

“Consumption Over the Life Cycle.” Gourinchas, Pierre-Olivier and Jonathan A. Parker. *Econometrica* Vol. 70, No. 1 (2002): 47-89.
<http://doi.org/10.1111/1468-0262.00269>

The copyright to this article is held by the Econometric Society, <http://www.econometricsociety.org/>. It may be downloaded, printed and reproduced only for personal or classroom use. Absolutely no downloading or copying may be done for, or on behalf of, any for-profit commercial firm or for other commercial purpose without the explicit permission of the Econometric Society.

All other permission requests or questions (including commercial purposes or on behalf of any for-profit entity) should be addressed to:

Wiley Permissions
www.wiley.com/go/rightslicensing
Then select Copyright & Permissions

For requests for any content not containing a ‘request permission’ link please contact permissions@wiley.com .

CONSUMPTION OVER THE LIFE CYCLE

BY PIERRE-OLIVIER GOURINCHAS AND JONATHAN A. PARKER¹

ABSTRACT. This paper estimates a structural model of optimal life-cycle consumption expenditures in the presence of realistic labor income uncertainty. We employ synthetic cohort techniques and Consumer Expenditure Survey data to construct average age-profiles of consumption and income over the working lives of typical households across different education and occupation groups. The model fits the profiles quite well. In addition to providing reasonable estimates of the discount rate and risk aversion, we find that consumer behavior changes strikingly over the lifecycle. Young consumers behave as buffer-stock agents. Around age 40, the typical household starts accumulating liquid assets for retirement and its behavior mimics more closely that of a certainty equivalent consumer. Our methodology provides a natural decomposition of saving and wealth into its precautionary and life-cycle components.

1. INTRODUCTION

THIS PAPER ESTIMATES a dynamic stochastic model of the life-cycle saving behavior of households. We focus on estimation of structural preference parameters and upon characterizing optimal behavior when households face exogenous, stochastic, labor income processes. We are motivated by two observations.

First, better methodology, data, and creative use of natural experiments are leading to significant evidence against consumption smoothing as an accurate description of household-level behavior. Despite generally poor-quality data on household consumption, recent tests often find that consumption responds to predictable changes in income. At lower frequencies, there is some evidence that consumption tracks expected and unexpected income changes across groups of households.²

Second, recent theoretical work has improved our understanding of optimal household consumption and saving behavior under uncertainty (Ayagari (1994), Zeldes (1989), Deaton (1991), Carroll (1997)). Usual log-linearized Euler equations can provide a poor approximation to marginal utility smoothing as well as unreliable estimates of the preference parameters in uncertain environments (Ludvigson and Paxson (1997)). Uncertainty can generate a positive correlation between expected

income changes and the conditional variance of consumption growth (and higher moments) through precautionary saving, typically subsumed in the error term of the approximation. In general, the presence of uncertainty invalidates most tests using a linearized form of the Euler equation.

In this paper, we use household survey data and simulation techniques to estimate a structural model of inter-temporal consumption choice with realistic levels of income uncertainty. We measure, exploit and analyze the systematic age-pattern of income and consumption. The estimated model is used to re-interpret life-cycle consumption and asset accumulation behavior. We find substantial age-heterogeneity in consumption behavior that results from the interaction between, and relative strengths of, retirement and precautionary motives for saving at different ages.

We proceed in two steps. Using weak identifying assumptions, we construct average total consumption and income profiles across the working lives of households of five different educational attainments and four different occupational groupings, using high quality household-level data on consumption and income from a sample of roughly 40,000 households from the Consumer Expenditure Survey (CEX) from 1980 to 1993. Consumption and income age profiles are both significantly hump-shaped, despite controlling for family composition and cohort effects, and consumption tracks income reasonably well early in life.

Second, using labor income uncertainty measured from the Panel Study of Income Dynamics (PSID) and our constructed profiles, we estimate a canonical stochastic life-cycle model of consumer behavior. For any value of the key parameters, we can solve numerically and recursively for the household optimal behavior and then aggregate to generate a simulated life-cycle consumption profile. By matching this simulated profile to its empirical counterpart, we estimate the parameters of the consumption problem using the Method of Simulated Moments.

To the best of our knowledge, this represents the first structural estimation of consumption functions that incorporates precautionary saving.

In our model, the optimal choice of consumption depends not only on lifetime resources, the real interest rate, and the discount factor, but also on the expected growth rate of income, so that consumer behavior may vary systematically as households age (Hubbard, Skinner, and Zeldes (1994), Carroll (1997)). When expected income growth and the discount rate are low relative to the interest rate, consumers' behavior remains similar to that predicted by the certainty-equivalent

life-cycle hypothesis model (CEQ-LCH henceforth). If, on the other hand, expected income growth or the discount rate are large relative to the interest rate, consumers behave as buffer-stock agents — wishing to consume large amounts out of future resources but instead saving enough to weather bad income draws. Given that expected income profiles are hump-shaped, the model can potentially produce concave average consumption profiles.

The paper yields four main findings.

First, the fitted model matches the correlation between consumption and income at young ages and the general concavity of the profile that is observed in the data. By appropriately constructing the inputs to the household problem and by estimating preference parameters, we improve on the previous studies based mostly on simulation of a limited number of scenarios.

Second, we find reasonable estimates of the preference parameters. The average household has a discount rate of 4.0 – 4.5 percent, and, less robust across specifications, a coefficient of relative risk aversion varying between 0.5 and 1.4. We also estimate consumption behavior at retirement and find a credible marginal propensity to consume out of liquid assets of 6 – 7 percent. On the other hand, the model provides a poor estimate of the marginal propensity to consume out of illiquid wealth, near zero. We evaluate the robustness of our method by estimating the model under several alternative assumptions.

Third, the paper contributes to the debate on the determinants of wealth accumulation. In our model, the relative shapes of the consumption and income profiles reveal the relative roles of precautionary and retirement motives for accumulating liquid assets. Solving for saving and wealth accumulation under complete markets, we then construct a measure of *precautionary* saving and wealth.³ Our fitted model indicates that wealth is accumulated early in life for precautionary reasons — were it not for income uncertainty households would instead borrow against future labor income. Households older than 40 instead save mostly for retirement and bequests. Thus observed saving patterns are consistent with forward-looking optimizing behavior in a life-cycle framework augmented to include income uncertainty. The importance of the precautionary motive early in life implies that between 60 and 70 percent of non-pension wealth is due to precautionary saving, according to our estimates and holding the real interest rate fixed.

Finally, and robustly, we find strikingly different consumption behavior for households at different ages: households behave like ‘buffer-stock’ consumers early in their working lives and more like CEQ-LCH households as retirement nears. We estimate that households make the transition from buffer-stock to CEQ-LCH behavior between age 40 and age 45. We conclude that, due to precautionary saving, a significant fraction of households behave as *target savers*, for whom the log-linearized form of the Euler equation is expected to fail.⁴ It is interesting to note that buffer-stock behavior arises early in life not from high levels of household impatience but rather from the steepness of the expected income profiles at young ages.⁵ Our results are not assumed: with a sufficiently low estimated discount rate, the average consumption profile would be very similar to that of the certainty equivalent case.

Our paper is related to three main strands of literatures. First, to construct our profiles we build on the studies of life-cycle behavior in Ghez and Becker (1974), Kotlikoff and Summers (1981), and Carroll and Summers (1991), using the techniques of Deaton (1985) and Browning, Deaton, and Irish (1985). Thus our work is closely related to Attanasio, Banks, Meghir, and Weber (1997) and Attanasio and Browning (1995) who argue that household level data can be explained by the CEQ-LCH once a more flexible treatment of preferences and or aggregation of commodities is allowed for. We argue instead that the data can be explained by variations in precautionary saving.⁶ Our methodology can be extended to nest preference heterogeneity and precautionary motives for saving.

Second, we build on previous studies that parameterize and simulate life-cycle consumption models with uncertainty. Hubbard, Skinner, and Zeldes (1994) and Carroll (1997) show that the optimal consumption choices of consumers can lead to profiles that are hump-shaped and track income over the early part of life for some parameterizations. Carroll and Samwick (1997) calibrate the discount rate of a structural model to replicate the sensitivity of asset holding to income uncertainty that is observed at young ages. Our approach goes beyond those studies by estimating a structural model of consumption.

Finally, a few papers estimate the optimal level of consumption as a function of a household’s fully specified environment. Palumbo (1999) uses individual consumption, income and asset data to estimate individual consumption levels for retirees. We choose to rely on average profiles because

we do not believe that the individual-level data are of sufficient quality to support the employed technique in general. In addition, there are now several papers that use a methodology derived from our paper to address asset holding (Gakidis (1998), Cagetti (1998)) and labor supply (French (1998)) decisions.

The structure of the paper is as follows. Section 2 lays out an empirically tractable model of consumer maximization and characterizes optimal behavior. Section 3 introduces the two-step Method of Simulated Moments methodology for estimating the model. The fourth section describes the first-stage estimation of income risk, retirement consumption behavior, and the real interest rate. The fifth section discusses the construction of life-cycle profiles of consumption and income and presents graphs of the profiles for various education and occupation cells. Finally, we present the results of the estimation and conclude. Appendices contain more detailed descriptions of the numerical optimization, the econometric procedure, and the CEX data.

2. CONSUMPTION BEHAVIOR WITH STOCHASTIC INCOME

2.1. The Canonical Model with Labor Income Uncertainty. Our starting point is the basic discrete-time, life-cycle model of household consumption behavior. Consumers live for N periods and work for $T < N$, where both T and N are exogenous and fixed. At every age $1 \leq t \leq T$, the consumer receives a stochastic income Y_t . There is one asset in the economy, with a constant, after-tax, gross real interest rate R . We assume that preferences take the standard additively separable expected utility form, with discount factor β :

$$(1) \quad E \left[\sum_{t=1}^N \beta^t u(C_t, Z_t) + \beta^{N+1} V_{N+1}(W_{N+1}) \right],$$

where C_t represents total consumption at age t , W_t represents total financial wealth and Z_t is a vector of deterministic household characteristics (e.g. family size). V_{N+1} represents the value to the consumer of any assets left at the time of death, capturing any bequest motive. The consumer maximizes equation (1) given an initial wealth level W_1 , the constraint that terminal wealth is non-negative $W_{N+1} \geq 0$, and the dynamic budget constraint:

$$W_{t+1} = R(W_t + Y_t - C_t).$$

We further assume that the felicity function $u(.,.)$ is of the Constant Relative Risk Aversion (*CRRA*) form, with intertemporal elasticity of substitution $1/\rho$, and multiplicatively separable in Z :

$$u(C, Z) = v(Z) \frac{C^{1-\rho}}{1-\rho}.$$

If income were certain, the solution to this program would be standard: the consumer would choose a consumption path such that

$$(2) \quad \frac{C_{t+1}}{C_t} = \left(\beta R \frac{v(Z_{t+1})}{v(Z_t)} \right)^{\frac{1}{\rho}}.$$

With constant individual characteristics, equation (2) implies a constant growth rate of consumption. Consumption increases (respectively decreases) over time when the interest rate is larger (respectively smaller) than the discount rate. The growth rate of consumption is independent of the income profile. The level of consumption is determined by the lifetime budget constraint and the terminal value function.

When individual characteristics vary over the life cycle, the growth rate of consumption may change accordingly. For instance, if the marginal utility of consumption increases with family size, consumption will grow faster as family size increases, and slower as children leave the household. These variations in individual characteristics may induce a positive correlation between consumption and income over the life cycle.

With individual income uncertainty and prudence, households hold liquid wealth to insure themselves against future contingencies. This precautionary saving motive has potentially far-reaching and striking implications. The main consequence of income uncertainty is to increase the slope of the consumption profile. Hubbard, Skinner, and Zeldes (1994) demonstrates that this uncertainty can lead to hump-shaped consumption profiles as households save for precautionary reasons early in life and run down these assets during retirement due to lower levels of uncertainty and an increased probability of death.

Zeldes (1989), Carroll (1992) and Deaton (1991) analyze the case in which consumers are also *impatient*: absent uncertainty, households would like to borrow in order to finance a high level of current consumption. Deaton (1991) and Zeldes (1989) impose liquidity constraints while Carroll

(1997) sets up a model in which consumers choose never to borrow. In either rendition, assets play the role of a *buffer stock* against bad income shocks. Consumers have a target level of liquid assets, above which impatience dominates and assets are run down, and below which the precautionary motive dominates and assets are accumulated. Thus the theory predicts a positive correlation between expected income growth and consumption growth.

This paper explicitly incorporates uninsurable idiosyncratic income uncertainty. We adopt Zeldes' (1989) formulation, and decompose the labor income process into a permanent component, P_t , and a transitory component, U_t :

$$(3) \quad \begin{aligned} Y_t &= P_t U_t \\ P_t &= G_t P_{t-1} N_t \end{aligned}$$

The transitory shocks, U_t , are independent and identically distributed, take the value 0 with probability $0 \leq p < 1$, and are otherwise log-normally distributed, $\ln U_t \sim \mathcal{N}(0, \sigma_u^2)$. The log of the permanent component of income, $\ln P_t$, evolves as a random walk with age specific drift $\ln G_t$. The shock to the permanent component of income, N_t , is independently and identically log-normally distributed, $\ln N_t \sim \mathcal{N}(0, \sigma_n^2)$.⁷ Thus $\Delta \ln Y_t \sim MA(1)$, income is a serially correlated process, with both permanent and transitory shocks, and a positive probability of zero income in every period.⁸

Two points are worth noting. First, innovations to the permanent component of income are only as permanent as the remaining length of the working life: all shocks are ultimately transitory, as consumers retire and die. As a consequence, the propensity to consume out of 'permanent' shocks will vary with age, a point emphasized by Clarida (1991). This property holds true for the CEQ-LCH also. Second, in this setup consumers will choose never to borrow against future labor income, a point shown by Schechtman (1976). This follows from (a) there being a strictly positive probability that labor income will be arbitrarily close to zero for the rest of the working life and (b) the Inada condition $\lim_{c \rightarrow 0} u'(c) = \infty$. To see this, suppose the household were to borrow in the next to last working period. Then, with strictly positive probability it would be left without any wealth in the last working period. The household would then have an infinite expected marginal utility. Backward induction implies that it will never be optimal to borrow in any period.

