

NBER WORKING PAPER SERIES

INTERTEMPORAL CHOICE AND THE
CROSS-SECTIONAL VARIANCE
OF MARGINAL UTILITY

Orazio P. Attanasio
Tullio Jappelli

Working Paper 6560
<http://www.nber.org/papers/w6560>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 1998

We would like to thank James Banks, Richard Blundell, Luigi Pistaferri and Guglielmo Weber for useful comments and suggestions, and Sarah Tanner for help with the UK data. Particular thanks are owed to Angus Deaton and Christina Paxson for interesting discussions and extensive exchanges which, we hope, improved the paper. We retain responsibility for any errors. This paper is part of a research project on "Structural Analysis of Household Savings and Wealth Positions over the Life-Cycle." Financial support has been provided by The Training and Mobility of Researchers Network Program (TMR) of the European Commission DGXII and by the Italian National Research Council (CNR). Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1998 by Orazio P. Attanasio and Tullio Jappelli. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Intertemporal Choice and the Cross-Sectional
Variance of Marginal Utility
Orazio P. Attanasio and Tullio Jappelli
NBER Working Paper No. 6560
May 1998
JEL No. E21

ABSTRACT

The theory of intertemporal choice predicts that the cross-sectional variance of the marginal utility of consumption is equal to its own lag plus a constant and a random component. Using general preference specifications and some assumptions about the nature of the random component, we provide an explicit test of this hypothesis. Our approach circumvents the necessity to identify a pure age profile of the cross-sectional variance of consumption and yields a well-specified statistical test. This test is applied to data from the United States, the United Kingdom and Italy. The results are remarkably consistent with the restrictions implied by the theory of intertemporal consumption choices.

Orazio P. Attanasio
Department of Economics
University College London
Gower Street
London WC1E 6BT
UNITED KINGDOM
and NBER
o.attanasio@ucl.ac.uk

Tullio Jappelli
Via Santa Lucia 173
Napoli 80132
ITALY
jappelli@synapsis.it

1. Introduction

One remarkable prediction of the permanent income hypothesis with certainty equivalence is that consumption inequality within a group of households with fixed membership should, on average and over long periods of time, increase with age. By this model, the change in individual consumption represents the annuity value of the revisions in labor income, which under rational expectations are unpredictable. Then, if income shocks are not perfectly correlated within the group, the cross-sectional variance of consumption will increase with age until retirement, i.e. until uncertainty is resolved. This prediction of the theory has recently been investigated by Deaton and Paxson (1994), who estimate the age-profile of the cross-sectional variance of consumption using average cohort data for the US, the UK and Taiwan and find that consumption inequality increases with age in all three countries.

If certainty equivalence is relaxed the permanent income hypothesis provides predictions about the marginal utility of consumption, not consumption itself. In particular, the model does not generate an explicit relation between age and consumption inequality or, for that matter, between consumption and income. We thus propose to focus directly on the Euler equation for consumption, and consider the time-series properties of the cross-sectional variance of an approximation of the marginal utility of consumption rather than consumption. Under the identifying assumptions discussed in Section 2, the theory implies that in a regression of the cross-sectional variance of the marginal utility of consumption on a constant and its own lag, the coefficient of the latter is unity. It is this hypothesis that we test in this paper. Our procedure allows us to specify quite flexible preferences without relying on certainty equivalence, does not impose assumptions about the relation between age, time and cohort effects in estimating the age-profile of the cross-sectional variance, and delivers a simple statistical test of a well-specified null hypothesis that allows us to apply standard inference tools.

The benefits are not cost-free, however. The test we propose is based on assumptions about the properties of the residuals of the Euler equation. A rejection of the null hypothesis could therefore be due to the failure of these assumptions, and not necessarily of the permanent income model. Furthermore, since we relax certainty equivalence, we cannot exploit the relation between consumption, income innovations and age. For instance, we cannot check if the increase in consumption inequality slows down around retirement, as implied by the certainty equivalence case. Nor can we attribute the spreading of consumption inequality to permanent and transitory changes in uncertainty, as is proposed by Blundell and Preston (1997).

We use average cohort data to test our hypothesis. In the absence of long panel data on consumption, these data are particularly well-suited to the problem at hand. Although we cannot compute the cross-sectional variance of the same group of individuals over time, average cohort data allow us to track the variance of a representative sample of the same cohort. Another advantage of cohort data is that they allow us to apply an appropriate instrumental variables estimator.

The empirical analysis, presented in Section 3, uses three approximations to the marginal utility of consumption. Initially, we measure marginal utility as the log of expenditure on non-durable goods. We then take into account the life-cycle of family size and define marginal utility as the log of non-durable consumption expenditure per adult equivalent. Finally, we compute marginal utility by relying on available estimates of the parameters of a flexible utility function. These parameters have been estimated with the same data sets used in this paper. Having estimated the marginal utility for each household in the sample, we can then compute the cross-sectional variance of marginal utility of population groups defined by year of birth and test whether the coefficient of lagged on current variance is unity. The advantage of the third procedure is that marginal utility is allowed to depend on a full set of demographic and labor supply variables.

We use three sets of cohort data. The British Family Expenditure Survey (1974-1993) and the US Consumer Expenditure Survey (1980-1992) cover a sufficiently long time span and allow consistent estimation of an Euler equation for consumption. We also use data drawn from the Italian Survey of Household Income and Wealth (1987-

1993); since this is too short a time span to provide reliable estimates of the Euler equation, in this case we do not attempt to measure marginal utility by using a flexible utility function. Section 4 summarizes the results and relates them to the debate over welfare comparisons.

2. The cross-sectional variance of marginal utility

We take off from Deaton and Paxson and their exploration of the implications of the permanent income hypothesis for the evolution of consumption and income inequality. It is well known that if the utility function is quadratic and if the real interest rate equals the discount rate, optimal consumption for individual i follows a martingale:

$$c_{i,t+1} = c_{i,t} + \varepsilon_{i,t+1} \quad (1)$$

The cross sectional variance of the variables in equation (1) is:

$$\text{var}(c_{i,t+1}) = \text{var}(c_{i,t}) + \text{var}(\varepsilon_{i,t+1}) + \text{cov}(c_{i,t}, \varepsilon_{i,t+1}) \quad (2)$$

where variances and covariances are computed over a cross-section of households whose composition is constant over time. If aggregate variables are part of the information set of each agent, Deaton and Paxson show that the time average of $\text{cov}(c_{i,t}, \varepsilon_{i,t+1})$ is zero. Equation (2) then implies that in a group with fixed membership the variance of consumption of individuals aged t is stochastically dominated by the variance of consumption of the same individuals aged $t+1$. In other terms, in a stationary population, consumption inequality increases with age on average and over long periods of time. Since $\varepsilon_{i,t+1}$ represents the annuity value of the innovations in permanent income, consumption inequality increases until individuals face income shocks. Such dispersion should at least slow down for older households if earnings shocks dominate interest rate shocks.

