

NBER WORKING PAPER SERIES

CONSUMPTION OVER THE LIFE CYCLE
AND OVER THE BUSINESS CYCLE

Orazio P. Attanasio
Martin Browning

Working Paper No. 4453

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1993

This paper was begun whilst Browning was a visiting professor at the Department of Economics, Stanford. We thank the department and the Canadian SSHRC for support. We also thank Tom MaCurdy, Costas Meghir and Guglielmo Weber for useful discussions. We also wish to acknowledge the stimulating comments received from seminar audiences, two anonymous referees and the editor. This study employs UK FES data; this was made available to us by the Institute for Fiscal Studies, London. The data is collected by the UK Department of Employment which bears no responsibility for the analysis below. This paper is part of NBER's research program in Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

CONSUMPTION OVER THE LIFE CYCLE
AND OVER THE BUSINESS CYCLE

ABSTRACT

The main aim of this paper is to assess the validity of the life cycle model of consumption. In particular, we address an issue that has recently received much attention, especially in the macroeconomic literature: that of "excess sensitivity" of consumption growth to income growth. We do this using a time series of cross sections and a novel and flexible parametrization of preferences. The former allows us to address aggregation issues directly, while with the latter we can allow both the discount factor and the elasticity of intertemporal substitution ϵ_{is} to be affected by various observable variables and lifetime wealth.

The main findings can be summarized as follows:

(i) the excess sensitivity of consumption growth to labor income disappears when we control for demographic variables. This is true both at life cycle and business cycle frequencies.

(ii) estimation of a flexible specification of preferences indicates that the elasticity of intertemporal substitution is a function of several variables, including the *level* of consumption. The ϵ_{is} increases with the level of consumption, as expected.

(iii) the variables that change the ϵ_{is} are also important in explaining why we observe excess sensitivity over the business cycle.

(iv) we are able to reconcile our results with those reported both in the macro and micro literature.

(v) in our specification the elasticity of intertemporal substitution is not very well determined. This result, however, should be taken with care, as we have not made an effort to construct a 'preferred' specification, which would probably include additional controls for labor supply behavior.

The evidence presented shows that the life cycle model cannot be easily dismissed. Indeed, we believe that the model does a good job at representing consumption behavior both over the life cycle and over the business cycle.

Orazio P. Attanasio
Stanford University
Stanford, CA 94305
and NBER

Martin Browning
McMaster University
Hamilton, Ontario L8S 4M4
Canada

1. Introduction

The overwhelming majority of empirical 'structural' models of consumption behavior employ representative agent models on aggregate time series data to estimate the parameters of intertemporal allocation. Although the conditions under which we can recover 'structural' parameters from aggregate data are sometimes explicitly invoked, usually no formal justification is given for employing a representative agent model. This is somewhat unfortunate since the conditions under which we can infer something about micro behavior from aggregate data are very stringent (see Stoker (1984)) and are not likely to be met in practice (see Browning (1990) and Blundell, Pashardes and Weber (1992)). If these conditions are not met then none of the time series consumption estimates that one sees in the literature have any obvious interpretations.

In this paper we present some results on the life cycle model of consumption obtained employing a long time series of cross sectional data. These findings have some intrinsic interest and they also support our claim that estimates of structural models derived from aggregate time series data are likely to be very misleading. This is true for qualitative findings (for example, the result that consumption is excessively sensitive to anticipated changes in income) as well as for specific quantitative results (for example, estimates of the elasticity of intertemporal substitution). Our claims are made possible by the nature of the data we use, which allows us to concentrate on the aggregation process and its consequences.

Simple life-cycle models assume intertemporally additive preferences, perfect capital markets¹ and rational expectations. A consistent finding of models of intertemporal allocation estimated on aggregate time series data under the assumption of a representative consumer, is that such simple versions are rejected by the data. These rejections take the form of violations of the overidentifying restrictions implied by the model, 'excess sensitivity' of consumption growth to expected income growth, and implausible values for the structural parameters². Some of these results have been interpreted as indicating the presence of liquidity constraints which prevent a substantial fraction of the population from achieving an efficient intertemporal allocation. It is our belief, based on

¹ As Grossman and Shiller (1982) show, this perfect capital markets requirement can be weakened to the assumption that there is at least one financial asset that has the same borrowing and lending rate and that is used by everyone. We shall return to this below.

² The papers in this area are too numerous to be cited here. Some of the most influential include Flavin (1981), Hansen and Singleton (1982, 1983), Mankiw (1981), Hansen, Rotemberg and Summers (1985), Campbell (1987) and Campbell and Mankiw (1989,1991)

our own earlier work with other authors (see Blundell, Browning and Meghir (1989) and Attanasio and Weber (1992)) and the results presented in this paper, that these rejections may plausibly be attributed to aggregation bias.

There are so many consumption studies in the literature that a justification is needed for presenting one more. In our case the principal justification is that we use micro data. Following Blundell et al. (1989) this allows us to condition on household specific factors that may affect consumption decisions. The most important of these factors are liable to be labor supply and family composition. In what follows we show that such factors have indeed an important influence on intertemporal allocation decisions. Furthermore we show that if one ignores these factors, then the results look very much like the results found on aggregate time series data. We also show that excess sensitivity of consumption to income disappears once we control for the effects of family composition.

The other innovation presented in this paper is a flexible parametrization of the instantaneous utility function. This is of some interest for at least two reasons. First of all, we show that the elasticity of intertemporal substitution varies in a plausible way with observable variables, including consumption. We find that wealthier households find it easier to substitute consumption intertemporally. Second, using our parametrization we can show why controlling for demographic variation which exhibits very low frequency movements, can explain the excess sensitivity results found on macro data at relatively high (quarterly) frequency.

The data we use are drawn from the UK Family Expenditure Survey for the years from 1970 to 1986. After some selection we have 44,334 households. Since this is a relatively little used and rather inaccessible source of data for consumption studies we present a full data description in Section 2. In that section we present the salient features of the data and illustrate how we might draw informal inferences about the life-cycle model. First we estimate age profiles for consumption and income and show that the two are highly correlated. This finding has recently been interpreted as strong evidence against simple forms of the life cycle model³. We then show that this correlation over the life cycle largely disappears when we control for changing family composition and labor supply behavior. The informal evidence in this section anticipates the results from our structural model.

The specification of the utility function is presented and discussed in Section 3.1. This speci-

³ See Carroll and Summers (1990). See also Heckman (1974) and Ghez and Becker (1975).

fication controls in a flexible way for the influence of various household characteristics and allows for non linearities in Engel curves and for substitution amongst goods. As Blundell, Browning and Meghir (1989) emphasize, most consumption studies assume forms for preferences that are at odds with established features of demand; we do not. In section 3.2 we deal with the econometric issues that arise in our model.

