] who use novel questions on expectations from a subsample of Health and Retirement Study participants to show that households anticipate that their consumption will fall about 20% at retirement. Similarly, Laitner [2001a] points out that financial advisors have long said that retirees should plan for less consumption than working people: he quotes a TIAA–CREF brochure stating that "you'll need 60 to 90 percent of current income in retirement, adjusted for inflation, to maintain the lifestyle you now lead" and a popular press article writing that "many financial planners say it will take 70 to 80 percent of your current income to maintain your standard of living when you retire." These sources suggest that households anticipate a drop in consumption at retirement; so, such a drop need not be indicative of defective planning or news about retirement resources.

This paper shows that if one is willing to treat the systematic drop in household consumption at retirement as a consequence of purposeful behavior, the magnitude of the decline in consumption upon quitting work provides useful information for estimating life–cycle–model parameters. Consider the elasticity of intertemporal substitution in consumption. Suppose that consumption and leisure are complements (i.e., that we have intratemporal nonseparability) and that a household anticipates a discrete increase in leisure upon retirement. Then the household will prefer either to increase or decrease consumption discretely at retirement depending on its taste for smoothing utility over time, i.e., on its intertemporal elasticity of substitution. If a household's taste for intertemporal smoothing is sufficiently high, it will choose to decrease its consumption at retirement so that lost utility from consumption offsets the gains from additional leisure. A household with a low desire for intertemporal smoothing, on the other hand, might increase its consumption at retirement to take advantage of the complementarity of consumption and leisure. The

[1] E.g., Banks *et al.* [1998], Bernheim *et al.* [2001], Hurd and Rohwedder [2003], Haider and Stephens [2004], and Aguiar and Hurst [2004]. See the discussion below.

size, and sign, of the consumption change can pin down the elasticity of intertemporal substitution. In other words, one can use the magnitude of the decline in consumption at retirement to identify a parameter of great importance for policy analysis, which authors have otherwise sought to identify from changes in interest rates.

This paper implements the identification strategy described above using pseudo–panel data on expenditure from the U.S. Consumer Expenditure Survey 1984-2001 to estimate parameters from a life–cycle model with nonseparable consumption and leisure. As in previous studies, we find a substantial average drop in household expenditure at retirement. Embedding that average decline in a life–cycle framework, and using lifetime earnings profiles and retirement ages from Health and Retirement Survey panel data, we estimate parameters of the model and use those parameters to assess the impact of Social Security reforms on retirement decisions.

Related Literature

Our paper relates to three main literatures. The first studies the decline in consumption at retirement and the implications of this decline for the life–cycle model. Banks, Blundell, and Tanner [1998] use British data and find consumption declines at retirement of 22-35%. To explain the decline, Banks *et al.* argue that individuals may, in their pre-retirement planning, tend to overestimate their pension entitlements. Bernheim, Skinner, and Weinberg [2001] examine consumption declines using the Panel Study of Income Dynamics. Proxying for consumption with expenditure on food and housing rents, they estimate an average annual consumption drop of 14% in the first two years after retirement. Bernheim *et al.* suggest that this decline, particularly the symmetric decline in expenditure on food consumed inside and outside the home, is inconsistent with the life–cycle model. Although Haider and Stephens [2004] find that instrumenting for anticipated retirement with Health and Retirement Study expectations data reduces the size of the estimated decline in expenditure at retirement by approximately one-third, the drop remains both economically and statistically significant, and they argue that the decline thus remains, in large part, a puzzle for the standard life–cycle model. In contrast, Hurd and Rohwedder [2003], as noted above, show that households expect their consumption to decline about 20% at retirement. Moreover, they show that households that are still working actually anticipate, on average, a larger drop in consumption at retirement than that reported by those who have retired.[2]

Second, in investigating the effects of changes in the Social Security system on retirement behavior, our paper joins the large literature that employs the life–cycle framework to study public policy. In addition to the well–known work of Auerbach and Kotlikoff [1987], examples from this literature include, for instance, Hubbard *et al.* [1995] and Altig *et al.* [2001].

Third, this paper contributes to the literature aimed at estimating the structural parameters of life–cycle models, and in particular the intertemporal elasticity of substitution

[2] Aguiar and Hurst [2004] take a different tack: augmenting a standard model with home production, and using data on food expenditure, consumption, and time spent on food shopping and production, they argue that consumption remains constant while expenditure declines with retirement.

— e.g., Hall [1988], Campbell and Mankiw [1989], Attanasio and Weber [1993], and Barsky *et al.* [1997].

The organization of this paper is as follows. Section 2 considers the possible structure of simple life–cycle saving models that incorporate labor/leisure choices. Section 3 outlines our strategy for empirical analysis. Section 4 introduces data from the U.S. Consumer Expenditure Survey, using it to assess the shape of household life–cycle consumption profiles. Section 5 introduces data from the Health and Retirement Study and, in combination with regression results from Section 4, estimates the structural parameters of a life–cycle model with nonseparable consumption and leisure. Section 6 uses the new parameter estimates to simulate the impact of possible changes to the U.S. Social Security System on household decisions of when to retire.

2. Framework of Analysis

In this paper's model, every household chooses its saving and labor supply to maximize its utility subject to a lifetime budget constraint.[3] The analysis assumes that (1) the work day is indivisible, (2) private–pension formulas do not affect workers' retirement age, (3) the household utility function may be intratemporally nonseparable with respect consumption and leisure, and (4) perfect markets exist for annuities and health insurance. The following provides our rationale for assumptions (1)–(3).

<u>Assumption 1</u>: the work day is indivisible

Households in our analysis must either work full time or retire. While in practice employers do offer part–time jobs, the rate of pay is, on average, substantially lower than that for full–time work. Reasons for the wage penalty for part-time work include daily fixed costs of startup and shutdown, scheduling and coordination problems, employer concern for timely return on training investments, and the fixed–cost nature of some employee benefits (see, for example, Hurd [1996]). As Rust and Phelan [1997,p.786] write,

> The finding that most workers make discontinuous transitions from full–time work to not working, and the finding that the majority of the relatively small number of 'gradual retirees' reduce their annual hours of work by taking on a sequence of lower wage partial retirement 'bridge jobs' rather than gradually reducing hours of work at their full-time pre-retirement 'career job' suggests the existence of explicit or implicit constraints on the individual's choice of hours of work.

The indivisibility assumption is also consistent with the fact that U.S. data show little trend in male work hours or participation rates after 1940, except for a trend toward earlier

[3] The option to save (and dissave) contrasts to Stock and Wise [1990] and Rust and Phelan [1997], in which pre-retirement households live on their earnings, and post-retirement households live on social security and pension benefits. Note that Lumsdaine and Mitchell [1999], for example, point out that Health and Retirement Study respondents in 1992, all of whom are near retirement, have substantial median net worth beyond Social Security and pension benefits.

retirement from 1940-80 (e.g., Pencavel [1986], Blundell et al. [1999], and Burkhauser et al. [1999]).

Assumption 2: Pension formulas do not affect retirement ages

Although some analyses stress the importance of workers' private pension plans as determinants of retirement behavior (e.g., Gustman and Steinmeier [1986], Stock and Wise [1990]), we assume that a worker chooses an employer whose pension plan matches his requirements; thus, in this paper, private pensions form a part of private wealth accumulation and do not require separate attention. In fact, while arguments in favor of explicitly modeling the distinct features of private pensions seem most applicable to defined benefit plans, in practice defined contribution plans are rapidly becoming more important. Moreover, to the extent that defined benefit pension plans remain relevant, many of the features of these plans originated at union (i.e., worker) initiative or during collective bargaining, and thus reflect worker preferences.[4]

Assumption 3: Utility is intratemporally nonseparable in consumption and leisure

In our framework, consumption and leisure are intratemporally non-separable. Since non-separability is central to this paper's focus, we now examine its implications in detail. For the sake of simplicity, all of our analysis assumes household preference orderings are intertemporally separable.

Gustman and Steinmeier [1986], Anderson et al. [1999], and many others assume intratemporal separability of consumption and leisure. To see the restrictions imposed by this assumption, consider a specific example. A single–person household lives from $t = 0$ to $t = T$, choosing to retire at $t = R$. The household's time endowment at each age is 1; when the household works, its leisure falls to $\bar{\ell} \in (0, 1)$. According to assumption (1), indivisibilities force $\bar{\ell}$ to be a fixed parameter. The wage is w; the interest rate is r. A household's consumption c_t yields utility flow $u(c_t)$; its leisure ℓ_t yields utility flow $v(\ell_t)$; and, its assets (net worth) are a_t. The household's behavior solves

$$\max_{R, c_t, \ell_t} \int_0^T [u(c_t) + v(\ell_t)]\, dt \tag{1}$$

$$\text{subject to:} \quad \ell_t = \begin{cases} \bar{\ell}, & \text{for } t < R \\ 1, & \text{for } t \geq R \end{cases}$$

$$\dot{a}_t = r \cdot a_t + (1 - \ell_t) \cdot w - c_t \,,$$

$$a_0 = 0 = a_T \,.$$

As utility depends on the sum $u(c) + v(\ell)$, this is the separable case.

Provided $u(.)$ is concave, specification (1) predicts that, despite the discrete increase in leisure at retirement, there should be no abrupt change in consumption at that age.

[4] Note that this paper's strategy of identifying the parameters of a life-cycle model from post–retirement consumption ultimately does not depend on our treatment of private pensions.

To see this, note that along an optimal consumption path, the additional utility at date s from one extra dollar's consumption, $u'(c_s)$, must equal the additional utility from the dollar if it were saved until later date t, by which time it will have grown to an amount $e^{r\cdot(t-s)}$. That is,

$$u'(c_s) = e^{r\cdot(t-s)} \cdot u'(c_t)\,, \tag{2}$$

the so–called "Euler equation" for consumption. Letting c_{t-} be consumption the instant before t, and c_{t+} the instant after, condition (2) yields

$$u'(c_{t-}) = u'(c_{t+})\,. \tag{3}$$

With $u'(.)$ continuous, equation (3) is inconsistent with a jump in consumption at any t, including at $t = R$.

