

TESTING THE LIFE CYCLE MODEL OF CONSUMPTION: WHAT CAN WE LEARN FROM MICRO AND MACRO DATA?*

Orazio P. ATTANASIO

Stanford University, Università di Bologna, NBER and CEPR

Martin BROWNING

McMaster University

In this paper we analyze the implications of the life cycle model for consumption and consider the possibility of testing the model using macro and micro data. We conclude that the implications of the model cannot be tested with macro data. We provide some evidence from US and UK micro datasets which shows that the life cycle model, at a first glance, is not inconsistent with the data.

1. Introducción

The life cycle-permanent income hypothesis of consumption behavior has been one of the most influential models in modern economics. Countless theoretical and empirical papers have been devoted to its study. The appeal of the model lies in its apparent simplicity, its consistency with microeconomic theory and, at the same time, its apparently strong empirical predictions. Consumers are assumed to maximize life time utility. The problem of allocating resources between periods is similar to that of allocating total expenditure to different commodities in a given time period. Given the appropriate relative prices, the problem can then be tackled in a similar way.

As discussed below, the life cycle model was partly motivated by the necessity of explaining some macroeconomic «stylized facts» that were apparently

* This paper was prepared for the second International Macroeconomics Programme Meeting organized by the Centre Economic Policy Research (London) in Madrid, Spain, June 7 and 8 1991. The empirical part of the paper uses two large surveys: the US Consumer Expenditure Survey and the UK Family Expenditure Survey. The study of the former was financially supported by the Center for Economic and Policy Research (Stanford). We are grateful to the Institute for Fiscal Studies (London) and the UK Department of Employment for providing the FES data. The paper draws from and discusses material contained in a previous paper by the two authors as well as various papers coauthored with Richard Blundell, Costas Meghir and Guglielmo Weber. We thank Angus Deaton, Tom MaCurdy and an anonymous referee for helpful comments. All errors are our responsibility.

inconsistent with the simple Keynesian consumption function. This was done with a model that was also consistent with the evidence from «budget studies».

The life cycle model is based on individual optimization and therefore, as every microeconomic problem, can be studied on the basis of aggregate data only if very stringent aggregation conditions are satisfied. The early empirical students of the life cycle model were aware of these aggregation conditions (see Modigliani and Ando (1963), for instance). Since the 1970s, however, the emphasis of the empirical research on the life cycle model has shifted more and more towards aggregate time series data, without worrying too much about aggregation. This was partly motivated by the lack of micro data on consumption and partly by the shift in attention towards the new theoretical developments occurring in that period. The assumption of rational expectations was applied to the life cycle model by Hall (1978) who derived its implications within the framework of a representative consumer model with separable preferences facing an uncertain environment.

During the 1980s numerous tests of the first order conditions of a representative consumer maximizing expected utility were performed and, by and large, found strong violations of the overidentifying restrictions implied by the model. As a consequence, the opinion that the life cycle model is an inadequate representation of consumption behaviour became more and more widespread. At the same time a number of explanations for the «empirical failure» of the model were proposed.

In this paper we argue that aggregate time series (ATS) data cannot be used to test the life cycle model. In Section 2 we shall show that whether or not the ATS data appear to be generated by a maximizing «life cycle» representative agent has no connection with whether individual agents behave in this way. Thus the consistency of the ATS data with a life cycle model is neither necessary nor sufficient for consistency of individual behavior with that model.

The increasing availability of micro data sets on consumption makes it possible to test the life cycle model with individual data. In Section 3 we discuss various ways of using time series of cross sections to such a purpose.

Finally, in Section 4 we present some evidence based on micro data from the UK and the US, which indicates that the life cycle model is not inconsistent with the facts emerging from the data. The results reported are not formal tests of the life cycle model. They should rather be interpreted as a first check of the simplest implications of the model: namely that consumption expenditure is smoothed over the life cycle. Section 5 concludes the paper.

2. What can be learned from aggregate time series?

If individual consumption behavior is well described by the life cycle model, what are the properties of aggregate consumption? Under what conditions

will aggregate time series behave as if they were generated by a representative individual? Under what conditions is it possible to identify the parameters of individual behaviour from aggregate data?

The most common reaction of macroeconomists to these problems has been to sweep them under the carpet. Most of the empirical work on testing the life cycle model *assumes* the existence of a representative consumer. The purpose of this section is to address these issues directly.

2.1. Under what conditions are aggregate data consistent with the existence of a representative consumer?

As can be inferred from the list of questions in the first paragraph, there are several levels at which the problem of linking microeconomic behavior to macroeconomic data can be posed. The first question we will ask is the following: when can we say that an aggregate time series for, say, the growth in consumption of non durables and services, varies in time as if it was generated by a single («representative») agent behaving according to the life cycle model? In other words: when is it that the analysis of aggregate data will not reject the representative consumer-life cycle model?

An obvious answer to the questions above is the counterfactual one: «when consumers are actually identical». More interesting answers are however available. Grossman and Shiller (1982) have partly addressed this question and shown under what conditions it is possible to aggregate individual Euler equations, and how to interpret the parameters derived from these equations¹. However, those conditions are not necessary. In what follows we construct an example which will stress a more subtle point about the relationship between macro and micro data.

One of the main implications of the life cycle model is that the rate of growth in consumption depends on the real interest rate, which represents the relative price of present and future consumption, and should not be affected by expected changes in current income. Indeed, the sensitivity of the consumption growth aggregate time series to expected labor income has often been used as a test of the life cycle model. The question we want to ask is: suppose that aggregate data reveal the relationship between consumption growth and the interest rate predicted by the life cycle model and do not show sensitivity to expected labour income; what can we say about individual

¹ The condition we need on preferences to aggregate individual Euler equations, is linearity of marginal utility. In a discrete time model, this follows from a quadratic utility function. In a continuous time model more general functional forms are possible. As we discuss below, in a discrete time model with isoelastic utilities, we will need additional assumptions. The problem of aggregating Euler equation can be rephrased in terms of a standard aggregation problem. The results in Gorman (1953) and Muellbauer (1975) can then be used.

behavior? Unfortunately the answer is: «not much», as the following example will make clear².

Consider an economy with a very large number of agents who are born at the same time and live for T periods. Suppose that no one has initial assets and that there is no uncertainty. Suppose that current labor income is the sum of a macro component \bar{y}_t and an idiosyncratic component u_t^i which has zero mean across the consumers in the economy:

$$y_t^h = \bar{y}_t + u_t^h \quad [1]$$

We assume that both components of income are deterministic and known in advance.

The value of lifetime income of a generic agent h discounted to the beginning of his/her lifetime is:

$$W^h = \sum_{t=1}^T R_t y_t^h = \sum_{t=1}^T (\bar{y}_t + u_t^h) R_t \quad [2]$$

where R_t is a discount factor: $R_t = \prod_{j=1}^t (1 + r_j)^{-1}$, $t > 1$.