2.1 Relaxing certainty equivalence

Equations (1) and (2) can be generalized to more flexible preference specifications. Suppose that utility is intertemporally separable, and that the instantaneous utility function, defined over non-durable consumption, depends on a set of conditioning variables z (demographic and labor supply, let us say) and on an unobservable component v that captures unobserved heterogeneity. In this case it can be immediately shown that the Euler equation that applies to individual i is:

$$E_i \left[\mu(c_{i,t+1}, \theta' z_{i,t+1}, v_{i,t+1}) (1 + r_{t+1}^k) \beta \right] = \mu(c_{i,t}, \theta' z_{i,t}, v_{i,t}) \quad (3)$$

where $\mu(\cdot)$ is the marginal utility of consumption, r_{t+1}^k the rate of return on a generic asset k , θ a set of parameters of the utility function and β the discount factor. Note that both θ and β are assumed to be time-invariant and common to all individuals. Inspection of equation (3) reveals that if one abandons certainty equivalence, it is not necessarily the case that the cross-sectional variance of marginal utility increases with age.¹

If the instantaneous utility function is isoelastic, that is $U(c, z, v) = \frac{c_{i,t}^{1-\sigma}}{1-\sigma} \exp[-\sigma^{-1}(\theta' z_{i,t} - v_{i,t})]$, the Euler equation (3) can be expressed as a linear function of the logarithm of consumption:

$$\begin{aligned} \ln(c_{i,t+1}) + \theta' z_{i,t+1} - v_{i,t+1} &= \\ \sigma \ln \beta + \ln(c_{i,t}) + \theta' z_{i,t} - v_{i,t} + \tilde{\phi}_{i,t+1} + \sigma \ln(1 + r_{t+1}^k) + \tilde{\varepsilon}_{i,t+1} &= \\ \sigma \ln \beta + \ln(c_{i,t}) + \theta' z_{i,t} - v_{i,t} + \phi_i + \sigma \ln(1 + r_{t+1}^k) + \varepsilon_{i,t+1} & \end{aligned} \quad (4)$$

where σ is the elasticity of intertemporal substitution, $\tilde{\varepsilon}_{i,t+1}$ an expectational error and $\tilde{\phi}_{i,t+1}$ the higher conditional moments of the distribution of consumption in equation

¹A sufficient condition for increasing dispersion of marginal utilities is that the marginal utility is concave and the rate of time preference greater than the real interest rate. See Deaton and Paxson, 1994, p. 461.

(3).² The second equality in equation (4) follows by defining $\varepsilon_{i,t+1} = \tilde{\phi}_{i,t+1} - \phi_i + \tilde{\varepsilon}_{i,t+1}$, where ϕ_i indicates the time-series average of the conditional moments. The cross-sectional variance of the two sides of equation (4) is:

$$\begin{aligned} \text{var}\left[\ln(c_{i,t+1}) + \theta' z_{i,t+1} + v_{i,t+1}\right] = \\ \text{var}\left[\ln(c_{i,t}) + \theta' z_{i,t} + v_{i,t}\right] + \text{var}(\phi_i) + \text{var}(\varepsilon_{i,t+1}) + 2 \text{cov}\left[\ln(c_{i,t}), \phi_i + \varepsilon_{i,t+1}\right] \end{aligned} \quad (5)$$

Equation (5), as such, cannot be measured empirically. For one thing, the parameters θ are unknown; second, the variable v is by definition unobservable; and finally, ϕ includes unknown parameters and moments of a distribution which has not yet been specified. But it does highlight a sharp prediction of the theory, i.e. that the coefficient of the lagged cross-sectional variance of marginal utility equals 1.³ It is precisely because of this property that in the simple case of quadratic utility the cross sectional variance of consumption increases, on average, with age. In the remaining of this section we illustrate how we propose to test this important implication of intertemporal optimization.

To make equation (5) operational, we need to make a number of identifying assumptions. We deal with the first problem (that the θ parameters are unknown) below. For the present we assume that the parameter vector θ is known. To address the fact that v and ϕ are unobservable, one can write equation (4) as:

$$\ln(c_{i,t+1}) + \theta' z_{i,t+1} = \sigma \ln \beta + \ln(c_{i,t}) + \theta' z_{i,t} + \phi_i + \sigma \ln(1 + r_{t+1}^k) + \Delta v_{i,t+1} + \varepsilon_{i,t+1} \quad (6)$$

Assuming that the real interest rate is common across households, the cross-sectional variance of the two sides of equation (6) is given by the following expression:

²For instance, if the interest rate is constant and consumption growth is normally distributed, this term equals the conditional variance of consumption growth.

³The left hand side of equation (4) is not exactly the marginal utility of consumption, as we have divided by σ , the elasticity of intertemporal substitution. If constant across individuals, however, this term represents only a factor of proportionality that does not affect the arguments below. For simplicity we continue to refer to the left-hand side of equation (4) as the marginal utility.

$$\begin{aligned} \text{var}[\ln(c_{i,t+1}) + \theta' z_{i,t+1}] &= \text{var}[\ln(c_{i,t}) + \theta' z_{i,t}] + \\ &\text{var}(\phi_i + \Delta v_{i,t+1} + \varepsilon_{i,t+1}) + 2 \text{cov}[\ln(c_{i,t}) + \theta' z_{i,t}, \phi_i - \Delta v_{i,t+1} + \varepsilon_{i,t+1}] \end{aligned} \quad (7)$$

Only the first term on the right hand side can be measured empirically using cross sectional data. However, the second and third terms can be decomposed into a part that is constant over time (such as the variance of ϕ and its covariance with the other components) and a time-variant component. Equation (7) can then be re-written as:

$$\begin{aligned} \text{var}[\ln(c_{i,t+1}) + \theta' z_{i,t+1}] &= \omega_0 + E(\omega_{t+1}) + \text{var}[\ln(c_{i,t}) + \theta' z_{i,t}] + \omega_{t+1} - E(\omega_{t+1}) = \\ &\omega_0 + \text{var}[\ln(c_{i,t}) + \theta' z_{i,t}] + \omega_{t+1} \end{aligned} \quad (8)$$

where $E(\omega_{t+1})$ denotes the time series average of ω_{t+1} , ω_0 includes the time-invariant terms of the right-hand-side of equation (7):

$$\omega_0 = \text{var}(\phi_i) + 2 \text{cov}[\ln(c_{i,t}) + \theta' z_{i,t}, \phi_i] + 2 \text{cov}[\phi_i, \Delta v_{i,t+1}] + 2 \text{cov}(\phi_i, \varepsilon_{i,t+1});$$

and ω_{t+1} the time-varying components of the cross-sectional variance:

$$\omega_{t+1} = \text{var}(\Delta v_{i,t+1}) + \text{var}(\varepsilon_{i,t+1}) + 2 \text{cov}[\ln(c_{i,t}) + \theta' z_{i,t}, \Delta v_{i,t+1}] + 2 \text{cov}(\Delta v_{i,t+1}, \varepsilon_{i,t+1}).$$

The terms ω_0 and ω_{t+1} are implicitly defined by the second equality in equation (8).