The precautionary motive acts as a *self-imposed* liquidity constraint. It is important to note that this holds true even when p , the probability of strictly zero income, is set to zero. The key here is that the income process has a zero lower bound. This and the requirement that the consumer dies without debt almost surely, impose a natural borrowing limit equal to 0 (see Ayagari (1994)). With a strictly positive lower bound on income, the consumer could borrow only up to the present discounted value of certain future income.

Going from the model to the data, we need to make four assumptions. First, in order to solve the consumer's problem as stated, we need to specify both the nature of uncertainty during retirement and a bequest function. While there have been good attempts at modelling consumer behavior after retirement,⁹ we feel that we know too little about the form that uncertainty takes after retirement to use our methodology and draw inferences from post-retirement behavior. Uncertainty arises from different sources – e.g. medical expenses, the timing of death and asset returns. *Inter-vivos* bequests are also likely to be important. Although these sources of uncertainty are also present to some extent in the last working years, labor income uncertainty is the dominant source of uncertainty when young. Further, high quality information on household asset holdings, together with consumption and income, is not available. Given that investment income, social security, and pensions represent the main sources of income during retirement, it is currently difficult to establish consumption patterns as a function of total wealth. Lastly, even with a proper treatment of retirement issues, one would have to make a guess about the bequest function. Instead, we make use of Bellman's optimality principle, and truncate our problem at the date of retirement.

Second, we assume that age variations in $v(Z_t)$ are common across households of the same age t , deterministic, and come from changes in family size, so that the evolution of the consumer problem can be captured by a single state variable.¹⁰ This has the added advantage that we can adjust and report consumption and income profiles that maintain a constant effective family size across ages.

Third, the model imposes a single vehicle for precautionary and retirement wealth accumulation, since there is only one asset. In practice, much of net worth at retirement is accumulated in the form of illiquid wealth, only available after retirement.¹¹ This suggests that the relevant model of consumption behavior should incorporate an additional asset that is illiquid and accessible only after retirement. However, this would substantially complicate the problem by introducing another

control variable (how much to save in liquid versus illiquid assets) and state variable (illiquid assets). In order to keep our estimation procedure tractable, we instead assume that illiquid wealth accumulates exogenously, cannot be borrowed against, and that illiquid wealth in the first year of retirement, H_{T+1} , is proportional to the last permanent component of income, $H_{T+1} \equiv hP_{T+1} = hP_T$. These assumptions eliminate both illiquid assets as a state variable of the program, and contributions to illiquid assets as a control variable.¹²

Finally, we postulate a retirement value function which summarizes the consumer's problem at retirement time. Defining cash on hand X_{t+1} as total liquid financial resources, $X_{t+1} = R(X_t - C_t) + Y_{t+1} = W_{t+1} + Y_{t+1}$, we choose a functional form that maintains the tractability of the problem and is flexible enough to allow robustness checks:

$$V_{T+1}(X_{T+1}, H_{T+1}, Z_{T+1}) = \kappa v(Z_{T+1}) (X_{T+1} + H_{T+1})^{1-\rho},$$

for some constant κ . Under the CRRA assumption, this functional form is exactly correct if time of death and/or asset returns are the only sources of uncertainty after retirement (Merton (1971)). Under this assumption, consumption is linear in total wealth at retirement. Since most households have large amounts of illiquid wealth in Social Security and pensions at retirement, a linear approximation for the true consumption rule is likely to be a good approximation.¹³

Defining the value function for the household problem at time τ as V_τ , our problem becomes:

$$\begin{aligned} V_\tau(X_\tau, P_\tau, Z_\tau) = \\ \max_{c_\tau, \dots, c_T} E_\tau \left[\sum_{t=\tau}^T \beta^{t-\tau} v(Z_t) \frac{C_t^{1-\rho}}{1-\rho} + \beta^{T+1-\tau} \kappa v(Z_{T+1}) (X_{T+1} + hP_{T+1})^{1-\rho} \right] \\ s.t. \quad X_{t+1} = R(X_t - C_t) + Y_{t+1}, \quad X_{T+1} \geq 0, \text{ and given (3).} \end{aligned}$$

where income is defined as disposable income, net of Social Security taxes and saving in illiquid assets. The assumption that illiquid wealth cannot be borrowed against imposes the borrowing constraint that liquid wealth must be weakly positive at retirement.

2.2. Solving for Optimal Consumer Behavior. The setup of the problem combined with our particular choice of retirement value function makes the problem homogeneous of degree $(1 - \rho)$ in the permanent component of income P_t . We write the optimal consumption rule as a function

of age, t , and the ratio of cash on hand to the permanent component of income, $x_t \equiv X_t/P_t$. The budget constraint becomes:

$$(4) \quad x_{t+1} = (x_t - c_t) \frac{R}{G_{t+1}N_{t+1}} + U_{t+1}, \quad 1 \leq t \leq T.$$

where lowercase letters are normalized by the permanent component of income. The following Euler equation holds for $t < T$:

$$(5) \quad u'(c_t(x_t)) = \beta R E \left[\frac{v(Z_{t+1})}{v(Z_t)} u'(c_{t+1}(x_{t+1}) G_{t+1}N_{t+1}) \right],$$

where $c_t(x_t)$ represents the optimal consumption rule at time t (normalized), as a function of normalized cash on hand x_t . In the last working period, the Euler equation becomes:

$$u'(c_T(x_T)) = \max \left\langle u'(x_T), \beta R \frac{v(Z_{T+1})}{v(Z_T)} u'(c_{T+1}(x_{T+1})) \right\rangle$$

since the household cannot borrow against illiquid wealth H_{T+1} and $P_{T+1} = P_T$.

Given our retirement value function, optimal consumption is linear in total wealth at retirement:

$$C_{T+1} = \gamma_1 (X_{T+1} + H_{T+1})$$

where γ_1 represents the marginal propensity to consume out of wealth.¹⁴ Normalized consumption, c_{T+1} , is then linear in normalized cash on hand x_{T+1} :

$$(6) \quad c_{T+1} = \gamma_0 + \gamma_1 x_{T+1},$$

with $\gamma_0 \equiv \gamma_1 h$.

The solution to the consumer problem consists of a set of consumption rules $\{c_t(x_t)\}_{1 \leq t \leq T}$. The consumption rule in the last working period is the solution to the above Euler equations for all values of cash on hand, where we replace c_{T+1} using equation (6). Solving recursively generates consumption functions $c_T(x_T), \dots, c_1(x_1)$. A complete description of the solution method is provided in Appendix A.

2.3. Characterization of Individual Consumption Behavior. Figure 1 panel *A* shows consumption rules at various ages, for a typical household working from age 25 to 65 and retiring thereafter with a consumption rule at age 66 characterized by: $\gamma_0 = 0.001$, $\gamma_1 = 0.071$. These parameter values are those of our baseline estimates.¹⁵ If G and Z were constant, the finite horizon problem would converge to the infinite horizon one, as we move further away from retirement. Consumption is always positive, increasing, and concave in cash on hand. Early in life, households exhibit ‘buffer stock’ behavior: for low level of cash on hand, typically less than the permanent component of their income ($x \leq 1$), households consume most, but never all, of their financial wealth, and move to the next period with a very low level of liquid assets. At high levels of cash on hand, the precautionary motive is small and households consume more than the income they expect to receive (which equals 1) and so run down their assets. As we discuss later, our baseline retirement consumption rule implies either little illiquid wealth or a low propensity to consume from it. Thus as households age, they must save for retirement: consumption rules decline and households accumulate significant amounts of liquid wealth.

Figure 1 panel *B* reports consumption rules for a similar household with the alternative retirement consumption rule $\gamma_0 = 0.594$, $\gamma_1 = 0.0774$.¹⁶ The higher value of γ_0 implies either larger illiquid wealth at retirement or a larger propensity to consume from it. The two households have very similar consumption rules early in life, for the same reasons that convergence occurs in the stationary problem. Both households behave as buffer stock households in their youth. Later in life, the household in panel *B* does not need to accumulate as much liquid wealth since it can rely more on illiquid wealth to maintain consumption in retirement. Thus the consumption rules decline less as the household ages.¹⁷

What general conditions are necessary to generate buffer-stock behavior? In a stationary infinite-horizon model, buffer stock behavior arises when agents are sufficiently impatient not to accumulate wealth without bounds. Formally, the mapping from current wealth to expected wealth in the next period must be a contraction. In this case, Carroll (1999) and Deaton (1991) prove that a solution to the consumer problem exists if:

$$(7) \quad R\beta E [(GN)^{-\rho}] < 1.$$

Thus, a higher preference for the present, steeper growth of expected income, or a lower interest rate each make buffer stock behavior more likely. Under some additional conditions, one can also show that cash on hand admits an ergodic distribution.¹⁸

In a life-cycle model, no such simple characterization exists because income growth and behavior change with age. Instead, we adopt a practical approach and call a household a *buffer stock* household if it saves more for precautionary motives than for life-cycle (bequest or retirement) motives. A household whose saving decisions are motivated primarily by uncertainty will have the characteristics of buffer stock behavior just described.

3. METHOD OF SIMULATED MOMENTS ESTIMATION

We estimate our model using a two-step Method of Simulated Moments (MSM) as follows. Given the parameters of the consumer problem defined in the previous section (preferences, income process, etc.) we can solve numerically for the age-dependent optimal consumption rules. For a given set of consumption rules, we can numerically simulate the associated expected consumption as a function of age only. The estimation procedure then minimizes the distance between the simulated consumption profiles and empirical consumption profiles, which are described in section 5.

According to section 2, consumption at age t for individual i depends on normalized cash on hand $x_{i,t}$, the realization of permanent component of income $P_{i,t}$, and the parameters of the consumption problem, which we now denote by $\psi \in \Psi \subset \mathbb{R}^s$. Defining the vector of state variable $s_{i,t} = (x_{i,t}, P_{i,t})$, we postulate the following data-generating process for each age, t :

$$(8) \quad \ln C_{i,t} = \ln C_t(s_{i,t}; \psi) + \epsilon_{i,t} = \ln(c_t(x_{i,t}; \psi) P_{i,t}) + \epsilon_{i,t},$$

where $\epsilon_{i,t}$ is an idiosyncratic shock that represents classical measurement error in consumption levels. We are interested in estimating ψ and then making inference about consumption behavior. Without quality panel data on consumption, asset and income for individual households, direct estimation using (8) is not possible.

We do however observe the average of log-consumption at each age, defined as $\ln \bar{C}_t \equiv \frac{1}{I_t} \sum_{i=1}^{I_t} \ln C_{i,t}$ where I_t represents the number of observations at age t and $\ln C_{i,t}$ is defined by equation (8). This suggests that we can look directly at the *unconditional* expectation of log-consumption at each age:

$$(9) \quad \ln C_t(\psi) \equiv E[\ln C_t(s; \psi) | \psi] = \int \ln C_t(s; \psi) dF_t(s; \psi),$$

where the unconditional cumulative distribution of normalized cash on hand and permanent component of income at age t , $F_t(s; \psi)$ depends on age t and on the parameters ψ . We seek to estimate the model from the following moment conditions: $E[\zeta(\ln C_i; \psi_0)] = 0$, where ψ_0 is the true parameter vector, $\ln C_i = \{\ln C_{i,t}\}_{t=1}^T$ and $\zeta(\ln C_i, \psi) \in \mathbb{R}^T$, with t^{th} element:

$$(10) \quad \zeta_t(\ln C_i; \psi) = \ln C_{i,t} - \ln C_t(\psi).$$

This approach must address two further complications. First, it is difficult to estimate accurately all elements of ψ in one step. This is due in part to the computational complexity of the problem and in part to the loss of information that averaging entails. Instead, we employ a two-stage estimation procedure. We partition ψ into two sub-vectors $\theta \in \Theta \subset \mathbb{R}^k$ and $\chi \in \mathbb{R}^r$, where Θ is a compact set. We use additional data and moments to estimate χ in a first stage according to $E[\mu(\chi)] = 0$ where $\mu \in \mathbb{R}^r$.¹⁹ We discuss this step in the next section and denote the associated first-stage estimator by $\hat{\chi}$. For instance, rather than estimate the variance of permanent and transitory shocks to income from average consumption and income profiles, where identification might prove difficult in practice, we use time-series moment conditions and true household-level panel data on income from the PSID.

The second complication arises from the analytical complexity of the unconditional expectation (9). The unconditional distribution for the state variables at age t , $F_t(s; \psi)$, is extremely cumbersome to evaluate and depends on the parameters ψ . To overcome this difficulty, we use the *Method of Simulated Moments*, as developed by Pakes and Pollard (1989) and Duffie and Singleton (1993). Using the budget constraint (4), and the process for P_t , we define a measurable transition function $\mathfrak{T} : \mathbb{R}^2 \times \mathbb{R}^2 \times \Theta \times \mathbb{R}^r \rightarrow \mathbb{R}^2$ that describes the dynamics of the state variables

$s_{t+1} = (x_{t+1}, P_{t+1}) = \mathfrak{T}(s_t, \nu_{t+1}; \theta, \hat{\chi})$, where $\nu_{t+1} = (U_{t+1}, N_{t+1})'$. This transition function can then be used to rewrite the unconditional expectation (9):

$$(11) \quad \begin{aligned} \ln C_t(\theta, \hat{\chi}) &= \int \ln C_t(s, \theta, \hat{\chi}) dF_t(s; \theta, \hat{\chi}) \\ &= \iint \ln C_t(\mathfrak{T}(s, \nu; \theta, \hat{\chi}), \theta, \hat{\chi}) dF_{t-1}(s; \theta, \hat{\chi}) dF^\nu(\nu). \end{aligned}$$

where $F^\nu(\nu)$ denotes the cumulative distribution of ν .

