



Taylor & Francis
Taylor & Francis Group



Wealth Accumulation over the Life Cycle and Precautionary Savings

Author(s): Marco Cagetti

Source: *Journal of Business & Economic Statistics*, Jul., 2003, Vol. 21, No. 3 (Jul., 2003), pp. 339-353

Published by: Taylor & Francis, Ltd. on behalf of American Statistical Association

Stable URL: <https://www.jstor.org/stable/1392584>

REFERENCES

Linked references are available on JSTOR for this article:

https://www.jstor.org/stable/1392584?seq=1&cid=pdf-reference#references_tab_contents

You may need to log in to JSTOR to access the linked references.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

Taylor & Francis, Ltd. and American Statistical Association are collaborating with JSTOR to digitize, preserve and extend access to *Journal of Business & Economic Statistics*

Wealth Accumulation Over the Life Cycle and Precautionary Savings

Marco CAGETTI

Department of Economics, University of Virginia, Charlottesville, VA 22903 (*cacio@virginia.edu*)

This article constructs and simulates a life cycle model of wealth accumulation and estimates the parameters of the utility function (the rate of time preference and the coefficient of risk aversion) by matching the simulated median wealth profiles with those observed in the Panel Study of Income Dynamics and in the Survey of Consumer Finances. The estimates imply a low degree of patience and a high degree of risk aversion. The results are used to study the importance of precautionary savings in explaining wealth accumulation. They imply that wealth accumulation is driven mostly by precautionary motives at the beginning of the life cycle, whereas savings for retirement purposes become significant only closer to retirement.

KEY WORDS: Precautionary savings; Simulated method of moments.

1. INTRODUCTION

Two of the most important reasons to save are to finance expenditures after retirement (retirement or life cycle motive) and to protect consumption against unexpected shocks (precautionary motive). Households are subject to several sources of risk (in earnings, health, and mortality). Markets to insure those risks are often limited or do not exist. The main way to self-insure against them is to accumulate a buffer stock of wealth.

This article reports an evaluation of the quantitative importance of the precautionary motive for saving by structurally estimating a model of wealth accumulation over the life cycle. Using data from the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF), the rate of time preference and the degree of risk aversion are estimated by matching the median wealth holdings by age for various educational groups (college graduates, high school graduates, and high school dropouts), and the estimates are used to evaluate how much of the wealth accumulation can be attributed to the precautionary motive.

The estimates imply a low degree of patience, together with significant aversion to risk. The estimated coefficient of risk aversion is usually higher than 3, and often higher than 4. The rate of time preference decreases with education and is higher (often significantly so) than 5%–10% for lower educational groups.

As stressed by other works (in particular, Carroll 1997 and Carroll and Samwick 1997), such parameter configuration generates large precautionary savings. For instance, simulations based on the estimates herein show that for college graduates, the median wealth level at retirement in the model with precautionary motives is twice as high as that implied by the model without uncertainty. Moreover, saving by younger households can be explained in large part as precautionary, whereas life cycle saving becomes relevant only after age 45–50, close to retirement.

The estimates of the two parameters differ from some previous estimates, which usually implied lower risk aversion. The difference is relevant because it has several implications on how much people save and how they react to various policies. High risk aversion and high impatience generate low elasticity of savings to the interest rate. There has been much debate, for instance, about the effect of tax-preferred forms of

savings, which deliver a higher interest rate. Although some authors have claimed that such instruments significantly increase wealth holdings, the results of the present study present evidence for the opposite thesis, argued by, among others, Engen, Gale, and Scholz (1996), that tax-favored forms of investment have little impact on total savings and formally justify the parameters chosen in that work.

Although the model presented herein does not distinguish assets with different risk characteristics, the degree of risk aversion also affects the portfolio choice over the life cycle, as discussed by Campbell and Viceira (2002). Moreover, heterogeneity in preferences, above all in the degree of patience, generates wealth heterogeneity and thus can help explain the concentration of wealth and have consequences for aggregation and general equilibrium. For instance, Carroll (2000a) showed the limitations of the representative agent hypothesis under parameter configurations similar to those estimated in this article.

There are several estimates of the preference parameters in the literature, in particular of the coefficient of risk aversion. Most were obtained using log-linearized Euler equations and consumption data (e.g., Dynan 1993; Attanasio and Weber 1995), and they typically imply lower risk aversion. This approach has been criticized by, among others, Carroll (2001a) and Ludvigson and Paxson (2001), who showed that in the presence of precautionary savings it may be difficult to recover the value of the preference parameters using Euler equations. Moreover, Attanasio and Low (2000) showed that although they may in certain cases correctly estimate the risk aversion coefficient, log-linearized equations do not recover the rate of time preference, which is also necessary for simulations of life cycle models.

In this article, instead the model is estimated structurally, using simulation methods, as was first done by Gourinchas and Parker (2002) (and later done also in French 2000 and French and Jones 2001 to study retirement behavior). Samwick (1998) performed a similar exercise, backing out the rate of time preference from simulated wealth at retirement, but he did not apply any formal estimation method. Gourinchas and Parker (2002) focused on consumption and used consumption profiles in their

© 2003 American Statistical Association
Journal of Business & Economic Statistics
July 2003, Vol. 21, No. 3
DOI 10.1198/073500103288619007

estimation. In contrast, because the aim is to study the importance of precautionary savings for wealth accumulation, the present study looks at wealth profiles and uses wealth data, using a quantile-based estimator, more appropriate for a variable that is highly concentrated in the top of the distribution (and, as shown in Carroll 2000b, different models are necessary to study the rich.) Using mean wealth holding, as explained in the paper, would lead to biased estimates (more patience and less risk aversion). Relative to the study of Gourinchas and Parker (2002), the estimates in this article imply higher risk aversion and less patience.

This article is also related to a large literature on precautionary savings (more recently summarized in Carroll 2001b). In particular, Carroll and Samwick (1997) showed that wealth is higher for people with higher income variability, which also suggests (as confirmed in this article) that the precautionary motive is quantitatively relevant. However, their exercise does not recover the preference parameters, which, as argued earlier, are fundamental for simulations of life cycle models.

The article is structured as follows. Section 2 discusses the forces determining precautionary savings. Section 3 presents the model, and Section 4 presents the estimation strategy and the data. Section 5 describes the results and presents some simulations that measure the amount of precautionary wealth (i.e., the increase in wealth due to earnings uncertainty). Section 6 further analyzes implications of my estimates for life cycle models: the low elasticity of savings generated by these models, the effect of heterogeneity in time preferences for the distribution of wealth, and the impact of bequest motives.

2. PRECAUTIONARY SAVINGS

To understand the forces that determine precautionary savings, it is useful to consider a condition derived by Deaton (1991) and also applied by Carroll and Samwick (1997). Deaton (1991) considered a model with infinitely lived agents with constant relative risk aversion utility and liquidity constraints, and derived a condition ensuring that wealth does not accumulate without bounds,

$$R\beta E(g^{-\gamma}) < 1, \quad (1)$$

where R is the interest rate, β is the discount factor, γ is the coefficient of relative risk aversion, and g is the rate of growth of income. When this condition is satisfied, agents will tend to run down their assets when they become too large. The only reason for holding a positive amount of wealth is to maintain a buffer

stock of savings to self-insure against random fluctuations in income.

In a life cycle model, the condition does not apply directly. Because income falls at retirement, most households will eventually start saving for retirement and will tend to decumulate assets when old. However, it highlights the main forces that determine the strength of the precautionary motive relative to the retirement motive. When income is expected to increase (as captured by the term g), households would prefer to borrow or save very little, because more resources will be available in the future. But, fearing negative random shocks, they save to build up a buffer stock of wealth. Moreover, precaution is the main force driving the saving behavior when the degree of patience, β , is low relative to the interest rate, R .

The earnings profile is increasing at the beginning of the career. This implies that in a certainty equivalent model, young agents will dissave and accumulate little wealth early in life. However, microeconomic data suggest significant amounts of uncertainty in income. When the precautionary motive is important, as my estimates suggest, this uncertainty generates significant wealth accumulation.

Table 1 reports the answers to a question in the section about expectations and attitudes of the 1992 SCF about the most important reasons for saving. (The entry "future" corresponds to the answers "to get ahead, for the future, to advance standards of living," which is difficult to interpret. The entry "other" also includes "don't save.") The table is very suggestive. Precaution is a relevant reason for saving for all the age and educational groups. Concerns about retirement, although also present when young, become more important with age. But even for many households who are close to retirement, the main reason for saving is precautionary. Note that early in life two other motives are particularly important: savings for the children's education and to buy a house. The model presented in the next section attempts to take all of these possible motives into account.

3. THE MODEL

The life cycle model considered here captures in a simplified and parsimonious way some of the main determinants of the saving behavior: to finance consumption after retirement, to keep a buffer stock of wealth for precautionary reasons, to finance various expenditures at certain ages (such as for the children), and possibly to leave bequests.

Table 1. Main Reason for Saving, 1992 SCF

| | Education | House | Retirement | Precautionary | Future | Other |
|----------------|-----------|-------|------------|---------------|--------|-------|
| Age | | | | | | |
| 26-35 | 12.23 | 10.64 | 19.15 | 29.79 | 7.98 | 20.03 |
| 36-45 | 17.88 | 2.65 | 23.51 | 26.85 | 10.93 | 18.10 |
| 46-55 | 8.26 | 1.71 | 31.05 | 32.65 | 5.70 | 20.63 |
| 56-65 | .31 | 0 | 38.65 | 32.72 | 7.06 | 21.26 |
| Degree | | | | | | |
| No high school | 4.29 | 1.43 | 14.29 | 28.09 | 11.43 | 40.47 |
| High school | 8.40 | 4.07 | 25.47 | 30.54 | 9.49 | 22.03 |
| College | 10.03 | 2.47 | 32.69 | 30.82 | 6.59 | 17.40 |

3.1 The Decision Problem

The decision unit is the household. The baseline model is

$$\max_{W_{t+1}, C_t} E \left(\sum_{t=0}^{\tilde{T}} \beta^t e^{Z_t \theta} \frac{C_t^{1-\gamma} - 1}{1-\gamma} + \beta^{\tilde{T}+1} e^{Z_{\tilde{T}+1} \theta} S(W_{\tilde{T}+1}) \right), \quad (2)$$

subject to

$$W_{t+1} = RW_t + Y_t - C_t + B_t, \quad \forall t \quad (3)$$

and

$$W_{t+1} \geq 0, \quad \forall t. \quad (4)$$

Here W_{t+1} represents the amount of assets carried on to $t + 1$, constrained to be positive. The model does not distinguish between various types of assets, and it assumes that W has a non-stochastic (after tax) rate of return R (in the computations, 1.03 is used). Y_t is a stochastic and exogenous stream of earnings, and B_t is the bequest received at age t . The expected value reflects the expectations about the level and variability of future earnings, the life span, and the receipt of a bequest.