Note that most components of ω_{t+1} do not vanish. An exception is $\text{cov}(\phi_i, \varepsilon_{i,t+1})$, which is zero under the permanent income hypothesis.

2.2 The empirical test

Suppose now that we can identify one or more population groups whose membership is fixed over time. The cross sectional variances in equation (8) can then be estimated by computing the corresponding sample variances within groups at different time periods. This application of synthetic panel techniques yields consistent estimates of the corresponding population variances. We can then run the regression:

$$\text{var}\left[\ln(c_{i,t+1}) + \theta' z_{i,t+1}\right] = \bar{\omega}_0 + \pi \text{var}\left[\ln(c_{i,t}) + \theta' z_{i,t}\right] + \bar{\omega}_{t+1} \quad (9)$$

and test the hypothesis that $\pi=1$.

The use of OLS to estimate equation (9) would produce inconsistent estimates of π , however, for at least three reasons. First, the hypothesis that the cross-sectional covariance of $\Delta v_{i,t+1}$ with the observed component of marginal utility at time t is zero is likely to be violated. Second, the time varying-terms included in $\bar{\omega}_{t+1}$ might be correlated *over time* with the variance term in equation (8). Finally, in a finite sample one never observes the true population moments, but only an estimate of these moments which are therefore affected by measurement error. This induces a further source of bias in the estimate of π .

The issue is therefore that of finding an instrument for $\text{var}\left[\ln(c_{i,t}) + \theta' z_{i,t}\right]$ which is uncorrelated with $\bar{\omega}_{t+1}$ and with the sampling error. Synthetic cohort data suggest an ideal candidate for dealing with the potential bias arising from measurement error: if the samples used in estimating the variances in equation (9) are independent - as they typically are in repeated cross sections - so are the sample errors of subsequent periods. Therefore, $\text{var}\left[\ln(c_{i,t-1}) + \theta' z_{i,t-1}\right]$ is, as far as sampling error is concerned, a valid instrument for the corresponding variance at time t with which it is obviously correlated. To guarantee that the (twice) lagged variance is a valid instrument overall, however, one must assume that $\text{cov}\left\{\bar{\omega}_{t+1}, \text{var}\left[\ln(c_{i,t-1}) + \theta' z_{i,t-1}\right]\right\}=0$, i.e. that the twice lagged variance is uncorrelated with the deviation of $\bar{\omega}_{t+1}$ from its unconditional mean. This is a strong assumption that deserves further scrutiny.

To better understand the issues involved, it is worth to strip the model down to its simplest version, the one with certainty equivalence behind equations (1) and (2), where $\bar{\omega}_{t+1} = \text{var}(\varepsilon_{i,t+1}) + \text{cov}(c_{i,t}, \varepsilon_{i,t+1})$. If aggregate shocks are part of agents' information set, the second term is a genuine innovation and therefore its covariance (over time) with the lagged cross sectional variance is zero under the permanent income

hypothesis. The assumption that $\text{var}(c_{i,t})$ is uncorrelated with the deviations of $\text{var}(\varepsilon_{i,t+1})$ from its time mean, however, is *not* an implication of the permanent income hypothesis. Rather, its validity depends on the nature of the shocks $\varepsilon_{i,t+1}$, on agents' information set as well as on the market structure in which they operate.

Suppose that the expectational error of the Euler equation is described by $\varepsilon_{i,t+1} = \alpha_i \psi_{t+1} + \eta_{i,t+1}$ where ψ_{t+1} is an aggregate shock and $\eta_{i,t+1}$ an idiosyncratic shock assumed to be i.i.d. across individuals and over time and independent of α_i and ψ_{t+1} (notice that the aggregate shock is allowed to affect consumers in different fashion). This representation of the error term encompasses several models.⁴

Given the assumed structure of the error term, $\text{var}(\varepsilon_{i,t+1}) = \psi_{t+1}^2 \text{var}(\alpha_i) + \text{var}(\eta_{i,t+1})$. If $\eta_{i,t+1}$ is indeed i.i.d., then $\text{var}(\eta_{i,t+1})$ is constant over time. If ψ_t is also i.i.d., so that $E(\psi_{t+1}^2)$ is constant over time, $\text{cov}[\text{var}(\varepsilon_{i,t+1}), \text{var}(c_{i,t})] = 0$ and $\text{var}(c_{i,t-1})$ is obviously a valid instrument for $\text{var}(c_{i,t})$. If instead ψ_t is heteroskedastic and/or ψ_{t+1}^2 exhibits a considerable amount of persistence, then $\text{var}(c_{i,t-1})$ may *not* be a valid instrument for $\text{var}(c_{i,t})$. For instance, if the conditional time variance of ψ_{t+1} follows an *ARCH*(1) process, $\text{cov}[\text{var}(\varepsilon_{i,t+1}), \text{var}(c_{i,t})] = 0$ and our instrument is still a valid one; but if ψ_{t+1}^2 evolves according to an *ARCH* of order 2 or higher, or as any *GARCH* process, it is not. Similar arguments apply to $\eta_{i,t+1}$ and to its cross sectional variance.

Although we develop the previous example with quadratic utility, the same considerations apply to the more general model that relaxes certainty equivalence. In particular, when the innovations to the Euler equation can be decomposed into an aggregate shock and an individual-specific shock, the period $t-1$ variance of the marginal utility of consumption is a valid instrument for its period t variance if: (a) the aggregate shock is part of the agents' information set; (b) its second moments do not exhibit persistence over time; (c) the evolution of the cross sectional variance of the

⁴In the case of complete contingent markets, for instance, the error term has the factor structure $\varepsilon_{i,t+1} = \psi_{t+1} + \eta_{i,t+1}$ where $\eta_{i,t+1}$ is i.i.d and independent of ψ_{t+1} .

idiosyncratic shocks is not correlated with the cross sectional variance of consumption. These are not predictions of the permanent income- life cycle model. Thus, our test that $\pi=1$ is a joint test of the model and of these additional assumptions. A finding that $\pi \neq 1$ cannot therefore be taken necessarily as a contradiction of the model; but a finding that $\pi=1$ implies that the model *and* our additional assumptions are consistent with the data.