In Section 4 we present our estimation results. Our principal conclusion is that although very simple forms of preferences seem to display excess sensitivity of consumption to expected changes in income, this disappears when we allow for the effects of demographics. We present a reconciliation of our results with other micro and macro studies that address the issue of excess sensitivity. Section 5 concludes the paper.

2. The Raw Data

The data used in this study are drawn from the UK Family Expenditure Survey (FES) conducted between 1970 and 1986. In this survey households keep a two week diary of expenditure on all goods. As well as these expenditures a wide range of supplementary information is also gathered. In particular, the survey records the composition of the household and the labor supply of each member. More details on the data and summary statistics are provided in the data appendix.

In all we start with about 120,000 households. We select households that live in England, Scotland and Wales and that are headed by a married couple. We exclude any household in which the husband is self-employed to minimize the contamination of the expenditure data by non-consumption expenditures. For reasons that will become clear below we also exclude any household in which the husband was born before 1920 or after 1949. These exclusions leave us with 44,334 households in our sub-sample.

We define total nondurable expenditure to be the sum of the recorded two week purchases on food, alcohol, tobacco, fuel, clothing, transport (excluding the purchase of any vehicles), other goods and services. The principal excluded goods are durables, vehicles and housing. For each good we also have a monthly price index: from these we construct a household specific price index. This is computed as the weighted geometric mean of the individual good prices (a 'Stone' price index) that takes the household budget shares for weights. In this section we define consumption as total nominal expenditure divided by this price index; we shall return to this issue in the next section.

We identify five broad groups of influences on consumption: lifetime income effects, cohort effects, life cycle effects, cyclical effects and heterogeneity.

It is clear that the level, timing and riskiness of *lifetime income* may affect the level and timing of consumption over the life cycle. This has been so extensively discussed that no more need be said here.

The next set of effects are those that are common to people born in the same period: *cohort effects*. These capture the idea that systematic differences in the social environment in which people grow up may well result in different attitude to risk, discount factors and preferences over the lifetime path of consumption (see Ryder (1965) for the classic statement on the importance of cohort effects). It is entirely plausible that people who came to maturity during the 1930's have

different attitudes to saving to those born after 1945. These effects are obviously related to lifetime income if different cohorts enjoy, maybe because of productivity growth, different levels of welfare.

The third set of influences are *life-cycle effects*. One of the most obvious here is the effect of family composition on consumption: for instance families with older children need to spend more. Another important set of influences are those related to labor supply behavior. If, for example, there are significant costs of going to work then we would expect to see consumption falling on retirement. Less obviously, there may also be pure age effects. For example, if most people have a big celebration on their fiftieth birthday then we should see a small 'blip' in consumption at that age.

The fourth set of influences are *cyclical effects*. These are effects that are common to all agents in the same time period. The most obvious examples are movements in common variables like prices or interest rates. We also include other less well defined general 'shocks'. Although common, we must allow that the effects may differ across households and across cohorts. For example, an interest rate 'shock' has a different effect for indebted families than for wealthy ones.

The final feature of any micro data is the large *heterogeneity* evident in the level of consumption by families that are identical in all other observable characteristics. There is little we can do about this other than including a conventional error term to pick up some of this heterogeneity.

In this section we present some simple descriptions of our data and relate them to the influences described above. Our main aim is to establish if households smooth consumption in the face of an uneven income profile. Given that we do not observe the same households over time, we focus on average cohort data. We start by assigning households to 6 cohorts on the basis of the date of birth of the husband. Households are then assigned to either year-cells (for the analysis in this section) or to quarter-cells (for the analysis in section 4) on the basis of the date of the interview⁴. The exact definition of the 6 cohorts considered is given in Table 1 along with the ages of the various cohorts in 1970 and 1986 and the mean, minimum and maximum cell size of the quarter-cohort cells.

To begin our look at the various influences on consumption we return to the original 44334 observations. For these data we regress log consumption and log real net income on 102 year cohort

⁴ The quarterly time series for cohort consumption, income and the other variables used in section 4 is obtained by averaging the relevant variables over all the households belonging to a given quarter-cohort cell.

Table 1
Cohort Definition and Cell Size

Cohort	Year of Birth	Age in 1970	Age in 1986	Minimum cell size	Maximum cell size	Mean cell size
1	1920-24	46-50	62-66	78	147	112.3
2	1925-29	41-45	57-61	63	126	100.2
3	1930-34	36-40	52-56	64	148	103.2
4	1935-39	31-35	47-51	76	134	107.2
5	1940-44	26-30	42-46	66	154	115.1
6	1945-49	21-25	37-41	53	172	129.2

Consumption and Income Over the Life-Cycle




Figure 1

Unadjusted and Adjusted Consumption Over the Life-Cycle

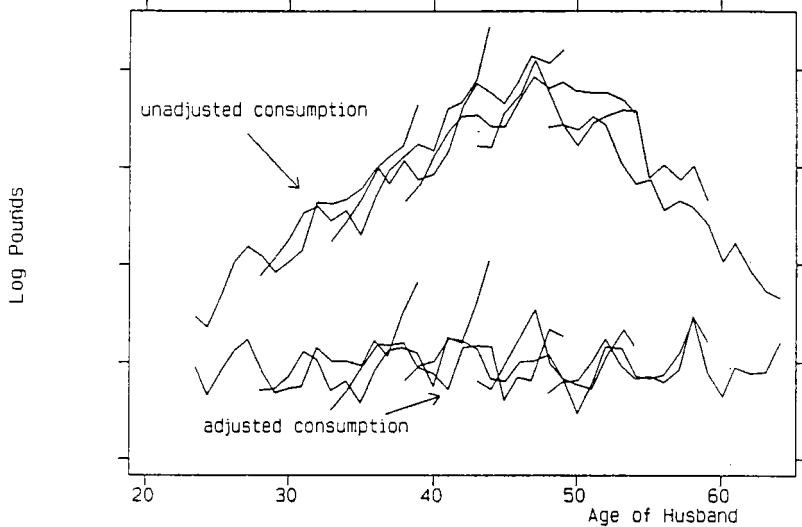


Figure 2

dummies (remember we have six cohorts and 17 years) and 3 quarterly dummies. The year cohort coefficients in this regression correspond to year-cohort means (with an adjustment to allow for seasonality).

These means are most conveniently presented in visual form. In figure 1 we plot the life cycle paths of consumption and income. In this figure each connected segment is the path over the 17 years of the survey for a particular cohort. In the graph we track part of the average consumption age profile for each of the cohorts considered. The life cycle paths given in Figure 1 are familiar: both consumption and income have an inverted U shape and they are highly correlated. This is consistent with the results presented by Carroll and Summers (1990) who interpret them as evidence of lack of consumption smoothing ⁵.