From (11), we approximate the theoretical unconditional expectation using Monte-Carlo integration. We generate income shocks for L households: an $\mathbb{R}^2 \times \mathbb{R}^T$ -valued sequence of random variables $\{\hat{\nu}_l\}_{l=1}^{l=L}$ where $\hat{\nu}_l = (\hat{\nu}_{l,1}, \dots, \hat{\nu}_{l,T})'$, that are identically independently distributed. From any initial distribution $F_1(s_1)$, candidate θ and estimated $\hat{\chi}$, we can generate the path of state variables according to:

$$\hat{s}_{l,t+1} = \mathfrak{T}(\hat{s}_{l,t}, \hat{\nu}_{l,t+1}; \theta, \hat{\chi}); \quad \forall 1 \leq t \leq T-1 \text{ and } 1 \leq l \leq L.$$

The unconditional expectation from the model $\ln C_t(\theta, \hat{\chi})$ is then simulated by:

$$\ln \hat{C}_t(\theta, \hat{\chi}) \equiv \frac{1}{L} \sum_{l=1}^L \ln C_t(\hat{s}_{l,t}, \theta, \hat{\chi}) \rightsquigarrow \ln C_t(\theta, \hat{\chi}),$$

where convergence occurs as $L \rightarrow \infty$.

For any parameter vector $\theta \in \Theta$, we can replace the theoretical expectation of consumption with its simulated counterpart in (10):

$$\hat{\zeta}_t(\ln C_i; \theta, \hat{\chi}) = \ln C_{i,t} - \ln \hat{C}_t(\theta, \hat{\chi}).$$

Estimation now simply proceeds by making the *simulated* empirical moments as close as possible to their theoretical values using sample averages:

$$\begin{aligned} g_t(\theta; \hat{\chi}) &= \frac{1}{I_t} \sum_{i=1}^{I_t} \hat{\zeta}_t(\ln C_{i,t}; \theta, \hat{\chi}) \\ &= \frac{1}{I_t} \sum_{i=1}^{I_t} \ln C_{i,t} - \ln \hat{C}_t(\theta, \hat{\chi}) \\ &= \ln \bar{C}_t - \ln \hat{C}_t(\theta, \hat{\chi}). \end{aligned}$$

Note that we do not require repeated observations on the same households. Our second-stage estimation procedure is then a method of simulated moments estimator (*MSM*) that minimizes over θ :

$$g(\theta; \hat{\chi})' W g(\theta; \hat{\chi}),$$

where $g(\theta; \hat{\chi}) = (g_1, \dots, g_T)' \in \mathbb{R}^T$ and W is a $T \times T$ weighting matrix. In the case where $W = I_T$, the identity matrix, the estimation procedure is equivalent to minimizing the sum of square residuals:

$$S(\theta; \hat{\chi}) = \sum_{t=1}^T \left(\ln \bar{C}_t - \ln \hat{C}_t(\theta; \hat{\chi}) \right)^2.$$

Note however, that we are minimizing the sum of squared residuals with a constant T and that asymptotic results still apply as long as I_t , the numbers of observation at age t , goes to infinity for each age.

Under the regularity conditions stated in Pakes and Pollard (1989) and Duffie and Singleton (1993), the *MSM* estimator $\hat{\theta}$ is both consistent and asymptotically normally distributed. However, one needs to correct the second-stage estimator of the variance-covariance matrix for the first-stage estimation.²⁰ Furthermore, an additional term is required to account for the fact that we have simulated moments rather than theoretical moments.

The variance-covariance estimator and its derivation is contained in Appendix B. We use two alternative weighting matrixes in our estimation. First, we choose a matrix that, while not fully optimal, does not depend on the fitted model. This is motivated by the observation that optimally weighting GMM estimators can worsen finite-sample bias (see for example West, Wong and Anatolyev (1998)). Second, we employ the optimal weighting matrix. Our 2-step procedure provides a test of the overidentifying restrictions in the second stage, which is derived in appendix B.

In practice, we simulate $\ln \hat{C}_t(\theta, \hat{\chi})$ by running $L = 20,000$ independent income processes for 40 years, and computing in each year the associated consumption and cash on hand. Once the optimum is found, the gradient of the moment vector is evaluated numerically and the variance-covariance matrix estimated.

4. FIRST-STAGE ESTIMATION

The first-stage parameters χ include the gross real after tax interest rate R ; the variances of the innovations to the permanent and transitory components, $p, \sigma_u^2, \sigma_n^2$; the initial distribution of liquid assets at age 26, which corresponds to the beginning of life in the empirical estimation of our model; and the income and family-composition profiles. The construction of profiles is discussed in the next section.

First, we estimate the gross real interest rate, R , from the average real return on Moody's AAA municipal bonds and the Gross Domestic Product implicit price deflator for personal consumption expenditures used to deflate the consumption and income data. Over the sample period from January 1980 to March 1993, the average real interest rate is 3.440 percent with a standard error of 0.281.

Second, we estimate the variance of the permanent and transitory components of shocks to income, σ_u^2 and σ_n^2 , using the PSID and the GMM methodology of Carroll and Samwick (1997).²¹ The GMM procedure and data employed are designed to estimate the parameters of exactly the income process we have specified. We redo their calculations since the definitions of occupation in the PSID and CEX do not exactly overlap. We aggregate the occupational groups in the PSID into categories that better match the groupings that we employ in the CEX.²² Table I displays the variances of the permanent and transitory shocks across occupation and education groups. Since measurement error in income may exaggerate these variances and since households may have access to some informal insurance mechanisms we study the impact of lowering these measures in section 6.3.

We calculate the probability of a truly zero income realization using the frequency of truly zero income in the PSID, as reported by Carroll (1992). Therefore, we set $\hat{p} = 0.00302$ with a standard error of 0.000764.

Finally, we require an initial asset level for households. We match both the typical level and the distribution in the population by estimating the ratio of liquid wealth to income for households between 24 and 28 years old. In simulating profiles, we assume that households are born with an initial level of assets relative to income drawn from a lognormal distribution with mean $\bar{\omega}_{26}$ and standard deviation $\sigma_{\omega_{26}}$ that we estimate from the CEX data. Table II displays our estimates of

the average log normalized liquid wealth, its standard deviation in the data and the implied ratio of average normalized wealth by education and occupation groups.

5. CEX CONSUMPTION AND INCOME PROFILES

5.1. Profile Construction Methodology. We estimate three profiles using household level data from age 26 to 65: a consumption age-profile $\{\bar{C}_t\}_{t=26}^{65}$, the average income profile $\{\bar{Y}_t\}_{t=26}^{65}$ from which we derive expected income growth $\{G_t\}_{t=27}^{65}$; and a profile for the typical shifts in marginal utility, $\{v(Z_t)\}_{t=26}^{65}$.

Assume that the data generating process for consumption is as defined in section 2, for each household. Define $\lambda_{i,t} \equiv v(Z_{i,t}) C_{i,t}^{-\rho}$ as the marginal utility of household i of age t . From the Euler equation (5) of section 2, we know that:

$$\lambda_{i,t} = \frac{1}{\beta R} \lambda_{i,t-1} \eta_{i,t},$$

where $\eta_{i,t}$ is the (multiplicative) innovation to the marginal utility of wealth and satisfies $E_{t-1}[\eta_{i,t}] = 1$.

Iterating backward until the birth year, and substituting for $C_{i,t}$, we obtain:

$$(12) \quad C_{i,t} = \left(\frac{v(Z_{i,t})}{v(Z_{i,26})} \right)^{\frac{1}{\rho}} (\beta R)^{\frac{t-26}{\rho}} \left(\prod_{l=27}^t \eta_{i,l}^{\frac{-1}{\rho}} \right) c_{26} P_{26} \quad , \quad 26 \leq t \leq 65.$$

The first term on the right hand side represents the effect of family size variation between age 26 and age t . The second term reflects the drift in marginal utility that depends upon the discount factor, the real interest rate and the intertemporal elasticity of substitution $1/\rho$. The third term reflects the effect of uncertainty and precautionary saving upon consumption at age t , through the past and current realizations of innovations to marginal utility, $\eta_{i,t}$, due to individual income innovations. The fourth term, c_{26} , represents variations in initial cash on hand across households. The final term, P_{26} , captures variations across households in the initial permanent component of income.

The household-level data from the CEX differs from the above $C_{i,t}$ in four ways. First, consumption data is mismeasured. Thus, we do not exploit the limited panel nature of the Survey,

but instead rely on data from repeated cross-sections. We assume that measurement error in consumption is classical and multiplicative.

Second, year-specific events, such as the stage of the business cycle, affect the average consumption of households in each year. These effects are not included in the model. One possible justification is that aggregate fluctuations account for a small share of individual uncertainty (Pischke (1995)). We assume that time-effects enter multiplicatively.²³

Third, the model of section 2 does not assume any variation in income across cohorts. In reality, households start life with different levels of wealth and permanent components of income and thus with different levels of consumption. Some of these differences are due to that fact that households born in later generations on average have higher values for their initial permanent component of income. This implies that there is a correlation between the effect of birth year and the effect of age on consumption. Households observed at age sixty, say, are born long before those we observe at young ages and so have on average lower lifetime resources, and lower levels of income and consumption at each age. Ignoring birth-year effects would lead to a negative bias in our estimate of the slope of income and consumption growth. We assume, as in much of the empirical literature that earnings have a time invariant age-profile so that cohort effects affect only the distance between the age-profiles of different cohorts. We decompose $P_{i,26}$ into a cohort effect, P_{26}^b , and an idiosyncratic component $\tilde{P}_{i,26}$.

Finally, the model refers to household consumption, adjusted so that all households have the same life-cycle pattern of family size $v(Z_t)$. We therefore estimate family size adjustments and apply them to the consumption data so as to control for within-age demographic variation.²⁴

Define $\tilde{C}_{i,t,\tau}$, as the observed household consumption in the CEX, where τ denotes the year. Taking logs, we can write $\ln \tilde{C}_{i,t,\tau}$ as:

$$(13) \quad \ln \tilde{C}_{i,t,\tau} = \frac{1}{\rho} \ln \left(\frac{v(Z_{i,t})}{v(Z_{i,26})} \right) + \frac{1}{\rho} E_{26} \left[\ln \left((\beta R)^{t-26} \left(\prod_{l=27}^t \eta_{i,l}^{-1} \right) \right) \right] + \ln \left(P_{26}^b \right) + \xi_\tau + \epsilon_i$$

where ξ_τ represents a time-effect and ϵ_i captures classical measurement error as well as household heterogeneity in initial conditions and innovations to marginal utility. The second term, captures

the amount of intertemporal substitution and expected precautionary saving and is just a function of age.

We estimate directly a linear regression based on (13) over households with male heads aged 26 to 65. As discussed in Deaton (1985), it is not possible to separately identify the linear component of the time, age, and cohort effects in equation (12).²⁵ We make the identifying restriction that time effects are due to the state of the regional economy and captured by the partial correlation of consumption with the regional unemployment rate.²⁶ Our empirical specification is then:

$$(14) \quad \ln \tilde{C}_i = f_i \pi_1 + a_i \pi_2 + b_i \pi_3 + \mathcal{U}_i \pi_4 + Ret_i \pi_5 + \varepsilon_i$$

where the π parameters are to be estimated, f_i is a set of family dummies, a_i is a complete set of age dummies, b_i is a complete set of cohort dummies (less the middle one), $\mathcal{U}_i$ is the Census region unemployment rate in year τ , Ret_i is a dummy for each group that is equal to 1 when the respondent is retired and ε_i is a residual that captures all the individual effects (measurement error, initial normalized cash on hand, initial income, etc.).²⁷

With these estimates, we first reconstruct household-level consumption data uncontaminated by cohort and time effects and removing within age variation in family size:

$$(15) \quad \ln C_i \equiv \bar{f}_i \hat{\pi}_1 + a_i \hat{\pi}_2 + \bar{\mathcal{U}} \hat{\pi}_4 + \hat{\varepsilon}_i.$$

Thus from here on, we work with household-level consumption that represents the consumption of the observed household with the typical age-dependent family size $\bar{f}_t$, facing the average unemployment rate $\bar{\mathcal{U}}$, born in the middle cohort, and not retired.²⁸ It is this version of consumption to which section 3 refers.

We construct average age-profiles of consumption by averaging this data across households. So that the role of precautionary saving is highlighted, we display most of our profiles and simulations for consumption per capita and income per capita. This profile is constructed by replacing f_t with $\bar{f}$:

$$\ln C_a \equiv \bar{f} \hat{\pi}_1 + a \hat{\pi}_2 + \bar{\mathcal{U}} \hat{\pi}_4.$$

We also construct three other sets of profiles. First, and in the same manner, we construct profiles for income and the typical household family size. Second, we build smooth profiles using fifth-order polynomials in age and year of birth instead of dummies. Lastly, we estimate similar profiles separately for various occupation and education subgroups of the population. The construction of these series is similar to that of consumption, and described in Appendix C.

The smoothed profiles for income and family size are used as inputs of the model. For instance, recall that $\ln Y_{i,t} = \ln P_{i,t} + \ln U_{i,t} = \ln G_t + \ln P_{i,t-1} + \ln N_{i,t} + \ln U_{i,t}$. After removing cohort and time effects, and averaging over a large number of households, I , with the same characteristics:

$$\frac{1}{I} \sum_{i=1}^I \ln Y_{i,t} = \ln G_t + \frac{1}{I} \sum_{i=1}^I \ln Y_{i,t-1} + \frac{1}{I} \sum_{i=1}^I \ln N_{i,t} + \frac{1}{I} \sum_{i=1}^I \ln U_{i,t} - \frac{1}{I} \sum_{i=1}^I \ln U_{i,t-1}$$

Applying the Law of Large Numbers, the probability limits of the last three term are all zero. Hence, we get

$$\text{plim} \left(\frac{1}{I} \sum_{i=1}^I \ln Y_{i,t} - \frac{1}{I} \sum_{i=1}^I \ln Y_{i,t-1} \right) = \ln G_t \equiv \ln \bar{Y}_t - \ln \bar{Y}_{t-1}.$$

Thus first differencing our log-average income levels gives expected income growth rates G_t .