Suppose that each individual consumes a fixed fraction α of his/her idiosyncratic income. In this sense individual consumption exhibits excess sensitivity to income. Let H^h be lifetime income net of these expenditures:

$$H^h = W^h - \alpha \sum_{t=1}^T R_t u_t^h \quad [3]$$

Suppose also that each agent consumes a random fraction q_t^h of H^h . Therefore, (discounted) total expenditure at time t for individual h is given by:

$$R_t c_t^h = q_t^h H^h + \alpha R_t u_t^h \quad [4]$$

where the q_t^h 's for $t < T$ are independent (over time and across agents) draws from a uniform distribution $[0, 1/2T]$. This implies that, on average, a fraction $1/2T$ of H^h is spent in each period³. Therefore, at the micro level we have both excess sensitivity to labor income and an unstable relationship of individual consumption to interest rates. However, if one aggregates equation [4] across individuals one gets:

$$c_t R_t = \bar{q}_t \bar{W} = \frac{1}{2T} \bar{W} \quad [5]$$

² The example is inspired by Becker (1962).

³ In the last period consumption is determined by the intertemporal budget constraint. With the distribution given, last period consumption is always positive.

where \bar{q}_t is the average realization of q_t^h . An analogous relation can be obtained for period $t+1$. Therefore, the rate of growth of aggregate consumption is given by:

$$\frac{c_{t+1}}{c_t} = (1 + r_{t+1}) \quad [6]$$

Equation [6] shows a stable relationship between aggregate consumption growth and the interest rate. Furthermore, this relationship does not present any «excess sensitivity» of consumption to income. This is because the «idiosyncratic» part of income (and consumption) is averaged out. And yet, at the micro level, we have both excess sensitivity and no reaction to the real interest rate. The example is obviously extreme, but makes the point that aggregate data cannot be used to test hypotheses about micro behavior.

2.2. *Under what conditions is it possible to estimate structural parameters from aggregate data?*

The construction of an example where aggregate data support the life cycle model while the individual behavior deviates substantially from it is necessarily artificial and elaborate. It is much easier to construct a model in which individuals behave according to the model, but aggregate data present, if interpreted within the representative agent framework, substantial violations of the restrictions implied by the model.

In this respect, an interesting question to ask is: provided that individual behaviour is in accordance with the model, under which conditions can macro data be used to infer the structural parameters of the model.

Suppose an individual household h maximizes life time utility, subject to an intertemporal budget constraint and a household technology.

$$\text{Max } E_t \sum_{j=0}^{T-t} \beta^j U(c_{t+j}^h) \quad [7]$$

$$\text{s.t. } A_{t+1+j}^h = y_{t+j}^h - C_{t+j}^h + (1 + r_{t+j}) A_{t+j}^h \quad [8]$$

$$c_{t+j}^h = G(z_{t+j}^h, C_{t+j}^h) \quad [9]$$

where c_t are consumption services, C_t is consumption expenditure, z_t is a vector of demographic and labor supply variables, A_t are financial assets, r is the interest rate and y is labour income. The function G represents «household technology». The z variables can include the number of household members, the number of children of various ages, seasonal dummies.

Under suitable assumptions on the functional forms involved in equations 7-9 it is possible to arrive at a specification which can be easily estimated on the

basis of individual observations. Blundell, Browning and Meghir (1994), Attanasio and Weber (1993) and Attanasio and Browning (1992), for instance, estimate the parameters of the utility function and test the model using average cohort data. Even though it is not necessary, it is convenient at this point to assume that the utility function is isoelastic, and the G function is homogenous of degree one in C_t .

Under these assumptions and the assumption of log-normality of the real interest rate and of consumption growth⁴, it is possible to derive a loglinear equation which relates consumption growth to the interest rate and other variables.

$$\Delta \log(C_{t+1}^h) = g' z_t^h + \sigma \log(1 + r_{t+1}) + \epsilon_{t+1}^h \quad [10]$$

where σ can be interpreted as the elasticity of intertemporal substitution in consumption and g is a vector of parameters. $g' z_t^h$ can be interpreted as a linear approximation to the function G in [9].

Two problems are apparent if we want to aggregate equation [10] and estimate it on a time series of National Account data. First, the equation is linear in the logs rather than the levels. Second, the z variables may not be observable at the macro level. We will analyze them in turn.

If there are H individuals in the economy, we can only observe $\Delta \log(\sum_{j=1}^H C_{t+1}^h)$ which is the log-difference of total personal expenditure. This may differ substantially from $\sum_{j=1}^H \Delta \log(C_{t+1}^h)$, which is what we get aggregating equation [10]. The difference between the log of the geometric mean and that of the arithmetic mean is sometimes referred to as the Theil entropy measure. This difference will induce a bias in the estimation of equation [10], unless it is constant or uncorrelated with the set of instruments used. This is the implicit assumption made in most studies estimating an equation like [10]⁵. Without the availability of a time series of cross sections it is not possible to test the validity of such an assumption. In Attanasio and Weber (1993) 68 quarters of the UK Family Expenditure Survey are used to construct a time series of the Theil entropy measure and study its properties⁶. By expanding such a measure in a Taylor series it is possible to decompose it into a weighted sum of central moments. Attanasio and Weber (1993) study the time series behavior of variance, skewness and kurtosis computed on each cross section. The three moments considered and in particular variance and skewness are characterized by substantial variability, autocorrelation and seasonality.

⁴ The assumptions of homothetic preferences and log-linearity are not necessary. Attanasio and Browning (1992) discuss the case in which the marginal utility of money is a linear function of the *log* rather than the *level* of consumption and possibly other variables. In this case the elasticity of intertemporal substitution is not constant.

⁵ See, for instance, Hall (1988).

⁶ See also Browning (1993) and Blundell, Pashardes and Weber (1993).

The variability of the cross sectional central moments may induce a rejection of the overidentifying restrictions implied by equation [10] when they are correlated with the set of instruments used. Attanasio and Weber (1993) find that when they aggregate the FES data using the arithmetic mean they obtain strong rejections of the overidentifying restrictions, which disappear when the geometric mean is used. A regression of the various components of the entropy measure on the instrument set shows that the second and third central moment of the cross sectional distribution are mainly responsible for the rejection, being correlated with the instruments.

From the discussion above it is clear that the problem of non-linearities can be interpreted as a missing variables problem. Indeed, Blundell, Pashardes and Weber (1993) show how a general aggregation problem can be expressed in terms of missing variables. For a large class of models, it is possible to express the difference between an aggregated and a micro relationship in terms of some function of cross sectional moments. The availability of long time series of cross sections allows the study of the dynamic properties of these moments and the possibility of testing the hypothesis that movements in this quantities can be safely ignored.

A similar interpretation can be given to the second problem we referred to above, namely the presence in [10] of a number of variables representing the «household technology» that transforms consumption expenditure into service flows, which enter the household utility function. Some of the variables in [10] are observable at the macro level: seasonality is the most obvious example⁷. Others (such as family composition, the composition of earners, education) are not. Omission of these variables can cause biases in the estimation of the structural parameters and rejections of the overidentifying restrictions. The problem becomes even more serious if the elasticity of substitution is not constant but a function of various demographic variables. In this case it will be necessary to introduce in equation [8] terms that interact the log of consumption with various demographic variables. The omission of these terms will certainly bias the estimation of the structural parameters of the model⁸.