Figure 1, however, does not control for either family composition or labor supply effects. To take into account the effect of demographic factors we regress the year cohort means plotted in figure 1 on the year cohort means of various demographic variables and plot the residuals of this regression in Figure 2 together with the unadjusted year-cohort means. The resulting graph should be interpreted as representing the life-cycle movements of consumption after removing the effect of the demographic variables in the regression ⁶. The demographic variables we consider in this regression are: number of children, number of adults, log of family size, and a dummy which equals one if there is at least one child in the household.

The age profile for consumption after removing the life cycle variation induced by changes in family composition is remarkably flat. While it is still possible to notice considerable business cycle year to year fluctuations, it is apparent that controlling in a simple way for changes in average family composition over the life cycle, eliminates completely the life cycle correlation of income and consumption. Figure 2 provides strong evidence in favor of the hypothesis that consumption is indeed smoothed over the life cycle.

Figures 1 and 2 are concerned with life-cycle allocation. The other focus of this paper is the smoothing of consumption over business cycle fluctuations. Before presenting a formal analysis

⁵ Carroll and Summers (1990) draw life cycle profiles for income and consumption for different occupational groups and notice that the shape of these profiles changes in similar ways across groups. It is interesting to notice that Ghez and Becker (1975) interpret the very same finding as evidence *in favor of* rather than against the model.

⁶ Performing the regression on the year cohort means rather than on the individual data allows us to remove the influence of household specific fixed effects.

Consumption Over the Business Cycle




Figure 3

of high frequency changes in consumption it is useful to plot the residuals of the last regression (which removes life cycle effects), against time rather than age. This is done in Figure 3 where the log of consumption (net of life cycle effects) of each cohort is plotted against time.

There are two notable features to figure 3. First, the consumption series are all very variable. For example, for all cohorts consumption rises by about 7% from 1970 to 1973 and then falls by about 9% to 1977. The aggregate data do not show anything like these swings (although they exhibit a similar cyclical pattern)⁷. There are several reasons for the greater variability of our measure of consumption relative to aggregate consumption data. First of all, our sample does not include the whole UK population. As stressed above, we eliminated all the households whose head was born before 1920 or after 1949 and other demographic and economic groups. Second, and more importantly, the average consumption figures we consider are derived as the *geometric* rather than *arithmetic* average of consumption. Finally, in Figure 3 we remove the part of consumption which is linked to demographic factors, factors that for the whole population are very slow moving (even though they exhibit strong life cycle variations).

The second notable feature of figure 3 is the synchronization in the movements over time for our cohorts. These large and synchronized movements in consumption are ripe for a macro explanation. The most obvious candidates for explanatory variables for these changes are changes in common discounted prices (that is, the real interest rate) and common shocks.

The evidence presented in this section indicates that at a first glance the life-cycle model is not inconsistent with the data. The hypothesis that households smooth consumption over the life cycle in the face of an uneven income profile, is not obviously rejected by the data, once we allow for the effects of changing family composition. Although this suggests that controlling for life cycle events may remove excess sensitivity, it is hardly conclusive. First, we have made no allowance for uncertainty. Second, although correcting for demographics seems to moderate the low frequency (life cycle) excess sensitivity it is by no means clear that we can account for the high frequency (business cycle) excess sensitivity documented by Campbell and Mankiw (1989,1991) among others, in the same way. After all aggregate movements in demographics are likely to be very sluggish so that it is not at all obvious how correcting for such movements will reduce high

⁷ The larger variability at business cycle frequencies of our consumption growth compared to the aggregate time series data has potential implications for the equity premium puzzle. See Deaton (1992) for a discussion.

frequency excess sensitivity. More powerful tests are needed to test the hypothesis that households smooth consumption over the business cycle. The hypothesis that we want to test is that of excess sensitivity of consumption growth to expected high frequency changes in income. These issues are best tackled in the context of a formal structural model. It is to this that we now turn.

3. Theory

3.1 A Structural Model

The conventional approach to structural modeling is to specify a utility function and then to derive the consumption function that it implies. Assume that life time utility is intertemporally separable and let $V(p, z, x)$ be the indirect utility function for each period, where p is a vector of within period prices, x is total expenditure and z is a vector of household characteristics. If p and z are given in discounted terms (using some nominal interest rate) then, under the assumption of the life-cycle model, agents try to keep the marginal utility of money (in this case $V_x(p, z, x)$) constant over time. This then gives a function for current x in terms of lagged (p, z, x) and current (p, z) (see Browning (1989)). The drawback with this procedure is that we usually have to start from very simple forms for the indirect utility function to ensure that the resulting consumption function is tractable.

Here we use a different procedure. We start with a consumption function that has desirable properties (flexibility and tractability) and then 'integrate back' to the indirect utility function to display all of the implications of the preferences that are thus implicitly chosen.

The first property we want is that intertemporal allocation depends on only one price index instead of many prices. Specifically, we model $c = \frac{x}{\alpha(p, z)}$, where $\alpha(p, z)$ is some linear homogeneous price index. The dependence on z allows for the fact that different types of households may be affected differently by relative price changes. For obvious reasons we shall call c 'consumption' and define the indirect utility function $V(p, z, x) = v(\frac{x}{\alpha(p, z)}, z) + \psi(p, z) = v(c, z) + \psi(p, z)$, where ψ is zero-homogeneous in p . Thus the marginal utility of expenditure is given by $V_x(\cdot) = v_c(\cdot) \frac{1}{\alpha(p, z)}$. This formulation does not require homotheticity or the existence of a Hicks composite commodity (see Gorman (1959)).

When using the first order condition for expected utility maximization in an intertemporal framework, we end up considering the first differences of discounted marginal utility of expenditure. It is therefore extremely convenient to model to model $v_c(\cdot)$ or $\ln(v_c)$ as a low order polynomial in known functions of c and z . As we discuss below, this is particularly so when using average cohort techniques since we can take means over period-cohort cells of known functions of the given variables.

There are advantages both in modeling $v_c(\cdot)$ and $\ln(v_c)$. In the former case we avoid the need

of assuming log normality of the residuals to get a linear relationship out of the Euler equation. In the latter case we do not have to impose any restrictions on the parameters to guarantee that marginal utility is positive and it is possible to obtain the much used isoelastic specification as a particular case of our more general formulations. In an earlier draft of this paper, we found that the empirical results were very similar in the two cases so that, for comparison with other studies, we report the results obtained modeling $\ln(v_c)$.

One of the most important parameters that characterizes intertemporal allocation is the elasticity of intertemporal substitution (eis). This gives the proportional change in consumption for an anticipated one percent increase in the discounted price of consumption (which equals the negative of the response to a one percent increase in the real interest rate). For a many good model it is defined as (see Browning (1989)):

$$\phi(p, z, x) = \frac{V_x}{xV_{xx}} = \frac{v_c}{cv_{cc}}$$

Since $V(\cdot)$ is increasing and concave in x this should be negative. This elasticity approaches zeros as agents become increasingly reluctant to substitute between periods. The bounding cases are Leontief preferences ($\phi = 0$) and linear preferences ($\phi = -\infty$).