Households start at age 25 ($t = 0$), and live up to age \tilde{T} . The life span \tilde{T} is random. While alive, households get utility from consumption, and after death they get some utility $S(A)$ from the wealth A that they bequeath. Households are assumed to live more than 65 and less than 91 years, and for older ages, the conditional probabilities of survival from the 1995 Life Tables (National Center for Health Statistics 1998) from age 65 to age s for women are used to compute the expected life span. The reason for using the probabilities for women (as also done in Hubbard, Skinner, and Zeldes 1995) is that women live longer, and savings decisions within the household take into account the utility of the survivor.

Following Attanasio, Banks, Meghir, and Weber (1999), the effect of demographic variables is introduced in the utility function through the discount factor. The discount factor between age t and $t + 1$ is given by $\beta e^{\Delta Z_t \theta}$, where β is the pure discount factor estimated in this article, and Z is a set of demographic variables that affect the marginal utility of consumption. This captures the idea that some expenditures (such as expenditures for the children) are higher at some ages, and households may want to save earlier in life to finance them. As explained in Appendix B, the present study uses the specification and the results provided by Attanasio et al. (1999), and assumes that the factors that most influence marginal utility are family size and leisure of the spouse.

3.2 The Earnings Process

The earnings process Y_t^i (earnings of household i at time t) is assumed to be given by

$$\log Y_t^i = G_t^i + u_t^i \quad (5)$$

and

$$G_t^i = gt + F(\text{age}^i, \text{education}^i), \quad (6)$$

where G is the deterministic component of log earnings, g is a common (and constant) productivity growth factor, and F is a

function of age and education that represents how income varies over the life cycle for different educational groups.

The stochastic component of earnings, represented by u_t^i , is a random walk with MA(1) innovations,

$$u_t^i = u_{t-1}^i + \eta_t^i - \psi \eta_{t-1}^i. \quad (7)$$

For simplicity, it is assumed that households retire after age 65, after which earnings become deterministic,

$$\log Y_t^i = gt + F(\text{age}^i, \text{education}^i) + u_{65}^i. \quad (8)$$

This formulation implies a constant replacement ratio for all households within each educational group.

Section 4.3 and Appendix A give more details about the earnings process and describe the estimates used in the article.

In addition to these earnings, households can receive gifts and bequests. The suggestion given by Laibson, Repetto, and Tobacman (1998) is used here. Using a probit model, these authors estimated the probability of receiving an inheritance in a given year as a function of age and of education using PSID data. Figure B.1 in Appendix B plots the estimates obtained using the gifts and bequests received between 1969 and 1989 reported in the PSID. As expected, the probability of receiving a bequest is highest around age 50–60 and is higher for college graduates than for the other groups.

Conditional on receiving an inheritance, the amount received should be related to parental income and wealth, but there is no variable relative to the parents in the model. This fact is captured by assuming that the amount received is proportional to the current earnings of the receiver (which are themselves correlated to parental earnings). From the PSID, the median ratio of gifts and bequests received to income in the year the bequest is received is estimated for the three different educational groups considered (no high school degree, high school degree, and college degree). It is assumed that in every period the household can receive this amount with the probability specified earlier. More details on the process for bequests can be found in Appendix B.

3.3 The Utility of Bequests

It is assumed that when a household dies, it may receive some utility $S(W_{T-t+1})$ from leaving a bequest,

$$S(W) = \alpha \frac{W^{1-\gamma} - 1}{1-\gamma}. \quad (9)$$

A similar, but more general, formulation was discussed by Carroll (2000b), who interpreted S not only as joy of giving, but also as utility from wealth for its own sake. For simplicity, S is assumed to be a function only of the amount of wealth W and to have the same coefficient γ of the utility function for consumption. The degree of altruism (or the utility from wealth) is dictated by the parameter α . When $\alpha = 0$, bequests are accidental, generated by the fact that the life span is uncertain. Note that when α is strictly positive, a household will always want to leave some bequests, whereas in reality many families do not leave any. In most of the simulations, it is assumed that $\alpha = 0$.

It is important to note that all the homogeneity assumptions (Constant Relative Risk Aversion utility, formulation for the bequest function, unit root in earnings) are useful when

solving the problem numerically, because they allow earnings to be dropped as a state variable, dramatically simplifying the numerical solution. A technical appendix available at <http://www.people.virginia.edu/~mc6se> gives more details about the algorithm used to solve the model, as well as the Fortran code.

4. THE ESTIMATION METHOD AND THE DATA

4.1 The Estimation Strategy

To estimate the parameters of the model, I adopt a strategy similar to that proposed by Gourinchas and Parker (2002), who simulated a life cycle model and matched simulated and empirical mean consumption. This article focuses on wealth accumulation, and because wealth is highly concentrated, median wealth profiles are matched instead of means. This section briefly summarizes the estimation strategy, which is described more extensively in Appendix C.

Given an earnings process Y , an initial distribution of wealth, and the parameters γ , β , and α , one can simulate the model for a large number of agents and compute the distribution of wealth for each age t and for different educational groups, $F_t^w(\beta, \gamma; Y)$. One way to estimate the preference parameters β and γ is therefore to choose those that generate a simulated distribution F^w that matches some aspects of the empirical distribution \hat{F} . In particular, two types of estimation are performed, one type considering the medians and the other using the means.

First, consider the condition on the median. Take α as given (the sensitivity to this parameter is discussed in Sec. 6.2) and estimate the rate of time preferences and of risk aversion that best match (in the sense specified later) the median. Let w_i^t be the wealth of agent i , who belongs to the age group t . Consider eight 5-year age groups, 26–30, . . . , 61–65. Let β_0 and γ_0 be the true parameter values. The median wealth, $m^t(\beta_0, \gamma_0)$, satisfies

$$E(.5 - \mathbf{1}(w_i^t \leq m^t(\beta_0, \gamma_0)) | t) = 0, \tag{10}$$

where $\mathbf{1}$ is the indicator function. As suggested by Powell (1994) (see Appendix C), an estimator based on the median can be constructed by noting that the median m^t also minimizes

$$\min_{m^t} E(\rho(w_i - m^t(\beta, \gamma)) | t) \tag{11}$$

and

$$\rho(y) = y(.5 - \mathbf{1}(y < 0)). \tag{12}$$

To estimate β and γ , one can compute the median $m^t(\beta, \gamma)$ implied by the model and choose the values $\hat{\beta}$ and $\hat{\gamma}$ that minimize the empirical counterpart to condition 11. Because no analytical expression exists for $m^t(\beta, \gamma)$, this quantity will be simulated.

The estimates obtained from condition 11 are also compared with those obtained by matching the means (although, as explained in Sec. 5, the estimates obtained from the means are biased). Let $W^t(\beta, \gamma)$ be the average wealth of the households of age t . The following moment conditions hold (one for each age group):

$$E(w_i^t - W^t(\beta_0, \gamma_0)) = 0. \tag{13}$$

Therefore one can apply the simulated method of moments (see, e.g., Duffie and Singleton 1993) and find the $\hat{\beta}$ and $\hat{\gamma}$ that minimize the criterion

$$\min_{\beta, \gamma} \hat{M}(\beta, \gamma)' T \hat{M}(\beta, \gamma), \tag{14}$$

where $\hat{M}(\beta, \gamma) = 1/N \sum (w_i^t - W^t(\beta, \gamma))$ is a column vector (one condition for each age group) and T is a (positive definite) weighting matrix. Again, because $W^t(\beta, \gamma)$ does not have an analytical expression, the simulated counterpart is used instead.

4.2 The Data

The two main sources of microeconomic data on wealth in the U.S. are the SCF and the PSID. The SCF collects detailed information about wealth for a cross-section of households (except for a small panel interviewed between 1983 and 1989). Because wealth has a highly skewed distribution, the SCF oversamples rich households by including, in addition to a national area probability sample (representing the entire population), a list sample drawn from tax records (to extract a list of high-income households).

The PSID is a panel, designed to study the income situation, above all for poorer households, a group that is oversampled. More recently, however, the PSID has started asking a few questions on wealth every 3 years.

Figure 1 plots the means and the quartiles of net worth by age group for the two datasets (limited to the households also used in the estimation). Figure 2 does the same for net worth excluding housing wealth. Because of the concentration of wealth at the top of the distribution, the means are much higher than the medians and are usually closer to the third quartile. For college graduates in particular, means are more than twice as high as medians. The means in the SCF tend to be higher than those in the PSID, because by oversampling the rich, the SCF is able

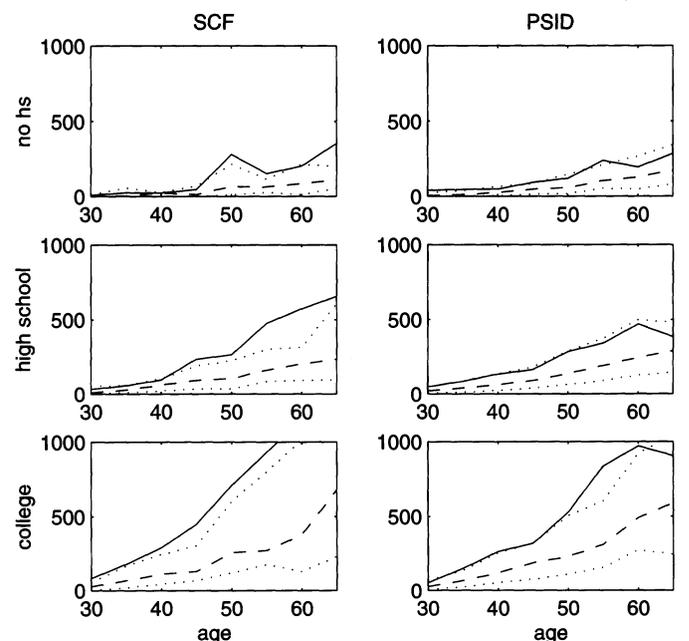


Figure 1. Means (straight line), Medians (dashed line), and First and Third Quartiles (dotted lines) of the Distribution of Net Worth, by 5-Year Age Groups, 1992 Dollars.