5.2. The Consumer Expenditure Survey and Our Use of it. We use the Consumer Expenditure Survey (*CEX*) to construct life-cycle profiles of consumption and income. The *CEX* contains information about consumption expenditures, demographics, income and assets, for a large sample of the US population. The Survey is conducted by the Bureau of Labor Statistics in order to construct baskets of goods for use in the bases for the Consumer Price Index, and has been run continuously since 1980. We use data from 1980 to 1993 from the family, member, and detailed expenditure files. The survey is known to have excellent coverage of consumption expenditures, to have reasonable data on liquid assets, and to have income information of moderate quality.²⁹ The survey interviews about 5500 households each quarter. In a household's first interview, the *CEX* procedures are explained to them and information is collected so that they can be assigned a population weight. They are then interviewed four more times (once every three months) about detailed consumption expenditures over the previous three months. In interviews two and five, income information is collected, and in the final interview asset information is collected. Families

rotate through the process, so that about 25% of households leave and are replaced in each quarter. About half of all households make it through all the interviews.

Each household contributes one data point to our sample. For each household we construct a measure of household income and consumption. Based on the characteristics of the male head of the household, we assign the household to an occupation group, an education group, a birth cohort, an interview year, and a Census region. In order to obtain a high quality sample that has the required information, we drop a significant portion of the data and make a series of adjustments. Further description of the data preparation is contained in Appendix D. We note here three major points.

First, we drop households that are classified as incomplete income reporters, that have any of the crucial variables missing, or that report changes in age over the course of the survey greater than one year or negative. We do not analyze the group of households with male heads holding less than 9 years of schooling due to very few younger households. Second, we drop all households with male heads younger than 26 or older than 70, given our focus on the working life. We are left with just under 40,000 households. Third, we use individual as well as family level information to correct most of the top-coding in household labor income.

We construct measures of income and consumption that match the concepts in the theoretical model. First, we define consumption as total household expenditures less those on education, medical care, and mortgage interest payments.³⁰ These categories of expenditure do not provide current utility but rather are either illiquid investments or negative income shocks. So that these expenditures do not incorrectly appear as liquid saving, they are also subtracted from income.

It should be noted that our model refers to total consumption at annual frequencies. In the data, we are averaging total expenditures, not consumption, across a large number of individuals and looking across one-year horizons. Hence, we suspect the distinction between durables and nondurables is not likely to matter much. Further, since the buffer stock model gives strong predictions about consumption tracking income, it is important not to break the consumption-income link when studying consumption. We will also report the results when consumption is restricted to nondurables and income is defined as disposable income after durable expenditures.

Our measure of disposable labor income is comprised of after-tax family income less Social Security tax payments, pension contributions, after-tax asset and interest income and, as noted,

those expenditures subtracted from consumption. The first two adjustments are saving in illiquid form and so are available to the household only after retirement. We remove asset income since the input to our theoretical model is a profile of income net of liquid asset returns. Consistent with the spirit of our model, all items removed from income involve a large amount of commitment and are hard to substitute intertemporally.

Finally, we put all data into real 1987 dollars using the Gross Domestic Product implicit price deflator for personal consumption expenditures.³¹

5.3. Life Cycle Profiles. Figure 2 presents consumption (raw and smoothed) and income profiles for our entire sample when the family-size is held constant over the life-cycle. Even after correcting for the effects of cohort, time, and family, both profiles are still hump shaped and track each other early in life. Consumption lies above income over the late twenties. Given that the CEX wealth data, and better household wealth surveys, show modest increases in liquid wealth over these ranges, this feature seems likely due to misreporting of income or consumption. One possibility is underreporting the assistance that is provided by intergenerational transfers early in life. After these first few years, consumption rises with income from age 30 to age 45, when consumption drops significantly below income. This tracking is however a lot less than is observed in profiles constructed by simply averaging cross-sections because we control for changes in family size and cohorts effects.

Figures 3 and 4 give some evidence that consumption and income track each other across subgroups of the population defined by education and occupation groups. These graphs are somewhat noisy. However, despite the noise, one can see that the occupation and education groups with the most pronounced humps in income present the most pronounced humps in consumption. Further, we can formally reject the null hypothesis that the consumption profiles are flat. This is essentially a now standard test of the linearized consumption Euler equation, as studied by Attanasio and Weber (1995), Lusardi (1996).

Our profiles differ slightly from the results of Attanasio and Browning (1995) and Attanasio and Weber (1995). These papers employ a larger set of preference shifters: once controlling for these, consumption is smoother and the CEQ-LCH is not rejected. In Attanasio and Weber (1995) and in the linear Euler-equation approach generally used in micro data, precautionary effects are

omitted so that preference shifters absorb, correctly or incorrectly, variation in consumption that we attribute to uncertainty. Clearly, allowing for enough preference variation can explain the life-cycle correlation without reference to substantial household-level uncertainty. The approach of this paper is to adjust only for the size of the household and examine how precautionary saving explains the observed consumption behavior. Our methodology has the potential to nest and test both explanations. It is also worth noting that our profiles represent total consumption expenditures rather than non-durable consumption.

6. RESULTS

We first estimate our model for the average household and discuss the implications of the fitted structural model for household behavior. We then turn to disaggregated results, by education and occupation groups, and to evaluating the robustness of our findings.

6.1. Findings from the Entire Sample. As a useful benchmark, we begin by asking what the standard Life-Cycle theory would predict, assuming away all uncertainty. To give the best chance to the CEQ-LCH, we perform first difference estimation, not asking it to fit the mean of the consumption profile. Under certainty, equation (2) holds, implying, after controlling for individual characteristics, a constant growth rate of consumption over the working period:

$$\Delta \ln \bar{C}_t = \frac{1}{\rho} \ln(\beta R).$$

We estimate $\frac{1}{\rho} \ln(\beta R)$ from the coefficient on age in a least-squares regression of the profile of consumption with family size, cohort, and time effects removed, as displayed in Figure 2. This procedure seems trivial only because of our earlier efforts to remove changing family-size and cohort effects. It is precisely this simplicity that gives the CEQ-LCH its power. From our estimate, we use the delta method to recover the discount factor and its standard error, using the real interest rate estimated in the previous section and a choice of the coefficient of relative risk aversion of unity.³² We estimate a discount rate of 3.44% with a standard error of 0.281%, after adjusting the estimate for both first order serial correlation and first-stage uncertainty. Since the CEQ-LCH consumption profile is flat and the data as shown in Figure 2 is hump-shaped, the certainty model performs poorly when it comes to explaining the dynamics of consumption across the life cycle.

Table III presents the results of estimating our structural model with income uncertainty. The associated consumption rules are displayed in Figure 1, panel A. The first column of Table III displays the results when the initial, robust weighting matrix is employed; the second displays the results with the optimal weighting matrix. The standard errors are calculated without (A) and with (B) the adjustment for uncertainty in the parameters $\hat{R}$, $\hat{p}$, $\hat{\sigma}_u^2$, $\hat{\sigma}_n^2$ and $\hat{\omega}_{26}$ from the first-stage estimation.³³ The standard errors typically increase, once the uncertainty in the first-stage parameters is accounted for.

The discount rate is estimated at just over four percent, which is within conventional significance levels of the real interest rate of 3.44 percent. It is worth noting that the discount rate is within a reasonable range. Using information on the elasticity of assets with respect to uncertainty, Carroll and Samwick (1997) estimate a discount rate in the vicinity of 10–15% and argue that even higher discount rates are needed to rationalize the findings of Hubbard, Skinner, and Zeldes (1994). Our lower discount rate, however, does not imply that households are not impatient enough to generate buffer stock behavior, as we will show shortly.

When we employ an optimal weighting matrix, the precision with which we estimate the discount rate declines significantly compared with the variance estimate that does not correct for uncertainty about the first stage parameters (A), but increases compared to the 2-step variance estimate (B). There is a growing literature that questions the small-sample validity of optimal weighting due to the correlation between parameter uncertainty and the weighting matrix. Optimal weighting can be more efficient; it can also be more biased. Thus we report both.

The intertemporal elasticity of substitution is estimated quite tightly, and is the sole parameter estimate that depends significantly on the weighting matrix employed. The estimated retirement rule suggests a marginal propensity to consume out of wealth at retirement (γ_1) between 6 and 7 percent, also quite reasonable. For instance, in the case of full certainty after retirement and no change in the utility shifter the marginal propensity to consume is given by

$$\left(1 - \hat{\beta}^{1/\hat{\rho}} \hat{R}^{1/\hat{\rho}-1}\right) / \left(1 - \left(\hat{\beta}^{1/\hat{\rho}} \hat{R}^{1/\hat{\rho}-1}\right)^{N-T}\right) = 7.05 \text{ percent},$$

given our estimates of β, ρ and R and setting death at age 88. Thus the estimate is very much in line with simple predictions of the model.

Finally, under our assumptions, the ratio γ_0/γ_1 provides an estimate of the ratio of illiquid wealth to the permanent component of income at retirement. The point estimate is extremely small, around 2%. The first thing to note is that this ratio is imprecisely estimated and we cannot reject more reasonable values. An alternative interpretation for a low γ_0 is that households have a low propensity to consume out of illiquid wealth at retirement.³⁴

It should be noted that both estimation methods reject the overidentifying restrictions at the 5% level. The 95% critical value for a χ^2 (36) is 50.71 and the chi-square always exceeds 150. This is not entirely surprising, given the number of moments we use (40) and the few parameters of the model. The estimated model should still be taken seriously however. As we now discuss, the model does much better in an economic sense than the CEQ-LCH model with which this section begins.

With our estimates in hand, we can address how well the stochastic model fits the life-cycle consumption profile. The first panel of Figure 5 plots the simulated and actual consumption data along with the income profile. The stochastic life-cycle model does a much better job at fitting the consumption profile than the consumption profile with constant growth rate of $\frac{1}{\rho} \ln(\beta R)$ that would obtain under the certainty-equivalent. The consumption profile from the fitted model tracks income until around age 40 and then falls as households start increasing their liquid wealth holding in preparation for retirement.

In two places however, the model fit is not good. First, actual consumption exceeds simulated consumption early in life. Since, as we discussed in the preceding section, in the data consumption exceeds income early in life, this may reflect a weakness of the data, rather than the model. Second, the actual consumption profile is slightly flatter and peaks slightly later. Apart from these features however, the tight structure imposed by the model produces good predictions in terms of consumption dynamics.

Why are we able, within the context of our model to obtain such tight estimates? The second panel of Figure 5 plots various simulated profiles for values of β , one percent away from the point estimate. This corresponds to choices of $\beta = 0.9498$ and $\beta = 0.9698$, or equivalently, to a discount rate between 4.18 and 5.28 percent. It is obvious that the profiles are very sensitive to small variations in the discount factor, holding ρ constant. With a higher discount factor, the household is willing to save more and earlier for retirement purposes. The consumption path exhibits less

of a hump shape, and may even be increasing over the entire working life, as is the case in this Figure. On the other hand, for more impatient households, consumption parallels income until much later in life and then falls more precipitously to build assets for retirement. This implies a stronger concavity of the consumption profile. Thus, our method yields tight estimates of the discount factor because the discount factor is an important determinant of the hump shape in the consumption profile.

We now turn to the question of how household behavior changes over the life cycle. Define the target level of cash on hand at age t , $\bar{x}_t$, as the level at which cash on hand is expected to remain unchanged from one year to the next (Carroll (1997)):

$$\bar{x}_t = E_t [x_{t+1} | x_t = \bar{x}_t]$$

For large initial levels of cash on hand ($x_t > \bar{x}_t$), households choose high levels of consumption and thus future levels of cash on hand are on average lower. For low levels of cash on hand ($x_t < \bar{x}_t$), households consume less than the income they expect to receive and so on average accumulate cash on hand.

Figure 6 presents the target level of cash on hand for consumers aged 26 – 42, from the fitted model. One can see from the graph that the target level of cash on hand –including current income– remains small early in life, around 1.2 times the permanent component of income. Around age 40, this target increases substantially, as consumers begin to build their wealth for retirement. This Figure shows a dramatic change in behavior. When the target level of liquid wealth is small, households behave as ‘buffer stock’ consumers. Good income shocks are consumed away and bad income shocks are imperfectly smoothed. Household consumption closely follows household income. Starting around age 38, agents begin to accumulate assets for retirement. As households move into their forties, they build significant amounts of liquid wealth. This retirement wealth allows households to smooth high frequency movements in income so that their behavior more closely mimics that of certainty-equivalent consumers.

We next decompose total saving and wealth at each age into that driven by life-cycle considerations and that additional amount driven by the presence of uninsurable risk. Our previous discussion might lead the reader to think that agents have no concern for retirement when they are

young and no concern for labor income uncertainty later in life. This is incorrect since consumers are rational and perfectly foresee their retirement needs. First, we define saving in liquid wealth as the discounted variation in financial wealth from one period to the next:³⁵

$$S_{i,t} = (W_{i,t+1} - W_{i,t}) / R = (R - 1) / R W_{i,t} + Y_{i,t} - C_{i,t}.$$

Saving is equal to investment income –liquid and illiquid– plus labor income minus consumption. From our empirical profiles in Figure 5, it follows that households save relatively little and consume roughly their income on average early in life. Second, at the estimated parameters, we compute the consumption path, $\{C_t^{LC}\}$, that would occur if all income risk were pooled, so that for all households $Y_t^{LC} = E_{26}[Y_t]$, but the household’s environment otherwise remains unchanged.³⁶ Finally, we define life-cycle saving as the difference between total income and life-cycle consumption:

$$S_t^{LC} = (W_{t+1}^{LC} - W_t^{LC}) / R = (R - 1) / R W_t^{LC} + Y_t^{LC} - C_t^{LC}$$

Precautionary saving is the complement of life-cycle saving.