We model $\ln(v_c)$ as a restricted polynomial in $\ln(c)$ and z :

$$(1) \quad \frac{1}{\sigma} \ln(v_c(c, z)) = \beta' z + \delta \ln(c) + (\gamma' z) \ln(c) + \eta (\ln(c))^2$$

The implied eis in this case is given by the following expression:

$$(2) \quad \phi(p, z, x) = \frac{\sigma}{\delta + \gamma' z + 2\eta \ln(c)}$$

In order to identify the parameters of this equation we use the normalization $\delta = -1$. Notice that if $\gamma = 0$ and $\eta = 0$ we obtain the usual isoelastic case. It is easy to show that integrating the marginal utility function in equation (1) one gets, in this case, an isoelastic indirect utility function. More generally, the indirect utility function is given by the following expression:

$$V(p, z, x) = \frac{1}{\alpha(p, z)} \int_0^x v_y(z, \frac{y}{\alpha(p, z)}) dy + \psi(p, z)$$

where $\psi(p, z)$ is a 'constant' of integration (constant as far as x is concerned, that is). Notice that this specification allows for complex substitution patterns across commodities within each period, and yet guarantees that intertemporal allocation is determined by the single price index $\alpha(\cdot)$.

Notice also that, if η is positive, the absolute value of the elasticity of substitution in equation (2) is an increasing function of the level of consumption, implying that wealthier families find it easier to substitute consumption across periods. Finally, note that the value of the elasticity of intertemporal substitution can vary with family characteristics z ; Blundell et al. (1989) found this to be important.

To derive an equation to estimate we consider the usual Euler equation derived under the assumption that all agents have access to a financial instrument that has the same lending and borrowing nominal interest rate i_{t+1} between periods t and $t + 1$. In that case we have:

$$(3) \quad E_t[\lambda_{t+1}(1 + i_{t+1})] = \lambda_t$$

where E_t denotes the expectations operator conditional on information available at time t and λ is the marginal utility of expenditure $V_x(p, z, x) = v_c(c, z)/\alpha(p, z)$.

Equation (3) can also be written in terms of v_c and the real interest rate:

$$(4) \quad v_{c,t+1}R_{t+1} = v_{c,t}\bar{\epsilon}_{t+1}$$

where $R_{t+1} = \frac{\alpha(p_{t+1}, z_{t+1})}{\alpha(p_t, z_t)}(1 + i_{t+1})$ and $E_t[\bar{\epsilon}_{t+1}] = 1$.

Combining equation (1) and equation (4) and using the normalization $\delta = -1$, we have:

$$(5) \quad \Delta \ln(c_{t+1}) = \text{constant} + \beta' z_{t+1} + \gamma' \Delta(z_{t+1} \ln(c_{t+1})) + \eta \Delta(\ln(c_{t+1}))^2 + \sigma r_{t+1} + \epsilon_{t+1}$$

where $r_{t+1} = \ln(R_{t+1})$ is approximately equal to the real interest rate between time t and $t + 1$ and $E_t[\epsilon_{t+1}] = 0$. The constant term includes various conditional moments of $\bar{\epsilon}_{t+1}$. Under the assumption of log-normality, it contains only second moments. We assume that these moments are constant over time or uncorrelated with the instruments used in estimation.

Equation (5) is the relationship that we will estimate and test in the next section. Notice that this equation is linear in the parameters and that if $\eta = 0$ and $\gamma = 0$ it reduces to the sort of equation which is usually estimated in the literature.

One last issue that remains to be discussed before we move to the econometric problems is that of the presence of the 'true' price index $\alpha(p, z)$ in the definition of the real interest rate that appears in equation (5). There are three approaches to dealing with this. One is to estimate the demand system associated with the given preference structure. This allows us to construct $\alpha(p, z)$ for each household and is the route followed by Blundell, Browning and Meghir (1989) for a different preference structure. This would take us too far away from our principal focus. The next alternative is to parametrize $\alpha(p, z)$ and to estimate these parameters along with (σ, β, γ) . This is likely to result in rather imprecise estimates of the $\alpha(p, z)$ parameters at the cost of increasing quite a bit our parameter size. The final approach is to replace $\alpha(p, z)$ with a known price index. We follow this route and use a geometric mean of prices with the household shares as weights. Blundell et al. (1989) present results indicating that this does not lead to significant bias. Using household specific weights allows for the dependence of $\alpha(\cdot)$ on household characteristics. We denote this index P .

3.2 Econometric Issues

Equation (5) is valid for a single household. Unfortunately we only observe each household once. To estimate equation (5) we have to aggregate across all the households belonging to a given cohort. In this respect the linearity in the parameters of equation (5) is crucial. At the same time the non-linearity of the equation in the variables is not a problem: we can take the cohort means of any known non-linear transformation of the data, we can interact any observable variable, etc. Of course this is not possible with aggregate time series data.

Equation (5) can be aggregated in a number of different ways, as long as the population from which the samples in various time periods are drawn is homogeneous. This is to say that the definition of cohorts is quite arbitrary. We chose to keep 6 separate cohorts for several reasons. First and most importantly, this allows us to isolate possible rejection of the model and impute them to a given cohort, for instance that observed in the early part of the life cycle. Second, it is possible that some variables, in particular demographic variables, do not vary much over time for the population at large, but exhibit strong variations over the life cycle. Therefore the analysis of many tightly defined cohorts allows us to identify more precisely parameters that might be potentially important. Finally, it is possible that the degree of heterogeneity within each cohort is different, inducing, as discussed below, heteroscedasticity. By explicitly allowing for this we may

obtain more efficient estimators.

It should be clear that the consideration of many cohorts simultaneously does not add genuinely new observations that would lead toward the consistency of our estimator. As stressed by Chamberlain (1984), consistency in this framework is only obtained when T , the number of time periods, goes to infinity. The reason for this is that the residuals for the same time period for different cohorts are correlated and do not have zero mean across cohorts.

The consideration of cohort means for the estimation of equation (5) with a cell size less than the population induces MA(1) residuals. This is because the cohort means are computed on the levels, so that taking the first difference of variables affected by measurement errors will induce first order serial correlation. The residuals of the estimated equation (5) will therefore be given by the sum of an expectational error (presumably white noise) and an MA(1) residual with a unit root. The resulting residuals will be (for each cohort) an MA(1) with the sign and size of the first order autocorrelation depending on the relative variance of the two components. In our experience the measurement error component tends to dominate so that we get negative first order autocorrelation.