We also experiment with an alternative instrument for $\text{var}\left[\ln(c_{i,t}) + \theta z_{i,t}\right]$. Theory suggests that in the case of certainty equivalence the cross sectional variance of consumption is an increasing function of age, so that age is likely to be correlated with the dependent variable even relaxing certainty equivalence. However, for age to be a valid instrument, one must assume that it is uncorrelated with the cross sectional variance of $\eta_{i,t}$ and with the time variance of ψ_t . The latter assumption is quite plausible. The former, however, is questionable, especially when one considers periods of the life cycle during which uncertainty changes in a dramatic and yet systematic way. The most obvious example is retirement. For this reason in the empirical section we use age both in addition to and in place of the lagged variance as an instrument for the current variance. We also check if our estimates of π are sensitive to the inclusion of the retired in our samples.

2.3 Measuring marginal utility and its cross sectional variance

The problem that the parameter vector θ is unknown can be handled in several ways. The simplest strategy is to drop the z variables, and to regress the variance of log consumption on the lagged variance. A more satisfactory measure of marginal utility is to assume that z includes only family size and to scale the consumption data using adult equivalent weights. This is equivalent to assuming that the utility function is defined on consumption per adult equivalent, rather than on consumption,

$$U(c, P, N) = \frac{N_{i,t}}{1 - \sigma^{-1}} \left(\frac{c_{i,t}}{P_{i,t}} \right)^{1 - \sigma^{-1}} \text{ and estimate:}$$

$$\text{var} \left[\ln \left(\frac{c_{i,t+1}}{P_{i,t+1}} \right) + \sigma \ln \left(\frac{P_{i,t+1}}{N_{i,t+1}} \right) \right] = \bar{\omega}_0 + \pi \text{var} \left[\ln \left(\frac{c_{i,t}}{P_{i,t}} \right) + \sigma \ln \left(\frac{P_{i,t}}{N_{i,t}} \right) \right] + \bar{\omega}_{t+1} \quad (10)$$

where $P_{i,t}$ is the number of adult equivalents, $N_{i,t}$ the number of household's members and σ the elasticity of intertemporal substitution.⁵

A third possibility is to posit more complex and flexible preference specifications that explicitly allow the utility function to depend on labor supply and demographic variables. Since the preference parameters are unknown, they have to be estimated from the data. We propose the following approach. In a first stage, we use the same cohort data and exploit the orthogonality conditions implied by the Euler equation to estimate the parameters of the utility function. In a second stage, we use the estimated parameters to compute the cross sectional variance in equation (9) and to test the hypothesis that $\pi=1$. In Section 3, rather than performing the first stage, we rely on previous estimates. In particular, we use preference parameters estimated by Euler equations that have been fitted to the same data sets used in this paper.⁶

A final econometric problem is that we estimate equations (9) and (10) for all cohorts simultaneously. Since the population groups (cohorts) might be characterized by different variances in ϕ and v , we always check if the results are sensitive to the introduction of dummies for cohort-specific intercepts.

⁵The second term in the square brackets in equation (2.10) appears if utility is defined in terms of consumption per adult equivalent, but is then summed over the number of *family members* rather than adult equivalents. With log utility, the expression simplifies to one where consumption *per capita* enters the marginal utility. In practice, the inclusion of the second term does not make much difference for the results we obtain.

⁶A more efficient strategy, which we do not pursue here, might be to consider equations (6) and (9) as a simultaneous equations system. This would be equivalent to estimating the parameters of the utility function exploiting the restrictions that the theory implies on the time series of both first and second cross sectional moments. One could then use the over-identifying restrictions (including $\pi=1$) to test the model. This procedure would also have the advantage of explicitly taking into account the fact that the vector θ is unknown.

2.4 Comparison with Deaton and Paxson methodology

Average cohort data lend themselves very well to the problem at hand. By definition, birth cohorts are groups with fixed membership.⁷ The most intuitive way of testing the implications of equation (5) is to track the cross-sectional variance of the marginal utility as the cohort ages. Indeed, in their seminal contribution, Deaton and Paxson constructed average cohort data for the US, the UK and Taiwan, considered the graphical representation of the cross-sectional variance and tested the hypothesis that the variance of (log) consumption increases with age.

Their methodology therefore involves identifying an age profile of the cross-sectional variance. This amounts to disentangling age, time and cohort effects. Since these variables are perfectly collinear ($time = cohort + age$), Deaton and Paxson normalize the time effects to zero, so that the evolution over time of the cross sectional variances of each cohort is explained only by a combination of age and cohort effects.⁸ Without structural and/or out-of-sample information, this normalization is indispensable. Any time trend can be written as combination of age and cohort effects and any age effect can be written as a combination of year and cohort effects.

The normalization is essentially an issue of interpretation when describing average cohort profiles of consumption or wages. In the case at hand, however, the implication of the theory being tested is of a structural relationship between age and the cross sectional variance, and the assumption that all variance trends are explained by cohort and age effects is not particularly appealing.

Our procedure has the advantage of circumventing this issue because it provides a direct test of the central implication of the theory (and of the auxiliary assumptions needed to obtain consistent estimates), namely that the coefficient π in equation (9)

⁷We are abstracting here from differential mortality and immigration. Since our samples are not restricted to couples, we also abstract from differential divorce rates.

⁸A slightly less restrictive assumption would be that the average of the time dummies in the sample period is zero and that the time dummies are orthogonal to a time trend. This is in fact equivalent to assuming that all the trends in the data can be interpreted as a combination of age and cohort effects and therefore, by definition, predictable. As Angus Deaton has pointed out to us, this normalization is similar to the normalization used when accounting for seasonality, when the coefficient of one of the seasonal dummies is set to zero. In our opinion, the only

equals unity. A second advantage is that it tests a well-defined null hypothesis against which standard inference tools can be applied.

The corresponding drawback, of course, is that we do not confront the theory with any particular shape of the age-profile of the variance of marginal utility. In principle, the shape of this profile allows one to discriminate among competing models,⁹ and by assuming certainty equivalence, Deaton and Paxson find a structural relation between the variance of consumption and that of earnings and study how this relation changes over the life-cycle. We are not able to specify such a relation: this is the price we pay for relaxing quadratic utility and using more general preference specifications.¹⁰

3. Empirical results

Our results are presented in three parts. We first describe the data sets used, then provide a graphical illustration of the main feature of the cross-sectional variance of consumption and marginal utility, and finally present the econometric results. Some of the details of the data treatment are given in the appendix.

3.1 Data description

We estimate equations (9) and (10) using three sets of average cohort data. The UK Family Expenditure Survey is the largest dates, including 20 annual surveys (1974-1993). The US Current Expenditure Survey is currently available for 13 years (1980-1992). The Italian Survey of Household Income and Wealth is available from 1984 to 1995; but given the characteristics of this survey, we use only data from 1987 to 1995.

difference is that in the seasonality example one normalizes the *level* of the variable under study, while in the present context one normalizes *its level and changes*.