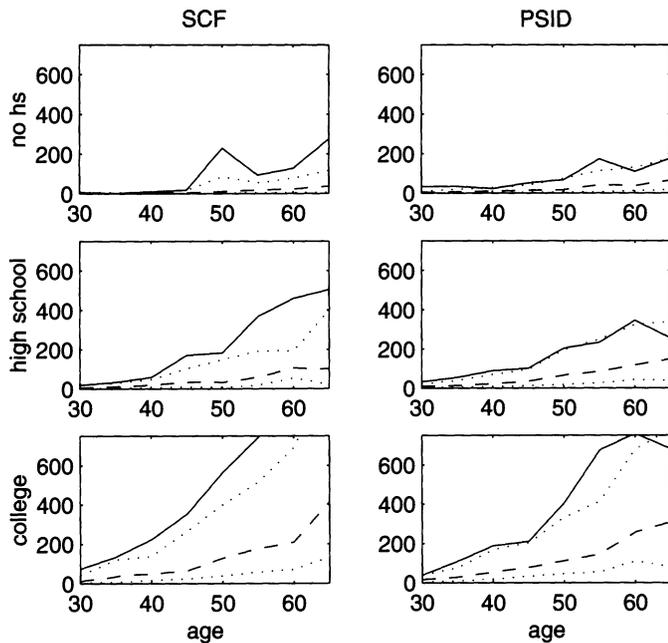


Figure 2. Means (straight line), Medians (dashed line), and First and Third Quartiles (dotted lines) of the Distribution of Net Worth Excluding Main Housing, by 5-Year Age Groups, 1992 Dollars.

to give a better description of the top 5% of the distribution. However, as also shown in detail by Juster, Smith, and Stafford (1999), for lower quantiles the two datasets give similar information. The PSID is thus also a good dataset for this article, which focuses on the medians.

The estimation uses the following data. The SCF 1989, 1992, and 1995 waves are used, deflating the values to 1992 dollars using the consumer price index (CPI) for urban consumers. Each observation is weighted according to the weight provided in each wave of the survey. For the PSID, all of the available waves (i.e., 1984, 1989, and 1994) are used. The census and the Latino sample are excluded, and equal weight is assigned to each observation in the remaining sample, which is representative of the whole population. Because the PSID is a panel, some households appear in more than one wave. Each household is counted as a different observation each time it appears in a wave.

Only households age 26–65 (age and education refer to the age and education of the head) are used. Younger households are not included, because educational choice is not modeled. The older households are excluded because several features that are not modeled explicitly (e.g., the survival of one spouse, certain types of medical shocks) may be relevant for their behavior. Only households composed at least by a head and a spouse are considered—singles were excluded (many singles are concentrated among the very young and the old, which were already excluded). Singles include a large number of single mothers who are on welfare. Asset-based welfare programs can have relevant influences on the savings decisions of the recipients (Hubbard et al. 1995) which are not fully captured in the present setup. Because these elements are not modeled, observations of singles are dropped. After also deleting the households with missing relevant demographic variables, the PSID sample includes 1,039, 3,353, and 1,924 observations for households

without a high school degree, with a high school degree, and with a college degree, whereas the SCF sample includes 212, 1,100, and 1,963.

Two measures of wealth are considered. The first measure, net worth, includes most assets [e.g., stocks, bonds, housing, and vehicles, as well as individual retirement accounts (IRAs) and thrift accounts] minus liabilities (e.g., mortgages, credit card debt). Pension wealth and capitalized social security are not part of net worth. In the model assets can be freely traded (or, equivalently, it is possible and costless to borrow against them at the constant rate of return R), and their payoff is the stream of interests that they give. Pensions and social security do not have these characteristics, because they are typically nonfungible and usually provide a stream of earnings after retirement. Therefore, they are accounted for in the earnings measure, as explained in Section 4.3. Many assets included in net worth are also not perfectly fungible; for instance, there are penalties in drawing money from IRAs, and markets that allow to borrow against them do not exist. However, because of the computational burden of accounting for these characteristics, IRAs and similar accounts are simply added to the other, more fungible forms of wealth. (Incidentally, note that the PSID, unlike the SCF, does not distinguish whether an asset is part of an IRA or of other more liquid forms.)

In particular, the inclusion of housing wealth may be problematic. Housing wealth does, at least in part, constitute wealth for precautionary purposes. A house can be sold in case of particularly bad shocks, and ownership of a house per se provides consumption insurance, because one does not have to pay rent, which is an important part of total consumption expenditures. However, selling a house can be expensive, and households rarely borrow against housing wealth for consumption purposes. Therefore, I also consider a second measure of wealth, net worth (as defined earlier) excluding primary housing. Housing is assumed to enter the utility function separately from other forms of consumption, and it does not affect savings in other assets.

4.3 The Earnings Data

As explained in Section 3.2, the earnings process has three components: a common productivity growth factor, g ; a seniority component, $F(\text{age, education})$; and a stochastic component, u . The term g represents productivity growth in post-World War II United States. It is assumed to be constant over time and is common to all educational groups. Using aggregate data on workers' compensation from the Bureau of Labor Statistics, it is calibrated to $g = .016$.

The age–education earnings profiles are estimated using data from the Consumer Expenditure Survey (CEX); see Appendix A for details on the exact equation. Two measures of earnings, corresponding to the two measures of wealth, are used.

As explained in the previous section, net worth does not include pensions and social security wealth. Rather, pension and social security wealth is assumed to generate an income stream that is included in earnings. Savings in the form of pension and social security contributions is not part of the savings that increases the measure of net worth used, and thus these contributions should not be counted in the earnings. The income measure used is thus labor earnings, net of taxes, social security,

and pension contributions, and inclusive of pensions, social security, and other transfer income (e.g., unemployment benefits). Moreover, educational and medical expenditures are excluded, which can be considered committed expenditures that decrease income but do not give utility. CEX datasets are used rather than other datasets, such as the PSID, precisely because they contain detailed information about such expenditures.

The second set of estimates uses wealth net of primary housing. For simplicity, it is assumed that housing wealth accumulates exogenously. Therefore, expenditures for housing, (i.e., mortgage down payments, interest payments, and housing additions) are also excluded from the previous definition of income.

For the stochastic component, the estimates of Carroll and Samwick (1997b), who used PSID data on before tax labor earnings, are used. Carroll and Samwick estimated a slightly different equation for u , which is, however, equivalent to the one used in this article. Appendix A explains the mapping. The resulting variances of the innovations η , σ_η^2 , equal .121, .090, and .070 for households with no high school degree, with high school, and college graduates, whereas ψ equals .59, .45, and .54. This implies that the variance of earnings decreases with education.

It should be noted that, although various types of transfer income (e.g., unemployment benefits, Aid to families with Dependent Children) are taken into account in constructing the mean income profiles, minimum income levels or consumption floors are not explicitly modeled. As already pointed out in Section 4.2, when these features are explicitly introduced (as shown in Hubbard et al. 1995), the number of people with zero wealth increases. This may affect my results regarding the left tail of the wealth distribution, above all for lower education groups. However, only median wealth holdings are used in the present estimation. Moreover, only households composed by a head and a spouse are used, whereas most of these programs are targeted to single-parent families.

Note also that some types of risks are also not explicitly included. Medical expenditures are excluded from the earnings profiles, and their variability is not considered. They may be an important source of risk for the old (who are, however, not considered in the estimation). The aggregate uncertainty coming from the business cycle is also not present, although typically the variability of the aggregate shocks is much smaller than the idiosyncratic component (see, e.g., Storesletten, Telmer, and Yaron 1999).

5. THE RESULTS

This section presents the estimates of β and γ (Sec. 5.1), and their implications for precautionary savings (Secs. 5.3 and 5.4).

5.1 Estimates of the Rate of Time Preference and Risk Aversion

Tables 2 and 3 show the estimates of β and γ , obtained using the condition on the medians (10) and on the means (14). Table 2 uses net worth, whereas Table 3 uses net worth excluding housing wealth. A 3% after-tax real interest rate and no voluntary bequests ($\alpha = 0$) are assumed.

First, consider the column for the medians, for the estimates obtained using total net worth. The results for all groups imply

Table 2. Estimated β and γ

| | Median | | | | Mean | | | |
|----------------|----------------|---------------|----------------|---------------|----------------|----------------|-----------------|----------------|
| | SCF | | PSID | | SCF | | PSID | |
| | β | γ | β | γ | β | γ | β | γ |
| College | .988 (.006) | 4.05 (.65) | .989 (.002) | 4.26 (.45) | 1.08 (.047) | 6.99 (.417) | 1.022 (.030) | 6.21 (2.00) |
| | | | | | (1.08) | | (4.39) | |
| High school | .923 (.005) | 3.70 (.27) | .952 (.007) | 3.27 (.15) | .933 (.026) | 4.14 (.927) | .869 (.020) | 5.57 (.109) |
| | | | | | (4.31) | | (2.73) | |
| No high school | .940 (.034) | 2.57 (.69) | .948 (.008) | 2.74 (.21) | .898 (.068) | 3.37 (1.16) | .920 (.003) | 3.58 (.166) |
| | | | | | (6.77) | | (1.26) | |

NOTE: Standard deviations are in parentheses, and χ^2 results (for the means) are in on the third rows.

a relatively high coefficient of risk aversion (greater than 3, except for the lowest educational group) and, for the groups without a college degree, a low rate of time preference—a configuration suggesting that precautionary savings are quantitatively very relevant, as shown in the following sections. The estimates obtained using the PSID and those using the SCF are relatively close—not surprisingly, because, as explained in Section 4.2, the two datasets imply similar median wealth profiles. Note incidentally that the SCF has few observations for the group without a high school degree, and the estimates have large standard errors; thus it is better to examine those from the PSID only.