In the macro literature, violations of the overidentifying restrictions and excess sensitivity of consumption to labour income have often been interpreted as a sign of the presence of liquidity constrained households, or of myopic behavior. Very few papers seriously consider the possibility of misspecification of the utility function. Browning (1989) derives a test for the life-cycle model that is independent of functional form (that is, a «revealed preference» or «non-parametric» test). He finds very few violations for the conditions of this test on ATS data for Canada, the UK and the US. Given that these tests make no allowance for surprises, measurement error or omitted variables, this

⁷ See for instance, Miron (1986).

⁸ This line of research has been pursued by Blundell, Browning and Meghir (1994) and by Attanasio and Browning (1992) and will be discussed in the next section.

suggests quite strongly that the rejections by parametric tests are plausibly due to functional form assumptions.

An obvious culprit for excess sensitivity would be non separability between leisure and consumption⁹. A number of macro papers do consider preferences which are non separable between leisure and consumption: see, for instance, Bean (1986) and Mankiw *et al.* (1985). However, without strong assumptions, it is impossible to obtain an equation which is estimable with aggregate data on employment and/or hours. In particular, as discussed below, the composition of earners in the households as well as the phase of the life cycle in which labour supply decisions are made, can be important. These variables are not observable at the macro level.

The possibility that a sizable fraction of the population is subject to liquidity constraints has received a considerable amount of attention in the macro literature¹⁰. The relevance of these phenomena can be addressed in a much more direct way with the help of micro data. On the one hand it is possible to use observable variables (like wealth) to isolate households that are less likely to be constrained (see Zeldes, 1989; Runkle, 1991; Keane and Runkle, 1992). On the other hand one can isolate households that are in a phase of their life cycle in which they are unlikely to have the necessity of borrowing (see Attanasio and Weber, 1993).

A conceptually similar problem is that of possible misspecifications of equation [10] arising from changes in preferences in the last part of the life cycle. The textbook derivation of the life cycle model is usually done assuming a fixed and known length of life. The model can easily be extended to consider an uncertain life time as long as the probability of death is constant or changes in a fairly regular way. However, we know that the probability of death changes dramatically after a certain age; this could cause dramatic changes in preferences. Furthermore, changes in family sizes due to mortality as well as failing health can also change preferences in a way that is substantially different from the simple specification assumed in equation [9]. The same applies to changes in preferences due to retirement. It may therefore be important to isolate households that are less likely to experience dramatic changes in preferences (see Attanasio and Weber, 1993). Once again, this cannot be done with aggregate time series data.

2.3. *Is it useful to study aggregate time series of consumption?*

From the discussion in the previous two subsections, it should be clear that, in our opinion, aggregate time series data cannot be used to test the

⁹ On the excess sensitivity of consumption to expected labor income see Flavin (1981), Campbell (1987), and Campbell and Mankiw (1989). For the argument about non separability between consumption and leisure see, for instance, Heckman (1974).

¹⁰ With few exceptions, no attempt has been made at modeling the behavior of liquidity constrained individuals. See Jackman and Sutton (1982) and, more recently, Deaton (1991).

life cycle model, nor to estimate structural parameters. This does not mean, however, that the analysis of aggregate time series data is completely useless, or that it should be done in a completely a-theoretical way. We do believe that the life cycle theory can be conceptually useful in the econometric modeling of consumption time series data. After all, the introduction of the life cycle- permanent income hypothesis was partly originated in the attempt to understand the failure of the simple Keynesian consumption function to explain some basic features of macroeconomic time series. Friedman (1957), for instance, in the introduction to his famous book, notices that «Estimates of savings... revealed no rise in the percentage of income saved during the past half century despite a substantial rise in real income. The corresponding ratio of consumption expenditure to income - the constancy of which means that it can be regarded as both the average and the marginal propensity to consume- is decidedly higher than the marginal propensity to consume that has been computed from either time series or budget data. ... the savings ratio... after WWII was sharply lower than the ratio that would have been consistent with findings on the relation between income and savings in the interwar period. This experience dramatically underlined the inadequacy of a consumption function relating consumption of savings solely to current income» (pp. 3-4). Similar comments can be found in Modigliani and Brumberg (1954) and Modigliani and Ando (1963).

The basic stylized facts that emerge from the analysis of aggregate consumption data are:

- a) Per capita consumption is smoother than disposable income.
- b) The long run marginal propensity to consume out of disposable income seems to be higher than the short run propensity and equal to the long run average propensity to consume.
- c) The short run average propensity to consume (unlike its long run equivalent) seems to be a decreasing function of income.

These basic facts are all consistent with the life cycle theory. Indeed, their observation motivated the development of the theory.

The characterization of the aggregate behavior of consumption expenditure can be useful for several reasons, ranging from policy analysis to econometric forecasts, to business cycle analysis. The study of individual behavior can be profitably integrated with the time series analysis of aggregate consumption expenditure. For such a purpose it is necessary to develop a statistical model of aggregate consumption expenditure which, while not pretending to test the life cycle theory, is flexible enough to be consistent with the stylized facts described above and with the main implications of the theory. In this respect the life cycle model can be useful in several respects.

One of the main implications of the life cycle model is that a transitory change in income should have a substantially different effect on consumption

from a permanent change. Furthermore, the theory suggests that, if utility is non separable between consumption of non durables and services and durables services, variables such as the relative price of durables and the stock of existing durables should probably be relevant. The same applies to other variables that might capture labor supply effects. While it is impossible, as discussed above, to estimate a fully fleshed structural model, it is desirable, while modelling the time series behaviour of aggregate consumption, to take into account these considerations.

First, it is possible to construct carefully time series for disposable income that isolate in the data episodes known to be reflecting transitory changes in tax law. Furthermore, it is possible, to a certain extent, to change the definition of consumption to make it more consistent with the definition relevant for life cycle behavior. These changes can be extremely important if we are interested in the analysis of different policy measures. While there is no way around the Lucas critique other than the estimation of a structural model, a careful analysis of particular episodes and of the time series behavior of consumption and income can lead us a long way in understanding the implications of different policy measures for consumption.

Second, it is possible to keep the dynamic specification of our equation flexible enough to reflect the differences between long and short run behaviour evident in the data and implied, to a certain extent, by the life cycle theory. In this respect, the *error correction* specification advocated by David Hendry and his collaborators can be extremely useful. As an aside, the *error correction* dynamic specification also deals in a direct way with the problem of non-stationarity which afflicts most aggregate time series. As is well known, the approach allows the study of co-integrated time series and the separation of long run relationships from the short run dynamics.

Third, variables the theory suggests to be important can be inserted in the regression equation and their effect on aggregate consumption studied within a fairly flexible specification.

An example of the methodology that we are suggesting is the paper by Blinder and Deaton (1985) where post-war US data are analyzed with care¹¹. Blinder and Deaton consider various tax changes episodes and remove from measured labor income transitory changes that were clearly perceived as such.

Before concluding this section a brief digression on the data used in macroeconomic analysis is necessary. In the US, the majority of consumption studies use seasonally adjusted data. The work of Miron (1986) has stressed the necessity of looking at seasonally unadjusted data. This is obviously crucial when trying to estimate structural parameters. However, even if we want only

¹¹ Blinder and Deaton also attempt to decompose the effects of expected and unexpected changes in various variables. This is only possible under strong identifying assumptions.

to describe the dynamic behavior of aggregate time series, it is probably better to start from seasonally unadjusted data and allow explicitly for seasonal factors at the estimation stage.