The presence of MA(1) residuals poses some problems for the choice of the instruments and the computation of the standard errors. In the absence of measurement error, variables known at time t would be valid instrument for the estimation of equation (5). To allow for MA(1) residuals we take instruments dated $t - 1$ and earlier.

The use of second or earlier lags has an important bearing on our test for excess sensitivity. In any period we take a (relatively small) sample of households. If we happen to sample a number (or even just one) with unusually high lifetime income then we will observe high income and high consumption. This is not, of course, inconsistent with the simple life-cycle hypothesis but it will introduce a spurious relationship between changes in income and consumption. By projecting onto variables that are lagged at least twice we purge these variables of their (correlated) measurement error. We shall demonstrate below that despite the presence of considerable sampling errors our predicting equations do track most of our variables quite well.

Strictly speaking the Euler equation (5) holds for each cohort so that we could apply the GMM orthogonality conditions for each cohort separately. On the other hand these cohort specific conditions also imply that the whole sample will satisfy the orthogonality conditions, so that we could just impose these conditions for the pooled sample. Effectively the choice here comes down to

whether we are going to use exogenous and lagged twice variables to predict each set of endogenous variables separately for each cohort or whether we are going to impose the same reduced form. Put another way, the instrument matrix can be formed as a block diagonal matrix of cohort instruments or as a stacked matrix. The 'block' procedure is weaker but has the disadvantage that we can only use a relatively small number of instruments.

Another problem with the consideration of many cohort simultaneously is the computation of the standard errors. The error structure of the pooled equations is fairly complicated. In computing the standard error we allow for first order autocorrelation within each cohort with a different variance and autocorrelation coefficient and for simultaneous correlation across cohorts. While we do allow for different variances (and first order autocorrelation) across different cohorts, we ignore the possibility of within cohort heteroscedasticity induced by variable cell size. This is justified by the fact that most of the variability in cell size is across cohorts (see Table 1).

4. Results

In this section we present our empirical results for the structural model discussed in section 3. In the first subsection we describe the specifications estimated. In the second subsection we present the estimated equations with a minimum of discussion. In subsection 4.3, we discuss in detail the implications of our estimates for tests of the life cycle hypothesis. We consider the robustness of our results and compare them to results obtained both with macro and micro data in the literature. In particular, we explain how ignoring demographics can lead to apparent excess sensitivity in the aggregate time series data, even though aggregate demographic variables display very little high frequency variability. In subsection 4.4 we discuss briefly the inclusion of labor supply variables in the Euler equation for consumption.

4.1 Variable and instrument choice

As stated above the consumption definition used in this paper includes consumption of non durables and services. To estimate equation (5), we also need to specify the z variables and the nominal interest rate.

The z variables fall into three distinct groups. First, there are variables like seasonal and cohort dummies, age and age squared that can all be taken as exogenous and measured without error. Next we have a set of demographic variables that describe the composition of the family. We consider the following variables: the number of children (nch), the number of adults (persons aged over 16) ($adult$), a dummy set to one if the number of children is greater than zero ($cdch$) and the log of total family members ($lnmem$). In one of the regressions we also include two labor supply variables: a dummy set equal to one if the husband works ($dphusb$) and another set to one if the wife works ($dpwife$). When included, the labor supply variables are considered as endogenous and therefore instrumented. While we also tried other demographic and labor supply variables, these did not affect in a substantial way the qualitative results discussed in this section.

In the specification above our z variables (demographics, age and labor supply) enter in two ways (see equation (5)): in levels (that is, as modifiers of the discount rate) and crossed with log consumption (that is, as modifiers of the intertemporal substitution elasticity). We term these two sets of variables 'levels' and 'crossed' respectively. All the interacted variables are considered endogenous and are instrumented.

The nominal interest rate chosen is the rate on Building Society deposits at the end of period

t. We chose this interest rate for two reasons: first, a large number of households have Building Society deposits or mortgages; second, given the tax deductibility of interest payments there is almost no difference between lending and borrowing rates.

The real interest rate (computed as the nominal rate minus the cohort specific inflation rate in the price index for non durables), the change in log consumption and in its square, the change in real disposable income and all the demographic labor supply variables are considered as endogenous and therefore instrumented.

In addition to the variables considered as exogenous we include in the list of instruments the second and third lag of: changes in log consumption and its square, changes in log real disposable income, the real interest rate and the inflation rate. We also used the second lag of all the z variables in levels and interacted with the log of consumption and the same for the labor supply variables.

Given the fairly large number of variables to be instrumented the choice between 'stacked' and 'blocked' instruments was resolved in favor of the former. However we also include in the instrument set cohort dummies and cohort dummies interacted with lagged twice consumption growth to pick up any cohort specific effects.

4.2 Estimation Results

We start by presenting the standard Euler equation for consumption with isoelastic preferences and no controls for demographics and/or labor supply variables. These estimates are reported in the first column of table 2. The results are consistent with the evidence from macroeconomic time series data (see, for instance, Campbell and Mankiw (1989,1991)). The estimated coefficient on expected income growth is large and significant (0.4, s.e. = 0.03), while that on the real interest rate is small and not significantly different from zero. The only difference is the non rejection of the over-identifying restrictions ⁸.

In the second column of Table 2 we add to the previous equation the change in the square of consumption. This turns out to be extremely significant (0.25, s.e.=0.02) indicating that there are substantial deviations from the isoelastic specification normally used in the literature. Notice that

⁸ Attanasio and Weber (1992) show that this is due to the use of the log of the geometric rather than that of arithmetic mean of consumption. Another possibility is that the use of a large number of instruments decreases the power of the test in small samples, if some of them are very noisy and uncorrelated with expected income.

Table 2

Effects of demographics and labour supply variables on the Euler equation for consumption

	1	2	3	4	5	6
Real interest rate	0.075 (0.047)	-0.005 (0.071)	0.211 (0.062)	0.137 (0.082)	0.062 (0.082)	0.080 (0.081)
$\Delta \ln(y)$	0.397 (0.029)	0.413 (0.030)	0.125 (0.027)	-0.056 (0.036)	0.086 (0.073)	0.188 (0.036)
$\Delta \ln(c)^2$	-	0.252 (0.022)	0.360 (0.035)	0.126 (0.045)	0.202 (0.070)	0.104 (0.066)
age and age ²	-	-	1553.6	223.2	-	-
χ^2 and (dof)	-	-	(3)	(3)	-	-
demographics ^a	-	-	-	374.0	1403.0	1198.9
χ^2 and (dof)	-	-	-	(8)	(8)	(8)
labour supply ^b	-	-	-	-	95.5	-
χ^2 and (dof)	-	-	-	-	(4)	-
Sargan test	30.266	30.391	29.945	25.017	24.344	30.2
d.o.f.	49	48	45	37	36	40

Notes: Numbers in parentheses are standard errors. $\Delta \ln(y)$ is the rate of growth in before tax family non capital income.