College graduates have a higher degree of patience, β , than the other two groups, that in turn have similar β . For college graduates, β is around .98, whereas for the others it drops to around .95–.93. The relation between education and patience has been already studied by, for instance, Lawrance (1991), who estimated the degree of patience from Euler equations on consumption and found that it increases with education. Note, however, that the risk-aversion coefficient is also decreasing with education, which suggests that not only different time preference, but also different attitudes toward risk can explain different saving behaviors across groups.

The estimates obtained excluding housing wealth (Table 3) present a similar picture, with some differences. Housing

Table 3. Estimated β and γ , Net Worth Excluding Housing

| | Median | | | | Mean | | | |
|----------------|----------------|---------------|----------------|---------------|-----------------|----------------|-----------------|----------------|
| | SCF | | PSID | | SCF | | PSID | |
| | β | γ | β | γ | β | γ | β | γ |
| College | .985 (.015) | 3.49 (.33) | .977 (.005) | 4.00 (.24) | 1.14 (.068) | 8.13 (1.05) | 1.041 (.027) | 7.01 (.233) |
| | | | | | (.96) | | (9.04) | |
| High school | .857 (.005) | 4.29 (.17) | .864 (.010) | 4.45 (.12) | .917 (.0052) | 4.45 (.518) | .856 (.0064) | 5.57 (.121) |
| | | | | | (4.18) | | (18.03) | |
| No high school | .898 (.034) | 2.40 (.73) | .843 (.014) | 4.04 (.21) | .839 (.037) | 3.22 (.485) | .781 (.079) | 5.45 (.360) |
| | | | | | (2.41) | | (7.33) | |

wealth tends to be a large fraction of households' wealth, particularly for the relatively less affluent. Correspondingly, even correcting the earnings as explained in the previous section, there is a large drop in the estimated degree of patience for the lowest two educational group, with β dropping to values around .84–.86. The drop of around .01 in the discount factor for the college graduates is instead smaller. The coefficient of risk aversion remains high (greater than 3) for all groups, although there is now little difference in the value between groups.

Regarding the difference in preferences between groups, Hubbard et al. (1995) have argued that social programs with asset-based means testing (i.e., transfers to people who have very low wealth holdings) induce many households with low lifetime income (most of them with low education levels) to hold little or no wealth. As previously mentioned, most of the households who are most affected by these programs are excluded, and then only the median households are considered, with no attempt to match the lower quantiles. These corrections thus help control for the effect discussed by Hubbard et al. (1995). However, at least in part, the lower degree of patience and the lower degree of risk aversion of the households without a high school degree may reflect not only genuine differences in preferences, but also the impact of these programs. Including them (along the lines of, e.g., Hubbard et al. 1995) could therefore be an important extension of the model. However, such extension would come at a very great computational cost, because the model would not be homogeneous in earnings anymore. Hubbard et al. (1995), moreover, showed that such programs have almost no impact on the median wealth holding of the college graduates. Therefore, the estimates for this group should not be affected.

5.2 Matching Means

Consider now the estimates obtained using the mean condition (the right most columns in Tables 2 and 3). They imply either a much higher coefficient of risk aversion or a higher degree of patience (or both), above all for of college graduates.

The reason for this is that the means are affected by the right tail of the distribution of wealth, whereas the medians are not. Because of the high concentration of wealth, the richest households hold a disproportionate share of total wealth. However, the model presented in this article is not able to adequately capture the behavior of the very rich, because it does not consider the earnings levels and above all the large bequests received by the very rich. To match the observed mean wealth, the estimates simply assign high patience and high risk aversion to all households. The effect is strong, particularly for the college graduates, because the richest households tend to be among this group.

For this reason, the estimates based on the means are biased upward and are inappropriate for the purpose of drawing inference about most of the population. Incidentally, this exercise suggests that the parameters of the utility function used to calibrate macroeconomic models of aggregate capital may differ from those found from microeconomic estimates. The behavior of the very rich has a large impact on the aggregates, and this group may be characterized by preferences and face different sources of risks than the rest of the population, on which many microeconomic estimates (including this one) are based. Thus the following sections concentrate on the first set of estimates, those based on median wealth.

5.3 Precautionary Savings

The parameter estimates presented earlier imply that buffer stock wealth is a large component of the total wealth of the median household. It is difficult to exactly measure the contribution of precautionary savings to total wealth, because wealth accumulated for precautionary purposes is then available also for retirement. To get a rough idea, the following experiment is performed. For a given set of preference parameters, one can compute the wealth profile when there is earnings uncertainty (as done until now) and when earnings are certain. The latter is called *profile life cycle wealth*. The difference between the simulated wealth profiles with and without earnings uncertainty can be considered the contribution of precautionary savings to total wealth.

In particular, three wealth profiles are computed: the one predicted by the precautionary saving model, the one predicted by the same model without earnings uncertainty and with borrowing constraints (so that households must hold a nonnegative amount of wealth), and the one predicted by the model without uncertainty and with the possibility of borrowing, so that households can borrow in certain periods and repay later. (This case can be considered the case of complete markets for the earnings uncertainty.) Note incidentally that longevity risk (i.e., about the age of death) is still present in all of the simulations. Only earnings are deterministic.

The implied wealth paths (for college graduates, $\alpha = 0$) are depicted in Figure 3. Figure 3(b) and (d) shows the empirical paths (+) and the wealth paths computed from the model (dashed lines), with and without uncertainty. The top line is the precautionary saving model, whose parameters have been estimated to match the data. When the precautionary motive is absent, households save much less. The estimates suggest that

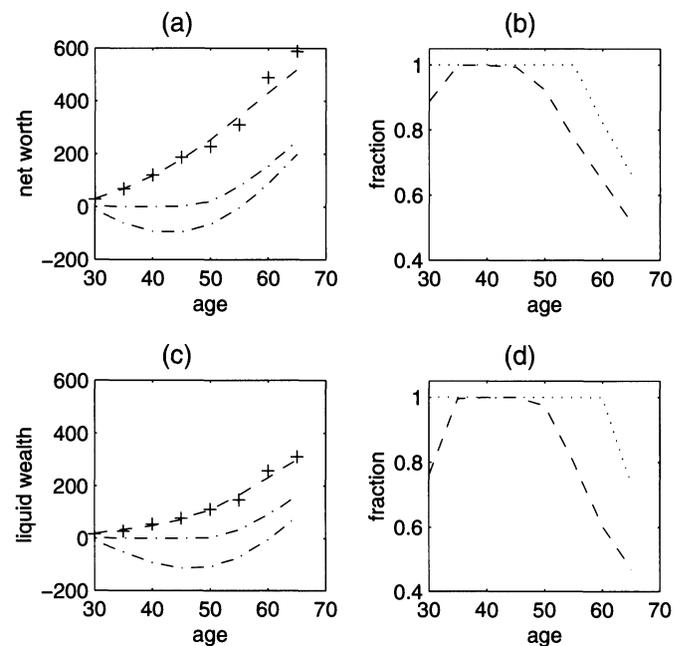


Figure 3. Simulated Wealth Paths, With and Without Uncertainty, for (a) Net Worth and (c) Liquid Wealth, and Proportion of Wealth Attributable to Precautionary Savings, for (b) Net Worth and (d) Liquid Wealth, College Graduates.

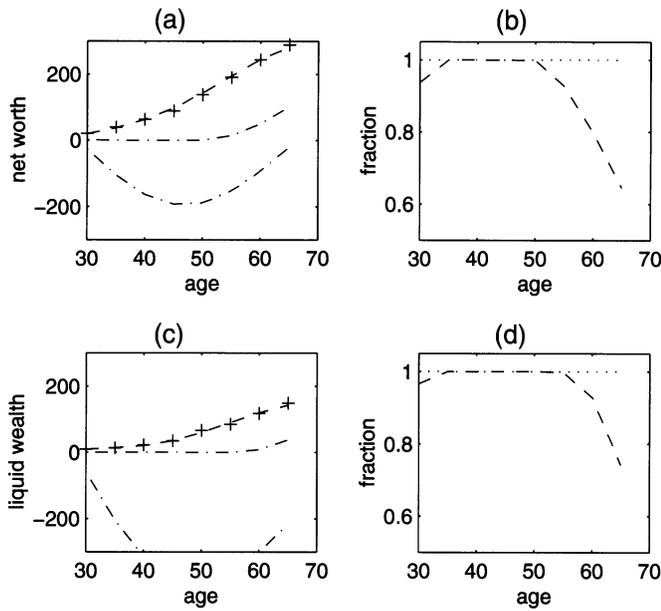


Figure 4. Simulated Wealth Paths, With and Without Uncertainty for (a) Net Worth and (c) Liquid Wealth, and Proportion of Wealth Attributable to Precautionary Savings for (b) Net Worth and (d) Liquid Wealth, High School Graduates.

without the precautionary motive persons will not save early in life, because they expect higher income in the future. If they can borrow, they will do so until around age 40 (bottom line), and they will repay their debt later. If they cannot borrow (middle line), then they will keep little or no wealth early in life. Eventually, households start saving for retirement, and the wealth profile becomes positive and rises until retirement at age 65. But the switch to retirement saving happens late, around age 45–50. These estimates thus confirm Carroll’s (1997) sugges-

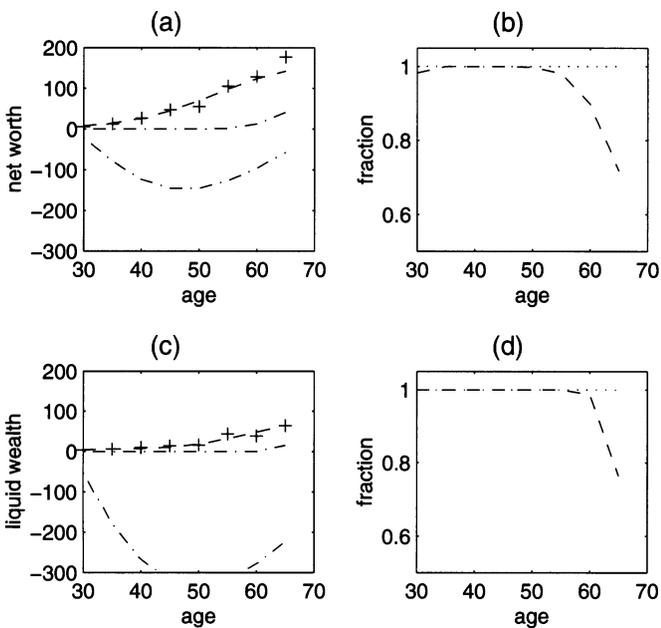


Figure 5. Simulated Wealth Paths, With and Without Uncertainty, for (a) Net Worth and (c) Liquid Wealth and Proportion of Wealth Attributable to Precautionary Savings for (b) Net Worth and (d) Liquid Wealth, No High School Degree.