3. What can be learned from micro data?

The life cycle-permanent income (LC-PI) theory of consumption behavior originated in the seminal work of Modigliani and Brumberg (1954), Friedman (1957), Modigliani and Ando (1963). In these studies, the desirability of testing the theory on micro data is mentioned several times. Modigliani and Ando (1963), for instance state: «... Friedman's formulation of the hypothesis is fairly well suited for testing against cross section data, ... Friedman's model, on the other hand, does not generate the type of hypotheses that can be easily tested against time series data»¹². Friedman (1957), considering the aggregation of the individual consumption functions over individuals that differ with respect to the real interest rate i , the ratio of non human wealth to permanent income and a taste factor u asserts: «The assumption that ... the distribution of consumers units by income is independent of their distribution by i , w , and u , is obviously false in a descriptive sense. The variable u , for example, covers such factors as age, size of family, perhaps education, and these are all known to be connected systematically with the distribution of income...».

Modigliani and Brumberg (1954), Modigliani and Ando (1960) and especially Friedman (1957) present a considerable amount of evidence from cross section data that shows a substantial support for the LC-PI hypothesis. Indeed, one of the most stunning successes of the theory was its « ... consistency with budget studies» (Friedman, 1957). Since then, however, almost all empirical research on the life cycle model was executed on aggregate time series data.

This is probably explained by the absence of good microeconomic data on consumption. The only US micro data sets available in the 1970's which contained some information on consumption were the 1960-1961 and 1972-1973 Consumer Expenditure Survey (CEX) and the PSID. The former are pure cross sections and therefore inadequate to test an intrinsically dynamic theory such as the life cycle model, while the latter only contained information on food expenditures.

The CEX was used by Ghez and Becker (1975) who identified dynamic effects by looking at individual of different ages. They were therefore forced to strong assumptions on the nature of cohort effects.

The PSID was used by Hall and Mishkin (1982), Zeldes (1989), Runkle (1991) and Keane and Runkle (1992), among others, to estimate and test Euler equa-

¹² In their paper, Modigliani and Ando (1963) do estimate and test the life cycle model on aggregate data. However, they spell out the assumptions necessary for these tests to be valid. These included: « ... the constancy in time ... [of] the age structure of the population and the relative distribution of income, of expected income, and of net worth over the age groups».

tions derived from a stochastic version of the life cycle model. The fact that only food consumption is available in the PSID has constituted a major problem with these studies: one of the assumptions for the formal tests performed in this literature is that utility is isoelastic in food and separable between food and the rest of consumption expenditure. Both assumptions are very strong.

If the lack of suitable data was the main explanation for the low proportion of studies based on micro data relative to those that used aggregate data, this ratio should increase in the next few years. Since 1980 the Consumer Expenditure Survey, which contains detailed and exhaustive information on consumption, has been collected on a regular basis. As discussed below, it can profitably be used to test the life cycle model.

In the UK, the Family Expenditure Survey has been collected on a regular basis since 1970 and has been extensively analyzed for almost a decade (see Atkinson and Micklewright, 1983). Only recently, though, has it been used in various ways to assess the empirical validity of the life cycle model¹³.

In the previous section we stressed the necessity of using individual data to evaluate the life cycle theory and estimate the structural parameters. In this section we wish to address the following questions. What use should be made of the available data sets? What can be and what cannot be learned from them?

The life cycle model is an intrinsically dynamic theory in that it describes the allocation of resources over the life time of an individual. The first step in assessing the empirical relevance of the theory is therefore likely to be the estimation of a life cycle profile for consumption and income which in turns requires the observation of individuals over time and, therefore, the use of panel data. Unfortunately, as yet, no consumption panel data is available: in most consumption data sets households are interviewed only once or over a very short time period.

If individuals born at different ages are identical and experience an identical history, one could use cross sectional variation to estimate the age profile for consumption and income. Unfortunately, the absence of cohort effects is highly questionable: in the presence of productivity growth the amount of resources available to different cohorts grows over time. Therefore, even in the absence of more complicated effects, the estimation of age profiles on the basis of single cross sections can be extremely misleading. Shorrocks (1975), in a paper studying the life cycle accumulation of wealth, constructed a simple example in which individual wealth is a linear function of age, while the analysis of a single cross section gives the illusion of a «hump shaped» age profile because of the presence of productivity effects¹⁴.

¹³ See Browning, Deaton and Irish (1985), Attanasio and Weber (1989, 1993), Browning and Meghir (1991), Attanasio and Browning (1992), Blundell Browning and Meghir (1994).

¹⁴ Ghez and Becker (1975) circumvent this problem by assuming smooth cohort effects exemplified in linear productivity growth.

A similar point applies to the estimation of Euler equations using a single cross section. The residual of an Euler equation is typically an expectational error which has zero mean conditional on lagged information. This implies a set of orthogonality conditions that can typically be used to estimate the parameters of the individual optimization problem. The expectation is, however, taken over time, not across individuals. In general the expectational errors of different individuals are going to be correlated: there is no reason why averaging across different individuals the error term should converge to zero. As a consequence, Euler equation estimates obtained using cross sectional variation are going to be inconsistent¹⁵.

In Subsection 3.1 we will discuss methods to describe and characterize the life cycle profile of consumption (and other variables) while in Subsection 3.2 we analyze more structural approaches more apt to analyze high frequency fluctuations.

3.1. Describing life cycle behavior using a time series of cross sections

The problem of estimating a dynamic relationship without observing the same individual over time can be solved if a time series of repeated cross sections is available. The basic idea, which was pioneered by Browning, Deaton and Irish (1985) and developed by Deaton (1985), is to follow over time the average behavior of homogeneous groups of households. Groups are formed on the basis of the year of birth (or cohort) of the household head (rather than their age). Thus we take means over all individuals aged, say, 25 in the 1980 survey and over individuals aged 26 in the 1981 survey and treat these means as observations on the same individual¹⁶. In this way, even though we do not observe individuals over time we can analyze the behavior of individuals homogeneous with respect to the most important factor in terms of life cycle behavior: their age at any point in time. The idea is that by averaging over individuals belonging to the same cohort, we will be able to estimate the life cycle behavior of an «average» individual. The implicit assumption behind this procedure is that it is possible to decompose individual behavior into the sum of two components: one that reflects life cycle behavior and the reaction to common time effects, and the other that reflects individual heterogeneity orthogonal to the previous effects that can therefore be safely ignored.

Of course, individuals differ in other interesting ways than their year of birth. Indeed, in some situations, ignoring within cohort heterogeneity may strongly bias our estimates of life cycle profiles (an example is given below). However, as long as enough individual information is available, we may hope to control, to a certain extent, for within cohort heterogeneity. To make this point clear, suppose we are interested in modelling the life cycle behavior of

¹⁵ This point is discussed in Chamberlain (1983).

¹⁶ See also Moffitt (1993).

a variable X_t^{ach} where the indexes a , c , h , t stand for age, cohort, individual and time¹⁷.

$$X_t^{ach} = \delta_i^c + \gamma'w^{ch} + \epsilon_i^{ach} \quad [11]$$

where w^{ch} are observable variables that do not change over the life cycle.