Other authors assume intratemporally nonseparable preferences. A well–known example is Auerbach and Kotlikoff [1987].[5] We can easily modify example (1) to accommodate their alternative. Let a household have a constant returns to scale "neoclassical" production function $f : R^2 \mapsto R^1$ that combines current consumption and leisure to generate a flow of services, the latter yielding a flow of utility, say, $u(f)$. The household then solves

$$\max_{R, c_t, \ell_t} \int_0^T u\big(f(c_t, \ell_t)\big)\, dt \tag{4}$$

$$\text{subject to:} \quad \ell_t = \begin{cases} \bar{\ell}, & \text{for } t < R \\ 1, & \text{for } t \geq R \end{cases}$$

$$\dot{a}_t = r \cdot a_t + (1 - \ell_t) \cdot w - c_t\,,$$

$$a_0 = 0 = a_T\,.$$

As before, post–retirement leisure is 1 and pre-retirement leisure is $\bar{\ell} < 1$. Since a bivariate constant–returns–to–scale neoclassical production function has $f_{12}(.) > 0$, inputs are complementary in the sense that more leisure (consumption) raises the marginal product of consumption (leisure). If $u(.)$ were linear, this complementarity would make the household want to increase its consumption at retirement: consumption would jump up after the discrete rise in leisure at retirement because the marginal product of consumption increases with leisure. If, on the other hand, $u(.)$ is sufficiently concave, the household would desire a level flow of consumption services at different ages. Since the household produces such services abundantly from leisure during retirement, it would choose more consumption prior to $t = R$ to produce a relatively even annual flow of utility. Hence, rational behavior may lead to an age profile of consumption which discontinuously changes in either direction at retirement.

A specific parameterization with nonseparable consumption and leisure can illustrate the connection between the drop in consumption at retirement and structural parameters

[5] See also King, et al. [1988], Hurd and Rohwedder [2003].

of the life–cycle model. Let the intratemporal household production function $f(.)$ be Cobb–Douglas:

$$f(c,\ell) = [c]^{\alpha} \cdot [\ell]^{1-\alpha}, \quad \alpha \in (0,1)\,.$$

Let $u(.)$ have the familiar isoelastic form

$$u(f) = \frac{[f]^{\gamma}}{\gamma}, \quad \gamma < 1\,.$$

Recalling that ℓ changes from $\bar{\ell}$ to 1 at retirement, condition (3) at retirement date $t = R$ is

$$\begin{aligned}
&([c_{R-}]^{\alpha} \cdot [\bar{\ell}]^{1-\alpha})^{\gamma-1} \cdot \alpha \cdot [c_{R-}]^{\alpha-1} \cdot [\bar{\ell}]^{1-\alpha} = ([c_{R+}]^{\alpha})^{\gamma-1} \cdot \alpha \cdot [c_{R+}]^{\alpha-1} \iff \\
&\quad [c_{R-}]^{\alpha\cdot\gamma-1} \cdot [\bar{\ell}]^{(1-\alpha)\cdot\gamma} = [c_{R+}]^{\alpha\cdot\gamma-1} \iff \\
&\quad [c_{R-}] \cdot [\bar{\ell}]^{-\frac{\gamma\cdot(1-\alpha)}{1-\alpha\cdot\gamma}} = [c_{R+}]\,. \qquad (5)
\end{aligned}$$

Equation (5) shows that given the technology for household production summarized by α, the change in consumption at retirement identifies the intertemporal elasticity of substitution parameter γ. It is straightforward to show that the change in consumption at retirement is strictly increasing in γ. In particular, if γ is nearly 1, $u(.)$ is nearly linear, and consumption jumps up at retirement. More generally,

$$\begin{aligned}
c_{R-} < c_{R+} \quad \text{if} \quad & [\bar{\ell}]^{-\frac{\gamma\cdot(1-\alpha)}{1-\alpha\cdot\gamma}} > 1 \iff \\
& -\frac{\gamma\cdot(1-\alpha)}{1-\alpha\cdot\gamma} < 0 \iff \\
& \gamma > 0\,.
\end{aligned}$$

So whenever $0 < \gamma$, consumption discontinuously rises at retirement. On the other hand, the preceding algebra indicates that

$$c_{R-} > c_{R+} \iff \gamma < 0\,.$$

Thus, whenever $\gamma < 0$, the model predicts a discontinuous drop in consumption at retirement and, because $(c_{R+} - c_{R-})$ is increasing in γ, the decline in consumption at retirement grows larger as the taste for intertemporal smoothing increases (γ grows more negative).

To summarize, the preceding analysis illustrates that an abrupt adjustment in consumption at retirement is consistent with rational behavior. Equation (5) also shows that data on the size of change in consumption can help us to estimate the taste for intertemporal substitution. We now turn to the latter task.

3. Empirical Model

This section presents the model upon which we base our empirical analysis. We incorporate several new features into the framework above: to facilitate comparisons with data, we allow changes in household size and composition that may influence households' preferences for consumption; we also consider the possible role of liquidity constraints. Then we outline our estimation strategy.

Specification. In practice households gain and lose members over their life spans, and households presumably desire greater consumption at ages when their membership is larger. Let the number of "equivalent adults" per household be n_t. Let a household's head constitute one "equivalent adult." For a married household, let the spouse constitute ξ^S additional equivalent adults. Although ξ^S might be 1, it could also be substantially less if, for example, there are scale economies to household size or if preferences for consumption depend on marital status. If at age t the household head has a spouse, set $n_t^S = 1$; otherwise, set $n_t^S = 0$. Similarly, let n_t^C be the number of children in a household when the head's age is t, and ξ^C be the adult equivalency of each child.[6] Set

$$n_t = 1 + \xi^S \cdot n_t^S + \xi^C \cdot n_t^C \, . \tag{6}$$

In a model without a labor/leisure choice, Tobin [1967] suggests a utility–flow model

$$n_t \cdot u(\frac{c_t}{n_t}) \, .$$

The idea is that a single–member household with consumption c^1 and the same household at a different age with n equivalent adults and consumption c^n achieve the same per capita current utility flow when $c^n = n \cdot c^1$, and that a household weights $u(.)$ by n because the household values the per capita utility flows of all members equally. Based on the same reasoning, this paper's utility–flow specification is

$$n_t \cdot u\big(f(\frac{c_t}{n_t}, \ell_t)\big) \, . \tag{7}$$

Our idea is that equivalent adult scale (6) allows us to convert a household's total consumption to per capita units. Children do not work in our model, and we do not attempt to assess the value of their leisure. Think of the household head as allocating current consumption expenditures to household members according to their relative "needs," as measured by the equivalent adult scale. If the head's leisure is ℓ_t, his/her utility flow is $u\big(f(c_t/n_t, \ell_t)\big)$. The head assumes that the equivalent–adult–based consumption allocation equates other member's utility flow to his/her own; thus, the head employs (7) as the household's criterion. In the case of a married couple, we assume that the man works in the market until retirement, and the woman works in the market or at home and retires at the same date as her husband. With either the man or his wife serving as the household's "head," expression (7) remains valid.

[6] We experiment below with more general forms that allow different weights on children of different ages.

Traditionally financial markets have been wary about extending loans without collateral, although credit–card debt may have mitigated that constraint in recent years. In addition, *inter vivos* gifts and other family transfers my lift or alleviate liquidity constraints for many households. If we, nevertheless, add a liquidity constraint $a_t \geq 0$ for every age t to our life–cycle model to capture collateral requirements, our consumption Euler equation only holds for ages at which the constraint is not binding.[7] At an age when the constraint binds, household net worth will be zero and household consumption will equal earnings, and hence grow as fast as earnings, and faster than the consumption Euler equation would dictate. Since life–cycle households build net worth to help finance their retirement, liquidity constraints are most likely to bind, if at all, at the beginning of life. In fact, our regression outcomes below provide some evidence of constrained behavior at young ages. Thus, as described below, we adapt our estimation strategy to account for the possibility of liquidity constraints early in life.

The life–cycle maximization model upon which we base our empirical analysis is then as follows: for household i,

$$\max_{R_i, c_{it}, \ell_{it}} \int_0^T e^{-\rho \cdot t} \cdot n_{it} \cdot u\big(f(\frac{c_{it}}{n_{it}}, \ell_{it})\big)\, dt \tag{8}$$

$$\text{subject to:} \quad \ell_{it} = \begin{cases} \bar{\ell}, & \text{for } t < R_i \\ 1, & \text{for } t \geq R_i \end{cases}$$

$$\dot{a}_{it} = r \cdot a_{it} + (1 - \ell_{it}) \cdot e_{it} \cdot w \cdot (1 - \tau - \tau^{ss}) + ssb_{it} \cdot (1 - \tau/2) - c_{it} \,,$$

$$a_t \geq 0 \quad \text{all} \quad t \,,$$

$$a_{i0} = 0 = a_{iT} \,,$$

where ρ is the subjective discount rate, and equivalent adults, n_{it}, come from (6). Household i supplies e_{it} "effective hours" in the labor market per hour of work time; hence, if w is the economy wide average wage rate, the household's earnings are $e_{is} \cdot w$ per hour of market work at age s. The latter include both the earnings of the head, and those of the spouse, if present and working in the labor market. We assume a proportional income tax τ on earnings, interest, and one half of Social Security benefits, ssb_{it}. The real interest rate r is already given in net–of–tax terms. There is also a proportional Social Security tax τ^{ss}.[8] Our simulations derive Social Security benefits, ssb_{it}, from the year-2000 U.S. formula. As before, we take $f(.)$ to be Cobb–Douglas and $u(.)$ to be isoelastic.

---

[7] Existing analyses of potentially binding liquidity constraints include Mariger [1987], Zeldes [1989], and Hubbard and Judd [1986]. See especially the discussion in Mariger.

[8] Our simulations below assume the Social Security tax applies only up to the statutory maximum.