^a Demographic variables are the changes in log of family size, number of children, number of adults, dummy for the presence of children and changes in the same variables crossed with $\ln(c)$.

^b Labor supply variables are the changes in the dummies for husband and wife in employment as well as the changes in these dummies crossed with $\ln(c)$.

Table 3

Experiments on Base Specification

	Drop levels variables	Drop crossed variables	Small Instrument set
Comparison column in Table 2	4	4	3
Coefficient on expected income growth (se)	-0.037 (0.025)	0.197 (0.052)	0.278 (0.209)

the estimate of the coefficient on $\ln(c)^2$ implies that the elasticity of intertemporal substitution increases with the level of consumption (holding everything else constant). The inclusion of this variable does not affect the coefficient on income growth but it does change that on the real interest rate.

In the third column we approximate the life cycle movements in family composition with a simple quadratic polynomial in age. We therefore add three variables: the first differences of age squared and age and age squared crossed with log consumption (the change in age is simply a constant). In table 2 we report only the value of the χ^2 test for the hypothesis that these additional coefficients are jointly equal to zero ⁹. The noticeable features are the strong significance of the new variables, the fact that the point estimate of the coefficient on income growth, albeit still significant, is less than one third of the estimate in the previous column (0.13, s.e.=0.02). and that the coefficient on the real interest rate is now strongly significant (0.21, s.e.=0.06). The coefficient on $\ln(c)^2$ is still strongly significant.

In column 4 we add to the previous specification the four demographic variables listed above; note that this means estimating eight new coefficients since each variable is included in 'levels' and 'crossed' terms. This has the effect of removing completely the excess sensitivity of consumption to income. The point estimate of the coefficient on income growth is now negative and insignificantly different from zero. The demographic variables are strongly significant.

In column 5 we remove the age and age squared variables and add two dummies indicating whether the husband and wife are in employment. The coefficient on income growth, again, is estimated to be very small and insignificantly different from zero. The labor supply variables are strongly significant. Their significance is discussed in the subsection 4.4 ¹⁰.

Finally, in column 6 we report a specification which excludes both labor supply and age variables. The coefficient on income growth is now sizeable and significant (0.19, s.e.=0.04). The coefficient on $\ln(c)^2$ is reduced in size and is not strongly significant (0.10, s.e.=0.07). It should be stressed that column 6 has been obtained by excluding two sets of significant variables.

We now turn to the discussion and interpretation of these results.

⁹ Full estimation results for this and the following equations are reported in Table A2 in the result appendix.

¹⁰ A specification with both labor supply and age variables indicated that both sets of variables are significant. However, both the size and the significance of the other coefficients were unaffected.

4.9 Robustness of results and comparison with the existing literature

The results in the previous sub-section suggest that controlling for demographics and age or labor supply removes all the apparent excess sensitivity in the data (compare columns 4 and 5 with column 2). On the other hand, if we control for age or demographics alone excess sensitivity is considerably reduced but not removed completely (see columns 3 and 6). We turn now to a consideration of how these results compare to those found using macro and micro data. We shall restrict attention to the specification that does not use labor supply variables (that is, column 4 of Table 2) since the use of these as conditioning variables in a consumption equation is unusual; we discuss column 5 in the next subsection.

To facilitate these comparisons we first present the results of three experiments on our data. In our first two experiments we drop, in turn, the 'levels' and the 'crossed' variables. In the third experiment we investigate the effect of using a smaller set of instruments on our results. As we shall see other studies use a much smaller instrument set than the one we have. Specifically, in the 'small' version we use as instruments: the seasonals; age variables; second lags of the real rate of interest; second lags of the first differences of log consumption squared; second lags of first differences of hours of work and participation dummies for husband and wife and second and third lags of income and consumption growth¹¹. Thus in this experiment we exclude the crossed and demographic variables from our instrument set.

The results of these experiments on our test of excess sensitivity can be seen in Table 3. In the first two columns we drop the demographic and age levels and crossed variables respectively from the specification in column 4 of Table 2; it should be noted that these variables are significant in that specification (the $\chi^2(5)$ and $\chi^2(6)$ statistics for excluding them are 44.5 and 469.8 respectively). As can be seen there is no evidence of excess sensitivity if we drop the levels variables but this changes dramatically if we replace the levels variables and remove the crossed variables: the coefficient on expected income is now large and significant.

The final column in Table 3 shows that if we drop the demographic and crossed variables from the instrument set then the coefficient on expected income becomes very ill determined; the standard error rises from 0.027 in Table 2, column 3 to 0.209. The obvious explanation here is that without the variables excluded from the instrument set the predictions for income growth are

¹¹ The reason for the choice of these variables will become clear below.

not very good.

For the purposes of comparison with previous results using micro data we shall consider the studies of Hall and Mishkin (1982) (hereinafter, HM); Altonji and Siow (1987) (AS); Zeldes (1989) and Runkle (1991). These papers provide a representative cross-section of results concerning excess sensitivity on the PSID: HM and Zeldes find evidence of excess sensitivity whilst AS find only weak evidence and Runkle finds none. Of course there are many differences between these studies and between any of them and our specification ¹² but we shall concentrate on the differences in the use of demographics.

When considering tests of excess sensitivity and the use of demographics there are two off-setting considerations. First, allowing for the dependence of tastes on demographics is likely to reduce the life-cycle correlation between income and consumption in the way already explained. On the other hand, demographics add significantly to the predictive power for income changes and so make rejections of excess sensitivity more likely.

All four papers referenced take some account of age and/or family composition but these are generally treated as nuisance variables. Thus both HM and AS use residuals from regressions of income and consumption (strictly, food) changes on a variety of demographic variables. Zeldes uses an adult equivalence scale to deflate food expenditures. Finally, Runkle includes an age variable in his first-differenced specification and reports in a footnote that using extra demographics didn't affect his results. The important point here is that none of these studies includes demographics crossed with consumption. As can be seen from Table 3 it seems that it is these variables that remove the excess sensitivity; thus we can reconcile our results with those of HM, AS and Zeldes by noting this difference.

The Runkle study presents a very different problem for reconciliation: Runkle finds 'no excess sensitivity' even though he uses only age variables. As can be seen from column 3 of Table 2 we cannot remove all the excess sensitivity with just these variables. We believe that the difference here may be associated with how well we are predicting income changes. Runkle uses age, and lagged asset variables, hours of work and the real rate as instruments. Thus he does not use demographics to predict income changes; in the other cited studies the coefficients on the demographic variables in the income change equation are highly significant when they are used (see, for example, HM

¹² Principally: we use cohort mean data for the U.K. on all non-durable consumption and we use a more flexible preference specification.

and Zeldes). Thus the experiment reported in column 3 of Table 3 above: when we exclude demographics and crossed variables (but not lagged levels of labor supply variables) from the instrument set then age variables alone 'take out' the excess sensitivity. This is similar to Runkle's result.