The first panel of Figure 7 plots the precautionary and life cycle-liquid saving of the average household. Given the estimated discount rate and the profile of expected income, young consumers facing no income risk would like to borrow large amounts, so life-cycle saving is negative early in life. Young households in fact hold a positive buffer stock of wealth in response to income risk, so that precautionary saving is positive early in life. In the early to mid forties, in accordance with our previous discussion, life-cycle saving becomes larger than precautionary saving. Households begin to build their liquid wealth for retirement purposes. As asset levels increase, the expected variance of consumption declines, decreasing the precautionary saving motive. Since households that face income uncertainty save more early in life due to risk, they are able to consume more and save less when older, leading to negative precautionary saving late in life. This discussion highlights the importance of the income profile. Absent income growth, households would start saving for retirement early in life. This liquid wealth would double as a buffer stock. Thus there would be no additional precautionary saving motive and household behavior would be well approximated by the CEQ-LCH.

Panel B of Figure 7 displays average total, life-cycle, and precautionary wealth, obtained by integrating the corresponding saving profiles. We observe first that total wealth always exceeds life cycle wealth, reflecting the presence of uninsurable income risk. Second, life-cycle wealth is negative early in life, representing the desire of households facing no risk to borrow as just discussed. Third, since young households actually hold small positive amounts of wealth on average, precautionary wealth early in life is substantial, peaking at \$110,000 or 3.6 times current income by age 42. Finally, as households get older, life-cycle wealth becomes very substantial, reaching \$229,000 or 6.7 times current income at the age of retirement, while precautionary wealth slowly declines, leaving total wealth at \$312,000, or about 9 times income at retirement.³⁷

An important caveat to this exercise is that our estimation does not try to match the model to wealth data, and household accumulation of liquid wealth is larger than observed liquid asset wealth at retirement. According to the 1998 Survey of Consumer Finances, mean total (excluding Social Security wealth but including pension and housing) wealth of 55 – 64 year old households is 7.4 times total pre-tax income, where our model predicts a factor of 9 times, excluding pension wealth.³⁸ This feature of the data is the same feature that drives the finding of low γ_0 .

Finally, we apply the same partial-equilibrium decomposition to aggregate wealth per effective household in our economy. We assume that there is no uncertainty after retirement and certain death at age 87. According to our definition, precautionary wealth represents 62 percent of all wealth holding with no population growth. This share rises to 71 percent under the assumption that the population grows at one percent per year. These shares are reasonable close to Caballero (1991), which reports findings for a calibrated CARA model, and Cagetti (1998), which employs our methodology to fit wealth profiles.³⁹

6.2. Findings from Different Education and Occupation Groups. We next estimate the model separately on subgroups of the population defined by education levels and occupation. This allows us to exploit variation in the environments of different agents and allows different subgroups of the population to have different preferences. Technically, we simply follow the procedure described above on each cell, using first-stage estimates already presented by subgroup in sections 4 and 5. Each group’s consumption profile is matched with different estimates using a different income profile (Figures 3 and 4), income uncertainty (Table I), and initial average log liquid wealth

(Table II). The results are summarized in Table IV and the simulated consumption profiles from the fitted model are presented in Figures 3 and 4.

The estimated discount factors are close to that obtained using the aggregate profile. The discount rate lies between 3.94 percent (Some High School) and 5.71 percent (Managerial and Professional). Interestingly, there is no clear pattern associated with the level of education. If anything, the discount rate is slightly higher for more highly educated households. By contrast, there is much wider dispersion in the estimated coefficients of relative risk aversion, from a low of 0.282 (Some High School) to a high value of 2.290 (Graduate School). The coefficients of the consumption rule at retirement are also within a reasonable range of the aggregate estimates. The marginal propensities to consume out of liquid wealth range between 5 and 7% despite being imprecisely estimated. On the other hand, we find extremely low and imprecise estimates of γ_0 for most cells, implying, as before, either an unreasonably low ratio of illiquid wealth to permanent income at retirement, or a very low marginal propensity to consume out of illiquid wealth.

Perhaps more interestingly, the overidentification tests statistics are much smaller, indicating that the model does a better job at fitting the actual consumption profiles. Indeed, looking back at Figure 3 and Figure 4, we see that the fitted consumption profiles are very much in line with the unconstrained smoothed consumption profile estimated from the data. For all cells, the consumption profile is hump-shaped and the simulated profiles indicate that households start accumulating for retirement purposes between 35 and 40 years. The behavioral implications of this set of results are the same as those just discussed for estimation that did not allow for cross-group heterogeneity.

6.3. Robustness. To provide further evidence on the strengths and weaknesses of the model, we present the results of five alternative procedures: (i) estimating the initial level of liquid wealth in the second stage; (ii) estimating the retirement consumption rule directly from micro data in the first stage; (iii) estimating by matching age-specific growth rates rather than levels; (iv) adjusting the second-stage inference for income and family profile estimation; and (v) reducing the variance of the income shocks to account for measurement error, (vi) matching the consumption profile for expenditures on non-durable goods only.

First, we estimate in the second stage the average initial cash on hand $\bar{\omega}_{26}$ with which households begin life. While the model fits the data better with this extra parameter, the second stage no longer

provides precise or sensible identification of the structural parameters. The results are presented in columns 1 and 2 in Table V and in Figure 8. This alternative simulated profile does a better job than the baseline model in terms of fitting household consumption during the first 5 years of life, when measured consumption lies above measured income. The procedure attains this by assigning an implausibly high value of initial cash on hand: the mean level of assets is estimated to be 12.6 times current income!⁴⁰ With high initial wealth, the typical household can easily buffer labor income shocks and enjoy high consumption early in life. To fit the observed high correlation between income and consumption prior to middle age, the estimation procedure guesses that agents are very impatient (so that they run this initial wealth down) and are quite unwilling to substitute inter-temporally. The discount rate is estimated as 14% and the coefficient of relative risk aversion at 5.2.

While it is possible that this high wealth comes from inheritances, this is at odds with all household surveys of wealth holding. In our minds, this is enough to reject this alternative scenario, however it informs us about the model. From this and other experiments, we conclude that identification hinges in part on providing the second-stage with more information than is contained in the profiles alone. In particular, from the typical consumption profile, one can accurately and sensibly estimate the preference parameters only by fixing either the retirement consumption rule or the initial value of assets. Without these parameters estimated from other sources in a first stage there is not enough information in the profiles to provide accurate estimation of the model.

In our second evaluation of the model, we estimate γ_0 and γ_1 in the first stage from consumption and asset measures around the age of retirement. We use the PSID to construct measures of liquid assets in 1989, income from 1989 to 1994, and total consumption from 1989 to 1994, from active saving, wealth and income measures. We then annualize the data and estimate a crude consumption function as:

$$C_i = \gamma_0 P_i + \gamma_1 X_i + \gamma_2 f_i + \varepsilon_i$$

across households. We obtain:

$$\hat{\gamma}_0 = 0.594 \quad s.e. (0.158); \quad \hat{\gamma}_1 = 0.0774 \quad s.e. (0.0125).$$

This estimated marginal propensity to consume at retirement, 0.0774, is very close to the one in our baseline results 0.071, while the ratio of illiquid wealth to the permanent component of income, at retirement, is equal to $\frac{\gamma_0}{\gamma_1} = 7.67$, a more realistic number.⁴¹

The third column of Table V shows that the estimated discount rate is slightly lower than in the baseline estimation and the coefficient of relative risk aversion is estimated to be quite low. While the parameters appear very tightly estimated, the overidentification restriction increases dramatically. The chi square value is 1690. As shown in Figure 8, the simulated profiles are substantially higher than the actual profile in most years. This accounts for the overidentification rejection. Given the above parameters, households receive 7 times their permanent income at retirement. This decreases saving in liquid assets over the working life. This alternative scenario highlights the difficulty the model has in fitting both the large amount of liquid saving implied by the CEX profiles and the implicit large amount of illiquid wealth made available at retirement.

Since there may be some problems in the CEX measuring the relative levels of consumption and income, we next estimate the model in first-differences, so as to try to match the growth of consumption over the life-cycle, not its level. This eliminates one moment condition. The resulting parameter estimates (column (4) of Table V and Figure 8) are substantively and statistically similar to our benchmark results. More importantly, the test of overidentification does not reject the model. Since it is likely that all household consumption expenditures are not measured in the CEX, this mismeasurement may well be the reason that the data reject the baseline model.

Fourth, we adjust the standard deviations of the second-stage estimates to account for the statistical uncertainty in the estimation of the profiles of income and utility shifters, as estimated in section 5. Formally, the second-step correction employs the variance-covariance matrix for the fifth-order polynomial coefficients from the smoothed income and utility shifter profiles. The results are reported in column (5) of Table V and demonstrate extremely large standard errors. That is, given the sampling uncertainty with which a fifth-order polynomial estimates the income profile, the consumption profile does not provide an accurate measure of preference parameters. Since the age-pattern of income determines the age-pattern of consumption, the income profile is important in identifying the consumption profile.

Fifth, we address the possibility that our measures of uncertainty overstate the true risk faced by households. If there were mismeasurement in labor income, this would bias upwards our estimates of income uncertainty. With serially uncorrelated measurement error, mismeasurement inflates only our measure of the variance of transitory income. Serially correlated mismeasurement could however bias both of our measures. Additionally, households may have access to informal mechanisms that mitigate the impact of income risk on consumption. If we have overestimated permanent and transitory income uncertainty, the reader might be concerned that this is generating some of our results, particularly our low estimate of risk aversion.⁴² A reasonable guess might be that roughly a third of the variance of measured income growth is due to mismeasurement and that most of this is transitory (see Bound and Krueger (1991) and Bound (1994)). However, given our concern about unmodeled sources of insurance, we experiment with an extreme alternative. We re-estimate our model halving the variance of transitory income shocks and reducing by one third the variance of permanent shocks. Column (6) of Table V reports our findings. Most parameter estimates are unchanged with a tightly estimated discount rate of 4.1% and a relative risk aversion of 0.48, close to the benchmark estimate of 0.51. Thus, our procedure estimates a relatively low risk aversion even for substantially lower estimates of individual income uncertainty.

Lastly, we estimate the same set of parameters using information on nondurable expenditures only. Strictly speaking, our model refers to total consumption, the sum of expenditures on non-durables and the flow consumption on durables. However, since households do not rent all durables, a more complete model would accommodate both sorts of goods separately. We stop short of doing so and investigate here what our model implies for nondurable expenditures assuming durable expenditures are mandated expenditures and subtracting them from our measure of income and consumption.⁴³ Since we are altering the level of the profiles, we also estimate the model in first-differences and adjust the parameters of the consumption rule at retirement so as to match the ratio of nondurable to total consumption at retirement. The results (column (7)) indicate a substantially lower coefficient of relative risk aversion, at 0.15. The discount rate is unchanged at 3.84%. Neither parameter is estimated very tightly which accounts for the relatively low value of the overidentification test which does not reject the model.

7. CONCLUSION

This article contributes to the analysis and understanding of household consumption and saving behavior.

We develop a new method for estimating household consumption behavior. We model consumer behavior in the presence of realistic levels of uninsurable income uncertainty and estimate preference parameters and household consumption behavior using the Method of Simulated Moments. The model fits well and yields tight estimates of the discount rate and intertemporal elasticity of substitution. This methodology is now being used to estimate expanded models that address portfolio choice, labor supply and retirement behavior (Gakidis (1998), Cagetti (1998) and French (1998)).

Our results indicate that small holding of liquid assets by young households is an optimal response to expected income growth and the riskiness of future labor income over the life cycle. Until their early forties, household consumption behavior, while fully optimal, appears short-sighted within the context of the CEQ-LCH. Despite this apparent impatience, we estimate that households discount the future at modest rates and are not particularly risk averse. Our results also imply that older households save actively for retirement purposes and behave in a manner more consistent with the CEQ-LCH. These two phases of consumer behavior turn out to be quite distinct and are at the heart of our identification procedure. The neoclassical representative-agent model of aggregate consumption is incorrect precisely because of this changing behavior over the working life.

We see two interesting avenues for future research. First, our method could be fruitfully extended to richer representations of household intertemporal choices. We suspect there is much to be learned from the study of portfolio choice, durable goods, retirement, labor supply, or the effect of tax policy and social insurance programs. Second, our approach could be applied to estimate the aggregate implications of micro heterogeneity. For instance, estimated age-variation in consumption behavior that we find in the data may explain Campbell and Mankiw (1989) finding that roughly 40% of all agents are ‘hand to mouth’.

Dept. of Economics, Princeton University, Fisher Hall, Princeton, NJ 08544-1013, U.S.A.;
pog@Princeton.edu, <http://www.Princeton.edu/~pog/>,

and

Dept. of Economics and Woodrow Wilson School, Princeton University, Roberston Hall, Princeton, NJ 08544-1013, U.S.A.; jparker@Princeton.edu, <http://www.Princeton.edu/~jparker/>.

Appendix

A. Numerical Solution to the Consumer Problem

This appendix describes our approach to solving numerically the consumer problem.

A.1 Gauss-Hermite quadrature

The algorithm exploits the recursive structure of the consumer problem by solving the Euler equation for the optimal consumption rule. Assume for the time being that we know how to compute $c_{t+1}(\cdot)$ for all possible values of cash on hand. One can rewrite the Euler equation (5) using the Intertemporal budget constraint (4) as:

$$(16) \quad u'(c_t(x)) = \beta R \frac{v(Z_{t+1})}{v(Z_t)} \left(p E \left[u' \left(c_{t+1} \left((x - c_t) \frac{R}{G_{t+1}N} \right) G_{t+1}N \right) \right] + \right. \\ \left. (1 - p) E \left[u' \left(c_{t+1} \left((x - c_t) \frac{R}{G_{t+1}N} + U \right) G_{t+1}N \right) \middle| U > 0 \right] \right).$$

The first problem consists in evaluating the expectation in (16). Since N and U are log normally distributed, a natural way to evaluate these integrals is to perform a two-dimensional Gauss-Hermite quadrature:

$$\begin{aligned} E [u'(c_{t+1}(x_{t+1}) G_{t+1}N) | U > 0] &= \int u'(c_{t+1}(x_{t+1}) G_{t+1}N) dF(N) dF(U) \\ &= \int_{-\infty}^{\infty} f_t(n, u) e^{-n^2} e^{-u^2} du dn \\ &\approx \sum_{i,j} f_t(n_i, u_j) \omega_{ij}, \end{aligned}$$

where $f_t(n, u) = \frac{1}{\pi} u' \left(c_{t+1} \left((x - c_t) \frac{R}{G_{t+1}} e^{\sqrt{2}\sigma_n n} + e^{-\sqrt{2}\sigma_u u} \right) G_{t+1} e^{-\sqrt{2}\sigma_n n} \right)$. The weights ω_{ij} and nodes n_i, u_j are tabulated in Judd (1998). In practice, we performed a quadrature of order 12.