Equation [11] decomposes the variability X into a part that is accounted for by changes in the mean for the cohort, and a part that reflects individual heterogeneity. This last part can in turn be decomposed in heterogeneity linked to observable variables (w^{ch}), and in unobservable heterogeneity. Examples of variables probably related to differences in consumption behavior within a cohort are variables such as race, the region of residence, the gender of the households head, and so on.

If we assume that ϵ_i^{ach} is uncorrelated with w^{ch} , equation [11] imposes a fair amount of structure on the way in which within cohort heterogeneity is modelled. Ideally, one would like to analyze separately groups that are homogeneous with respect to the w variables. Unfortunately, this procedure would often leave us with extremely small cells. Equation [11] could therefore be a reasonable compromise.

If the w variables are uncorrelated with δ_i^c , their introduction in equation [11] is only going to improve the precision with which the δ_i^c are estimated. On the other hand, if there are systematic changes in the w^{ch} in our sample, failing to control for the w 's is going to introduce a bias in the estimation of the δ_i^c 's. An example of such a situation is the bias in cohort analysis induced by differential mortality. It is well known that longevity is positively correlated with wealth. This implies that, if we follow a cohort over the last part of its life cycle, such a cohort is going to become, on average, «richer» as it ages. Controlling for variables that are correlated with wealth (like education levels) can help in reducing this problem.

The δ_i^c 's in equation [11] represent the behavior of the average household belonging to cohort c in year t . By following this average household over a long time interval we can study the properties of the age profile for a particular cohort. Two caveats are, however, in order. First, even if a relatively long time series of cross sections is available, it is unlikely that we will have enough years to cover the whole life cycle for a given cohort. Therefore, to cover the whole life cycle, it will be necessary to consider several cohorts simultaneously. Second, the macroeconomic environment changes through time and this will be reflected in changes in the δ_i^c which will interact with life cycle movements and will probably affect different cohorts in different ways. These considerations imply that movements in the estimated δ_i^c will reflect life cycle, business cycle and cohort effects which

¹⁷ The notation here is slightly superfluous because of the linear relationship that links age, cohort and time: $t = c + a$. The exposition here draws from Attanasio (1993).

are, in general, impossible to disentangle without the help of some «structure». The fundamental identification problem here arises from the equality linking year of birth, age and time¹⁸. An example will make this point clear. Suppose that a certain cohort is observed between ages 39 and 40 in a given year, and suppose that we observe a marked increase in consumption for that particular group in that year. By studying the pattern of the estimated δ_i^c it will be impossible to distinguish between the following hypotheses: a) consumption of 40 year old's is always «high» because everybody celebrates his or her 40th birthday (a pure age effect); b) consumption for that particular cohort has been high because they have become eligible for a particularly advantageous tax scheme (a pure cohort effect); c) there has been a change in interest rates which has affected consumption (over and above the pure life cycle effect) for the whole population (a pure time effect). We will further discuss these issues in the next section.

The life cycle behavior of a variable like total household consumption is likely to be linked to a number of demographic and economic variables which have a pronounced life cycle pattern. The most obvious example is family composition. In the early part of a household life cycle family size will increase, it will typically peak around age 40-45 and it will start declining as children start leaving the household. Another example is labor force participation.

It might therefore be useful to decompose life cycle movements in δ_i^c 's into the part that can be linked to movements in a vector of variables z_i^{ach} and the remainder. To do so we can re-specify equation [9] as follows:

$$X_i^{ach} = \tilde{\delta}_i^c + \gamma'w^{ch} + \pi'z_i^{ach} + \tilde{\epsilon}_i^{ach} \quad [12]$$

Suppose we can model the z variables as we modelled the X in equation [11]:

$$z_i^{ach} = \alpha_i^c + \beta'w^{ch} + \eta_i^{ach} \quad [13]$$

Substituting equation [13] into [2] we see that there is a relationship between the $\tilde{\delta}_i^c$'s and the δ_i^c 's:

$$\delta_i^c = \tilde{\delta}_i^c + \pi'\alpha_i^c \quad [14]$$

The $\tilde{\delta}_i^c$ will be interpreted as the average level of X_i^{ach} after having removed the effect of the z variables. These averages can be of interest for at least two reasons.

First, it can be useful in testing one of the main implications of the life cycle model: the smoothing of consumption over the life cycle. While in the

¹⁸ For a discussion of these issues see Heckman and Robb (1987), MaCurdy and Mroz (1991), Attanasio (1993).

textbook presentation of the model this implication follows directly from the concavity of the utility function of an individual consumer, if we want to test it with actual data we will have to modify the model to take into account, among other things, the fact that most consumption decisions are taken by households rather than by single individuals and that household composition obviously affects the utility derived from a given amount of expenditure. An implication of the model will therefore be that consumption per adult equivalent is smoothed over the life cycle. An equation like [14] can therefore be used to remove the effects of changing family composition: the hypothesis relevant for the life cycle model would then be that the $\bar{\delta}_i^c$ rather than the δ_i^c are «smooth» over the life cycle.

The $\bar{\delta}_i^c$ can be estimated by regressing the estimated δ_i^c 's on the estimated α_i^c 's and evaluating the residuals of such a regression. These will be orthogonal, by construction, to the estimated α_i^c , which represent the average behavior of the z 's over the life cycle for different cohorts. If one is interested in testing the hypothesis that consumption «tracks» income over the life cycle, and is found that for $\bar{\delta}_i^c$ such a hypothesis is rejected, one could argue that income is highly correlated with the z 's, so that considering the $\bar{\delta}_i^c$ removes most of the correlations one was trying to detect. While this argument is valid and could be justified for labor supply variables, it is hard to argue that changing family composition removes *all* the tracking of income and consumption because income and family composition are strongly correlated. This objection, however, is valid and it should be stressed that this method does not allow to distinguish between the hypotheses that consumption per adult equivalent is being smoothed versus that consumption tracks income which in turns is correlated with family composition. These issues cannot be addressed without imposing more structure on the problem.

Second, it might be of interest to establish whether an observed shift in behavior can be accounted for by a change in the life cycle profile of one of the z variables or is due to something else. Suppose for instance that we observe that over a particular time period consumption of the cohorts that happen to be in their 30s and 40s is lower than expected from the projected age profile. It might be interesting to know if such a downward shift can be explained, for instance, by a decrease in the average number of children or similar factors.

The estimation of the δ_i^c can only be a first step in the evaluation of the life cycle model. The following steps can go into two different directions. On the one hand, one can try to isolate the pure life cycle profile for consumption. On the other hand one can try to fit a structural model to the data, test it and estimate the structural parameters.

As far as the first alternative is concerned, the identification of the age profile for consumption cannot be done without the help of a number of very strong assumptions, *a priori* information, and/or additional data. From a purely statistical point of view, if we assume that the measured δ_i^c are a complex function of age, time and cohort effects, it is immediately clear that we cannot identify the linear terms of this relation, but only the coefficients on the non

linear terms. This implies that we will be able to say something about the shape of the age profile for the first differences in a given variable but not about its levels. These difficulties can be overcome in different ways. For instance we can assume something about the way in which time effects affect the level of the variable for different cohorts or about the level of a given variable at a certain age for different cohorts. This is the procedure that is behind the smoothing procedure used in the next section. Alternatively we can use information on a different variable which is linked by some sort of relationship to the variable of interest. For instance, in Attanasio (1993), the shape of the age profile for savings is pinned down by looking at information on wealth. Estimating δ_t^i 's for wealth it will be possible to identify the shape of the age profiles for changes in wealth, which is obviously related to savings.