Although not conclusive this suggests that our results can be reconciled with other results found using micro data. What of studies that use aggregate data? How can demographics, which are not very volatile in aggregate, 'explain' the finding of high frequency excess sensitivity on aggregate data (see, for example, Campbell and Mankiw (1989,1991))? Once again, the crossed variables play a central role. Consider the dummy variable for the presence of a child in the household; denote this z_t^h where h stands for household and t for time period. The aggregate of this over the population is $Z_t = \sum_h z_t^h$ which is simply the number of households with children. It is reasonable to assume that Z_t changes very little from period to period; consider the extreme case where Z_t is constant over the sample period. Omitting such a variable from the aggregate relationship will not cause any bias. More generally, if aggregate demographics are evolving slowly over time then they can probably be captured by simple trend variables (see Blinder (1975) and White (1978)).

Thus it seems we can ignore the levels variables when we use aggregate data. The same is not true of the corresponding crossed variables. For the dummy discussed in the last paragraph this variable is $\ln(c_t^h)z_t^h$. The aggregate of this variable is $\sum_h (\ln(c_t^h)z_t^h)$ which is the sum of $\ln(c_t^h)$ over the households with children. There is no reason to believe that changes in this variable are more or less volatile than, say, the geometric mean over the whole population of consumption or income. Moreover, this variable is probably serially correlated and correlated with expected income changes. Thus the widespread rejection of the orthogonality condition on lagged variables and the specific finding of excess sensitivity on aggregate data can be seen as due to omitted variable bias. Of course, with aggregate data there is no way to construct the omitted variable but column 2 in Table 3 above suggests strongly that for cohort mean data (which is somewhere between true micro data and aggregate data) there is such bias and that accounting for it removes all the apparent excess sensitivity.

4.4 The inclusion of labor supply variables

In column 5 of table 2 we included two labour supply variables and showed that in the absence

of pure age variables they help to remove the apparent excess sensitivity in the unconditional specification (that is column 2). This result can be interpreted in two different ways.

Within the theoretical framework of section 3, the inclusion of labor supply variables is justified if preferences over consumption and leisure are non additively separable. This non-separability has long been recognized as potentially important (see Heckman (1974) and Ghez and Becker (1975)). Moreover Browning and Meghir (1991) strongly reject the hypothesis that preferences over individual commodities are separable from labor supply. This suggests that preferences over total consumption and labor supply are not additively separable.

A very different interpretation of the comparison of columns 2 and 5 of Table 5 is the following. Suppose that consumption changes are correlated with expected income changes (for example, because of the presence of liquidity constraints) but that the latter are measured with a good deal of error. Suppose also that changes in employments status can be better predicted by the econometrician than can changes in income. Then it might be that including anticipated employment status variables would 'drive out' expected income whether or not the labor supply variables are relevant for consumption changes.

The results in Table 2 dispose us against the second of these interpretations. After all, we can remove all the excess sensitivity with just age and demographic variables (see column 4). On the other hand, someone who *a priori* ruled out pure age and labor supply effects could argue that the relevant comparison is between columns 2 and 6 and that this shows that there is some excess sensitivity, albeit less than usually found. This argument rests on the noiseiness of the construction of the income expectations process. As we have seen the results we have obtained regarding excess sensitivity are changed if we use poor instruments (see column 3 of table 3).

5. Conclusion

The main aim of this paper is to assess the validity of the life cycle model of consumption. In particular, we addressed an issue that has recently received much attention, especially in the macroeconomic literature: that of "excess sensitivity" of consumption growth to income growth. We do this using a time series of cross sections and a novel and flexible parametrization of preferences. The former allows us to address aggregation issues directly, while with the latter we can allow both the discount factor and the elasticity of intertemporal substitution to be affected by various observable variables and lifetime wealth.

The main findings can be summarized as follows:

(i) the excess sensitivity of consumption growth to labor income disappears when we control for demographic variables. This is true both at life cycle and business cycle frequencies.

(ii) estimation of a flexible specification of preferences indicates that the elasticity of intertemporal substitution is a function of several variables, including the *level* of consumption. The *eis* increases with the level of consumption, as expected.

(iii) the variables that change the *eis* are also important in explaining why we observe excess sensitivity over the business cycle.

(iv) we are able to reconcile our results with those reported both in the macro and micro literature.

(v) in our specification the *eis* is not very well determined. This result, however, should be taken with care, as we have not made an effort to construct a 'preferred' specification, which would probably include additional controls for labor supply behavior.

The evidence presented shows that the life cycle model cannot be easily dismissed. Indeed, we believe that the model does a good job of representing consumption behavior both over the life cycle and over the business cycle.

Data Appendix

The UK Family Expenditure Survey is a continuous survey of households. Each household is interviewed only once. Each fortnight about 140 new households are selected. In total we have about 1,750 households per quarter, which gives us about 120,000 households over the 17 years of the Survey.

During the two-week interview period the adult members of the household each keep a diary of their expenditures which is then checked by the interviewer at the end of the survey fortnight. As well, a wide range of information on various household characteristics is collected. Table A1 reports summary statistics for the data used .

Result Appendix

In Table A2 we report the detailed results of the GMM estimates discussed in section 4 in the text.

References

- Altonji, J. and A. Siow (1987) : "Testing The Response of Consumption to Income with (Noisy) Panel Data" , *Quarterly Journal of Economics* , 102, 293-328.
- Attanasio, O.P., and Weber, G. (1992): "Consumption Growth, the Interest Rate and Aggregation" , *Review of Economic Studies*, forthcoming.
- Blinder, A. (1975) : "Distribution Effects and the Aggregate Consumption Function" , *Journal of Political Economy* , 98, 447-75.
- Blundell, R., Browning, M. and C. Meghir , (1989) : "A Microeconomic Model of Intertemporal Substitution and Consumer Demand" University College London DP. 89-11.
- Blundell, R., Pashardes, P. and G. Weber, (1992) : "What Do We Learn About Commodity Demands From Micro Data?" , *American Economic Review* , forthcoming.
- Browning, M. (1989) : "The Intertemporal Allocation of Expenditure on Non-Durables, Services and Durables" *Canadian Journal of Economics*, 22, 22-36.
- Browning, M. (1990) : "Estimating Micro Parameters from Macro Data: Some Pessimistic Evidence" , Mimeo, McMaster University.
- Browning, M., Deaton, A. and M. Irish (1985): "A Profitable Approach to Labor Supply and Commodity Demands over the Life Cycle" , *Econometrica* , 53, pp.503-44.
- Browning, M., and C. Meghir (1991): "Testing for Separability of Commodity Demands from Male and Female Labor Supply" , *Econometrica* , 59 925-952.
- Campbell, J.Y. (1987) "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis" , *Econometrica* , 55, 1249-73.
- Campbell, J.Y., and Mankiw, N.G. (1989): "Permanent Income, Current Income, and Consumption" , *NBER Macroeconomic Annual*, Cambridge.
- Campbell, J.Y., and Mankiw, N.G. (1991): "The Response of Consumption to Income: A Cross-Country Investigation" , *European Economic Review*, 35, 715-21.
- Carroll, C.D. and L.H. Summers (1991): "Consumption Growth Parallels Income Growth: Some New Evidence" , in: B.D. Bernheim and J.B. Shoven (eds.) : *National Savings and Economic Performance*, Chicago.
- Chamberlain, G. (1984) : "Panel Data" , in: *Handbook of Econometrics* , North Holland.
- Cochrane, J. (1991) : "A Simple Test of Consumption Insurance" , *Journal of Political Economy* , 99, 957-81.