One can then find the root of the Euler equation at any point x using standard methods. We constrain the root to be positive and less than x , the current value of cash on hand. As discussed in the paper, this restriction is always satisfied when there are no illiquid assets.

A.2 Consumption rules

We initialize the algorithm with the consumption rule at retirement. One can show that the consumption rules for this problem are continuously differentiable as long as there are no liquidity constraints. However, in the presence of liquidity constraints, the consumption rules may exhibit

a kink. See Deaton (1991) and Ayagari (1994). We effectively impose a liquidity constraint by not allowing the household to borrow against illiquid assets. This indicates that smooth approximation methods are inappropriate. Instead, we use a standard discretization method: we specify an exogenous grid for cash-on-hand: $\{x^j\}_{j=1}^J \subset [0, x^{\max}]$. In order to capture the curvature of the consumption rule at low values of cash on hand, the grid will be finer for $x \in [0, x^{int}]$. In practice, for each value of cash on hand on the grid, x^j , we find the associated consumption, c^j , that satisfies (16). In choosing the size and coarseness of the grid, we face the usual trade-off between precision and computing time. Adding points on the grid gives a finer approximation of the consumption rules. Since the consumption rule at age $t + 1$ is the input necessary to get the consumption rule at age t , imprecisions could compound over time. On the other hand, the Euler equation is the innermost loop of the entire algorithm. With 100 points on the grid and 40 time periods, we must solve (16) 4000 times. This takes approximately 1.5 minutes on a Sun Sparc Enterprise10000. We also face a decision regarding the range of cash on hand, $x^{\max}$. For small values, cash on hand in sample is likely to move out of the grid. Consumption will then be evaluated using extrapolation methods, much less precise than interpolation. On the other hand, increasing the range for a fixed number of grid points implies less precise estimation of the curvature. One solution consists in endogenizing the grid so that, for instance, cash on hand remains within the grid with probability 0.95. We adopted a simpler approach consisting in checking that cash on hand, in the simulations, remains strictly inside the grid. In practice, we took $x^{\max} = 40$, $x^{int} = 2$ and $J = 100$, with 50 points between 0 and x^{int} . We checked the quality of the approximation by solving the stationary infinite horizon problem and checking the rate of convergence to the fixed point of the functional Bellman equation. Note that even if simulated x remains strictly inside the grid approximation errors can arise from the Euler equation since evaluating expected future marginal utility may require looking outside the grid.

A.3 Simulations

We simulate consumption profiles by generating a sequence of 20000 fictitious income processes over 40 years. For given parameters of the consumption problem, we solve for the optimal age-dependent consumption rules as explained above. Lastly, we calculate the consumption decisions

of households receiving these income processes. We then construct life cycle consumption profiles by log-averaging across age.

B. Asymptotic Variance Covariance Matrix and Test of Overidentifying Restrictions

We follow Newey and McFadden (1994) and use an expansion method. Define the first-stage sample moments that correspond to the theoretical moments as $m(\chi) = \frac{1}{J} \sum_{j=1}^J \mu_j(\chi)$, where J is the number of observations for the first stage. $\hat{\chi}$ is asymptotically normally distributed with a theoretical variance covariance matrix $V_\chi \equiv M_\chi^{-1} \Omega_m M_\chi^{-1'}$ where $\Omega_m = E[\mu(\chi_0)\mu(\chi_0)']$ and $M_\chi = E\left[\frac{\partial \mu(\chi_0)}{\partial \chi'}\right]$. Consistent estimation of V_χ uses the empirical counterparts to M_χ and Ω_m as is standard in *GMM*. Since the first-step is exactly identified, the choice of a weighting matrix is irrelevant here. $\hat{\chi}$ is an asymptotically linear estimator with influence function $\psi_j = -M_\chi^{-1} \mu_j(\chi_0)$ where $\sqrt{J}(\hat{\chi} - \chi_0) = \frac{1}{\sqrt{J}} \sum_{j=1}^J \psi_j + o_p(1)$.

The first order condition for the second stage estimator is:

$$g_\theta'(\hat{\theta}, \hat{\chi}) W g(\hat{\theta}, \hat{\chi}) = 0,$$

where $g_\theta(\hat{\theta}, \hat{\chi}) = \frac{\partial g(\hat{\theta}, \hat{\chi})}{\partial \theta'}$. Expanding $g(\hat{\theta}, \hat{\chi})$ around θ_0 and rearranging yields:

$$\sqrt{I}(\hat{\theta} - \theta_0) = - \left(g_\theta(\hat{\theta}; \hat{\chi})' W g_\theta(\bar{\theta}; \hat{\chi}) \right)^{-1} g_\theta(\hat{\theta}; \hat{\chi})' W \sqrt{I} g(\theta_0; \hat{\chi}),$$

where $\bar{\theta}$ denotes a consistent mean value. Expanding the same term further, this time around χ_0 , we obtain:

$$\begin{aligned} \sqrt{I}(\hat{\theta} - \theta_0) = & - \left(g_\theta(\hat{\theta}; \hat{\chi})' W g_\theta(\bar{\theta}; \hat{\chi}) \right)^{-1} g_\theta(\hat{\theta}; \hat{\chi})' W \times \\ & \left[\sqrt{I} g(\theta_0; \chi_0) + \sqrt{\frac{I}{J}} g_\chi(\theta_0; \bar{\chi}) \frac{1}{\sqrt{J}} \sum_{j=1}^J \psi_j + o_p(1) \right]. \end{aligned}$$

By the Slutsky and central limit theorems, $\sqrt{I}(\hat{\theta} - \theta_0)$ converges in distribution to a mean zero normal distribution with asymptotic covariance matrix:

$$V_\theta = (G'_\theta W G_\theta)^{-1} G'_\theta W E \left[\left(g(\theta_0; \chi_0) + \sqrt{\frac{I}{J}} G_\chi \psi \right) \left(g(\theta_0; \chi_0) + \sqrt{\frac{I}{J}} G_\chi \psi \right)' \right] W G_\theta (G'_\theta W G_\theta)^{-1},$$

where $G_\theta = E[\frac{\partial \zeta(\theta_0, \chi_0)}{\partial \theta'}]$ is $T \mathbf{x} k$ and $G_\chi = E[\frac{\partial \zeta(\theta_0, \chi_0)}{\partial \chi'}]$ is $T \mathbf{x} r$. As the mean values converge to the true ones the sample gradients converge to their theoretical counterparts.

Since our first-stage and second-stage estimator use different structural models and mostly different data, we assume that the first-stage and second-stage moments are uncorrelated. In that case, the formula simplifies to:

$$\begin{aligned} V_\theta &= (G'_\theta W G_\theta)^{-1} G'_\theta W E \left[g(\theta_0; \chi_0) g(\theta_0; \chi_0)' + \frac{I}{J} G_\chi \psi \psi' G'_\chi \right] W G_\theta (G'_\theta W G_\theta)^{-1} \\ (17) \quad &= (G'_\theta W G_\theta)^{-1} G'_\theta W \left[\varsigma \Omega_g + \vartheta G_\chi V_\chi G'_\chi \right] W G_\theta (G'_\theta W G_\theta)^{-1}, \end{aligned}$$

where $\Omega_g = E[h(\theta_0, \chi_0) h(\theta_0, \chi_0)']$ is $T \mathbf{x} T$, $\varsigma = \lim_{I \rightarrow \infty} (1 + \frac{I}{L})$ and $\vartheta = \lim_{I \rightarrow \infty} \frac{I}{J}$. This makes clear that as the relative precision of the first stage increases (i.e. $\vartheta \rightarrow 0$), the correction for the first stage disappears. ς controls for the simulation error. As the simulation becomes more and more precise, $\varsigma \rightarrow 1$. To construct an estimate of (17), we replace the theoretical concepts with their empirical counterparts, as is done in the usual GMM framework.

We use two different choices of the weighting matrix in our estimation. First, we choose a matrix that, while not fully optimal, does not depend on the fitted model. This is motivated by the observation that optimally weighting GMM estimators can worsen finite-sample bias (see for example West, fu Wong, and Anatolyev (1998)). We set the weighting matrix to Ω_g^{-1} . This is the optimal weighting matrix when the first stage correction does not matter, i.e. when $G_\chi = 0$ or $\vartheta = 0$. This yields:

$$V_\theta = (G'_\theta \Omega_g^{-1} G_\theta)^{-1} G'_\theta \Omega_g^{-1} \left[\varsigma \Omega_g + \vartheta G_\chi V_\chi G'_\chi \right] \Omega_g^{-1} G_\theta (G'_\theta \Omega_g^{-1} G_\theta)^{-1}.$$

While Ω_g^{-1} is not the optimal weighting matrix given the first-stage correction, this case is easy to implement since an estimate of Ω_g can be obtained directly as the sample counterpart of

$E[(\ln C_i - E(\ln C_i))(\ln C_i - E(\ln C_i))']$. Since we do not observe households across years, this is equivalent to assuming that Ω_g is diagonal with $\hat{\Omega}_{gt} = \frac{1}{I_t} \sum_{i=1}^{I_t} (\ln C_{i,t} - \ln \bar{C}_t)^2$. In the case of difference estimation, we use an identity weighting matrix since different households are observed in different years.

Second, we also employ the optimal weighting matrix $W = (\varsigma \Omega_g + \vartheta G_\chi V_\chi G_\chi')^{-1}$. In this case, the variance formula simplifies to:

$$V_\theta = (G_\theta' [\varsigma \Omega_g + \vartheta G_\chi V_\chi G_\chi']^{-1} G_\theta)^{-1}.$$

This requires using the estimates from our first case and constructing an estimate of W using sample counterparts.

There are two other interesting special cases. First, if the second stage is exactly identified ($T = k$), the covariance matrix simplifies to $V_\theta = G_\theta^{-1} [\varsigma \Omega_g + \vartheta G_\chi V_\chi G_\chi'] G_\theta^{-1'}$. Second, if there is no cross derivative, so that $G_\chi = 0$, the first stage uncertainty does not affect the second stage. In this case the formula collapses to the familiar one in which first stage estimation does not appear: $V_\theta = \varsigma (G_\theta' W G_\theta)^{-1} G_\theta' W \Omega_g W G_\theta (G_\theta' W G_\theta)^{-1}$.

This methodology also provides a useful test of the overidentifying restrictions in the second stage. If the model is correct and $G_\chi = 0$, the statistic:

$$\chi_{T-2} = \frac{I}{\varsigma} g(\hat{\theta}; \hat{\chi})' \hat{\Omega}_g^{-1} g(\hat{\theta}; \hat{\chi}),$$

is distributed asymptotically as Chi-squared with $T - 2$ degrees of freedom.

Under the correction for the first stage, the statistic

$$\chi_{T-2} = I g(\hat{\theta}; \hat{\chi})' [\varsigma \Omega_g + \vartheta G_\chi V_\chi G_\chi']^{-1} g(\hat{\theta}; \hat{\chi})$$

is distributed asymptotically as Chi-squared with $T - 2$ degrees of freedom.

C. Construction of Life Cycle Profiles

We construct profiles of consumption over the working life both smoothed and un-smoothed. Un-smoothed profiles are constructed by averaging our reconstructed measure of consumption by age. Smooth profiles are constructed by estimating an equation similar to (14) that fixes π_2 at the value estimated on the un-smoothed data, replaces the age and cohort dummies by fifth order

polynomials, and extends the highest age to 70 to avoid some of the endpoint problems commonly encountered with polynomial smoothers.

We generate a profile for per-household-equivalent consumption for a constant family size, by replacing the typical family size, $\bar{f}_t$, for each age with its sample average, 2.8. Finally, we construct the typical profile of the shift in marginal utility caused by changes in family size for our typical household as

$$\widehat{v(Z_{i,t})}^{\frac{1}{\rho}} = k \exp \left(\frac{1}{I_t} \sum_{i=1}^{I_t} f_i \hat{\pi}_2 \right).$$

using $\widehat{\pi}_2$ estimated from equation (14), where I_t is the number of households observed at age t where k is a constant.⁴⁴ The smoothed version of this is constructed by predicting $\bar{f}_t \widehat{\pi}_2$ using a fifth-order polynomial.

We employ the exactly same procedure for income, except that the adjustment for family size heterogeneity, π_2 , is set equal to its value from the consumption regression. $\bar{Y}_t = \exp[\frac{1}{I_t} \sum_{i=1}^{I_t} \ln Y_{i,t}]$ is an estimate of average income by age. To estimate separate profiles by education group or by occupation group, we add interaction terms between the categories considered and age and retirement status.

D. The Consumer Expenditure Survey

We use the CEX family, member and detailed expenditure files for years 1980 to 1993, as provided by the NBER. Most of our information about the CEX is obtained from Bureau of Labor Statistics (1993, and years 1980-1992) and conversations with BLS statisticians. Households are discarded if they are missing any of the information necessary for the regressions, if they report changes in age from the second to fifth interview of more than a year or less than zero years, if they are classified as incomplete income reporters, or if their reporting implies less than \$1000 in annual income or consumption.