⁹For instance, full risk sharing suggests that the variance is constant over time, dynastic models that the variance increases over time, and finite horizon models that it increases only up to retirement.

¹⁰In our most general specification we also consider labor supply as a determinant of the marginal utility of consumption. This implies that consumption and leisure are not separable in the utility function so that one cannot consider the income process as exogenous. In this case, even with quadratic utility one cannot explicitly relate the variance of consumption to that of earnings.

Our choice to focus on the US, the UK and Italy was guided mainly by the availability of cohort data for these countries.

The FES

The main motivation for the collection of the Family Expenditure Survey (FES) on the part of the UK Department of Employment is the computation of the weights for the Retail Price Index. In recent years the data set has been used extensively to describe the behavior of UK households and to estimate structural models of consumption behavior. The sample includes around 7,000 households per year (the survey has no panel element). Each household stays in the sample for two weeks, during which it compiles a diary with its expenditure. At the end the diary is collected and further information on various expenditure items during the previous three months (typically durable goods and utilities) is gathered. The survey also takes information on several economic and demographic household characteristics, which range from labor supply to household composition. The quality of the consumption data, thanks especially to the use of the diary method rather than retrospective interviews, is remarkably good.¹¹

As each household stays in the sample for just two weeks, consumption figures are heavily affected by seasonality. This problem is particularly serious for the cross sectional variances we compute. For instance, part of the difference in consumption between a respondent interviewed in December and someone interviewed in August will certainly be due to seasonal effects, not to genuine cross sectional variability. To account for this, we deseasonalize the individual consumption observations before computing the cross sectional variance. The seasonal adjustment is described in the appendix.

The CEX

The Consumer Expenditure Survey (CEX) is run by the US Bureau of Labor Statistics for the same reasons as the FES, namely to compute the Consumer Price

¹¹A recent study edited by Banks and Johnson (1997) compares expenditure and income data in the FES with those of the National Accounts. It finds that controlling for differences in definitions (especially of housing services), reference population, and sampling structure, the FES data track very well their counterparts in the national accounts.

Index, and the size of the sample is roughly similar (around 7,000 households per year). There are, however, two important differences between the two. First, in the CEX there is a short panel element, in that each household is interviewed 5 times over 15 months. The first interview is only a contact interview on which no information is disclosed in the public use tape; in each of the subsequent four interviews the household reports detailed information on expenditures incurred during each of the three previous months. Potentially, that is, each household provides twelve monthly expenditure accounts. Second, the households in the CEX do not compile a diary but answer retrospective questions. Like the FES, the CEX also collects data on income, demographics and several other variables.¹²

The households in the CEX do not complete their cycle of interviews simultaneously. Since interviews are performed every month, households interviewed in different months report information on different periods. For example, a household that completes its last interview in January 1990 provides information on each of the 12 months of 1989; while a household that completes its cycle of interviews in June 1990 covers the period from June 1989 to May 1990. One twelfth of the sample is replaced each month. And of course not all households complete the cycle of 5 interviews.

This sampling structure complicates the computations of the cross sectional variances we are interested in. We cannot compute the cross sectional variance in a given year using the monthly observations as we do in the FES because several observations would refer to a single household. Variation across monthly observations would therefore reflect time series (seasonal or cyclical) rather than cross sectional variability. Estimating annual consumption for each household is complicated by the overlapping structure of the sample: the “year” over which we have information for each household would be different depending on the month in which the household completes its cycle of interviews. Annual estimates are further complicated by the fact that not all households complete the cycle.

¹²Some questions are asked only in the last or in the first and last interview. In addition, the timing of some variables is different from that of expenditure. Most notably, the consumption information refers to each of the three months preceding the interview, while the income and labor supply questions (asked only in the second and fifth interviews) refer to the 12 months preceding the interview.

We consider two possible methods of constructing the cross-sectional variances. The first is to estimate the “calendar year” consumption figure for each household on the basis of the months available, with seasonal adjustment. By this method, most households would provide figures to be used in the computation of the variance in two different years, and this should be taken into account in the choice of instruments. The second procedure involves considering only one monthly observation per household; after seasonal adjustment one can compute the cross sectional variance for all the households in a given quarter or year.

We use the second procedure. Its drawback is that by limiting itself to one observation per household, it wastes useful information; on the other hand, it is very simple to apply and yields a sample whose structure is very similar to that of the FES. To minimize the effect of non-random attrition, the monthly observation we select is the first available. The procedure for seasonal adjustment is the same one used for the FES (detailed in the appendix).

The SHIW

The primary purpose of the Bank of Italy Survey of Household Income and Wealth (SHIW) is to collect detailed data on demographics and households' income and wealth. It uses a representative sample of the Italian resident population; probability selection is enforced at every stage of sampling. Like the FES and the CEX, it samples the household, defined as all persons residing in the same dwelling and related by blood, marriage or adoption. The SHIW consumption data are collected by retrospective questions on broad categories of durable and non-durable consumption during the previous year. This implies that seasonal adjustment is not necessary. Given the different purpose of the survey, the detail and quality of the consumption data is not as good as in the FES or the CEX. Nonetheless, the SHIW consumption data match the trends in the national accounts data reasonably well. The sample size is slightly larger than that of the other two surveys (about 8,000 households).

Until 1984 the age of the household head is available only in wide bands precluding the construction of year-of-birth cohorts. In 1986 the data for expenditures on durable goods is not distinguishable from non-durable consumption. Since 1987 the

survey was conducted every other year. For our purpose, therefore, the only usable surveys are the five surveys taken from 1987 to 1995.

3.2 Measures of marginal utility

In all three surveys we define consumption as total household expenditure on non-durable goods. This excludes expenditures on durable goods, health, education and housing, except that in the case of the SHIW, where it excludes durable expenditure only. Current values are transformed into constant values using a consumer price index. As discussed in Section 2, we use three different approximations of the marginal utility of consumption. The first is simply the logarithm of consumption. The second is the log of consumption per adult equivalent. The third filters the consumption data using the preference parameters estimated by Attanasio and Weber (1993) for the UK and by Attanasio and Weber (1995) for the US. The definition of adult equivalent used in the second procedure assigns a weight of 1 to the first adult, 0.8 to any additional adult and 0.25 to any child. To compute the expressions in equation (10) it is also necessary to know the value of the elasticity of intertemporal substitution. Those reported in the tables below correspond to a value of σ of 0.8. We experimented with alternative specifications, both for the definition of adult equivalents and for the assumption about the elasticity of intertemporal substitution, obtaining very similar results.¹³

The SHIW allows a straightforward estimation of the variance of annual log consumption. The sample size and structure of the FES and CEX allow us to construct quarterly data. The results for the UK and the US using annual observations are similar to those presented below and can thus be omitted for brevity. The only notable difference is that the standard errors for the US based on annual observations are slightly larger than those based on quarterly observations.