We assume a constant real interest rate of 4% per year. Since we set $\tau = .25$, we have $r = .03$. Adults work 40 hours per week until retirement and 0 hours per week after retirement. With 16×7 waking hours per week, we set[9]

$$\bar{\ell} = \frac{16 \times 7 - 40}{16 \times 7} = .6429\,.$$

Identification. We estimate two equations based on model (8). The first is the Euler equation from utility maximization with respect to consumption; the second stems from utility maximization with respect to retirement age. The parameters of interest form a vector

$$\theta \equiv (\alpha, \gamma, \rho, \xi^S, \xi^C)\,.$$

Consider first the Euler equation for consumption. Suppose the liquidity constraint binds only for $s < S$. Then the Euler equation yields

$$c_{is} = \begin{cases} c_{iS} \cdot e^{\frac{r-\rho}{1-\alpha\cdot\gamma}\cdot s} \cdot n_{is}\,, & \text{if } S \leq s < R_i, \\ c_{iS} \cdot e^{\frac{r-\rho}{1-\alpha\cdot\gamma}\cdot s} \cdot n_{is} \cdot [\bar{\ell}]^{\frac{-\gamma\cdot(1-\alpha)}{1-\alpha\cdot\gamma}}\,, & \text{if } s \geq R_i. \end{cases} \tag{9}$$

The first right–hand side term, c_{iS}, sets the level of the consumption path for household i. Determinants of this level include the household's earning ability, its taste for leisure, and the number of its children. The remaining right-hand side terms in (9) reflect relative changes in consumption with respect to age, arrival and maturity of children, marriage and death of spouse, and retirement. On the basis of (6),

$$\ln(n_{is}) \approx \xi^S \cdot n_{is}^S + \xi^C \cdot n_{is}^C\,.$$

Our empirical specification takes logarithms of both sides of (9) and appends a term ν_{is} to the right side to reflect measurement errors in $\ln(c_{is})$. Writing (9) at $s \geq S$ and again at $s+1$, and differencing the logs, we have

$$\ln(c_{i,s+1}) - \ln(c_{is}) = \frac{r-\rho}{1-\alpha\cdot\gamma} + \xi^S \cdot [n_{i,s+1}^S - n_{is}^S] + \xi^C \cdot [n_{i,s+1}^C - n_{is}^C] - \frac{\gamma\cdot(1-\alpha)}{1-\alpha\cdot\gamma} \cdot \ln(\bar{\ell}) \cdot [\chi_{s+1}(R_i) - \chi_s(R_i)] + \nu_{i,s+1} - \nu_{is}\,, \tag{10}$$

where we use the indicator function

$$\chi_s(R) \equiv \begin{cases} 0, & \text{if } s < R, \\ 1, & \text{if } s \geq R. \end{cases}$$

Differencing eliminates the c_{iS} terms of (9) from (10).

[9] See also Cooley and Prescott [1995] — who, on the basis of time–use studies, determine that households devote 1/3 of waking hours to work.

Below we estimate (10) from a regression. A household's retirement age depends upon its number of children, earning power, etc.; thus, R_i is almost certainly jointly determined with c_{iS}. Nevertheless, we have just noted that differencing removes the c_{iS} term from (10). So, there is no reason to expect R_i to be related to ν_{is}. Hence, our regressors are independent of our error term.

On the other hand, the coefficients from difference equation (10) are composites of the structural parameters of the underlying model, and regression coefficients from (10) alone cannot separately identify α, γ, and ρ. Previous studies (e.g. Hall [1988], Banks et al. [1998]) have typically used changes in expected interest rates to extract the intertemporal elasticity of substitution parameter γ. We turn instead to a second equation based on utility-maximizing retirement ages. Our regression of (10) estimates

$$\frac{r-\rho}{1-\alpha\cdot\gamma}, \quad \xi^S, \quad \xi^C, \quad -\frac{\gamma\cdot(1-\alpha)\cdot\ln(\bar{\ell})}{1-\alpha\cdot\gamma} \equiv -\xi^R. \tag{11}$$

Letting $x \equiv -\xi^R/\ln(\bar{\ell})$, the fourth term yields

$$x = \frac{\gamma-\alpha\cdot\gamma}{1-\alpha\cdot\gamma} \Longleftrightarrow \gamma = \frac{x}{(x-1)\cdot\alpha+1}\,. \tag{12}$$

If our second equation can identify α, (12) will yield γ — and the first term in (11) will yield ρ. Fortunately, retirement age is related to α. For example, if $\alpha = 0$, leisure but not consumption yields utility; so, households would choose never to work at all. If $\alpha = 1$, consumption but not leisure yields utility; so, households would choose never to retire.

More precisely, our second estimating equation works as follows. For any θ, we can solve model (8) for a household's retirement age. Denote with $g(\theta)$ the optimal retirement age. Given our functional forms, even if households have different wage rates, reflecting different inherent earning abilities, their desired retirement ages are the same. Nevertheless, retirement ages could differ in practice because households may anticipate differences in longevity due to heredity; households may have differing numbers of children; occupations differ in their physical and emotional stress, leading to occupational differences in e_{it} at advanced ages; households have different health status; households may have differing tastes for leisure; etc. To model these differences in preferences and endowments, we can modify the lifetime utility function of (8) to

$$\int_0^T e^{-\rho\cdot t}\cdot n_{it}\cdot u\big(f(\frac{c_{it}}{n_{it}}, \ell_{it})\big)\,dt + \mu_i\cdot\int_{R_i}^T dt$$

where μ_i is positive (negative) if household i has an unusually strong (weak) preference for leisure. The optimal retirement age is then given by

$$R_i = g\big(\theta(\alpha)\big) + \eta_i\,, \tag{13}$$

where R_i is the household's actual retirement age; η_i is a random error reflecting deviations from average in taste for leisure and health status; and, given (11)-(12), we can write $\theta = \theta(\alpha)$. We use (13) together with (10) to estimate θ.

4. The Consumption Euler Equation

We have argued that the regressors in (10) are uncorrelated with the error. After transforming the data to diagonalize the covariance matrix, we therefore run an OLS linear regression to determine the parameter composites of (11). This section describes our data and results.

CXS Data. We have two primary data sources. One is the U.S. Consumer Expenditure Survey (CXS). This is the most comprehensive source of disaggregate consumption data for the U.S. Using it, we can chart variations in household consumption with age.

The CXS obtains diary information on small purchases from one set of households; for a second set of households, it conducts quarterly interviews that catalog major purchases. The survey also collects demographic data, data on current income, and data on value of the respondent's house. The sample is large. The survey was conducted at multi–year intervals prior to 1984, and annually thereafter. This paper uses the CXS surveys from 1984-2001.[10]

Tables 1-2 compare National Income and Product Account (NIPA) personal consumption for 1985, 1990, 1995, and 2000 with weighted totals from the CXS.[11] We omit household expenditures on pensions and life insurance from the CXS: the former constitute saving rather than consumption, and our concept of earnings is net of insurance. Looking at the last row of Table 2, total consumption measured in the CXS is only about 50-70 percent as large as the NIPA equivalent, with the discrepancy larger in later years.[12] This paper assumes that the NIPA numbers are accurate; that item–nonresponse on the survey makes the CXS totals too low; and, that the relative magnitude of survey errors does not systematically vary with age. Thus, for each year we scale CXS consumption categories, uniformly across ages, to match NIPA amounts. Appendix 1 describes our procedure for assembling the CXS data into categories for Table 2.

We make three special adjustments. First, we subdivide "shelter" into "services from own house" and "other." We scale the latter as with other categories, but we drop the CXS "services from own house," which consists of expenditures on mortgage payments, repairs, etc., and impute a substitute which divides the annual NIPA total service flow from residential houses in proportion to CXS house values by age.

Second, CXS medical expenditures are far too low (one important reason being that they omit employer contributions to health insurance). The Department of Health and Human Services provides a breakdown of U.S. health care spending.[13] For each year,

---

[10] The web site http://stats.bls.gov.gov/csxhome.htm presents aggregative tables, codebooks, etc., for the CXS. This paper used the raw data from the ICPSR archive, and we gratefully acknowledge the assistance of the BLS in providing stub files of the changing variable definitions.

[11] We abstract from the empirical difference between consumption and expenditure. See Aguiar and Hurst [2004] for evidence on the possible importance of this difference.

[12] The reader may notice a particularly large gap for "apparel" in 1985. The 1984 and 1985 data files omit a number of apparel subcategories. We assume this does not create biases with respect to age — so that our scaling procedure below eliminates the problem.

[13] See http://www.cms.hhs.gov/statistics/burden-of-health-care-costs/table01.asp. The

Table 1. National Income and Product Accounts Personal Consumption for 1985, 1990, 1995, and 2000 (billions of current dollars)[a]

Category	1985	1990	1995	2000
food[b]	498.5	677.9	802.5	1027.2
apparel	188.3	261.7	317.3	409.8
personal care	37.6	53.7	67.4	87.8
shelter:	406.8	585.6	740.8	960.0
own home	279.3	410.7	529.3	704.9
other	127.5	174.9	211.5	255.1
household operation	344.0	433.6	555.0	723.9
transportation	372.8	455.4	560.3	768.9
medical care	367.4	619.7	888.6	1171.1
recreation	187.6	284.9	401.6	564.7
education	53.8	83.7	114.5	164.0
personal business:	188.1	284.7	406.8	632.5
brokerage fees	130.9	192.1	284.4	449.3
other	57.2	92.6	122.4	183.2
religious & welfare	60.4	97.1	134.9	190.1
total[c]	2705.1	3837.8	4989.7	6699.8

a. Source: http://www.bea.doc.gov/bea/dn/nipaweb/AllTables.asp, Section 2, Table 2.4.
b. Includes tobacco and alcohol.
c. Omits foreign travel.

we normalize CXS expenditures on private health insurance to the HHS total (including employer and employee premiums), we do the same for premiums paid by individuals to the Medicare Supplementary Medical Insurance Trust Fund, and we do the same for out–of–pocket health spending.

Third, the NIPA "personal business" category includes bank and brokerage fees, many of which are hidden in the form of low interest on saving accounts, fees deducted from mutual fund income, etc., and hence absent from expenditures which CXS households perceive. This paper assumes that bank and brokerage fees make their way into the life–cycle model in the form of lower–than–otherwise interest rates on saving (see below); therefore, we normalize annual personal business expenditures measured in the CXS to match the corresponding NIPA amount less bank and brokerage fees, and omit bank and

annual figures cover 1987-2000. We extrapolate to 1984-86 and 2001 using the growth rate of NIPA total medical consumption.

Table 2. Consumer Expenditure Amount ÷ NIPA Amount (percent)[a]				
Category	1985	1990	1995	2000
food	73.5	69.6	64.9	62.3
apparel	22.0	60.0	55.4	49.5
personal care	73.7	65.8	61.7	70.2
shelter:	82.9	82.4	81.4	81.0
own home	74.1	69.7	73.0	71.4
other	102.1	112.2	102.5	107.6
household operation	76.0	82.6	78.6	71.4
transportation	111.7	109.0	110.7	105.4
medical care	27.6	23.2	20.1	19.3
recreation	61.8	55.5	50.9	45.0
education	65.1	61.2	58.8	57.4
personal business:	14.8	12.2	9.9	6.8
brokerage fees	0.0	0.0	0.0	0.0
other	48.5	37.5	33.0	23.4
miscellaneous	120.0	80.0	68.0	67.5
total	66.7	64.3	59.7	56.0

a. Source: see text.

brokerage fees from our measure of consumption.

For all other rows in Table 2, we scale, year by year, CXS expenditures across all ages in each given category to the corresponding NIPA amount.