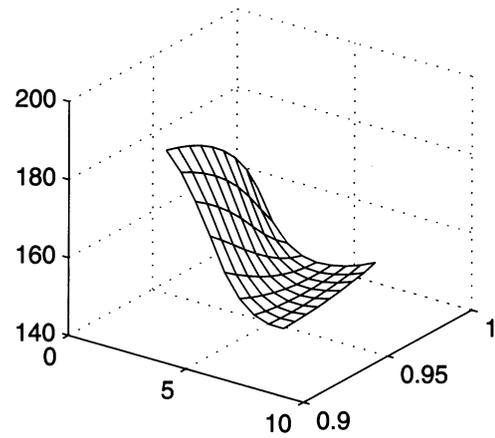


Figure 6. Value of the Objective Function (10), College Graduates, PSID. The minimum is .989 and 4.26.

tion about the relative importance of the motives for saving at different ages.

Precautionary wealth is thus an important component of total wealth. To stress this point even more, Figure 3(a) and (c) shows the fraction of total wealth attributable to precautionary savings. This is computed as total wealth minus life cycle wealth over total wealth (with the ratio set to 1 when total wealth is negative). The ratio is close to 1 until age 45–50, meaning that most of the wealth can be attributed to the precautionary motive. The ratio then decreases as households approach retirement, but it still remains high. Even at retirement, the level of wealth implied by the model with precautionary savings is twice as high as the level implied by the model without uncertainty. The results for high school graduates (Fig. 4) and for households without a high school degree (Fig. 5) are even starker, since these two groups are even more impatient than the college graduates.

5.4 Patience Versus Risk Aversion

Figure 6 shows the value of the objective function (10) for the case of college graduates in the PSID. The plot is relatively flat along one direction; if γ is fixed and only β is estimated, then β decreases as we increase γ . The results of this exercise are reported in Table 4.

When risk aversion is increased, households accumulate more assets for precautionary purposes. Therefore, to match the

Table 4. Estimated β , PSID (1984–1994), Using the Condition on the Medians (10), for Various γ (Columns)

| | Net worth | | Excluding housing | |
|----------------|------------------|-----------------|-------------------|-----------------|
| | 1 | 6 | 1 | 6 |
| College | 1.009 (.0004) | .950 (.003) | .998 (.0005) | .917 (.0038) |
| High school | .996 (.0003) | .820 (.0035) | .991 (.0002) | .750 (.0034) |
| No high school | .988 (.0005) | .771 (.0066) | .982 (.0009) | .680 (.009) |

NOTE: Standard deviations are in parentheses.

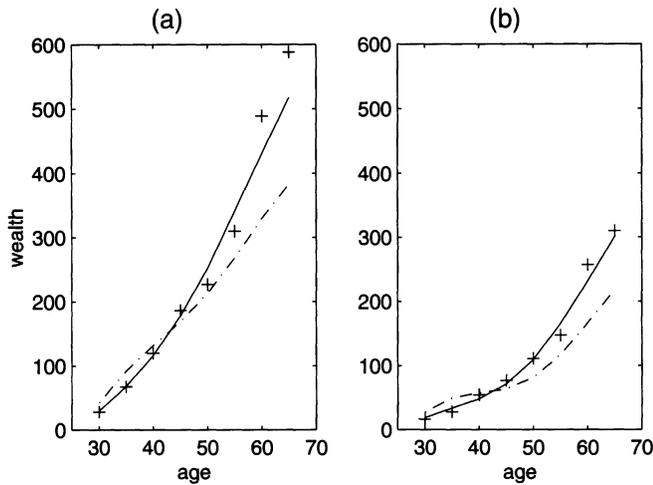


Figure 7. Wealth Paths, Net Worth (a) and Net Worth Excluding Primary Housing (b), for College Graduates. Data (+), $\gamma = 1$ (dashed and dotted line) and $\gamma = 6$ (dotted line) and β as in Table 4, Estimated β and γ (dashed line).

same median amount of wealth requires decreasing β and assuming less patience. The opposite occurs when γ is decreased. However, the life cycle profiles are different for different combinations of the two parameters.

Figure 7 plots the implied wealth paths for $\gamma = 1, 6$ and the corresponding estimates of β shown in Table 4, and for the jointly estimated β and γ from Section 5.1 (that imply a γ around 4). The estimates are from the PSID.

The profiles for γ around 4 (the estimated value), and $\gamma = 6$ are very close, which suggests that it may be difficult to identify the two parameters for such values of γ . The profile for $\gamma = 1$ is, however, somewhat different from the other two; it implies relatively large savings very early in life, when agents save for the anticipated expenditures (described in Sec. 3.1) occurring around age 40, and when they quickly build up their buffer stock of wealth. The profile is rather flat from 35 to 50, when consumption for demographic purposes (represented by the discount factor corrections) is highest. Less risk-averse households will save very little, or dissave, during this period, whereas more risk-averse ones will tend to save more to keep a higher buffer stock of wealth. Finally, after age 50, there is again accumulation of wealth for retirement purposes. When

γ is higher, the required buffer stock of saving is higher, and there is more wealth accumulation also at later ages and during periods of high consumption needs (such as around age 40).

5.5 Consumption Profiles

One of the reasons for the development of models with precautionary savings was the observation that the profile of consumption over the life cycle is hump-shaped and tends to track that of income, a feature not replicated by the flat profiles of certainty equivalent models. Figure 8 plots empirical and simulated median profiles for consumption and income for all three groups; the graphs refer to the experiment with housing wealth. The (smoothed) empirical consumption profiles are obtained from the CEX (excluding the same type of households as for wealth and earnings).

The simulated profiles tend to reproduce some aspects of the behavior of the empirical series. The consumption paths are very steep at the beginning of the life cycle, as is the income profile. Consumption then falls significantly below income, as the household saves for retirement. And it exhibits the hump shape (except for college graduates, whose degree of patience is high—.989—and whose consumption peaks at retirement).

Relative to the data, the simulated profiles tend to show a little more savings early in life (until age 35) for college graduates and high school graduates, but the opposite for people without a high school degree, who have the lowest risk aversion and time preference. This is due to the high estimates of risk aversion; young households are saving for precautionary motives, even if income will be much higher later in life.

This may explain the difference between the estimates of the present study and those of Gourinchas and Parker (2002), who looked at consumption profiles and found lower estimates of the coefficient of risk aversion (between .5 and 1.5, depending on the group). The consumption paths that these authors generated imply consumption higher than income for households younger than 30–35 years, which is consistent with consumption data. Wealth data instead tend to show much higher savings at that age, which the model explains as precautionary savings.

6. FURTHER IMPLICATIONS OF THE ESTIMATES

The previous section discussed the implications of the estimates of β and γ on the magnitude of the precautionary motive

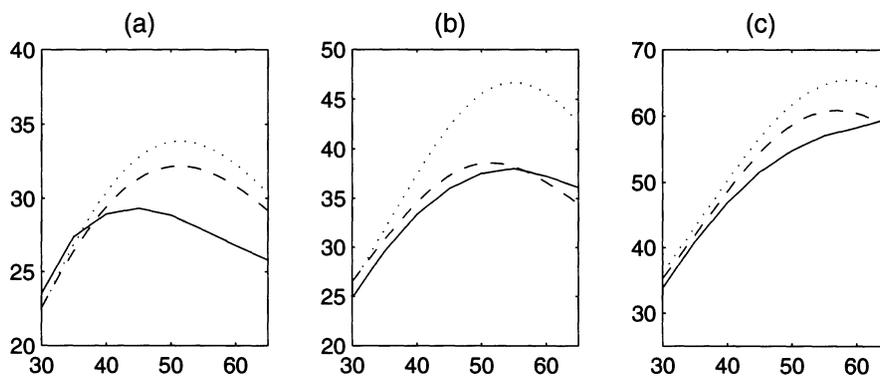


Figure 8. Median Consumption Simulated (solid line) and Data (dashed line), and Income (dotted line) Paths for the Three Educational Groups: (a) No High School, (b) High School; (c) College.

for saving. In this section some further implications are examined, to highlight the fact that a careful choice of these parameter is also important for other economic questions.

Simulations of life cycle models are often used to study how households increase their savings in response to various taxes and incentive schemes. For instance, Summers (1981) showed that in a life cycle model without uncertainty, savings are very interest-elastic, and thus small changes in taxes on capital can generate large increases in capital accumulation. This may not be true in a precautionary savings model, however. If precautionary motives are important, then people want to accumulate a buffer stock of wealth, but changes in interest rates do not increase their savings much beyond that point. Savings are thus very interest-inelastic. As shown by Cagett (2001), the value of the elasticity depends on the combination of β and γ chosen to simulate the model.

This point has been made by others, including Carroll (1997), and is the basis for the results of Engen et al. (1996) on the impact of saving incentive schemes. There is a large literature examining the effects of IRAs. Some authors (e.g., Poterba, Venti, and Wise 1996) found large effects. However, a model with parameter values similar to those of the present study (such as that of Engen et al. 1996) implies very small effects. The results in this paper thus confirm the choice of their parametrizations.

Of course, even if the median household is very interest-inelastic, a considerable proportion of total capital is concentrated in a small fraction of the population, who may have different savings behavior. Therefore, capital accumulation may be very interest-elastic in the aggregate, if this small fraction of households is responsive to changes in interest rates.

6.1 Heterogeneity in the Discount Factor

Heterogeneity in the degree of patience is often considered an important factor in explaining the dispersion of wealth. For instance, Krusell and Smith (1998) reproduced the variability of wealth holding (in an infinitely lived agent model) by assuming two types of agents, who have slightly different discount

factors. In the absence of precautionary savings, savings are extremely sensitive to the choice of β . In the standard, infinitely lived agent model with quadratic utility (the certainty equivalence model), one must assume that $\beta R = 1$. If individuals were more patient, they would accumulate wealth without bounds.