We use information about the reference person to assign the household to cells, unless the reference person is female. In this case we use the spouses information. If there is no spouse, or his information is missing, the household is discarded. When this cut was made it eliminated 20% of the sample. All information besides individual labor income and consumption is taken from the family files. Values are assigned to a household based on information gathered in the fifth interview,

otherwise information is used from the second interview, or, if it is not available, the household is discarded. Households should not be matched across 1985 to 1986, and are not. Care is taken to assure consistency in our data despite variable classification changes through time, and across reference person and spouse. Information was provided by the Division of the CEX in the Bureau of Labor Statistics about various issues including the matching of occupation codes from 1980-81 to later years.

Pension contributions, income, Social Security contributions, and all asset income all refer to the past twelve months. Our definition of pension contributions is the sum over the CEX subcategories and thus includes private pensions, public pensions, Railroad Retirement pensions, and self-employed, IRA, and Keogh plans. If the after-tax family income variable is topcoded, reference person and spouse labor incomes are subtracted and we add, for each, the variable created by multiplying the earnings in last paycheck by the fraction of the year the pay period covers. These labor income variables are the sole variables from the member files used. Assets and asset income refers to the sum over savings accounts, checking accounts, bonds, and stocks, as of the time of interview. Each household is assigned to a year based on the midpoint between the first and fifth interview if both data are available; otherwise simply the single interview date is used. Age is the average of both interviews if both are available, otherwise it is the single one available. Due to some extreme reports, we reset reported tax rates above 50% back to 50%, and below zero percent to zero. We perform a similar exercise for Social Security contribution rates and pension contribution rates, using 25% as the upper bound.

While topcoding is very infrequent in consumption information, the household annual income variable reflects summation over a topcoded item for roughly half a percent of our households. Since, in most years, topcoding occurs at \$100,000 in income subcategories, reported individual annual labor income is the source of almost all income topcoding problems. However, households are also asked the gross amount of their paycheck and what length of time-period this paycheck covers. By multiplying these two variables together, we construct a second measure of annual labor income. Topcoding on this variable occurs only for a few cases. We correct our measure of after-tax family income by replacing the reported annual labor income in family income with

our constructed measure whenever the family income variable is topcoded. We are able to correct almost all topcoding.

Consumption data is compiled from the detailed expenditure files as all expenditures by a household except for those for health care, mortgage interest, and education. The consumption level is then the average monthly expenditure times twelve. Five percent of households have consumption data for 4, 7, 10, 13, or 14 months and these households' consumption are treated as if they were over 3, 6, 9, and 12 months. That is the recall interview period extended beyond the basic three months and some expenditures are recorded in a later month. BLS statisticians recommend treating these expenditures as if they occurred in the preceding month. Those covering 1 or 2 months (one percent of the sample) were dropped.

One might be concerned about the coverage of income and consumption and thus about the relative levels of the profiles. We get reasonable relative levels of consumption and income. Also, we have checked a total income profile and a food consumption profile against similar profiles constructed from PSID data. In the case of food, the level and shape of the two profiles were nearly identical. For total income profile, the shapes of the profiles were less similar, but still quite close. Finally, in our results section, we present evidence on estimation that instead uses only information from the changes in income and consumption. See the working paper version for further details.

The unemployment rates merged to the CEX are the regional unemployment rates for civilian population from the Household survey conducted and published by the Bureau of Labor Statistics in "Employment and Earnings." The GNP IPD PCE is from Council Of Economic Advisors (1995).

REFERENCES

- ABOWD, J., AND D. CARD (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57, 411–45.
- ATTANASIO, O., J. BANKS, C. MEGHIR, AND G. WEBER (1997): "Humps and Bumps in Lifetime Consumption," Manuscript, IFS, University College London.
- ATTANASIO, O., AND M. BROWNING (1995): "Consumption over the Life Cycle and over the Business Cycle," *American Economic Review*, 85, 1118–1137.
- ATTANASIO, O., AND G. WEBER (1995): "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103(6), 1121–57.
- AYAGARI, R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Savings," *Quarterly Journal of Economics*, 109, 659–684.
- BOUND, J. (1994): "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data," *Journal of Labor Economics*, 12(3), 345–68.
- BOUND, J., AND A. B. KRUEGER (1991): "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right," *Journal of Labor Economics*, 9(1), 1–24.
- BRANCH, E. R. (1994): "The Consumer Expenditure Survey: a comparative analysis," *Monthly Labor Review*, pp. 47–55.
- BROWNING, M., A. DEATON, AND M. IRISH (1985): "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle," *Econometrica*, 53(3), 503–43.
- BROWNING, M., AND A. LUSARDI (1996): "Household Saving: Micro Theories and Macro Facts," *Journal of Economic Literature*, 34(4), 1797–1855.
- BUREAU OF LABOR STATISTICS (1993, and years 1980-1992): *Consumer Expenditure Survey, 1993: Interview Survey*. Washington, DC: US Department of Commerce, Bureau of the Census, Ann Arbor, MI: Inter-university Consortium of Political and Social Research.
- CABALLERO, R. J. (1991): "Earnings Uncertainty and Aggregate Wealth Accumulation," *American Economic Review*, 81(4), 859–71.
- CAGETTI, M. (1998): "Estimating Preferences from Wealth Data," Manuscript, University of Chicago.

- CAMPBELL, J. Y., AND N. G. MANKIW (1989): "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence," in *N.B.E.R. Macroeconomics Annual*, ed. by O. J. Blanchard, and S. Fischer, pp. 185–215. Cambridge: MIT Press.
- CARROLL, C. D. (1992): "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 2, 61–135.
- (1994): "How Does Future Income Affect Current Consumption?," *Quarterly Journal of Economics*, 109, 111–147.
- (1997): "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 107(1), 1–56.
- (1999): "The Buffer Stock Saving: Some Theory," Manuscript, The Johns Hopkins University.
- CARROLL, C. D., AND A. A. SAMWICK (1997): "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40(1).
- CARROLL, C. D., AND L. H. SUMMERS (1991): "Consumption Growth Parallels Income Growth: Some New Evidence," in *National Savings and Economic Performance*, ed. by B. Bernheim, and J. B. Shoven. Chicago University Press.
- CLARIDA, R. (1991): "Aggregate Stochastic Implications of the Life-Cycle Hypothesis," *Quarterly Journal of Economics*, 106, 851–867.
- COUNCIL OF ECONOMIC ADVISORS (1995): *The Economic Report of the President*. Washington, DC: Government Printing Office.
- DEATON, A. (1985): "Panel Data from Time Series and Cross Sections," *Journal of Econometrics*, 30, 109–26.
- (1991): "Saving and Liquidity Constraints," *Econometrica*, 59(5), 1121–1142.
- (1992): *Understanding Consumption*. Oxford:Clarendon Press.
- DUFFIE, D., AND K. SINGLETON (1993): "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, 61, 929–952.
- FLAVIN, M. (1981): "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy*, 89, 974–1009.

- FRENCH, E. (1998): "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior," Working Paper, University of Wisconsin.
- FRIEDMAN, M. (1957): *A Theory of the Consumption Function*. Princeton, Princeton University Press.
- GAKIDIS, H. E. (1998): "Stocks for the Old? Earnings Uncertainty and Life-Cycle Portfolio Choice," Working Paper, M.I.T.
- GHEZ, G. R., AND G. S. BECKER (1974): *The Allocation of Time and Goods over the Life Cycle*. NBER, New York.
- GOURINCHAS, P.-O., AND J. A. PARKER (1997): "Consumption Over the Lifecycle," Social Science Research Institute Working paper 9722.
- (2001): "The Empirical Importance of Precautionary Saving," *American Economic Review (Papers and Proceedings)*, 91(2), 406–12.
- HUBBARD, G., J. S. SKINNER, AND S. P. ZELDES (1994): "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," in *Carnegie-Rochester Conference Series on Public Policy*, ed. by A. H. Meltzer, and C. I. Plosser, vol. 40, pp. 59–125.
- HURD, M. (1989): "Mortality Risk and Bequests," *Econometrica*, 57(4), 779–813.
- JAPPELLI, T., J.-S. PISCHKE, AND N. SOULELES (1998): "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," *Review of Economics and Statistics*, 80(2), 251–62.
- JUDD, K. L. (1998): *Numerical Methods in Economics*. Cambridge, MA, M.I.T. Press.
- KOTLIKOFF, L. J., AND L. H. SUMMERS (1981): "The Contribution of Intergenerational Transfers in Aggregate Capital Accumulation," *Journal of Political Economy*, 86, 706–32.
- LUDVIGSON, S., AND C. PAXSON (1997): "Approximation Bias in Linearized Euler Equations," mimeo, Princeton University.
- LUSARDI, A. (1996): "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets," *Journal of Business and Economic Statistics*, 14(1), 81–90.
- MEGHIR, C., AND G. WEBER (1996): "Intertemporal Nonseparability or Borrowing Restrictions? A Disaggregate Analysis Using a U.S. Consumption Panel," *Econometrica*, 64(5), 1151–81.

- MERTON, R. (1971): "Optimum Consumption and Portfolio Rules in a Continuous Time Model," *Journal of Economic Theory*, 3, 373–413, and erratum, *Journal of Economic Theory*, 6, pp 213–214, (1973).
- NEWKEY, W., AND D. L. MCFADDEN (1994): "Large Sample Estimation and Hypothesis Testing," in *Handbook of Econometrics, Volume 4*, ed. by R. F. Engle, and D. L. McFadden, pp. 2113–2245. Elsevier.
- PAKES, A., AND D. POLLARD (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 57, 1027–1057.
- PALUMBO, M. (1999): "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle," *Review of Economic Studies*, 66(2), pp. 395–422.
- PARKER, J. A. (1999): "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, 89(4), 959–73.
- PISCHKE, J.-S. (1995): "Individual Income, Incomplete Information, and Aggregate Consumption," *Econometrica*, 63(4), 805–40.
- SCHECHTMAN, J. (1976): "An Income Fluctuation Problem," *Journal of Economic Theory*, 12, 218–41.
- SCHECHTMAN, J., AND V. ESCUDERO (1977): "Some Results on 'An Income Fluctuation Problem'," *Journal of Economic Theory*, 16(2), 151–66.
- SHEA, J. (1995): "Union Contracts and the Life Cycle-Permanent Income Hypothesis," *American Economic Review*, 85, 186–200.
- WEST, K. D., K. FU WONG, AND S. ANATOLYEV (1998): "Feasible Optimal Instrumental Variables Estimation of Linear Models with Moving Average Disturbances," Social Science Research Institute WP No. 9901, University of Wisconsin.
- ZELDES, S. P. (1989): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 104(2), 275–98.

NOTES

¹We wish to thank Olivier Jean Blanchard, Ricardo Caballero, Chris Carroll and Angus Deaton for detailed comments and encouragement. We also thank Daron Acemoglu, Orazio Attanasio, David Card, Charles Fleischman, Eric French, Jang-Ting Guo, Gregori Kosenok, Franco Modigliani, Steve Pischke, Nicholas Souleles, Robert Shimer, anonymous referees, and seminar participants at the Board of Governors of the Federal Reserve, Chicago, Columbia, Harvard, Michigan, MIT, Princeton, Stanford, Tillburg, Wisconsin, and Yale. Parker gratefully acknowledges financial support from the Sloan Foundation and National Science Foundation grant SES-0096076.

²Some of the most recent papers are Carroll and Summers (1991), Lusardi (1996), Meghir and Weber (1996), Parker (1999), Shea (1995), and Jappelli, Pischke, and Souleles (1998). See Browning and Lusardi (1996) or Deaton (1992) for surveys.

³We do not enter the debate on the relative importance of retirement versus bequest wealth. Implicitly, our retirement wealth measure will include both.

⁴This is, in part, a confirmation of Carroll and Samwick (1997) and Carroll (1994) that argue, based on asset data, that buffer-stock models apply only to households before ages 45 to 50.

⁵Carroll (1997) argues that this is likely to be true.

⁶Attanasio, Banks, Meghir, and Weber (1997) shows that the residuals from a regression of consumption on family composition and labor supply variables are uncorrelated with age. However, if precautionary saving is part of the reason for the initial hump shape in consumption over the life cycle, this regression suffers from an omitted variable bias, which will incorrectly assign the hump to changes in demographics. Similarly, if we are fitting a model with too little preference heterogeneity to the data, then we will incorrectly assign too much of the shape to variations in precautionary saving. We allow for heterogeneity, in a limited manner, by controlling for heterogeneous family size at a given age and estimating separately across occupation and education groups.

⁷The permanent component of income is that level that would obtain without transitory shocks, as in Friedman (1957), not the annuity value of the present discounted sum of future income streams, as in Flavin (1981).

⁸While Abowd and Card (1989) found that change in labor income was best characterized by an $MA(2)$ process, they also found little gain in moving from an $MA(1)$ to an $MA(2)$.

⁹See Hubbard, Skinner, and Zeldes (1994), Palumbo (1999), and Hurd (1989) .

¹⁰We directly control for heterogeneous family size at a given age when constructing consumption and income profiles. The alternative is to model explicitly the stochastic process for changes in family size.

¹¹Social security wealth is definitely illiquid and is only available as annuities after retirement. Early withdrawal of pension and saving vehicles targeted for retirement purposes, such as IRA's, 401k plans and Keogh, is often restricted and fiscally penalized, if allowed at all. One might also consider a substantial part of housing wealth as illiquid wealth (at least during the time period that we study).

¹²We recognize that this assumption is problematic in some ways. For instance, two households with different income trajectories but the same permanent component of income at retirement are predicted to have identical accumulated illiquid wealth.

¹³To the extent that there is curvature in the relevant range of the consumption rule our assumption of linearity will bias us away from finding important effects of uncertainty on consumption.

¹⁴In the case of full certainty after retirement and no change in the utility shifter, $\gamma_1 = \frac{1 - \beta^{1/\rho} R^{1/\rho - 1}}{1 - (\beta^{1/\rho} R^{1/\rho - 1})^{N-T}}$.

¹⁵Other relevant parameters are $\beta = 0.960$, $\rho = 0.514$, $R = 1.0344$, typical income uncertainty, family size, and expected income profile, as discussed in sections 4 and 5.

¹⁶This rule is estimated directly from micro data. See section 6.3.