Deflating with the NIPA personal consumption deflator, we derive an adjusted consumption amount, say, $\bar{c}_{it}$, for each age i and year t. Due to the construction of the CXS from separate interview and diary surveys, we do not have consumption figures for individual households; nevertheless, we have a synthetic panel of average household consumption for each age and year. (Panel analysis of individual households at yearly frequencies would be infeasible in any event because the CXS only follows each given household for four quarters.[14]) The number of interviewed households per cell that we use below varies from 127 to 981. The left–hand side variable for our Euler equation is

[14] Sabelhaus, see http://www.nber.org/data/ , derives CXS consumption for individual households by utilizing only the interview surveys. Sabelhaus, however, reduces his sample size by limiting attention to households which participate in the interview survey for four quarters. In contrast, our construction makes use of all data available in each quarter.

$$\Delta \ln(\bar{c}_{it}) \equiv \ln(\bar{c}_{i+1,t+1}) - \ln(\bar{c}_{it})\,.$$

We organize the data so that a household's age is the age of the wife in the case of a married couple, and the age of the single household head in all other cases.[15]

In addition to ages for a household's husband and wife, the CXS provides information on whether the household is married, the number of children of age 0-17 in the household, and whether the husband (or the single adult) currently participates in the labor market. The first provides our regressor n^S; the last provides $\chi(R)$. To obtain more flexibility we construct our own measure of children per household as follows. Using Census data on births per woman at age i, i =15,...,49, in year t, t = 1920,..., 2001, we simulate the number of children of age 0-17 and of age 18-22 for women of each age i in 1984,...,2001.[16] Notice that the natality surveys are not samples; they provide a complete census of births. As stated, we have reorganized our CXS cells so that "age" is age of the woman in all but single male households; thus, we append numbers of children to each cell on the basis of the ages and birth dates of women.

Regression Results. Tables 3-4 provide our regression results. We use households of age 20-80 for 1984-2001, so that our differences cover ages 20-79 and years 1984-2000. Each of our regressions includes separate time dummies for 1984, 1985, ... , 1999. As stated, we transform the data to diagonalize our covariance matrix. The other independent variables are a constant, presence of a spouse (n^S), and retirement status ($\chi(R)$); some specifications include number of minor children, aged 0-17, n^{MC}; and, all regressions include number of children 18-22, n^C.

Table 3 presents estimates of a specification with separate regressors for children 0-17 and 18-22. We consider household ages as low as 20. The estimates of these specifications seem unsatisfactory in that they are sensitive to the starting age. In particular, the constants ξ^S and ξ^{MC} change considerably from column to column. What is more, the coefficient on minor children steadily declines and is negative for starting ages of 30 or more. For the middle column of Table 3, Chart 1 plots average residuals for each age. The chart suggests that liquidity constraints may invalidate our Euler equation at early ages: Chart 1 shows a sequence of positive residuals through age 30, which would be consistent with binding liquidity constraints — implying unusually fast consumption growth.

[15] This differs from conventional BLS tables and the original conception of the survey: in the case of married couples, the CXS randomly chooses one parent as the household's "reference person," and the CXS household age is the age of that person. Notice also that in our data some variables are top coded.

[16] The natality statistics for 1920-40 come from
http://www.cdc.gov/nchs/products/pubs/pubd/vsus/1963/1963.htm,
Vital Statistical Rates In The U.S. 1900-1940. Statistics for 1940-1999 come from
http://www.cdc.gov/nchs/datawh/statab/unpubd/natality/natab99.htm,
Table 1-7: Total Fertility Rates and Birth Rates, By Age of Mother and Race: US, 1940-99. Data for 1999-2001 come from
http://www.cdc.gov/nchs/births.htm.

Table 3. Estimated Coefficients for Consumption Euler Equation:[a, b]
Elaborate Specification of Children's Consumption Role;
Consumer Expenditure Survey Data 1984-2000 (see text)

Para– meter	CXS Sample Ages				
	Age 20-79	Age 22-79	Age 24-79	Age 26-79	Age 28-79
Constant	.0257	.0238	.0231	.0220	.0197
(Std Err)	(.0010)	(.0009)	(.0010)	(.0011)	(.0013)
T-Stat	26.6341	25.6131	23.2779	20.0968	15.6608
Spouse	.7482	.5901	.5437	.5107	.4742
(Std Err)	(.0463)	(.0468)	(.0499)	(.0520)	(.0539)
T-Stat	16.1704	12.6131	10.8985	9.8189	8.8037
Retire	-.2565	-.2679	-.2637	-.2455	-.2058
(Std Err)	(.0327)	(.0310)	(.0313)	(.0321)	(.0340)
T-Stat	-7.8540	-8.6388	-8.4243	-7.6376	-6.0546
Child18-22	.2152	.2277	.2324	.2383	.2436
(Std Err)	(.0161)	(.0153)	(.0154)	(.0155)	(.0157)
T-Stat	13.3925	14.8901	15.1339	15.3618	15.5408
Child 0-17	.1036	.0847	.0741	.0563	.0247
(Std Err)	(.0104)	(.0102)	(.0110)	(.0126)	(.0151)
T-Stat	9.9687	8.3387	6.7604	4.4759	1.6398
Summary Statistics					
R^2	71.9370	66.1434	60.5917	55.3782	50.1876
Obs	1020	986	952	918	884
Error Mean Sq	.0120	.0108	.0108	.0109	.0110

a. Time dummies 1984, 1985, ... , 1999 not reported.
b. Note: coefficient of Child 0-17 is negative for starting ages 30 and above.

CHART 1. YEAR-AVERAGE RESIDUALS;
REGRESSION STARTS HOUSEHOLD AGE 24;
REGRESSORS INCLUDE CHILD 0-17

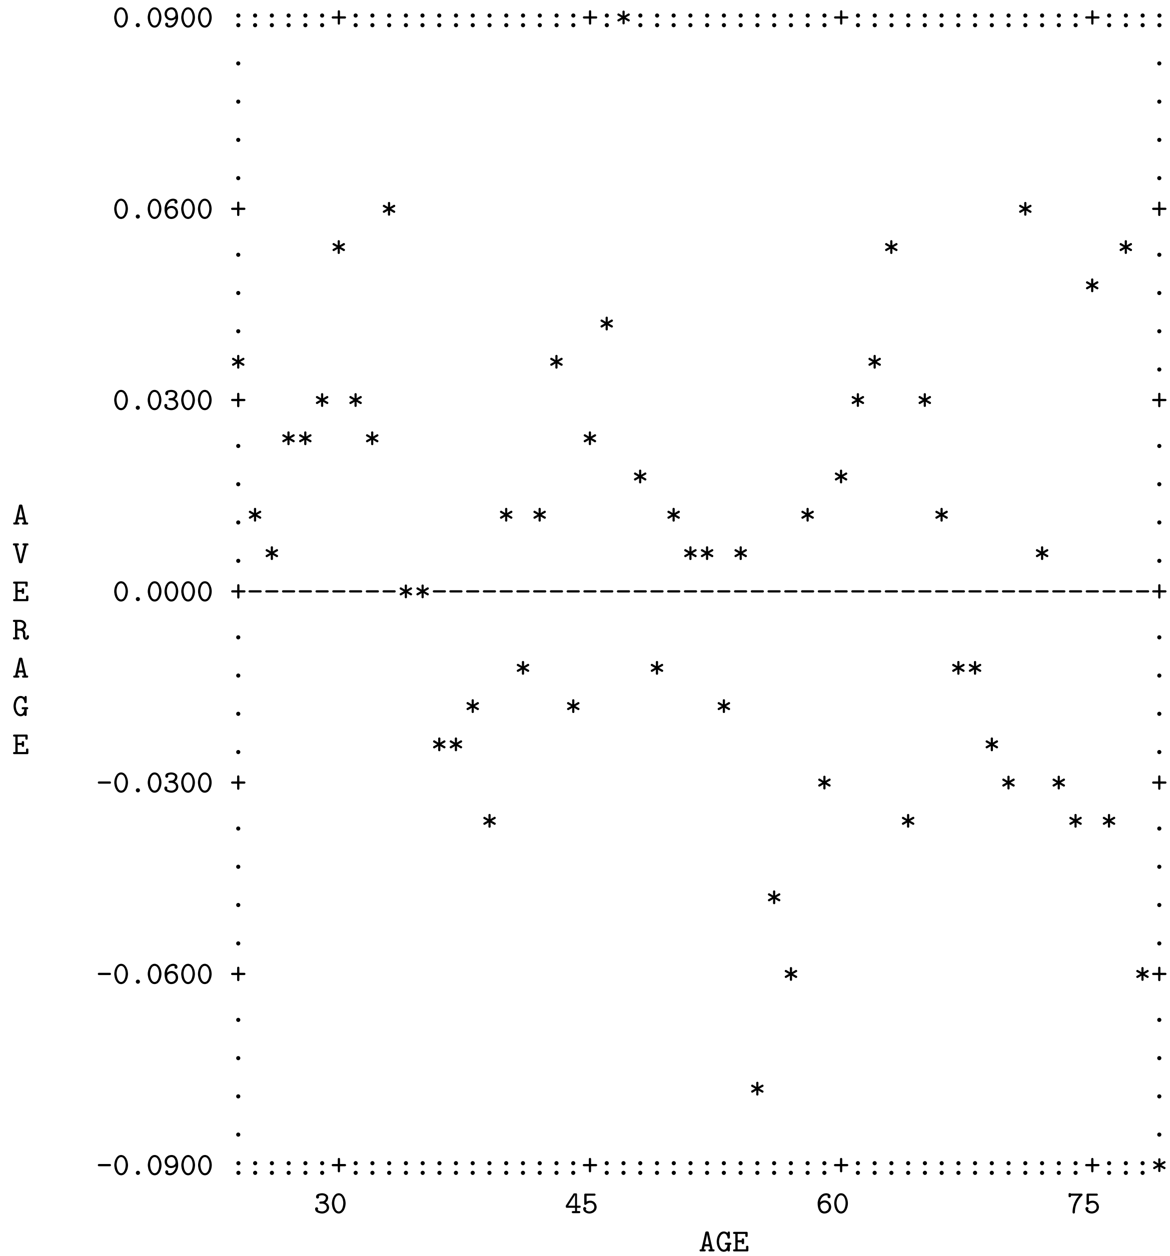

Tables 4a-b drop the regressor for minor children — in Table 3, the estimated coefficient of this regressor seems substantially smaller than that of the variable for children 18-22, and its sign is implausible for starting ages beyond 28.