But, when precautionary savings are important, larger differences in patience are needed to generate large dispersion in wealth. Figure 9 shows how the median wealth (net worth) profile changes when the discount factor is changed while keeping γ fixed. Figure 9(a) is for the estimated $\gamma = 4.26$, and $\beta = .969, .989$, and 1.09 . Figure 9(b) shows a similar experiment, but when precautionary savings are less important ($\gamma = 1$, $\beta = .98, 1.00, 1.02$). When $\gamma = 1$, changes in the discount factor have a large impact on the amount of wealth held at retirement, but the impact is lower when $\gamma = 4.26$, above all when moving β away from 1.

The present estimates, however, do suggest also large differences in the rate of time preference among groups, sometimes greater than .05. Carroll (2000a) showed that with precautionary savings and differences in β of the magnitude of those between my estimates for different educational groups, one can generate significant dispersion in wealth holdings (and interesting macroeconomic dynamics not captured by the representative agent model). This article points out, however, that differences in risk aversion can be equally important in explaining differences in behavior across households.

6.2 Bequest Motives

It has been assumed until now that there are no bequest motives ($\alpha = 0$). When $\alpha > 0$, households have an additional motive for saving, and, for a given level of β and γ , they will accumulate more wealth. This section examines how the wealth profile change with $\alpha > 0$.

Here α is calibrated to two other values, chosen so that the marginal propensity to consume in the last period is .5 and .1. That is, a person who is sure to die would leave 50% and 90% of his own wealth to the children ($\alpha = 0$ obviously implies a

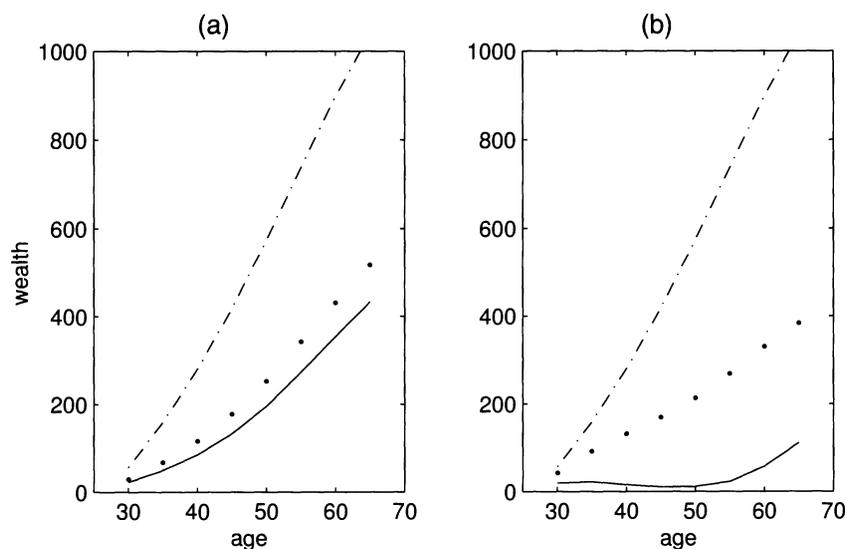


Figure 9. Wealth Paths (net worth), for $\gamma = 4.26$ (a) and $\gamma = 1$ (b) for College Graduates. Low β (solid line), medium β (dotted line) and high β (dashed and dotted line).

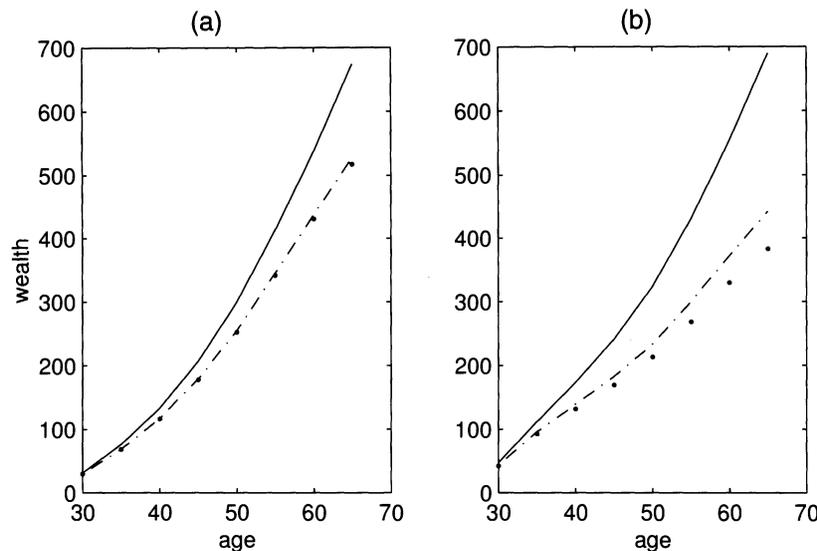


Figure 10. Wealth Paths (net worth) for $\gamma = 4.26$ (a) and $\gamma = 1$ (b), Various α , for College Graduates.

marginal propensity to consume of 1, because the person would consume everything and leave no bequests). Figure 10 shows the wealth profiles before retirement for the three choices of α , (a) for the estimated $\beta = .989$ and $\gamma = 4.26$, and (b) for $\gamma = 1$ and the corresponding $\beta = 1.00$. Note that the actual parameters α chosen are different in the two panels, because the marginal propensity to consume in the last period also depends on β and γ . The actual parameter values are 1 and 10,000 for 10(a) and 1.07 and 7.7 for 10(b).

For the estimated parameter values, there is almost no difference in switching to a marginal propensity to consume of .5, whereas the wealth path for the case of a propensity of .1 is 15% higher at retirement. This suggests that α has a relatively small effect on wealth profiles before retirement, unless a very high degree of altruism is assumed. For the other parameter configuration (low risk aversion and high patience), the differences in the wealth path are a somewhat higher (for instance, the wealth at retirement increases by 50% when the altruism parameter is at its highest value). The reason for the difference is that the utility from bequest is obtained only at the time of death and thus is more heavily discounted by more impatient households, such as the ones in Figure 10(a). Bequest motives, therefore, do not much affect the behavior of impatient households, but matter more for more patient ones.

These results show that it is difficult to identify bequest motives from (median) wealth data before retirement; thus in most of the article α is set to 0. To identify the presence of bequest motives, it may be necessary to look at data on the very old (that has been excluded from this analysis), and on whether, and if so, how they run down their wealth. Moreover, the bequest motive can be quantitatively important in explaining the wealth holding of a subgroup of (very rich) households, as illustrated by De Nardi (2000), who showed that bequest motives can account for part of the wealth concentration in the hands of the top 1% of the population. It is also important to note that the results depend on the particular way of introducing bequest motives. These motives may enter the utility function in a different way—in particular, providing utility not only at the time of death, but also during the lifetime.

7. CONCLUSIONS

This article has assessed the importance of precautionary savings for the accumulation of wealth by formally estimating a life cycle model of consumption and saving. Age profiles of wealth holdings were constructed and the preference parameters (time preference and risk aversion parameters) found that generate simulated profiles closest to those constructed from microeconomic data (PSID and SCF). The observed life cycle profiles are consistent with a low degree of patience, a high degree of risk aversion, and differences in the degree of patience among various educational groups. Under such a configuration, saving is dictated by precautionary motives early in life, and concerns for retirement are important only for older households. Precautionary motives explain a large fraction of the wealth of the median household. The amount of wealth at retirement implied by the model with precautionary savings is twice as high as that implied by a pure life cycle model without uncertainty.

The model and the estimation methods presented in this article can be applied to more complex settings. First, various models of wealth accumulation have been simulated to examine portfolio composition over the life cycle (e.g., Campbell and Viceira 2002), or the effects of different social security and pension schemes (e.g., Samwick 1998). The results of these experiments depend on the extent of precautionary savings; it is therefore fundamental to choose appropriate preference parameters used to perform such exercises. Second, most models, including this one, fail to account properly for the extreme concentration of wealth among the richest households. A richer setup, involving a better characterization of the income and returns to entrepreneurial activities, may be helpful in analyzing such phenomenon. Finally, the analysis presented here suggests that borrowing constraints, preference heterogeneity, and life cycle behavior can interact in ways not captured by the simplified, infinitely lived representative agent models common in the macroeconomic literature. Some contributions on the importance of heterogeneity and market incompleteness for general equilibrium have begun to appear (Browning, Hansen, and Heckman 1999; Carroll 2000a), but the problem deserves further study.

ACKNOWLEDGMENTS

The computer codes and the data are available on request or at <http://www.people.virginia.edu/~mc6se/>. I received suggestions from seminar participants at many institutions. I would like to thank in particular Francisco Ciocchini, Mariacristina De Nardi, Eric French, James Heckman, Manuel Lobato, Annamaria Lusardi, Jonathan Parker, Annette Vissing-Jørgensen, my advisor Lars Peter Hansen, the associate editor, and two anonymous referees. Financial support from the Margaret Reid Fund is gratefully acknowledged.

APPENDIX A: THE EARNINGS PROCESS

Figure A.1 displays the deterministic component of earnings. The dotted line represents college graduates; the straight line, households with a high school degree; and the dashed line, for those without. As mentioned in the text, the productivity growth factor is calibrated from total workers' compensation data. The age and education component is estimated from a regression of CEX data about log earnings on a fourth-degree polynomial in age (separately for each educational group). The two panels correspond to the two measures of earnings described later. The drop at age 65 is generated by the assumption that households retire after that age. The regression is estimated separately for working and retired household. A household is defined as retired if the head works for less than 500 hours and is older than 60.

Income and expenditure data from the Consumer Expenditure Survey, in the extracts prepared by Sabelhaus, available on the National Bureau of Economic Research website at http://www.nber.org/ces_cbo.html are used. As explained in the main text, because marital risk is not considered, only households with a head and a spouse are considered. For retired people, however, all households are considered, to capture the fact that one of the spouses may die, but savings decisions may be made considering the welfare of the survivor. Observations with

missing demographic information, those with an incomplete report for any quarter, and those with income of less than \$500 1992 dollars are excluded. Apart from very few cases of business or financial losses, zero earnings most likely represent inaccurate data, given the inclusion of various transfer income in the measure of earnings.