¹³For brevity, these are not reported.

3.3 Cohort definition and descriptive analysis

In each survey we restrict the sample to households headed by individuals born between 1910 and 1959 and define 10 groups: cohort 1 includes those born in 1915-19, cohort 2 those born in 1920-24, and so on, up to cohort 10, 1955-59. We define the “age” of the cohort as the median age of each cohort in a given sample period (for instance, the age of cohort 1, born in 1910-14, is 68 in 1980, 69 in 1981, and so on).

Table 1 reports the cohort definition and, for each survey, mean and minimum cell size by cohort. The assumption behind our procedure is that the cross sectional variance refers to the same group of individuals at different points in time. Even though we do not need to observe the same individuals over time, the test requires that the sample composition be constant. The positive correlation between survival probabilities and wealth implies that rich households are over-represented in the older cohorts. The correlation between wealth and young headship (young working adults living independently are likely to be wealthier than average) suggests that rich households may also be over-represented in the youngest cohorts. Accordingly, we exclude households where the head is older than 80 years or younger than 20.

The construction of cohort data involves taking means of the marginal utility within each of the cells over successive time periods. Aggregation, however, is not an issue in this context. Each of the methods of measuring the marginal utility of consumption first defines an index of marginal utility at the household level; the cross-sectional variance of this index can then be readily computed within each cell.

Before presenting the regression results, we provide a graphical exposition of the data. In Figures 1, 2 and 3 we plot the cross-sectional variance of the logarithm of consumption for the UK, the US, and Italy, respectively. Each connected segment refers to a different cohort, observed over time at different ages. The cohort numbers, let us recall, run from 1 (the oldest) to 10 (the youngest). To facilitate the comparison across countries, the graphs are plotted on the same scale. These graphs are similar to the raw

variances produced by Deaton and Paxson (see their Figures 2 and 3 for the US and the UK, respectively).¹⁴

As we stress in Section 2, interpreting these graphs is not easy because cohort, age and time are perfectly collinear variables. Their distinct effects are not identified: for instance, any interpretation of the data in terms of age and cohort effects (normalizing time effects to zero) can be recast in terms of an alternative decomposition in terms of age and time effects (normalizing cohort effects to zero). Nonetheless, some features of the raw data are highly suggestive.

Interestingly, the absolute level of the variance is considerably higher in the US (0.4 on average), lowest in Italy (0.25) and intermediate in the UK (0.3), reflecting very different patterns of consumption inequality in these countries. In the UK there is some evidence of cohort effects, especially for households born after the Second World War (cohorts 7 to 10 in Figure 1). In the US and Italy cohort effects are absent or weak at best. In the UK and Italy the variance of log consumption increases up to age 60, and diminishes for older households, which Deaton and Paxson interpret as evidence in support of models with finite horizon, where income shocks cease after retirement. The graph for the US, by contrast, is much flatter over the entire life-cycle.

In Figures 4, 5 and 6 we plot the cross-sectional variance for consumption defined in terms of adult equivalent, assuming an intertemporal elasticity of substitution of 0.8 (see equation 10). The behavior of these variances is rather different than in the previous figures in all three countries. First of all, the variance is now considerably smaller, especially in the UK. Second, the age pattern is much flatter than in the previous set of figures suggesting that variation in household size and composition over the life-cycle explains a good part of the age-profile of the cross-sectional variance of consumption. In the UK and the US there is no evidence of the spreading of consumption over the life-cycle. In Italy the cross-sectional variance declines slightly with the age of the household head, reversing the pattern in Figure 3.¹⁵

¹⁴As mentioned, in the estimation for the UK and the US we use the variance of quarterly (log) consumption. To have legible pictures, we plot the variance only for the first quarter of each year and for five cohorts (2, 4, 6, 8, and 10).

¹⁵Using a flexible specification of preferences to impute marginal utility further reduces the cross-sectional variance in the UK and the US. There is again weak evidence of age effects in consumption inequality only in the UK. For brevity, these figures are not reported.

The last set of figures summarizes the age effect in the cross-sectional variance according to the different measures of marginal utility available in each of the three countries. The age effect is obtained by regressing the cross-sectional variance against a fifth-order age polynomial, a full set of cohort dummies and a set of time dummies constrained to sum to zero and to be orthogonal to a time trend.¹⁶ The estimated age-profiles show that consumption inequality increases with age only for the variance of log consumption. Even in this case, however, the age effects show little spreading of consumption. In the UK, the variance increases only from 0.2 to 0.4 (see Figure 7), in the US from 0.35 to 0.5 (Figure 8) and in Italy from 0.15 to 0.30 (Figure 9). If instead consumption is defined in terms of adult equivalent, or if we directly impute marginal utility, we find very flat, or even declining, age profiles in each of the three countries.

The pattern of the estimated age effects contrast with the findings of Deaton and Paxson, particularly in the case of the US. Using the CEX, they find that the age effects in the US increase from about 0.2 to 0.5 for both the variance of the logarithm of consumption and the variance computed defining consumption per adult equivalent. Part of the difference may derive from differences in the treatment of the data, but not by the different sample periods.¹⁷ They also use quarterly cells, but limit the analysis to the 1980-90 period. However, computing the age effects dropping the 1991-92 years does not change the profiles in Figure 10.

Only under special assumptions does the theory delivers a closed form solution for consumption and allows to test the hypothesis that consumption inequality increases with age on average and over long periods of time. Furthermore, such a prediction of the theory is testable only by arbitrarily normalizing the time effects. Using Deaton and Paxson's methodology, this prediction of the theory is rejected in each of the three data sets when we take into account the fact that part of the inequality in consumption over the life-cycle is explained by variation in household size and composition. We now turn to our alternative - and in some respect more general - test of the implication of

¹⁶This procedure differs slightly from that used by Deaton and Paxson because we use a polynomial in age rather than age dummies. This, however, does not affect our results in any substantive way.

¹⁷Before computing the cross sectional variance we adjust our consumption figures using a model that allows monthly seasonal effects. Furthermore, to avoid complicated MA structures

intertemporal choices for the evolution of the cross sectional variance of the marginal utility of consumption.

3.4 Regression results

We regress the variance of log consumption on its own lag (see equation 9) for the three countries for which we have data. Table 2 reports the point estimates, the standard errors and the p-value of a test that the coefficient of the lagged variance equals one. The basic specification for the variance of marginal utility also included a full set of cohort dummies. However, an F-test for the significance of the cohort dummies was never statistically different from zero, regardless of the measure of marginal utility. The cohort dummies are therefore dropped in the final specifications.