Table 4a repeats the starting ages from Table 3. Results resemble Table 3: the starting age has a disquietingly large effect on the magnitude of the first two coefficients. Likewise, Chart 2 plots average residuals by age and shows a sequence of positive values through age 30. Again, results suggest that at early ages binding liquidity constraints are invalidating

Table 4a. Estimated Coefficients for Consumption Euler Equation:[a]
Simplified Specification of Children's Consumption Role;
Consumer Expenditure Survey Data 1984-2000 (see text)

Para– meter	CXS Sample Ages Age 20-79	Age 22-79	Age 24-79	Age 26-79	Age 28-79
Constant	.0234	.0217	.0208	.0197	.0185
(Std Err)	(.0010)	(.0009)	(.0010)	(.0010)	(.0010)
T-Stat	23.8357	23.1921	21.8035	20.1472	18.3091
Spouse	.9928	.7522	.6333	.5385	.4739
(Std Err)	(.0410)	(.0440)	(.0492)	(.0522)	(.0539)
T-Stat	24.2142	17.0834	12.8626	10.3196	8.7882
Retire	-.1295	-.1695	-.1826	-.1845	-.1777
(Std Err)	(.0315)	(.0297)	(.0296)	(.0294)	(.0294)
T-Stat	-4.1104	-5.7117	-6.1684	-6.2710	-6.0463
Child18-22	.2428	.2506	.2533	.2533	.2490
(Std Err)	(.0166)	(.0155)	(.0154)	(.0153)	(.0153)
T-Stat	14.6428	16.1493	16.4568	16.5586	16.2304
Summary Statistics					
R^2	69.1454	63.7038	58.6572	54.3816	50.0324
Obs	1020	986	952	918	884
Error Mean Sq	.0132	.0116	.0113	.0112	.0111

a. Time dummies 1984, 1985, ... , 1999 not reported.

Table 4b. Estimated Coefficients for Consumption Euler Equation:[a]
Simplified Specification of Children's Consumption Role;
Consumer Expenditure Survey Data 1984-2000 (see text)

Para–meter	CXS Sample Ages				
	Age 30-79	Age 32-79	Age 34-79	Age 36-79	Age 38-79
Constant	.0173	.0165	.0160	.0167	.0163
(Std Err)	(.0011)	(.0011)	(.0012)	(.0014)	(.0017)
T-Stat	16.4165	14.8512	13.2865	12.3042	9.7462
Spouse	.4175	.4061	.3893	.4091	.3944
(Std. Err)	(.0556)	(.0572)	(.0594)	(.0615)	(.0659)
T-Stat	7.5036	7.1022	6.5489	6.6566	5.9879
Retire	-.1704	-.1612	-.1594	-.1642	-.1628
(Std. Err)	(.0294)	(.0299)	(.0305)	(.0313)	(.0332)
T-Stat	-5.7921	-5.3981	-5.2348	-5.2461	-4.9035
Child18-22	.2388	.2278	.2210	.2380	.2301
(Std Err)	(.0156)	(.0164)	(.0183)	(.0222)	(.0286)
T-Stat	15.3293	13.8959	12.0019	10.5734	8.0361
Summary Statistics					
R^2	44.5853	39.1733	32.0862	27.0867	19.7622
Obs	850	816	782	748	714
Error Mean Sq	.0110	.0112	.0115	.0118	.0122

a. Time dummies 1984, 1985, ..., 1999 not reported.

CHART 2. YEAR-AVERAGE RESIDUALS
REGRESSION STARTS HOUSEHOLD AGE 24;
OMIT REGRESSOR CHILD 0-17

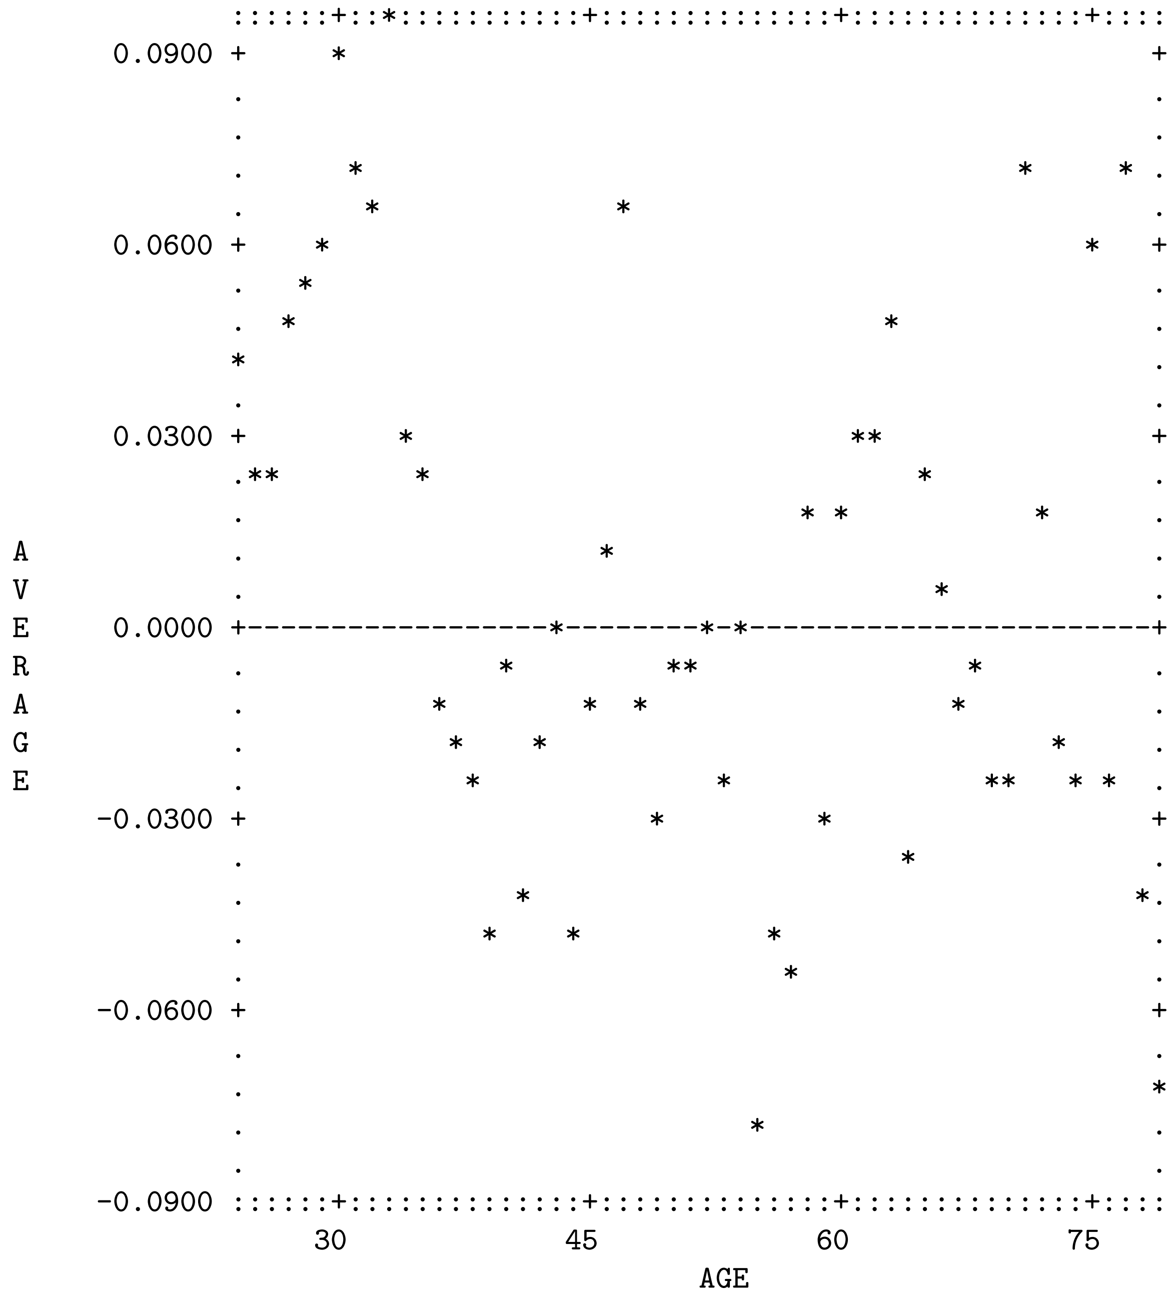

```
                CHART 3. YEAR-AVERAGE RESIDUALS
              REGRESSION STARTS HOUSEHOLD AGE 36;
                   OMIT REGRESSOR CHILD 0-17
             ::::::+::::*:::::::::+::::::::::::::+:::::::
      0.0750 +                                          +
             .                                        * .
             .                                  *       .
             .                                          .
             .                                      *   .
      0.0500 +                                          +
             .                          *               .
             .                        *                 .
             .                         *                .
             .      *  *                                .
      0.0250 +                     * *                  +
             .   *                        *             .
             .        *                                 .
A            .           * *   *                 *      .
V            .     *        **                          .
E     0.0000 +------------------------------------------+
R            . *              *            *            .
A            .*                                         .
G            *    *       *                 **          .
E            .                                    *     .
     -0.0250 +       *              *         **     *  +
             .  *                                       .
             .                                     *    .
             .                   *       *             *.
             .                                          .
     -0.0500 +                    *                     +
             .                                          .
             .                                          .
             .                                          .
             .                  *                       .
     -0.0750 ::::::+::::::::::::::+::::::::::::::+::::::*
                  42             57             72
                                 AGE
```

the Euler equation.

Table 4b considers starting ages 30-38. Here the coefficient estimates are quite stable. All have large T-statistics. Chart 3, which uses a sample with starting age 36, shows the typical outcome for residuals in Table 4b: the initial run of positive values evident in Charts 1-2 is absent. If liquidity constraints bind at early ages, estimates in Table 4a may be inconsistent; if constraints never bind, all of the estimates of Tables 4a-b are consistent, though Table 4b sacrifices some observations. To estimate the structural parameters of the life-cycle model, we use the subsamples of Table 4b.

5. Retirement Age

As described in Section 3, we want to use the utility–maximizing retirement age, R_i, from model (8) to estimate our underlying parameters. The Euler equation of the preceding section provides estimates of the composites in (11) but cannot separately identify our structural parameters. In terms of Section 3, we want to estimate alpha from (13); given α, we can deduce γ and ρ from the Euler–equation coefficients.