Two definitions of earnings are used, depending on the experiment. For the experiment with net worth, earnings represent all sources of nonasset income, net of taxes. The CEX reports federal, state, and social security taxes paid by the household. The ratio λ of labor income over total income (setting asset income to 0 if it is negative) is computed, and such fraction of the taxes is assigned to labor income. Then property taxes, pension contributions, medical, and educational expenditures, are subtracted from this measure.

Health expenditures should be considered a negative shock to income. For instance, Hubbard et al. (1995) explicitly modeled a random process for medical expenditures, which adds another dimension of variability to the stream of income. For simplicity, medical expenditures are subtracted when constructing the average profiles for income, but the randomness of such expenditures is not considered explicitly. Such randomness can be an important reason for precautionary savings, above all for the old. Educational expenditures are excluded because they are to some extent precommitted, and there is an important life cycle component in their timing; that is, they cannot be smoothed. Agents save at the beginning of the life cycle to finance these expenditures. Therefore, they are subtracted from income and considered an (anticipated) negative shock to earnings.

The second definition of earnings, that used when looking at net worth excluding primary housing, subtracts from the income measure explained above also all expenditures related to the accumulation of housing wealth, that is, mortgage payments, housing additions, and rents. Other expenditures on housing, such as heating and cleaning, are not excluded, because such expenditures can be considered current consumption. Note that mortgage down payments are also subtracted from income. Because of this, there are some households for

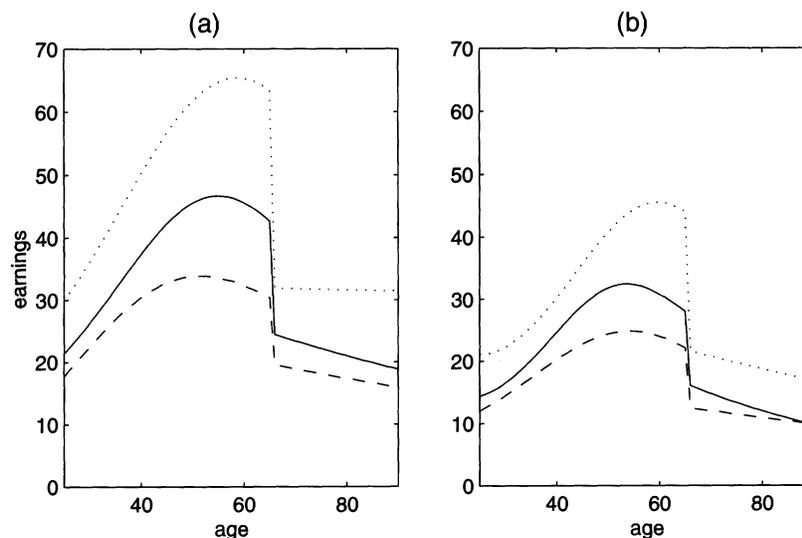


Figure A.1. Estimated Earnings Profiles for the Three Educational Groups, in Thousands 1992 Dollars, CEX 1980–1995. (a) Profile used in the estimation with net worth. (b) Profile used in the estimation with net worth excluding housing.

whom the implied measure of earnings is below \$500 or negative (approximately 6% of the sample). Because logarithms are taken, such observations are excluded.

For the stochastic component of earnings η_t , the results of Carroll and Samwick (1997) are used. They estimate a slightly different process,

$$u_t = u_{t-1} + \hat{\eta}_t + \hat{\epsilon}_t - \hat{\epsilon}_{t-1}.$$

Basically, the innovation in the present formulation is an MA(1) process, whereas Carroll and Samwick (1997) gave it a permanent ($\hat{\eta}_t$) and a purely transitory component ($\hat{\epsilon}_t$). The two formulations are observationally equivalent, in the sense that they generate the same autocovariances for u_t . If a process with a certain σ_η and ψ (the parameters that characterize the MA(1) formulation) is assumed, then an equivalent process of the form assumed by Carroll and Samwick (1997), that is, a process that for a certain σ_ϵ and $\sigma_{\hat{\eta}}$ generates the same variances and autocovariances, can be found. Given the estimates of σ_ϵ and $\sigma_{\hat{\eta}}$ of Carroll and Samwick (1997), σ_η and ψ thus can be recovered. To do so, the variance and the first-order autocorrelation of Δu_t (a system of two equations in the two unknowns σ_η and ψ) are equated.

To simulate the model, I need to initialize the cross-sectional distribution of earnings at age 25 must be initialized. For simplicity, this distribution is assumed to be lognormal, with a mean and a variance equal to the empirical ones; a grid of initial income levels (5 in the present simulation) is then chosen using Gauss–Hermite gridpoints and weights. For instance, for college graduates, this implies a weight of .53 at the mean of \$30,000, with the highest initial point being \$95,000, with weight .01. Given the initial level of income, the initial wealth-to-income ratio is then initialized at its observed distribution. Three initial values are chosen, corresponding to the quantiles .17, .5, and .83 of the wealth-to-income ratio.

APPENDIX B: BEQUESTS AND DEMOGRAPHIC CORRECTIONS

Figure B.1 plots the probability that a household receives a bequest as a function of age. The data is taken from reported

gifts and bequests in the PSID. This probability is modeled with a probit,

$$\Pr(\text{bequest}) = \Phi(\alpha_e * \text{educ} + P(t)),$$

where Φ is the normal cumulative distribution function, educ is a dummy variable for education, and $P(t)$ is a third-degree polynomial in age. This formulation implies that the life cycle profiles of the probabilities for different educational groups differ only by a multiplicative constant. An attempt was made to estimate education-specific age coefficients, but the estimates were very imprecise.

The median bequest-to-income ratio (conditional on receiving a bequest) was also estimated, and it was assumed that when receiving an inheritance, the household receives this proportion of his income. This ratio does not seem to depend significantly on age, so it is considered a constant. The ratios for the three educational groups are 1.16, 1.07, and .96. Note that the PSID does not capture the bequests left by very rich households, which, as noted by, for instance, Kotlikoff and Summers (1981), constitute a significant fraction of the total wealth in the economy. These bequests are fundamental in explaining the concentration of wealth in the right tail of the distribution, as shown by De Nardi's (2000) model of wealth accumulation with intergenerational linkages. However, this article focuses on the median wealth holding.

As suggested by Attanasio et al. (1999), the discount factor β is corrected to capture the effect of demographic variables on consumption. The correction has the exponential form, $e^{Z_t\theta}$. Z includes log family size and log hours of leisure of the spouse (computed as 5,000 minus the number of hours worked in the year). The estimate for θ is taken from Attanasio et al. (1999). The average profile of the two demographic variables over the life cycle, \bar{Z}_t , is computed and then the implied corrections, $e^{\bar{Z}_t\theta}$, shown in Figure B.2, are calculated. For simplicity, this discount factor is assigned to all households in the simulations instead of using different demographic profiles for each household. Thus there is no uncertainty coming from the discount factor.

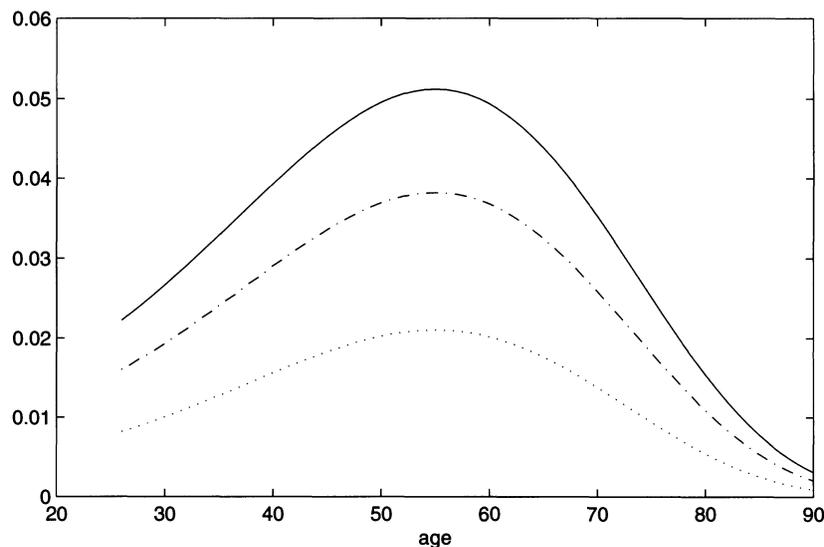


Figure B.1. Estimated Probability of Receiving a Bequest, PSID 1969–1989 (— college; - - - high school; no high school).

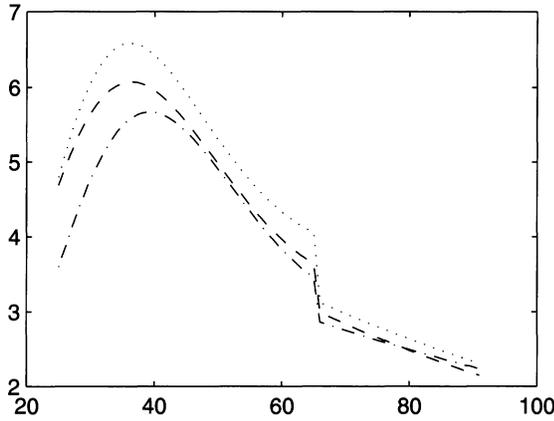


Figure B.2. Discount Factor Corrections (... no high school; --- high school; - · - · college).

Note that the shape differs between education groups only because the profile of the demographic variables is different (due mainly to the fact that college graduates have children later). There is no education dummy in θ .

APPENDIX C: THE ESTIMATION METHOD

To exploit the median restriction (10), the method described by Powell (1994) is used. Powell (1991) gave conditions for consistency and asymptotic normality.