As is shown in Section 2, OLS estimates of equation (9) are potentially subject to measurement error and endogeneity bias. Thus, for each specification we report three regressions, each using a different set of instruments. The first set includes only the second lag of the cross-sectional variance. The second replaces the lagged variance with the (median) age and age square of the cohort. The final set includes both age, age square and the second lag of the variance.

None of the regressions reported in Table 2 rejects the null hypothesis that the coefficient of the lagged variance equals one, with the exception of the last regression for Italy, where the test rejects the null hypothesis at the 5 percent level. The coefficient of the lagged variance are generally estimated with small standard errors, particularly in the case of the UK. Furthermore, the point estimates for the US and the UK are remarkably close to one and similar across different sets of instruments. The results are not affected by excluding the youngest and oldest cohorts or cohorts whose cells are not very large. For instance, excluding retired households (older than 57 years) does not affect the estimate of π , regardless of the instruments used in the estimation.

due to the overlapping structure of the sample, as well as biases induced by non random attrition, we only use the first interview completed by each household.

In Table 3 the dependent variable is the cross-sectional variance of log consumption per adult equivalent, as in equation (10).¹⁸ The estimates are not as precise as in Table 2, particularly for the US, but the general pattern is confirmed. The magnitude of the estimated coefficients for the UK and the US is hardly affected, while in Italy the point estimates are much closer to one. In no case can the hypothesis that $\pi=1$ be rejected at conventional significance levels. The results are again robust to the exclusion from the sample of the oldest cohorts.

In Table 4 we compute marginal utility using the specification $U(c, z) = \frac{c_{i,t}^{1-\sigma}}{1-\sigma} \exp(-\theta'z_{i,t})$, where z includes family size and indicator variables for blue-collar workers, heads out of employment, full-time working spouse, and more than two income recipients. This specification of the instantaneous utility function has been estimated by Attanasio and Weber (1993, 1995) on the same FES and CEX cohort data. Thus, we can readily use their estimates of θ to compute marginal utility.¹⁹ The actual values of the estimates and the details about the construction of marginal utility are described in the appendix. We do not attempt to fit a flexible Euler equation to Italian data. Due to the few surveys and the fact that we work with annual data, the number of observations is much lower than in the other two countries. The point estimate of π for the UK and the US in Table 4 are again very close to unity, regardless of the instrument set. Given also the small standard errors of the estimate, in no case we reject the hypothesis that $\pi=1$. The point estimates obtained by dropping the retired do not show again appreciable differences.

¹⁸ When the utility function is defined in terms of adult equivalent, marginal utility depends on the elasticity of intertemporal substitution. The results in the tables refer to the case in which the parameter equals 0.8. Different values yield similar results.

¹⁹ If the intertemporal elasticity of substitution σ is constant across households, it represents only a factor of proportionality which does not affect the regression results. Attanasio and Weber (1993) also consider seasonal shifts of preferences. Since these shifts are assumed to be

4. Conclusions

The life-cycle hypothesis predicts that the cross-sectional variance of the marginal utility of consumption is equal to its own lag plus a constant and a random component. Using fairly general preference specifications and auxiliary assumptions about the nature of the random component, we provide an explicit test of this hypothesis. Our approach circumvents the necessity to identify a pure age profile of the cross sectional variance of consumption and yields a well specified statistical test. On the other hand, the test we propose is only valid under specific assumptions about the shocks of the Euler equation for consumption.

Cohort data for the UK, the US and Italy provide strong support for the hypothesis that the coefficient of the lagged cross-sectional variance of marginal utility is equal to one, i.e., the joint hypothesis that the theoretical restrictions on the evolution of the cross-sectional variance and the auxiliary identification assumptions we use are consistent with the data.

As we have repeatedly stressed, our test does not assume certainty equivalence. Rather, we are able to confront with the data a flexible representation of preferences, which allows for departures from certainty equivalence, dependence on family size, labor supply and other demographic variables. This increased generality, however, has a price: we lose the ability to relate the evolution of the cross-sectional variance of consumption to the evolution of the cross-sectional variance of earnings. This relation has proved to be particularly useful in welfare comparisons among households.

Blundell and Preston (1997) study under which conditions consumption is a better indicator of welfare inequality than income. They note that in the life-cycle model, cross-sectional comparisons of consumption inequality are informative about welfare inequality only on restrictive assumptions concerning the form of the utility function. They also show that under certainty equivalence the evolution of the variance of consumption should reflect permanent income innovations; whereas the variance of earnings should reflect permanent as well as transitory changes in uncertainty. Our study is complementary to theirs in that it shows that one of the most general implications of multiplicative and constant across consumers, they do not affect the computation of the cross

individuals' intertemporal choices for the evolution of the cross sectional variance is borne out by the data, particularly in the case of the UK.

Dropping certainty equivalence, only simulation analysis can give an idea of the structural relation between the cross sectional variances of earnings and consumption (examples of such simulation models are in Hubbard, Skinner and Zeldes, 1995). In principle, these models could simulate the consumption behavior of a generation and compute the evolution of the cross sectional variance of marginal utility. By changing the stochastic properties of the income generating process it would then be possible to study the relation between income and consumption inequality. We regard this as an interesting topic for future research.

Appendix

1. Seasonal adjustment

In the FES and in the CEX consumption is observed at a frequency higher than annual. In the FES consumption refers to the two-week period over which the household is interviewed; in the CEX it measures the flow of expenditure in each of the 3 months preceding the interview. This feature of the data makes seasonal adjustment indispensable: we must avoid mixing seasonal variations with the cross sectional variability of consumption.

Since in the FES each household is interviewed only once, seasonal adjustment is easier. This is performed using a simple log-linear model of monthly seasonality. In the CEX, each household provides up to 12 observations. There are two pitfalls. On the one hand, we do not want to mix seasonal with genuine cross sectional variability; on the other, we want to avoid having the same household providing more than one observation within each period cell.²⁰ If the interviews were synchronized, it would be easy to construct annual (or quarterly) consumption figures and use those in computing means and variances by year. Given the overlapping nature of the sample, we rely on an alternative procedure. Among several alternatives, we report results for the simplest one: for each household we consider only the first monthly observation. We use the first rather than the last observation to minimize the influence of non-random attrition. The data are then adjusted for seasonality using the same model as in the case of the FES data.