We implement our second stage as follows. Because we need lifetime earnings to solve (8) for an optimal retirement age, we turn to data from the Health and Retirement Study (HRS). For a large set of households that reach retirement around the year 2000, the HRS collects data on retirement status, net worth, and earnings.[17] A majority of the households have signed a permission waiver allowing the HRS to link to their Social Security Administration (SSA) earnings history. Each history runs 1951-1991; the HRS itself covers 1992, 1994, 1996, 1998, 2000, and 2002. In this paper, we use the HRS and linked SSA records to estimate a representative lifetime earnings profile for men who reach age 65 in 2002 and for women who reach age 65 in 2000. We then construct a representative household with a husband, wife, and two children. The wife reaches age 65 in the year 2000; the husband is 2 years older; the husband starts work at age 22 and marries at 24, with the latter being the time at which our household begins; both children are born two years after the marriage; and, both children leave home at age 22. The husband lives through age 74; the wife lives through age 80. All of our calculations use 2000 dollars, with our price index being the NIPA consumption deflator. As stated above, we assume an after–tax real interest rate of 3% per year, a proportional income tax (on earnings and interest) of 25%, and a social security payroll tax — half of which employees pay — of 13% per year. We assume that one–half of social security benefits are subject to the income tax. The social security benefit formula, including the ceiling on taxable annual earnings, follows the history of the U.S. system.

Fix any $\alpha \in (0, 1)$. Determine $\theta = \theta(\alpha)$ from the coefficient estimates of Table 4b. Inserting our empirical earning profiles, for any R we can find the optimal consumption profile for (8).[18] Since the problem conceivably is not concave in R, we simply try every possible month of life as a retirement date and pick the one yielding maximal lifetime

[17] We use the original HRS survey cohort, consisting of households in which the respondent is age 51-61 in 1992.

[18] Without liquidity constraints, the problem is standard. With such constraints, we use Mariger's [1987] algorithm.

utility. HRS data shows a median male retirement age of about 62. In this draft, we treat (13) as nonstochastic and numerically determine the α that solves

$$62 = g\big(\theta(\alpha)\big)\,. \tag{14}$$

Appendix 2 provides details on our derivation of earnings profiles. Charts 4-6 plot, in 2000 constant dollars, the (net–of–tax) profiles that we use for males, females, and households. The graphs assume retirement age 62, and Chart 6 includes Social Security benefits.

Results. Tables 5-6 present our results.

As explained above, we favor the Euler–equation estimates in Table 4b, which employ data only for ages 30 and greater. Solving model (8) with a nonnegativity constraint on household net worth, we find that the constraint does indeed bind in every case until age 28-30. Since the latter is the husband's age, the constraint binds until the wife is 30-32. This is generally consistent with our Euler equation regression outcomes, which seem to show evidence of binding liquidity constraints until roughly age 30. To be precise, however, note that our HRS households were in their twenties from 1955 to 1965, whereas households in our CXS sample are of all ages over the span 1984 to 2001. We cannot be sure that the cohorts in our CXS data, which are relatively recent, will have the same lifetime earning profile shapes as the HRS cohort — in fact, we have no way of observing what the lifetime earnings of younger households in the CXS will be.

The estimates of gamma in Tables 5-6 vary from -.45 to -1.05; however, Table 6b presents our preferred estimates, and they fall in the range from -.55 to -.60. The intertemporal elasticity of substitution (IES) for consumption/leisure service flows is $1/(1-\gamma)$; so, our estimate of it from the age-36 column of Table 6b is .6346. Tables 5-6 provide 95% confidence intervals for gamma using the delta method and the covariance matrix from the CXS estimates. For a sample with starting age at least 36, the confidence interval for gamma is $[-.8211, -.3308]$. The corresponding confidence interval for the IES for consumption/leisure service flows is $[.5491, .7514]$. The intertemporal elasticity of substitution of consumption alone, on the other hand, is $1/(1-\alpha\cdot\gamma)$. Given our small estimate for alpha, the IES for consumption for the age-36 column data is closer to one: .8704.

Our results may be compared with estimates that have identified the IES from expected changes interest rates. Using aggregate consumption data Hall [1988], Cambell and Mankiw [1989], and Patterson and Pesaran [1992], for example, estimate the IES to be very nearly zero. Micro studies tend to estimate larger intertemporal elasticities. Banks, Blundell, and Tanner [1998], for example, estimate the average IES to be approximately 0.5. In another example, Attanasio and Weber [1993] estimate an IES of approximately 0.75 from micro data.[19] Thus, while we are using a very different source of variation to

[19] Barsky, et al. [1997] use answers to hypothetical questions to estimate the a distribution of IES in their sample. They find an average IES of 0.2 with less than 20% having an IES greater than 0.3. Others who have attempted to estimate a distribution of IES find evidence that the IES is increasing with wealth. See, e.g., Blundell, Browning and Meghir [1994]

CHART 4. NET-OF-TAX MALE EARNINGS:
2000 CONSTANT DOLLARS,
START WORK AGE 22,
RETIREMENT AGE 62,
REACHES 67 IN YEAR 2000

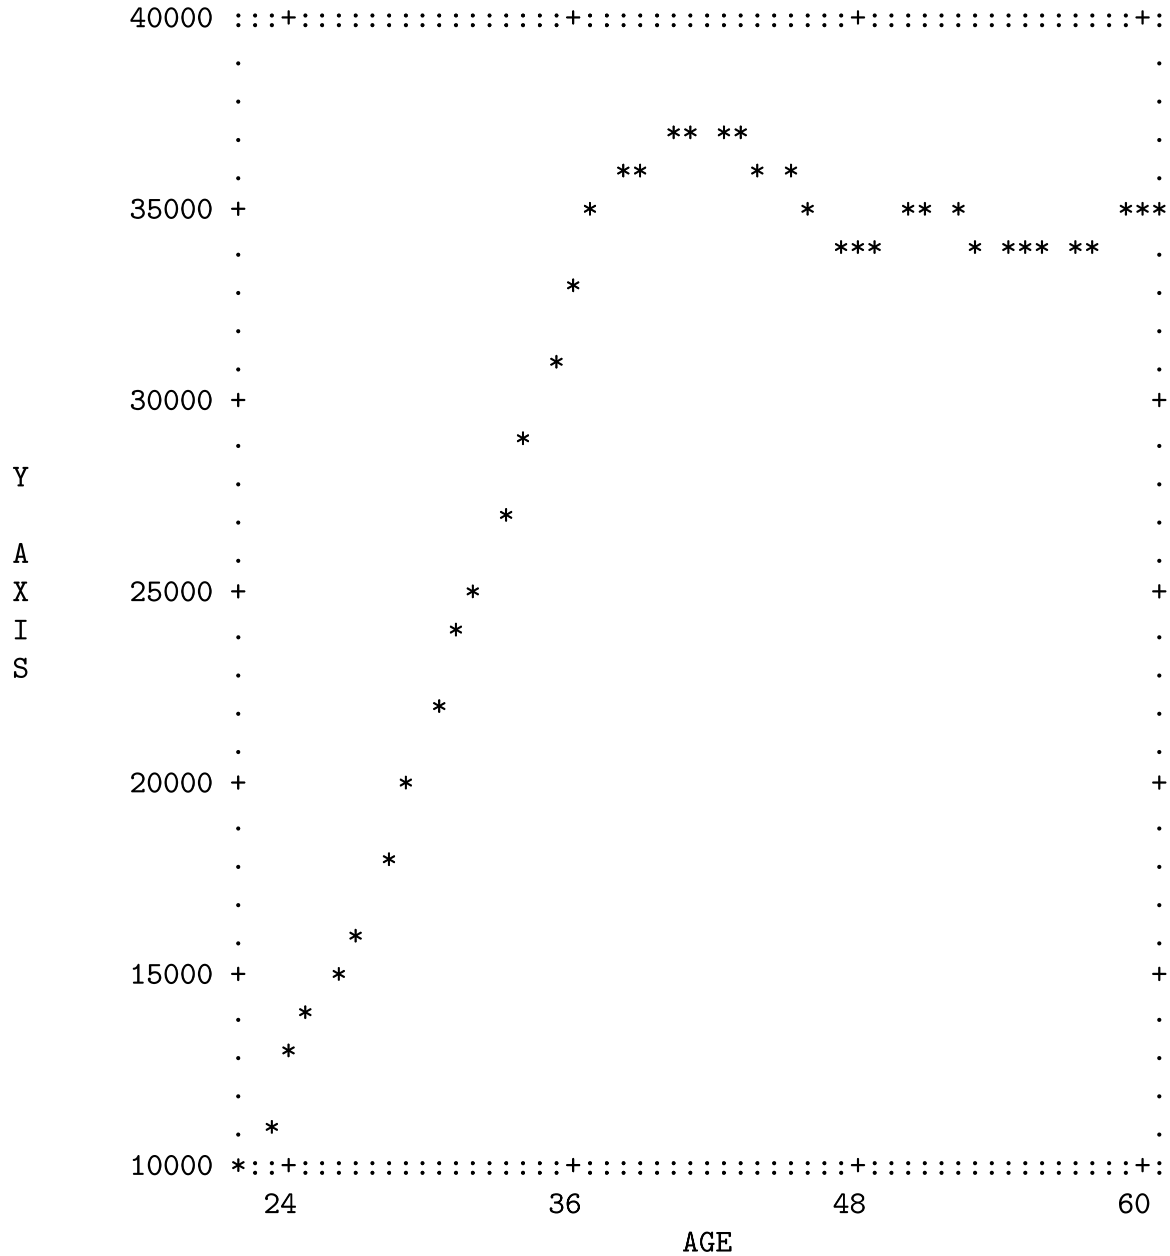

CHART 5. NET-OF-TAX FEMALE EARNINGS:
2000 CONSTANT DOLLARS,
START WORK AGE 22,
RETIREMENT AGE 60,
REACHES 65 IN YEAR 2000,
ABSCISSA IS HUSBAND'S AGE