Let $\xi = [\beta, \gamma]'$ be the set of preference parameters to be estimated. For each age group t (in this case, for each of the 5-year age groups), the π th quantile can be defined as

$$E(\pi - \mathbf{1}(w_i^t \leq m^t(\xi))|t) = 0, \tag{C.1}$$

where w_i^t is the wealth of an individual i belonging to group t , m^t is the π th quantile of the distribution of wealth for each age group, and $\mathbf{1}$ is the indicator function. m^t also solves the loss function

$$\min_m E(\rho_\pi(w_i^t - m^t(\xi))|t) \tag{C.2}$$

and

$$\rho_\pi(y) = y(\pi - \mathbf{1}(y < 0)), \tag{C.3}$$

and the true parameter ξ_0 solves

$$\min_\xi E(\rho_\pi(w_i^t - m^t(\xi))|t). \tag{C.4}$$

The previous restrictions are conditional on the age group t . The unconditional counterpart,

$$\min_\xi E(\rho_\pi(w_i^t - m^t(\xi))q(t)), \tag{C.5}$$

where $q(t)$ is some weighting function, can be formed. The sample analog is

$$\min_\xi \sum_{i=1}^N \omega_i \rho_\pi(w_i^t - m^t(\xi))q(t), \tag{C.6}$$

where t is the age group of agent i , N is the total number of observations, and ω_i is the weight of observation i in the entire population.

Because no analytic expression exists for m , m must be simulated; then the cumulative distribution function between simulated points is linearly interpreted. The variance of the estimator due to the simulation is not considered; however, the life cycle profiles are simulated for 10,000 households while each age group contains from 20 to 100 observations, so the ratio of observations to simulated points is extremely low.

An efficient choice of $q(t)$ is $f(m^t|t)$, that is, the density of the distribution of wealth of agents of age t at the median. With this choice, the estimator $\hat{\xi}$ has a distribution

$$\sqrt{N}(\hat{\xi} - \xi_0) \xrightarrow{d} N(0, \Omega) \tag{C.7}$$

and

$$\Omega = \pi(1 - \pi) \left(E \left(f^2(m^t) \frac{\partial m^t(\xi)}{\partial \xi} \frac{\partial m^t(\xi)}{\partial \xi'} \right) \right)^{-1}. \tag{C.8}$$

First, $q(t) = 1$ is used to get a consistent estimate $\hat{\xi}_1$ of ξ_j ; then $\hat{f}(m(\hat{\xi}_1)|t)$, where \hat{f} is the empirical density, is computed and this value is used in a second round of estimation. \hat{f} is a nonparametric estimate with Gaussian kernel.

Experiments using means (condition 13) instead of medians were also performed, using a simulated method-of-moments estimator. Let w_i be the observed wealth of individual i . $W^t(\xi)$ is the (theoretical) mean wealth of individuals of age t for the parameter value ξ , and $W_s^t(\xi)$ is the simulated average wealth level of individuals of the age group t , again for parameter value ξ . The moment conditions (one for each age group t) are

$$E(\mathbf{1}(t)(w_i - W^t)) = 0, \tag{C.9}$$

where $\mathbf{1}()$ is the indicator function, which has value 1 if t_i , the age of individual i , equals t and 0 otherwise. The expected value is not conditional on age (i.e., it is taken over all households).

Observations taken from the datasets are weighted differently (see Sec. 4.2); therefore, one must construct the empirical counterpart to the foregoing moments using weights. Let \hat{w}_i^t be the wealth observation on a household that has age t_i and a weight ω_i in the total sample (the weights are here normalized to 1). The empirical counterparts to the foregoing moments are

$$\hat{M}^t = \sum_i \omega_i (\hat{w}_i^t - W_s^t). \tag{C.10}$$

In the first stage, the identity weighting matrix is used to minimize

$$\hat{M}' I \hat{M} \tag{C.11}$$

with respect to ξ . Using the computed $\hat{\xi}$, an estimate $\hat{\Omega}$ of the covariance matrix can be constructed,

$$\Omega = \text{var}(\mathbf{1}(t)(w_i - W_s^t)). \tag{C.12}$$

The off-diagonal elements are 0. The elements on the diagonal are estimated by

$$\sum_i \omega_i (w_i^t - W_s^t)^2. \tag{C.13}$$

$T = \Omega^{-1}$ is the optimal weighting matrix, so in the second stage

$$\hat{M}' \hat{\Omega}^{-1} \hat{M} \tag{C.14}$$

is minimized. The distribution of the resulting estimate $\hat{\xi}$ is

$$\sqrt{N} (\hat{\xi} - \xi) \xrightarrow{d} N(0, Q), \quad (\text{C.15})$$

where N is the number of observations, and letting τ be the ratio of the number of observations to the number of simulated points,

$$Q = (1 + \tau) (D' \Omega^{-1} D)^{-1} \quad (\text{C.16})$$

and

$$D = E \left(\mathbf{1}(t) \frac{\partial W_s^t}{\partial \xi} \right), \quad (\text{C.17})$$

which is estimated by computing the numerical derivative

$$\hat{D}^t = \left(\sum_i \omega_i \right) \frac{\partial W_s^t}{\partial \xi}. \quad (\text{C.18})$$

To test the overidentifying restrictions, one can use

$$N \hat{M}' \hat{\Omega}^{-1} \hat{M} \xrightarrow{d} \chi_7^2. \quad (\text{C.19})$$

[Received October 2001. Revised October 2002.]

REFERENCES

- Attanasio, O. P., Banks, J., Meghir, C., and Weber, G. (1999), "Humps and Bumps in Lifetime Consumption," *Journal of Business and Economic Statistics*, 17, 22–35.
- Attanasio, O., and Low, H. (2000), "Estimating Euler Equations," Technical Working Paper 253, National Bureau of Economic Research.
- Attanasio, O. P., and Weber, G. (1995), "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence From the Consumer Expenditures Survey," *Journal of Political Economy*, 103, 1121–1157.
- Browning, M., Hansen, L. P., and Heckman, J. (1999), "Micro Data and General Equilibrium Models," in *Handbook of Macroeconomics*, Vol. 1A, eds. J. Taylor and M. Woodford, New York, NY: Elsevier, pp. 543–636.
- Cagetti, M. (2001), "Interest Elasticity in a Life-Cycle Model With Precautionary Savings," *American Economic Review*, 91, 418–421.
- Campbell, J. Y., and Viceira, L. (2002), *Strategic Asset Allocation*, New York: Oxford University Press.
- Carroll, C. D. (1997), "Buffer Stock Saving and the Life-Cycle Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112, 1–55.
- (2000a), "Requiem for the Representative Consumer? Aggregate Implications of Microeconomic Consumption Behavior," *American Economic Review*, 90, 110–115.
- (2000b), "Why Do the Rich Save So Much?," in *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, ed. J. B. Slemrod, Cambridge, MA: Harvard University Press, pp. 466–484.
- (2001a), "Death to the Log-Linearized Consumption Euler Equation," *Advances in Macroeconomics*, 1, article 6.
- (2001b), "A Theory of the Consumption Function, With and Without Liquidity Constraints," *Journal of Economic Perspectives*, 15, 23–45.
- Carroll, C. D., and Samwick, A. (1997), "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40, 41–71.
- Deaton, A. (1991), "Saving and Liquidity Constraints," *Econometrica*, 59, 1221–1248.
- De Nardi, M. (2000), "Wealth Inequality and Intergenerational Links," Working Paper n. 13, Federal Reserve Bank of Chicago.
- Duffie, D., and Singleton, K. J. (1993), "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, 61, 929–952.
- Dynan, K. (1993), "How Prudent Are Consumers?," *Journal of Political Economy*, 101, 1104–1113.
- Engen, E. M., Gale, W. G., and Scholz, J. K. (1996), "The Illusory Effect of Saving Incentives on Savings," *Journal of Economic Perspectives*, 10, 113–138.
- French, E. (2000), "The Effects of Health, Wealth and Wages on Labor Supply and Retirement Behavior," Working Paper WP-00-03, Federal Reserve Bank of Chicago.
- French, E., and Jones, J. B. (2001), "The Effects of Health Insurance, Self-Insurance and Liquidity Constraints on Retirement Behavior," mimeo.
- Gourinchas, P. O., and Parker, J. A. (2002), "Consumption Over the Life Cycle," *Econometrica*, 70, 47–89.
- Hubbard, R. G., Skinner, J., and Zeldes, S. P. (1995), "Precautionary Savings and Social Insurance," *Journal of Political Economy*, 103, 360–399.
- Juster, F. T., Smith, J. P., and Stafford, F. (1999), "The Measurement and Structure of Household Wealth," *Labour Economics*, 6, 253–275.
- Kotlikoff, L. J., and Summers, L. H. (1981), "The Role of Intergenerational Transfers in Aggregate Capital Accumulation," *Journal of Political Economy*, 89, 706–732.
- Krusell, P., and Smith, A., Jr. (1998), "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896.
- Laibson, D. I., Repetto, A., and Tobacman, J. (1998), "Self-Control and Retirement Savings," *Brookings Papers on Economic Activity*, 1, 91–172.
- Lawrance, E. (1991), "Poverty and the Rate of Time Preference: Evidence From Panel Data," *Journal of Political Economy*, 99, 54–77.
- Ludvigson, S., and Paxson, C. H. (2001), "Approximation Bias in Linearized Euler Equation," *Review of Economic and Statistics*, 83, 242–256.
- National Center for Health Statistics (1998), *Vital Statistics of the United States, 1995*, Vol. II, Mortality, Part A, Sec. 6, Life Tables, Hyattsville, MD: author.
- Poterba, J. M., Venti, S. F., and Wise, D. A. (1996), "How Retirement Saving Programs Increase Saving," *Journal of Economic Perspectives*, 10, 91–112.
- Powell, J. L. (1991), "Estimation of Monotonic Regression Models under Quantile Restrictions," in *Nonparametric and Semiparametric Methods in Econometrics and Statistics*, eds. W. A. Barnett, J. L. Powell and G. E. Tauchen, Cambridge, UK: Cambridge University Press.
- (1994), "Estimation of Semiparametric Models," in *Handbook of Econometrics*, Vol. 4, eds. R. F. Engle and D. L. McFadden, New York, NY: Elsevier, pp. 2443–2521.
- Samwick, A. A. (1998), "Discount Rate Heterogeneity and Social Security Reform," *Journal of Development Economics*, 57, 117–146.
- Storesletten, K., Telmer, C., and Yaron, A. (1999), "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," mimeo.
- Summers, L. H. (1981), "Capital Taxation and Accumulation in a Life Cycle Growth Model," *American Economic Review*, 71, 533–544.