¹⁷In Gourinchas and Parker (1997), we show that unexpected transitory shocks to income are better smoothed later in life, despite the fact that they then contain greater information about total resources for the remainder of the life. This change in consumption risk is what drives the cohort average consumption growth.

¹⁸Risk aversion and the income process must be bounded. These conditions are analyzed in more details in Deaton (1991), Ayagari (1994) and Schechtman and Escudero (1977).

¹⁹The first stage is exactly identified while the second stage is overidentified.

²⁰We thank the anonymous referees and David Card for suggesting that we move from calibration and robustness checks to formally integrating the first stage.

²¹We thank Andrew Samwick for providing us with their raw data.

²²We drop Service, Farming, Armed Forces and Self-Employed occupations due to small samples in the CEX.

²³The assumption that the time effects enter multiplicatively is only a convenient short-cut. A fuller model could include aggregate shocks as state variables. This is beyond the current paper.

²⁴It is worth noting that the assumption that family size is exogenous might in fact be reducing the importance that we attribute to precautionary saving. If the buffer stock model is correct and having children is a form of consumption, then the decision to have children is affected by precautionary saving that is closely related to the timing of lifetime income. By normalizing by a family size adjustment over the lifecycle, one is removing some portion of changes in consumption driven by expected income changes.

²⁵This follows from the annoying identity that interview year less age equals birthyear!

²⁶This assumes that the time effects are not important for the linear trend in consumption and that they are observed by the household. The allocation of the trend in consumption to cohorts and age effects may be sensitive to this assumption.

²⁷Dummy variables are constructed for family sizes 1 through 9 and then for whether family size is 10 or greater. We also experimented with exogenous family size adjustments –assuming f_i is simply family size raised to the power -0.7 . This led to profiles that were noisier and flatter early in life.

²⁸The middle cohort is 1941 and the average regional unemployment rate 7 percent.

²⁹See Lusardi (1996), and Branch (1994).

³⁰We are arguing that user cost of housing -repairs, maintenance, utilities, and housing services- captures the expenditures made for consumption on housing by homeowners. Note that we do not subtract down payments or payment of principal from consumption or income so that nearly all housing wealth is modelled as liquid.

³¹It is important not to use different deflators for income and consumption. This could break the relationship between cash on hand and consumption in nominal terms which is the relationship predicted by the buffer-stock theory.

³²Since $\hat{\beta}R \approx 1$, the choice of ρ matters little. The coefficient of relative risk aversion and the discount factor, are not separately identified because we do not have any variation in the desire to shift consumption across time.

³³We do not include all the first-stage parameters in the first-stage variance correction. In particular, we do not include the parameters that drive the family shifters and the expected income profile. We report results including uncertainty about these parameters in section 6.3 where we investigate robustness issues.

³⁴For instance, if the ratio of illiquid wealth to permanent income, h , were equal to 6, the marginal propensity to consume out of illiquid wealth would be a mere 0.35%.

³⁵The discount comes from the assumption that income is received and consumption occurs at the beginning of the period.

³⁶In order to do this, we input the consumption rule at retirement as estimated in our benchmark case. Our estimates imply that if households faced no risks after retirement, the age of death, N , is a reasonable 87 years. That is, if we set preferences at our estimated values and $N = 87$, the standard life cycle model with a certain date of death implies the same value function at retirement as we estimate (up to a constant).

³⁷These current income and wealth numbers refer to averages of levels (not logs) and so are not immediately comparable to the income profiles.

³⁸In favor of our model, the households we would like for this comparison are 65 years olds in 2006. On the other hand, some housing wealth is de facto illiquid in our model.

³⁹Gourinchas and Parker (2001) provide some additional calculations based on our model.

⁴⁰Recall that we estimate $\bar{\omega}_{26} = E[\ln w_{26}]$, and so the mean level of cash on hand is given by: $\exp(\bar{\omega}_{26} + \frac{1}{2}\sigma_{\omega_{26}}^2)$.

⁴¹Gourinchas and Parker (1997) present a similar scenario based on the *CEX* data.

⁴²We thank one of our referee for pointing this out to us.

⁴³Under this assumption, durable goods cannot be used to offset labor income uncertainty. While this is an extreme assumption, the common assumption of adjustment costs implies that durable goods are worth a fraction of their original value on second-hand markets.

⁴⁴This is correct up to a multiplicative constant.

TABLE I. VARIANCE OF INCOME SHOCKS

GROUP	VARIANCE OF PERMANENT SHOCK	VARIANCE OF TRANSITORY SHOCK
TOTAL	0.0212	0.0440
OCCUPATION		
Managerial and Prof. Speciality	0.0175	0.0341
Tech., Sales, and Admin. Support	0.0235	0.0361
Precision Prod., Craft, and Repair	0.0175	0.0432
Operators and Laborers	0.0299	0.0458
EDUCATION		
Some High School	0.0214	0.0658
High School Degree	0.0277	0.0431
Some College	0.0238	0.0342
College Graduate	0.0146	0.0385
Graduate School	0.0115	0.0500

Source: Carroll and Samwick (1997) and authors' calculations. Standard errors range from ten percent to thirty percent of the coefficient estimate, depending largely on the sample size. These data are available from the authors.

TABLE II. LIQUID WEALTH HOLDING AT THE BEGINNING OF LIFE

RATIO OF LIQUID ASSETS TO INCOME GROUP	LOG		LEVEL
	MEAN	S.D.	MEAN
	$\bar{\omega}_{26}$	$\sigma_{\omega_{26}}$	$\exp[\bar{\omega}_{26} + \frac{1}{2}\sigma_{\omega_{26}}^2]$
TOTAL	-2.794	1.784	0.300
OCCUPATION			
Managerial and Prof. Speciality	-2.261	1.649	0.405
Tech., Sales, and Admin. Support	-2.574	1.625	0.285
Precision Prod., Craft, and Repair	-3.312	2.051	0.299
Operators and Laborers	-3.172	1.654	0.164
EDUCATION			
Some High School	-3.772	1.886	0.136
High School Degree	-3.296	1.797	0.186
Some College	-2.756	1.679	0.260
College Graduate	-2.181	1.554	0.377
Graduate School	-1.887	1.540	0.495

Source: Authors' calculations based on the CEX measures of wealth and the constructed measure of income.

TABLE III. STRUCTURAL ESTIMATION RESULTS

MSM ESTIMATION	ROBUST WEIGHTING	OPTIMAL WEIGHTING
Discount Factor (β)	0.9598	0.9569
S.E.(A)	(0.0101)	
S.E.(B)	(0.0179)	(0.0150)
Discount Rate ($\beta^{-1} - 1$) (%)	4.188	4.507
S.E.(A)	(1.098)	
S.E.(B)	(1.949)	(1.641)
Risk Aversion (ρ)	0.5140	1.3969
S.E.(A)	(0.1690)	
S.E.(B)	(0.1707)	(0.1137)
Retirement Rule:		
γ_0	0.0015	$5.68 \cdot 10^{-6}$
S.E.(A)	(3.84)	
S.E.(B)	(3.85)	(16.49)
γ_1	0.0710	0.0613
S.E.(A)	(0.1215)	
S.E.(B)	(0.1244)	(0.0511)
χ^2 (A)	175.25	
χ^2 (B)	174.10	185.67

Note: *MSM* estimation for entire group. Standard errors calculated without (A) and with (B) correction for first stage estimation. Cell size is 36,691 households. The last row reports a test of the overidentifying restrictions distributed as a Chi-squared with 36 degrees of freedom. The critical value at 5% is 50.71. Efficient estimates are calculated with a weighting matrix $\hat{\Omega}$ computed from the robust estimates.

TABLE IV. ESTIMATES FROM THE STOCHASTIC MODEL BY EDUCATION AND OCCUPATION

GROUP:	Discount Factor (β)	Discount Rate ($\beta^{-1} - 1$)	Risk Aversion (ρ)	Retirement Rule		χ^2	N
	γ_0	γ_1					
EDUCATION							
Some High School	0.962 (0.082)	3.94 (8.92)	0.282 (1.481)	0.209 (5.04)	0.072 (2.360)	53.60	4,270
High School Graduate	0.949 (0.015)	5.30 (1.64)	0.869 (0.220)	$3.79 \cdot 10^{-3}$ (20.05)	0.059 (0.049)	59.12	12,445
Some College	0.960 (0.159)	4.15 (17.29)	0.394 (2.344)	0.351 (4.095)	0.043 (3.156)	84.21	9,653
College Graduate	0.930 (0.060)	7.48 (6.97)	2.290 (0.423)	$1.55 \cdot 10^{-8}$ (54.60)	0.049 (0.075)	111.70	6,350
Graduate School	0.944 (0.087)	5.93 (9.77)	1.694 (0.843)	$1.06 \cdot 10^{-7}$ (18.23)	0.057 (0.076)	87.26	5,973
OCCUPATION							
Managerial and Prof.	0.946 (0.060)	5.71 (6.69)	1.672 (0.524)	$5.20 \cdot 10^{-8}$ (22.78)	0.050 (0.067)	115.62	12,693
Tech., Sales, Admin.	0.953 (0.037)	4.90 (4.11)	1.059 (0.339)	$2.13 \cdot 10^{-7}$ (39.42)	0.049 (0.064)	64.02	6,548
Precision Prod., Craft	0.953 (0.333)	4.97 (36.77)	0.990 (3.895)	0.003 (18.49)	0.054 (0.997)	52.86	4,469
Operators, Laborers	0.953 (0.489)	4.90 (53.80)	0.867 (4.846)	$3.14 \cdot 10^{-6}$ (1365)	0.049 (2.35)	57.58	6,063


Note: *MSM* estimation in levels. Standard errors calculated with correction for first stage estimation. The next to last column reports a test of the overidentifying restrictions distributed as a Chi-squared with 36 degrees of freedom. The critical value at 5% is 50.71.

TABLE V. ROBUSTNESS CHECKS

ESTIMATION	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Consumption Measure:	Total	Total	Total	Total	Total	Total	NonDurable
OPTIMAL WEIGHTING	No	Yes	No	No	No	No	No
FIX $\bar{\omega}_{26}$	No	No	Yes	Yes	Yes	Yes	Yes
METHOD	Level	Level	Level	First Diff	Level	Level	First Diff
FIX (γ_0, γ_1)	No	No	Yes	No	No	No	Yes
INCOME CORRECTION	No	No	No	No	Yes	No	No
VARIANCE CORRECTION	No	No	No	No	No	Yes	No
Discount Factor (β)	0.8741	0.8724	0.9648	0.9643	0.9598	0.9606	0.9630
S.E.(A)	(0.0984)		(0.0002)	(0.0202)	(0.0101)	(0.0013)	(0.0962)
S.E.(B)	(0.1032)	(0.0197)	(0.0087)	(0.0206)	(1.1147)	(0.0033)	(0.1080)
Discount Rate ($\beta^{-1} - 1$) (%)	14.406	14.627	3.6412	3.6992	4.188	4.102	3.843
S.E.(A)	(12.881)		(0.0224)	(2.171)	(1.0981)	(0.1363)	(10.370)
S.E.(B)	(13.515)	(2.589)	(0.935)	(2.212)	(121.01)	(0.3597)	(11.635)
Risk Aversion (ρ)	5.1817	5.2823	0.1278	0.1585	0.5140	0.4803	0.1536
S.E.(A)	(0.5640)		(0.0004)	(0.5855)	(0.1690)	(0.0249)	(1.8082)
S.E.(B)	(0.6522)	(0.1195)	(0.0088)	(0.6669)	(16.592)	(0.2707)	(2.0180)
Retirement Rule:							
γ_0	5.37 10^{-5}	1.36 10^{-4}	0.594	0.001	0.0015	0.0001	0.2977
S.E.(A)	(20.70)			(1.668)	(3.849)	(5.510)	
S.E.(B)	(20.74)	(7.86)		(1.675)	(1221.)	(8.390)	
γ_1	0.0211	0.0203	0.0774	0.1455	0.0710	0.0832	0.0388
S.E.(A)	(0.1824)			(0.3163)	(0.1215)	(0.0533)	
S.E.(B)	(0.1847)	(0.0465)		(0.3184)	(38.976)	(0.0750)	
Initial Wealth ($\exp(lw_{25})$)	2.5715	2.5715	0.0612	0.0612	0.0612	0.0612	0.0612
S.E.(A)	(0.6472)						
S.E.(B)	(0.8578)	(0.4808)					
χ^2 (A)	91.89		1690	34.13	175.25	163.25	42.22
χ^2 (B)	91.83	91.78	1619	33.37	155.66	154.83	41.36

Note: *MSM* estimation in levels and first-differences. Standard errors calculated without (A) and with (B) correction for first stage estimation. The next to last column reports a test of the overidentifying restrictions distributed as a Chi-squared with 36 degrees of freedom. The critical value at 5% is 50.71.

Normalized consumption Panel A: $\beta = 0.960$, $\rho = 0.514$, $\gamma_1 = 0.071$, $\gamma_0 = 0.001$



Panel B: $\beta = 0.960$, $\rho = 0.514$, $\gamma_1 = 0.077$, $\gamma_0 = 0.594$




FIGURE 1. A Tale of Two Households

Thousands of 1987 dollars




FIGURE 2. Household Consumption and Income over the Life Cycle




FIGURE 3. Household Consumption and Income over the Life Cycle, by Education Group





FIGURE 4. Household Consumption and Income over the Life Cycle, by Occupation Group

Thousands
of 1987 dollars

Panel A: Baseline Estimation

$\beta = 0.960$, $\rho = 0.514$, $\gamma_1 = 0.071$, $\gamma_0 = 0.001$



Panel B: Various β




FIGURE 5. The Fitted Consumption Profile





FIGURE 6. Normalized Target Cash-on-Hand by Age

Thousands
of 1987 dollars

Panel A: Life Cycle and Buffer Saving



Panel B: Life Cycle and Buffer Wealth




FIGURE 7. The Role of Risk in Saving and Wealth Accumulation




FIGURE 8. Fitted Consumption Profiles under Different Model Assumptions