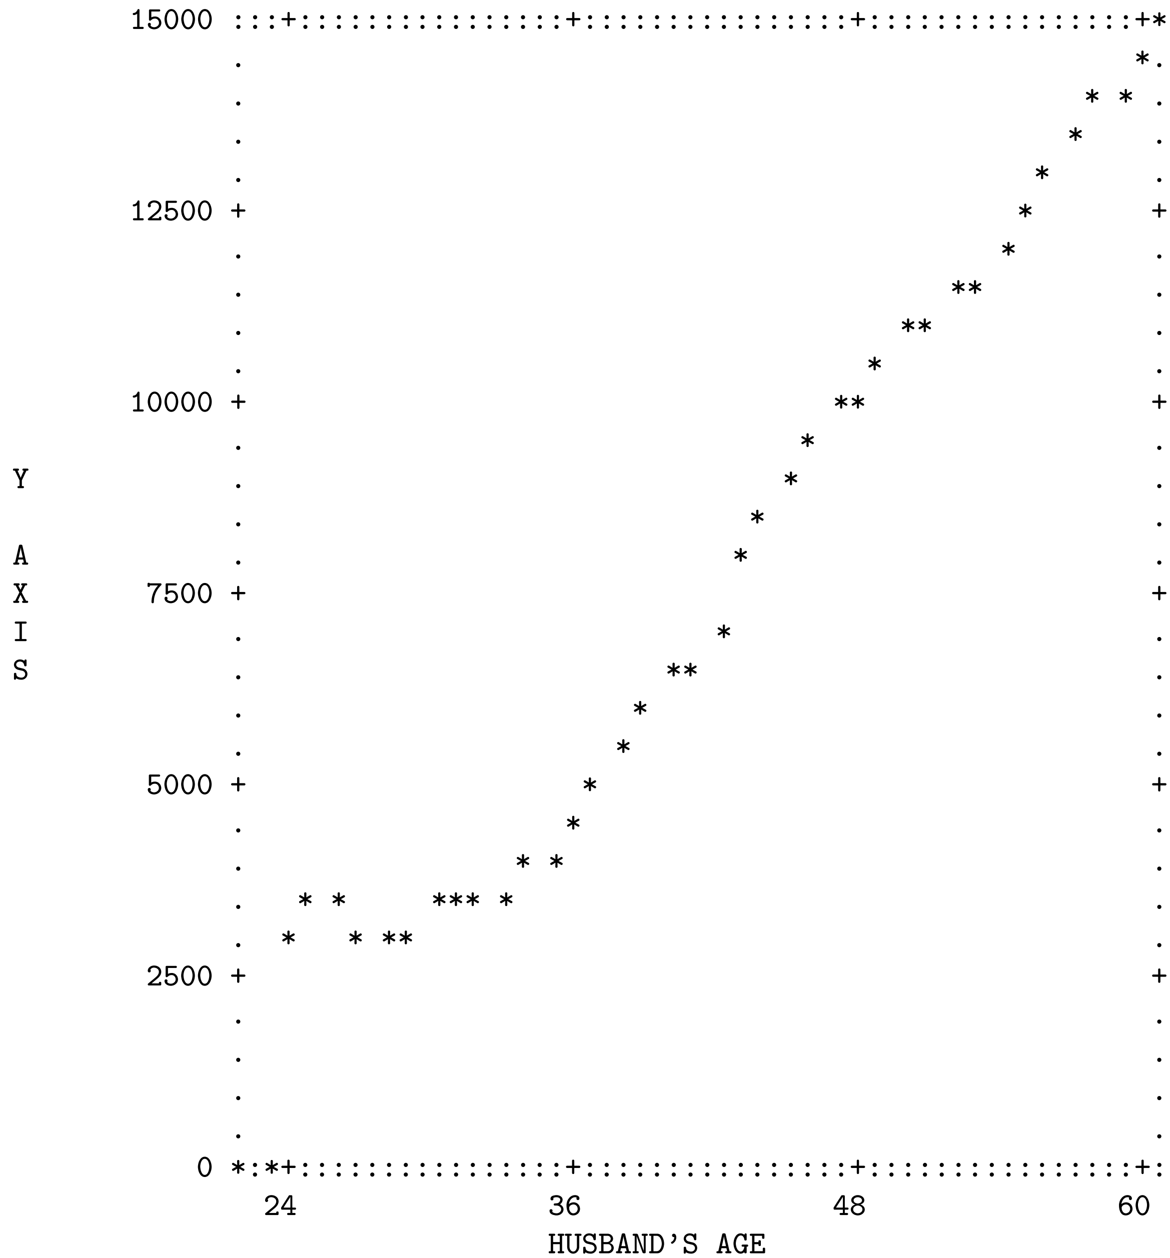

```
      CHART 6. NET-OF-TAX HOUSEHOLD EARNINGS PLUS SOCIAL SECURITY BENEFITS:
                         2000 CONSTANT DOLLARS,
                     HUSBAND 24 WHEN HOUSEHOLD STARTS,
                  COUPLE RETIRES WHEN HUSBAND REACHES 62,
            WIFE TWO YEARS YOUNGER, WIFE REACHES 65 IN YEAR 2000,
                         ABSCISSA IS HUSBAND'S AGE
             ::::::+::::::::::::::+::::::::::::::+*:::::::::::::+:::::::
             .                                  **                     .
       48000 +                                 *                       +
             .                            *****                        .
             .                        ****                             .
             .                 *******                                 .
             .               **                                        .
       42000 +             **                                          +
             .            *                                            .
             .                                                         .
             .           *                                             .
             .                                                         .
       36000 +          *                                              +
Y            .                                                         .
             .         *                                               .
A            .                                                         .
X            .        *                                                .
I      30000 +                                                         +
S            .       *                                                 .
             .                                                         .
             .      *                                                  .
             .     *                                                   .
       24000 +                                                         +
             .    *                                                    .
             .   *                                                     .
             .                                                         .
             .  *                                                      .
       18000 + *                                     ***********       +
             .*                                                        .
             *                                                         .
             .                                                         .
             .                                                         .
       12000 ::::::+::::::::::::::+::::::::::::::+:**:::::::::::********
                  30             45             60             75
                                HUSBAND'S AGE
```

Table 5. Estimated Structural Parameters for Life-Cycle Model: Euler Equation Coefficients from Table 3; Retirement Age 62; Health and Retirement Study Data[a]

Para–meter	CXS Regression Sample[b]				
	Age 20-79	Age 22-79	Age 24-79	Age 26-79	Age 28-79
Gamma (Std Err) [95% Confid.]	-1.010 (.1605) [-1.3247, -.6956]	-1.0536 (.1522) [-1.3519, -.7552]	-1.0287 (.1515) [-1.3257, -.7317]	-.9385 (.1500) [-1.2324, -.6445]	-.7545 (.1479) [-1.0444, -.4647]
Alpha (Std Err) [95% Confid.]	.2690 (.0023) [.2645, .2736]	.2641 (.0024) [.2594, .2689]	.2628 (.0025) [.2579, .2678]	.2621 (.0026) [.2570, .2672]	.2610 (.0020) [.2561, .2660]
Rho (Std Err) [95% Confid.]	-.0031 (.0020) [-.0070, .0008]	-.0013 (.0018) [-.0049, .0023]	-.0002 (.0019) [-.0038, .0035]	.0018 (.0019) [-.0019, .0056]	.0058 (.0020) [.0018, .0098]
Addendum: Household Age Last Binding Liquidity Constraint[c]					
Age (Yr, Mth)	(29,11)	(29,11)	(29,11)	(29,11)	(29,11)

a. See text.
b. See Table 3.
c. See text.

Table 6a. Estimated Structural Parameters for Life-Cycle Model: Euler Equation Coefficients from Table 4a; Retirement Age 62; Health and Retirement Study Data[a]

Para–meter	CXS Regression Sample[b]				
	Age 20-79	Age 22-79	Age 24-79	Age 26-79	Age 28-79
Gamma (Std Err) [95% Confid.]	-.4597 (.1254) [-.7055, -.2138]	-.6151 (.1247) [-.8595, -.3708]	-.6657 (.1258) [-.9122, -.4192]	-.6688 (.1240) [-.9119, -.4256]	-.6361 (.1213) [-.8739, -.3982]
Alpha (Std Err) [95% Confid.]	.2800 (.0021) [.2758, .2841]	.2718 (.0021) [.2677, .2759]	.2684 (.0022) [.2640, .2728]	.2649 (.0024) [.2603, .2696]	.2621 (.0025) [.2572, .2669]
Rho (Std Err) [95% Confid.]	.0029 (.0016) [-.0003, .0060]	.0042 (.0015) [.0012, .0072]	.0048 (.0015) [.0018, .0078]	.0061 (.0015) [.0031, .0091]	.0078 (.0015) [.0048, .0107]
Addendum: Household Age Last Binding Liquidity Constraint[c]					
Age (YR, MTH)	(28,11)	(28,11)	(29,11)	(29,11)	(29,11)

a. See text.
b. See Table 4a.
c. See text.

Table 6b. Estimated Structural Parameters for Life-Cycle Model: Euler Equation Coefficients from Table 4b; Retirement Age 62; Health and Retirement Study Data[a]

Para– meter	CXS Regression Sample[b]				
	Age 30- 79	Age 32- 79	Age 34- 79	Age 36- 79	Age 38- 79
Gamma (Std Err) [95% Confid.]	-.6023 (.1190) [-.8356, -.3691]	-.5639 (.1186) [-.7963, -.3315]	-.5556 (.1203) [-.7914, -.3199]	-.5759 (.1251) [-.8211, -.3308]	-.5695 (.0400) [-.6478, -.4912]
Alpha (Std Err) [95% Confid.]	.2595 (.0025) [.2546, .2644]	.2586 (.0026) [.2535, .2636]	.2576 (.0027) [.2523, .2630]	.2585 (.0029) [.2530, .2641]	.2579 (.0074) [.2433, .2724]
Rho (Std Err) [95% Confid.]	.0093 (.0015) [.0064, .0123]	.0105 (.0015) [.0075, .0135]	.0111 (.0016) [.0079, .0143]	.0102 (.0018) [.0066, .0138]	.0107 (.0066) [-.0022, .0236]
Addendum: Household Age Last Binding Liquidity Constraint[c]					
Age (YR, MTH)	(30,11)	(30,11)	(30,11)	(30,11)	(30,11)

a. See text.
b. See Table 4b.
c. See text.

identify the IES, our estimates are similar to those obtained in micro studies of the change in consumption growth with expected changes in interest rates.

6. Policy Simulations

This section uses our preferred structural parameter estimates — those of Tables 4b and 6b — to simulate the effect on the life–cycle model's utility–maximizing retirement age of various possible changes in the U.S. Social Security system.

In May 2001, the President appointed the bipartisan Commission to Strengthen Social Security. The Commission filed its report on December 21, 2001 (revised March 19, 2002).[20] In a section entitled "Fiscal Problems Facing Social Security," the report outlines the following scenario [p. 64]:

> Absent Congressional action, the Trust Funds will be exhausted in 2038. At that time, the Social Security system's dedicated revenue will be enough to cover only 74 percent of promised benefits. To pay full promised benefits would require an increase in the total tax rate from payroll and benefit taxation from the current 12.4 percent to 17.8 percent.