The second alternative is to use the average cohort technique outlined above to estimate and test a structural model. It is to this issue that we turn now.

3.2. Structural model of life cycle behavior

One of the big advantages of working with individual data is that the relationship to be estimated does not have to be linear and can depend on a number of different individual characteristics: aggregation within a cohort can be done over different transformations of the data and over different variables. An example will make this clear.

Suppose that we have a utility function of the form in equation [7] and that the household technology specializes so that we can rewrite [7] as follows:

$$\text{Max } E_t = \sum_{j=0}^{T-t} \beta^j F(C_{t+j}^h, z_{t+j}^h) \quad [15]$$

where:

$$F(C_t^h, z_t^h) = \frac{(C_t^h)^{1-\gamma} - 1}{1-\gamma} \Phi(z_t^h); \quad \Phi(z_t^h) = \exp(\alpha' z_t^h)$$

and the vector z_t^h includes demographic, labor supply and seasonal variables. It is possible to show that under some assumptions about the conditional distribution of innovations to the consumption and the interest rate, it is possible to obtain, for individual h , an Euler equation of the following form:

$$\Delta \log(C_{t+1}^h) = \text{constant} + \alpha' \Delta z_{t+1}^h + \sigma \log(1 + r_{t+1}) + \epsilon_{t+1}^h \quad [16]$$

where $\sigma = 1/\gamma$ and ϵ_{t+1}^h is an expectational error (for a derivation see, for instance, Deaton, 1992). Equation [16] can be aggregated across individuals belonging to the same cohort by summing over the relevant h 's. If all the variables included in the vector z are observable at the individual level, the equation can then be estimated using instrumental variables.

In principle, the definition of a cohort is arbitrary: we could define it as all the individuals born in the same year, or as all the individuals born in the same decade. Defining cohorts tightly allows us the possibility of considering more cohort simultaneously. It should be clear however, that this does not increase the number of independent observations. As stressed above, when estimating an equation like [16] consistency can only be achieved by a large number of *time* observations. The advantages of considering more cohorts are others. First, it is possible to isolate the behavior of cohorts that are observed over a part of their life cycle where the model is more (or less) likely to hold. Second, some variables, such as demographic over, do not exhibit strong variations when averaged across the whole population, while they vary substantially over the life cycle. The consideration of cohorts defined over a narrow interval, therefore, allows us to identify more precisely the effects of these variables. In this sense, therefore, such a procedure increases efficiency.

A few econometric problems can be discussed at this point. Notice that if we had a long panel, equation [16] could be estimated at the individual level and if time aggregation (as discussed in Hall, 1988) was not a problem, instruments dated t or earlier would be valid, under the null of correct specification. Using average cohort data, by aggregating [16] over individuals: we introduce a measurement error (if the sample is smaller than the population) in the *levels* of each variable. This implies the presence of an MA(1) residual when we consider the first difference of the same variable, as in [16]. In empirical applications, like Attanasio and Weber (1993) and Attanasio and Browning (1992) the residuals of equations similar to [16] indeed exhibit strong negative first order autocorrelation. This makes the use of instruments dated t invalid. If the residuals do not exhibit autocorrelation of order higher than one, instruments dated $t-1$ and earlier are valid¹⁹.

A nice interpretation of the average cohort technique used to estimate equation [16] is that provided by Browning and Meghir (1994) and also discussed by Moffitt (1991). The time series of average cohort data can be thought of as the projection of all the individual observations on the interaction of year and cohort dummies. Therefore one can interpret the use of average cohort data as an instrumental variable technique, where the year-cohort dummies are the instruments.

Equation [16] has been estimated by Attanasio and Weber (1993) with the z 's variables including various demographic and labor supply variables. Equation [16] can be generalized by relaxing the assumption of constant elasticity of intertemporal substitution. Blundell, Browning and Meghir (1994) and Attanasio and Browning (1992) allow γ to vary in time as a function of a vector of variables $z_{1,t}^h$. In this case it can be shown that if γ_t^h is given by:

$$\gamma_t^h = \theta_0 + \theta' z_{1,t}^h \quad [17]$$

¹⁹ The information contained in the cross sectional variation can in principle be used to improve the efficiency of the estimates of equation [14]. Deaton (1985) and Fuller (1985) suggest methods that correct for the presence of measurement error.

the Euler equation for individual h will be:

$$\log(1 + r_{t+1}) = \text{constant} + \alpha' \Delta z_{t+1}^h + \theta_0 \Delta \log(C_{t+1}^h) + \theta' \Delta (\log(C_{t+1}^h) z_{1,t+1}^h) + \epsilon_{t+1}^h \quad [18]$$

Notice that now, when we aggregate equation [18] over different individuals, we need not only to compute the mean of the log of consumption, but also of the product of consumption with each of the z_1 variables that affect the elasticity of intertemporal substitution. This is clearly impossible using National Accounts data but is straightforward using micro data. Attanasio and Browning (1992) show that these interaction terms are important on FES data.

Attanasio and Browning also offer a different interpretation of equation [18]. A method which is sometime used to derive Euler equations is the so called λ -constant or Frisch demand functions (see MaCurdy, 1981 and Browning, Deaton and Irish, 1985). This method involves assuming that the marginal utility of money is a function of consumption and a vector of other variables. It is however possible to assume that the marginal utility of money is a linear function of the *logarithm* of consumption and eventually other variables. If this is the case the Euler equation derived is similar, but not identical to equation [16]. Furthermore, if we go back to the specification in terms of the indirect utility function, it is possible to show that utility is not any longer isoelastic.

Estimation of Euler equations for consumption based on aggregate data have strongly rejected the restrictions imposed by the model²⁰. Those studies that have used average cohort techniques on time series of cross sections have reported a much better empirical performance of the model²¹.

Still much work needs to be done. The use of micro data in studying Euler equations for consumption has just begun. In concluding this section we just want to point out some promising areas of research.

While the use of average cohort techniques can easily accommodate, as we have discussed, any non linear transformation in the data, it still requires linearity in the parameters. An interesting challenge is to extend the method to allow the estimation of equations that are non linear in the parameters.

A problem that we have not touched in the discussion so far is that of consumer durables. The usual argument that is made to avoid the treatment of durables is that of separability between non durables and durable services. If the micro data contain information on the stock of durables by individual household, it is possible to condition on the value of that stock and on the relative price of durables, when estimating Euler equations for non durables. More generally, information on durable expenditure requires the study of

²⁰ See, for instance, Hansen and Singleton (1982), Mankiw *et al* (1985), Campbell and Mankiw (1989).

²¹ Attanasio and Weber (1993), Attanasio and Browning (1992).

discrete choice problems that are completely intractable with aggregate time series. While the analysis of discrete choice variables has a well established tradition, to the best of our knowledge, it has never been modelled simultaneously with non durables and services consumption choices.