2. Parameters of the Euler equations used to compute the marginal utility

The parameters used in the constructions of the marginal utility of consumption are taken from Attanasio and Weber (1993, 1995). If $U(c,z) = \frac{c_{i,t}^{1-\sigma^{-1}}}{1-\sigma^{-1}} \exp(\theta'z_{i,t})$, the Euler equation for consumption is:

$$\Delta \ln(c_{i,t+1}) = \text{constant} + \sigma(1 + r_{t+1}^k) + \theta' \Delta z_{i,t+1} + \varepsilon_{i,t+1}.$$

Further assuming that z includes demographic as well as labor supply variables, and using quarterly cohort data from the FES (1970-1987), Attanasio and Weber (1993) estimate

$$\begin{aligned} \Delta \ln(c_{i,t+1}) = & 0.747 * r_{t+1} + 0.478 * \Delta wch + 0.2201 * \Delta adults + 0.7086 * \Delta ma + \\ & 0.0063 * \Delta children + 0.0446 * \Delta young\ children - 0.9415 * \Delta oow + \\ & 0.2725 * \Delta ww + \text{seasonal dummies} \end{aligned}$$

²⁰This is because if a household provides observations in different cells, problems arise in the appropriate choice of instruments.

where *oow* is a dummy for the household head out of work, *ma* a dummy indicating more than two adults, *wch* a dummy for white collar workers, and *ww* a dummy for working spouse; *young children* are those of pre-school age (0-5). For details and standard errors, see Attanasio and Weber (1993, column 2 of Table 1 of the Appendix).

For the US, Attanasio and Weber (1995) use quarterly cohort data derived from the CEX from 1981:3 to 1990:4 and estimate

$$\Delta \ln(c_{i,t+1}) = 0.341 * r_{t+1} + 1.172 * \Delta family\ size - 0.539 * \Delta children - 1.551 * \Delta ww + \\ -2.578 * \Delta \ln wl - 2.239 * \Delta single + seasonal\ dummies$$

where *children* is the household members between the ages of 0 and 15, *ww* is a dummy for the wife working full-time and *wl* is the wife's annual hours of leisure (see Attanasio and Weber 1995, Table 3, column 3).

We ignore the real interest rate term, which is only a factor of proportionality and does not affect the regression results. We also ignore the seasonal dummies because they do not affect the computation of the cross sectional variance. The parameters above allow straightforward computation of $[\ln(c_{i,t}) + \sigma\theta' z_{i,t}]$.

References

- Attanasio, Orazio P. (1994), "A Cohort Analysis of Saving Behavior by US Households", in James Poterba (ed.), *International Comparison of Personal Saving*. Chicago: The University of Chicago Press.
- Attanasio, Orazio P. and Guglielmo Weber (1993), "Consumption Growth, the Interest Rate and Aggregation", *Review of Economic Studies* 60, 631-49.
- Attanasio, Orazio P. and Guglielmo Weber (1995), "Is Consumption Growth Consistent With Intertemporal Optimization? Evidence from the Consumer Expenditure Survey", *Journal of Political Economy* 103, 1121-57.
- Banks, James and Paul Johnson eds. (1997), *How Reliable Is The Family Expenditure Survey Data. Trends in Incomes and Expenditure over Time*. London: Institute for Fiscal Studies.
- Blundell, Richard, and Ian Preston (1997), "Consumption Inequality and Income Uncertainty". London, Institute for Fiscal Studies, Working Paper n. 97/15.
- Deaton, Angus (1992), *Understanding Consumption*. Oxford: Oxford University Press.
- Deaton, Angus, and Christina Paxson (1994), "Intertemporal Choice and Inequality", *Journal of Political Economy* 102, 437-467.
- Hall, Robert E. (1978), "Stochastic Implications of the Life-Cycle Permanent Income Hypothesis," *Journal of Political Economy* 96, 971-987.
- Hubbard, Glenn R., Jonathan Skinner and Stephen P. Zeldes (1995), "Precautionary Saving and Social Insurance", *Journal of Political Economy* 103, 360-399.

Table 1
Average and minimum cell size in the FES, the CEX and the SHIW

Cohort	Year of birth	Age in 1985	United Kingdom (FES)		United States (CEX)		Italy (SHIW)	
			Mean	Minimum	Mean	Minimum	Mean	Minimum
1	1910-14	71-75	133	55	65	29	696	582
2	1915-19	66-70	127	81	78	47	764	739
3	1920-24	61-65	160	125	90	60	937	874
4	1925-29	56-60	142	102	88	53	859	838
5	1930-34	51-55	134	106	86	55	918	856
6	1935-39	46-50	141	114	92	66	836	807
7	1940-44	41-45	151	118	114	76	799	766
8	1945-49	36-40	183	134	151	96	714	661
9	1950-54	31-35	156	58	174	128	390	298
10	1955-59	26-30	121	1	201	148	355	284

Table 2
Dependent variable: variance of log(consumption)

Country	Coefficient	Standard error	P-value	Instrument set	Number of observations
United Kingdom	0.945	0.037	0.14	lag2 variance	766
	0.996	0.039	0.91	age, age ²	776
	0.969	0.035	0.38	age, age ² , lag2 variance	766
United States	1.029	0.173	0.87	lag2 variance	500
	0.978	0.099	0.82	age, age ²	510
	0.977	0.101	0.82	age, age ² , lag2 variance	500
Italy	0.798	0.123	0.11	lag2 variance	28
	0.830	0.106	0.12	age, age ²	38
	0.773	0.108	0.05	age, age ² , lag2 variance	28

Table 3
Dependent variable: variance of log(consumption per adult equivalent)

Country	Coefficient	Standard error	P-value	Instrument set	Number of observations
United Kingdom	0.912	0.075	0.24	lag2 variance	766
	1.029	0.112	0.80	age, age ²	776
	0.942	0.070	0.41	age, age ² , lag2 variance	766
United States	0.941	0.480	0.90	lag2 variance	500
	0.975	0.229	0.91	age, age ²	510
	1.014	0.232	0.95	age, age ² , lag2 variance	500
Italy	1.121	0.362	0.74	lag2 variance	28
	0.930	0.190	0.71	age, age ²	38
	0.987	0.185	0.94	age, age ² , lag2 variance	28

Table 4
Dependent variable: variance of marginal utility

Country	Coefficient	Standard error	P-value	Instrument set	Number of observations
United Kingdom	0.924	0.068	0.26	lag2 variance	766
	1.001	0.073	0.93	age, age ²	776
	0.966	0.066	0.57	age, age ² , lag2 variance	766
United States	0.952	0.193	0.81	lag2 variance	500
	0.925	0.217	0.73	age, age ²	510
	0.948	0.163	0.75	age, age ² , lag2 variance	500




Figure 1: United Kingdom




Figure 2: United States




Figure 3: Italy

Var(log consumption), ad. equiv




Figure 4: United Kingdom

Var(log consumption), ad. equiv




Figure 5: United States

Var(log consumption), ad. equiv




Figure 6: Italy




Figure 7: Age effects, United Kingdom




Figure 8: Age effects, United States




Figure 9: Age effects, Italy