Our simulations accordingly consider a 5.4% increase in the Social Security payroll tax. All of our simulations are partial equilibrium in the sense that they maintain a 4%/year pre–tax real interest rate and HRS earnings profiles for men and women. We do, however, assume that the entire increase in the payroll tax, including the half that employers would pay, is promptly shifted to labor — so that aftertax earnings fall by the full amount. We also consider a 26% reduction in Social Security benefits, and, for illustrative purposes, a 50% cut as well. Finally, the earlier so–called Greenspan Commission recommended an increase in the Normal Retirement Age for computing Social Security benefits, and current statutes have that age rising from 65 in 2000 to 67 by 2022. Accordingly, we simulate an increase in the Normal Retirement Age from 65 to 67.

Results. Table 7 presents outcomes. The columns correspond to parameter estimates from the same column of Tables 4b and 6b.

The first row of Table 7 simulates an increase in the payroll tax of 5.4 percent. The tax increase has a substitution effect: the substitution effect tends to increase demand for leisure because it lowers the net–of–tax remuneration for work hours. The tax increase also has an income effect: the income effect tends to reduce the demand for leisure as household resources fall. With a Cobb–Douglas function producing household services from consumption and leisure, one might expect these two effects to offset one another exactly; however, while an exact offset would follow if earnings were the only source of household resources, in our framework households have Social Security benefits as well as earnings. The benefit formula for the U.S. Social Security System averages a household's earnings over its 37 top earning years — using a procedure that raises the relative importance of earlier years. A nonlinear, progressive formula then defines benefits as a function of the average. In the end, the link between one more year of earnings and additional benefits is tenuous, so that Social Security benefits are approximately lump–sum transfers. Thus, an increase in the payroll tax affects only part of households' lifetime resources, and the

[20] See http://www.commtostrengthensocsec.gov/reports/Final_report.pdf.

Table 7. Comparative–Static Results: Months Change in Retirement Age Following an Increase in Social Security Tax, an Increase in the Social Security Normal Retirement Age, or a Decrease in Social Security Benefits; Euler Equation Coefficients from Table 4b[a]

Policy Change	CXS Regression Sample Ages[b]				
	Age 30-79	Age 32-79	Age 34-79	Age 36-79	Age 38-79
Increase Social Security Tax by 5.4%[c]	-2	-2	-3	-2	-3
Increase in Normal Retirement Age to 67	+4	+4	+4	+4	+4
Cut Social Security Benefits to 74% of Current	+7	+7	+7	+7	+7
Cut Social Security Benefits to 50% of Current	+18	+19	+19	+19	+19

a. See text.
b. See Table 4b.
c. Assumes entire tax increment transferred to labor. See text.

substitution effect from a tax change dominates the income effect. Nevertheless, our table's substantial tax increase only reduces work by 2-3 months in the simulations.

The next three rows of Table 7 present benefit reductions. Given our assumption about life spans, for example, increasing the retirement age to 67 has the same effect as decreasing benefits by about 20%. To the extent that Social Security benefits are essentially a lump–sum transfer, we expect a pure income effect. In these three cases, the fall in resources leads to a negative income effect, tending to decrease leisure. Corresponding increases in work effort — i.e., increases in the retirement age — range from 4 to 19 months.

7. Conclusions

A number of recent papers have noted the typical fall in household consumption at retirement. Some authors find this to be a "puzzle." Section 2, however, argues that under plausible assumptions — that household preferences over consumption and leisure are (intratemporally) nonseparable, and that households essentially face a discrete choice of working full time or not at all — a drop in consumption at retirement is fully consistent with the life-cycle model of behavior. Furthermore, Section 2 shows how one can employ the empirical magnitude of the decline in consumption at retirement to help estimate key parameters of the model (e.g., equation (5)).

Sections 3-5 use a consumption Euler equation, regressed on CXS pseudo-panel data, and the life-cycle model's utility-maximizing retirement age, calibrated from HRS data (including linked Social Security earning records), to estimate key structural parameters for a nonstochastic life-cycle specification. Our identification strategy uses the size of the drop in consumption at retirement — as opposed to relying on time series variation in real interest rates. Our estimates of the intertemporal elasticity of substitution vary from one half to nearly one — depending on whether we are analyzing service flows or consumption.

One surprising result is how well our new elasticity estimates agree with those in the literature obtained from micro data, but using very different procedures. A second surprise, at this point, is the apparent importance of liquidity constraints for households up to the age of 30-32 (c.f., Mariger [1987]).

Section 6 presents policy simulations of the possible effect of changes in the U.S. Social Security System on retirement ages. Increases in the Social Security tax imply earlier retirement; reductions of Social Security benefits lead to later retirement. In our simulations so far, even fairly substantial policy changes affect retirement ages by less than one year.

Bibliography

[1] Altig, David; Auerbach, Alan J.; Kotlikoff, Laurence J.; Smetters, Kent A.; and Walliser, Jan, "Simulating Fundamental Tax Reform in the United States," *American Economic Review* vol. 91, no. 3 (June 2001): 574-595.

[2] Anderson, Patricia M.; Gustman, Alan L.; and Steinmeier, Thomas L., "Trends in Male Labor Force Participation and Retirement: Some Evidence on the Role of Pensions and Social Security in the 1970s and 1980s," *Journal of Labor Economics* vol. 17, no. 4, part 1 (October 1999): 757-783.

[3] Attanasio, Orazio; and Browning, Martin, "Consumption over the Life Cycle and Over the Business Cycle," *American Economic Review* 85, no. 2 (December 1995): 1118-1137.

[4] Auerbach, Alan J.; and Kotlikoff, Laurence J. *Dynamic Fiscal Policy*. Cambridge: Cambridge University Press, 1987.

[5] Banks, James; Blundell, Richard; and Tanner, Sarah, "Is There a Retirement–Savings Puzzle?" *American Economic Review* 88, no. 4 (September 1998): 769-788.

[6] Barsky, Robert; Kimball, Miles; Juster, F. Thomas; and Shapiro, Matthew, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics* vol. 112, no. 2 (May 1997): 537-579.

[7] Bernheim, B. Douglas; Skinner, Jonathan; and Weinberg, Steven, "What Accounts for the Variation in Retirement Wealth Among U.S. Households," *American Economic Review* 91, no. 4 (September 2001): 832-857.

[8] Blundell, Richard; and MaCurdy, Thomas, "Labor Supply: A Review of Alternative Approaches," in Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics: vol. 3A*. Amsterdam: Elsevier, 1999.

[9] Burkhauser, Richard V.; Dwyer, Debra; Lindeboom, Maarten; Theeuwes, Jules; and Woittiez, Isolde, "Health, Work, and Economic Well–Being of Older Workers, Aged Fifty–One to Sixty–one: A Cross–National Comparison Using the U.S. HRS and the Netherlands CERRA Data Sets," in James Smith and Robert Willis (eds.), *Wealth, Work, and Health: Innovations in Measurement in the Social Sciences*. Ann Arbor: The University of Michigan Press, 1999.

[10] Gustman, Alan L.; and Steinmeier, Thomas, L., "A Structural Retirement Model," *Econometrica* vol. 54, no. 3 (May 1986): 555-584.

[11] Hall, Robert E., "Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy* 86, no. 6 (December 1978): 971-987.

[12] Hall, Robert E., "Intertemporal Substitution in Consumption," *Journal of Political Economy* 96, no. 2 (1988): 339-357.

[13] Hubbard, R. Glen; and Judd, Kenneth, "Liquidity Constraints, Fiscal Policy and Consumption," *Brookings Papers on Economic Activity* 1 (1986): 1-50.

[14] Hubbard, R. Glen; Skinner, Jonathan; and Zeldes, Stephen P., "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," *Carnegie–Rochester Conference Series on Public Policy* 40 (June 1994): 59-125.

[15] Hubbard, R. Glen; Skinner, Jonathan; and Zeldes, Stephen P., "Precautionary Saving and Social Insurance," *Journal of Political Economy* 103, no. 2 (April 1995): 360-399.

[16] Hurd, Michael, "The Effect of Labor Market Rigidities on the Labor Force Behavior of Older Workers," in (David Wise ed.), *Advances in the Economics of Aging.* Chicago: The University of Chicago Press, 1996.

[17] Hurd, Michael; and Rohwedder, Susann, "The Retirement–Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement," NBER working paper 9586, http://www.nber.org/papers/w9586.

[18] King, Robert; Plosser, Charles; and Rebelo, Sergio, "Production, Growth and Business Cycles: I. The Basic Neoclassical Model," *Journal of Monetary Economics* vol. 21, no. 2 (1988): 195-232.

[19] Laitner, John, "Secular Changes in Wealth Inequality and Inheritance," *The Economic Journal*, vol. 111, no. 474 (October 2001a): 691-721.

[20] Lumsdaine, Robin L.; and, Mitchell, Olivia, "New Developments in the Economic Analysis of Retirement," in Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics: vol. 3C.* Amsterdam: Elsevier, 1999.

[21] Mariger, Randall P. *Consumption Behavior and the Effects of Government Fiscal Policy.* Cambridge: Harvard University Press, 1986.

[22] Mariger, Randall P., "A Life-Cycle Consumption Model with Liquidity Constraints: Theory and Empirical Results," *Econometrica* 55, no. 3 (May 1987): 533-557.

[23] Paterson, Kerry D.; and Pesaran, Bahram, "The Intertemporal Elasticity of Substitution in Consumption in the U.S. and in the U.K.," *Review of Economics and Statistics* 74, no. 4 (May 1992): 573-584.

[24] Pencavel, John, "Labor Supply of Men: A Survey," in Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics: vol. 1.* Amsterdam: Elsevier, 1986.

[25] President's Commission, "Strengthening Social Security and Creating Personal Wealth for All Americans," 2001.

[26] Rust, John; and, Phelan, Christopher., "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets," *Econometrica* 65, no. 4 (July 1997): 781-831.

[27] Stock, James H.; and, Wise, David A., "Pensions, The Option Value of Work, and Retirement," *Econometrica* 58, no. 2 (September 1990): 1151-1180.

[28] Tobin, J., "Life Cycle Saving and Balanced Growth," in W. Fellner (ed.), *Ten Economic Studies in the Tradition of Irving Fisher.* New York: Wiley, 1967.

[29] Zeldes, S., "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy* 97, no. 4 (April 1989): 305-346.