A similar issue applies to labor supply. It is very likely and has been shown that leisure and consumption are not separable²². It is therefore necessary to condition, when estimating Euler equations for consumption, on labor supply variables. An explicit modelling of such a variables is much more difficult, given the importance of corner solutions in the determination of labor supply.

Finally the level of aggregation across commodities is another important issue. Gorman (1959) derives the assumptions necessary to obtain that a singular price index is sufficient to determine intertemporal allocation. If these conditions do not hold, and there is substantial variability in relative prices, it is necessary to consider more commodities separately.

The consideration of conditional Euler equations, such as those that do not impose separability between leisure and consumption or those that consider more commodities separately, poses some conceptual problems. Consider for instance the problem of the excess sensitivity of consumption to income. If conditioning on labor supply the excess sensitivity disappears, it could be argued that the labor supply variables (hours worked) are strongly correlated with income and that such variables, because of possible measurement problems in income, actually pick up the symptoms of «lack of smoothing».

More generally, without any restrictions on preferences, it is possible to fit, with a flexible enough functional forms, virtually any data set. The question is then one of interpretation of the results: ultimately we are not interested in fitting the data. Rather we ask: given a model of optimizing behavior that fits the data, are the structural parameters economically plausible? If not, this constitutes evidence against the proposed model.

4. Empirical evidence

In this section we will report some simple evaluations of the life cycle model. These evaluations are not formal tests, but rather a simple check of the consistency of the model with the data. They focus on the hypothesis that consumption is smoothed over the life cycle in the face of an uneven income profile. In this sense they can be interpreted as a first step toward a more formal test of the implications of the model.

The main implication of the simplest version of the life cycle model is that consumers with a concave utility function will try to smooth consumption over

²² Browning and Meghir (1991) find that commodity demands are not separable from labor supply, which immediately implies that labor supply cannot be additive separable from consumption.

their life cycle. This does not mean that consumption should be constant: allocation will depend on the relative price of consumption at different dates (interest rates) and on preferences (the elasticity of substitution between different dates and impatience - the discount rate). If preferences are separable between consumption and leisure, there should be no relationship between consumption and current income.

A first test of the life cycle model can therefore be the comparison of the life cycle profiles for consumption and income. The question we will be asking is: to what extent do the estimated profiles resemble the textbook picture of a flat consumption profile and hump-shaped income profile?

If we look at the simple age profile for total family consumption and income, the answer to the question in the previous paragraph is: not at all. However the reasons for these deviations might be pretty obvious. Even without estimating a fully fleshed household technology model as in [9] or [13], it will be necessary to control for demographic changes, if we want the model to have any hope of fitting household data. We also believe that the assumption of separability between consumption and leisure is untenable, especially in the wake of dramatic changes like retirement. We will therefore try to control for labor supply effects in estimating the age profile for consumption.

Another possible interpretation of the results presented in this section is as an illustration of the techniques whose use was advocated in Section 3. The examples we present should make clear both the advantages and the disadvantages of using time series of repeated cross sections.

The data we use are taken from two major surveys: the US Consumer Expenditure Survey and the UK Family Expenditure Survey. Both surveys are collected to compute the weights of the consumer price index and both collect consumption information. The similarities, however, stop here. Each household in the CEX is interviewed 4 times over a one year period. In each interview the reference person is asked to recall the amounts spent in about 500 different commodities in the previous 3 months. Together with these data the survey supplies a substantial amount of information on various household variables including income, demographic and labor supply variables.

FES households stay in the sample only for two weeks and they are asked to keep a diary of all their purchases and to document all major purchases. The quality of the data seems to be higher for the FES. The CEX, however, contains richer information. The FES is available from 1969 to 1990, while the CEX is available from 1980 to 1990. For what we are concerned with in this paper, the much longer time horizon is probably the major advantage of the FES over the CEX. Both for the UK and the US we use expenditure on non durable and services as our definition of consumption.

Another advantage of using individual data is the possibility of computing household specific price indexes, using household expenditure shares. This was done both for the UK and the US.

4.1. Evidence from UK data

The first step in our study is going to be the estimation of equation [11] for income and consumption. The households in the survey are partitioned in cohorts according to their year of birth. We consider ten cohorts and follow them for the 22 years of the survey²³. The cohort definitions are reported in Table 1. The only w variables we controlled for are seasonal factors: the δ_i^c are therefore basically average cohort data. Estimates of the δ_i^c are obtained by OLS over the entire sample. In Figure 1 the δ_i^c obtained for log consumption and log income are plotted against age. Each connected segment represent the average for a given cohort, which is observed, over the sample period, at different ages. From this picture it is evident that consumption tracks income closely, as stressed by Carrol and Summers (1991).

The next step is to control for some demographic variables: equation [14] is estimated by using as z variables the number of adults and the number of children of different ages. The δ_i^c so obtained are plotted against age, along with the original data, in Figure 2. As can be seen from this picture, the consumption age profile looks now much «flatter», especially in the middle part of the life cycle.

TABLE 1
Cohort definition for the UK data

Cohort	Year of birth	Age in 1970	Age in 1986
1	1945-49	21-25	37-41
2	1940-44	26-30	42-46
3	1935-39	31-35	47-51
4	1930-34	36-40	52-56
5	1925-29	41-45	57-61
6	1920-24	46-50	62-66

Cohort definition for the US data

Cohort	Year of birth	Age in 1980	Age in 1988
0	1955-59	21-25	29-33
1	1950-54	26-30	34-38
2	1945-49	31-35	39-43
3	1940-44	36-40	44-48
4	1935-39	41-45	49-53
5	1930-34	46-50	54-58
6	1925-29	51-55	59-63
7	1920-24	56-60	64-68
8	1915-19	61-65	69-73
9	1910-14	66-70	74-78

²³ Some of the cohorts are only followed for part of the sample either because they are too young or too old.




Figure 1
Conditional Income and Non-durable consumption




Figure 2
Non-durable consumption controlling for demographics

In Figure 3 we also control for labor supply variables²⁴. Again the profile is much flatter than in the original case. The interesting aspect here is the fact that the slight decline in consumption in the last part of the life cycle has largely disappeared.

Finally, in Figure 4 we plot the same data presented in Figure 2 against time rather than age. What is apparent in this figure is the degree of synchronization among the various cohorts (with the possible exception of the youngest cohort, which is only observed in the last years of the survey).

The simple facts presented here show that, at least at a first glance, the life cycle model is not dramatically at variance with the data. This grants further investigations. In particular, without some kind of structural model it is not going to be possible to isolate business cycle from age and cohort effects.

4.2. Evidence from US data

The Consumer Expenditure Survey is the only US micro data set with complete information on consumption. Unfortunately the survey covers only 11 years. We consider, again, ten cohorts. The cohort definition is reported in Table 1.

Unlike in the UK, households stay in the sample for 4 consecutive quarters. We construct annual consumption by adding the quarterly figures, while income is taken from the last interview which asks for the amount earned in the twelve months before the interview. All figures are in real dollars and are deflated, as with the UK data, with a household specific Stone price index.

When estimating equation [1], we control for several sources of within cohort heterogeneity. The w variables included 2 education dummies, a dummy for black household head, and 3 regional dummies. The «average cohort data» so obtained are going to be the object of our analysis.