- Deaton, A. (1992) : *Understanding Consumption*, Oxford University Press, Oxford.
- Flavin, M. (1981) : " The Adjustment of Consumption to Changing Expectations About Future Income", *Journal of Political Economy* , 89, 974-1009.
- Gorman, W.M. (1959) : "Separable Utility and Aggregation", *Econometrica* , 21, 63-80.
- Grossman, S. J. and R. J. Shiller (1982) : "Consumption Correlatdness and Risk Measurement in Economies with Nontraded Assets and Heterogeneous Information", *Journal of Financial Economics*, 10, 195-210.
- Hall, R.E. (1978) : "The Stochastic Implications of the Life Cycle Permanent Income Hypothesis", *Journal of Political Economy* , 86, 971-87.
- Hall, R.E., (1988) : " Intertemporal Substitution in Consumption", *Journal of Political Economy* , 96, 339-57.
- Hall, R.E. and F. Mishkin, (1982) : "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households", *Econometrica* , 50, 461-82.
- Hansen, L.P. and K.J. Singleton (1982) : "Generalized Instrumental Variable Estimation of Non-linear Rational Expectations Models", *Econometrica* , 50, 895-916.
- Hansen, L.P. and K.J. Singleton (1983) : "Stochastic Consumption, Risk Aversion and teh Temporal Behavior of Asset Returns", *Journal of Political Economy* , 91, 249-65.
- Mankiw, N.G., Rotemberg, J.J and L.H. Summers (1985) : " Intertemporal Substitution in Macroeconomics", *Quarterly Journal of Economics* , 100, 225-53.
- Meghir, C. and G. Weber, (1991) : "Intertemporal Non-Separability or Borrowing Restrictions?", Institute for Fiscal Studies, Mimeo.
- Runkle, D.E. (1991) : "Liquidity Constraints and the Permanent Income Hypothesis: Evidence from Panel Data", *Journal of Monetary Economics* , 27, 73-98.
- Ryder, N.B. (1965) : "The Cohort as a Concept in the Study of Social Change", *American Sociological Review*, 30, 843-61.
- Stoker, T.M. (1984) : "Completeness, Distribution Restrictions and the Form of Aggregate Functions", *Econometrica* , 52, 887-907.
- White, B.B. (1978) : "Empirical Tests of the Life Cycle Hypothesis", *American Economic Review* , 68, 547-60.
- Zeldes, S. (1989): "Consumption and Liquidity Constraints: An Empirical Investigation", *Journal of Political Economy* , 97, 305-46.

Table A1
Data Summary

Variable	mean	min	max
$\Delta \ln(c)^1$ (x100)	0.486	-24.95	22.76
$\Delta \ln(y)^1$ (x100)	0.511	-22.15	19.57
$\ln mem^2$	0.517	0.057	0.822
$cdch$	0.611	0.011	0.975
nch	1.205	0.021	2.675
$adult$	2.381	2.0	2.929
$dphusb^3$	0.074	0.0	0.681
$dpwife^3$	0.429	0.16	0.820
$real\ rate^1$ (x100)	-0.746	-8.63	2.41
$inflation^1$ (x100)	2.54	-1.16	10.50

Notes: 1) First differences are all quarterly.

2) The participation dummies are set to 1 if the person is not employed.

3) $\ln mem = \ln(\text{number of household members}) - \ln(2)$

Table A2

Effects of demographics and labour supply variables on the Euler equation for consumption

	3	4	5	6
Real interest rate	0.211 (0.062)	0.137 (0.082)	0.062 (0.082)	0.080 (0.081)
$\Delta \ln(y)$	0.125 (0.027)	-0.056 (0.036)	0.086 (0.073)	0.188 (0.036)
$\Delta \ln(c)^2$	0.360 (0.035)	0.126 (0.045)	0.202 (0.070)	0.103 (0.066)
age^2	-0.179 (0.053)	-0.135 (0.066)	- -	- -
$\ln(c) * age$	-0.0103 (0.004)	0.011 (0.003)	- -	- -
$\ln(c) * age^2$	0.0028 (0.0001)	0.001 (0.0002)	- -	- -
$\ln mem$	- -	0.236 (0.373)	-0.024 (0.571)	0.562 (0.371)
$\ln(c) * \ln mem$	- -	2.411 (0.516)	3.197 (0.693)	2.579 (0.547)
nch	- -	-0.040 (0.078)	-0.040 (0.127)	-0.180 (0.086)
$\ln(c) * nch$	- -	-0.272 (0.092)	-0.452 (0.129)	-0.297 (0.112)
$adult$	- -	0.048 (0.139)	0.158 (0.199)	0.040 (0.135)
$\ln(c) * adult$	- -	-0.641 (0.138)	-0.796 (0.220)	-0.692 (0.172)
$cdch$	- -	0.101 (0.096)	0.346 (0.096)	0.272 (0.072)
$\ln(c) * cdch$	- -	-0.349 (0.225)	-0.734 (0.277)	-0.565 (0.212)

d_{husb}	-	-	0.283	-
	-	-	(0.109)	-
$\ln(c) * d_{phusb}$	-	-	0.783	-
	-	-	(0.121)	-
$dpwife$	-	-	-0.068	-
	-	-	(0.039)	-
$\ln(c) * dpwife$	-	-	0.233	-
	-	-	(0.085)	-
Sargan test	29.945	25.017	24.344	23.9
d.o.f.	45	37	36	40

Notes: Numbers in parentheses are standard errors. $\Delta \ln(y)$ is the rate of growth in before tax family non capital income. The column numbers correspond to that of table 2 in the text.

^a *Demographic variables are the changes in log of family size, number of children, number of adults, dummy for the presence of children and changes in the same variables crossed with $\ln(c)$.*

^b *Labor supply variables are the changes in the dummies for husband and wife in employment as well as the changes in these dummies crossed with $\ln(c)$.*