Figure 5 plots average log consumption and average log income against age for the ten cohorts we are considering. Once again, as with the UK data, we observe that the two profile track each other.

In Figure 6, we plot the original average for consumption along with those obtained controlling for demographic variables²⁵. In Figure 7, we control for the same labor supply variables considered for the UK. The story told by these two pictures is analogous to that told by Figure 2 and 3.

²⁴ The labor supply variables considered are: a dummy for working husband, a dummy for working wife, the number of hours worked by husband and wife, a white collar dummy.

²⁵ These variables included the number of children between 0 and 2, the number of males and females between 2 and 15 the number of males and females above 16, and the number of adults.




Figure 3
Consumption controlling for demographics and labor supply




Figure 4
Non-durable consumption over the business cycle




Figure 5
Conditional Income and Non-durable consumption




Figure 6
Non-durable consumption controlling for demographics




Figure 7
Consumption controlling for demographics and labor supply




Figure 8
Non-durable consumption over the business cycle

Finally in Figure 8, we plot average log consumption (after removing the effects of demographic and labor supply variables) against time. The degree of coordination among different cohorts is less pronounced than in the analogous figure for the UK. It should be considered that the sample period is much shorter and, unlike in the UK, barely covers one business cycle. The profile appears to be considerably flatter, even though not, in this case, we can still observe a peak in the middle part of the life cycle. Furthermore, the increase in the early part is more marked than in the UK.

From the evidence presented in this section it can be concluded that neither in the UK nor in the US does the life-cycle model seem to be at variance with the observed allocation of resources over the life cycle, once one controls for changing family composition and/or labor supply behavior.

5. Conclusions

In this paper we stressed the need to use microeconomic data to assess the empirical relevance of the life cycle model. We first show that evaluating the model on the basis of aggregate time series can be extremely misleading. This does not imply, however, that the study of macroeconomic time series is useless. Not only are they useful for forecasting and policy analysis, but an appropriate statistical modeling of their dynamic properties is going to be extremely useful to understand their long run properties. The life cycle model can be useful in this respect to suggest the variables that can be used in formulating the dynamic equation to describe the behaviour of aggregate consumption.

The dynamic specification should be flexible enough to reflect the main implication of the life cycle model. The observation that aggregate consumption reacts in different ways to anticipated and unanticipated and to transitory and permanent changes to income was one of the motivations for the development of the life cycle theory. A statistical description of aggregate consumption should allow for these properties.

What we object to is an evaluation of the life cycle model on the basis of checking the first order conditions of an *individual* optimization problem with aggregate time series data.

In Section 3, we showed that the analysis of a single cross section is of limited use in evaluating the life cycle model. The absence of panel data can, however, be overcome by the use of time series of repeated cross sections.

The analysis of repeated cross sections can be made at different levels: from a mainly descriptive one, to a more structural. In Section 4, we estimated simple age profiles for income and consumption and showed that, once we control for demographic and labour supply variables, the model is not dramatically at variance with the data.

More work is needed in this area, from the theoretical, but especially from the

empirical point of view. The analysis of micro data set on consumption has just started and an enormous amount of work remains to be done.

References

- Atkinson, A. B. and Micklewright, J. (1983): «On the Reliability of Income Data in the Family Expenditure Survey, 1970-77», *Journal of the Royal Statistical Society, A*, 146, pp. 33-61.
- Attanasio, O. P. (1993): «A Cohort Analysis of Saving Behavior by US Households», IFS Discussion Paper No. W93/4.
- Attanasio, O. P. and Browning, M. J. (1992): «Consumption over the Life Cycle and over the Business Cycle», Mimeo.
- Attanasio, O. P. and Weber, G. (1989): «Intertemporal Substitution, Risk Aversion and the Euler Equation for Consumption», *Economic Journal*, Supplement 99, pp. 59-73.
- Attanasio, O. P. and Weber, G. (1993): «Consumption Growth, the Interest Rate and Aggregation», *Review of Economic Studies* 60, pp. 631-649.
- Bean, C. R. (1986): «The Estimation of 'Surprise' Models and the 'Surprise' Consumption Function», *Review of Economic Studies* 53, pp. 497-516.
- Becker, G. (1962): «Irrational Behavior and Economic Theory», *Journal of Political Economy* 70, pp. 1-13.
- Blinder, A. and Deaton, A. (1985): «The Time Series Consumption Function Revisited», *Brookings Papers on Economic Activity* 2, pp. 465-521.
- Blundell, R.; Browning, M. J. and Meghir, C. (1994): «A Microeconomic Model of Intertemporal Substitution and Consumer Demand», *Review of Economic Studies* 61, pp. 57-80.
- Blundell, R.; Pashardes, P. and Weber, G. (1993): «What Do We Learn from Micro Data?», *American Economic Review* 83, pp. 570-597.
- Browning, M. J. (1989): «A Non-Parametric Test of the Life-Cycle Rational Expectations Hypothesis», *International Economic Review* 30, pp. 979-992.
- Browning, M. J. (1993): «Estimating Micro Parameters from Macro Data: Some Pessimistic Evidence», *Ricerche Economiche* 47, pp. 101-115.
- Browning, M. J.; Deaton, A. and Irish, M. (1985): «A Profitable Approach to Labor Supply and Commodity Demands over the Life Cycle», *Econometrica* 53, pp. 503-544.
- Browning, M. J. and Meghir, C. (1991): «Testing for the Separability of Commodity Demands from Male and Female Labour Supply», *Econometrica* 59, pp. 925-951.
- Browning, M. J. and Meghir, C. (1994): «Food, Children and Fixed Effects», in progress.
- Campbell, J. Y. (1987): «Does Savings Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis», *Econometrica* 55.
- Campbell, J. Y. and Mankiw, N. G. (1989): «Permanent Income, Current Income and Consumption», *NBER Macroeconomic Annual*, Cambridge.
- Carroll, C. and Summers, L. H. (1991): «Consumption Growth Parallels Income Growth: Some New Evidence». In Bernheim, B. D. and Shoven, J. B. (eds.): *National Saving and Economic Performance*, Chicago University Press, Chicago.
- Chamberlain, G. (1983): «Panel Data», in *Handbook of Econometrics*, North Holland.
- Deaton, A. (1985): «Panel Data from Time Series of Cross Sections», *Journal of Econometrics* 30, pp. 109-126.
- Deaton, A. (1987): «Life Cycle Models of Consumption: Is the Evidence Consistent with the Theory?», in Bewley, T. (ed.): *Advances in Econometrics*, vol. 2, Cambridge University Press.
- Deaton, A. (1991): «Savings and Liquidity Constraints», *Econometrica* 59, pp. 1221-1248.

- Deaton, A. (1992): *Understanding Consumption*, Oxford University Press.
- Flavin, M. (1981): «The Adjustment of Consumption to Changing Expectations about Future Income», *Journal of Political Economy* 89, pp. 974-1009.
- Friedman, M. (1957): *A Theory of the Consumption Function*, Princeton University Press.
- Fuller, W. A. (1987): *Measurement Error Models*, John Wiley, New York.
- Gorman, W. M. (1953): «Community Preference Fields», *Econometrica* 21, pp. 63-80.
- Gorman, W. M. (1959): «Separable Utility and Aggregation», *Econometrica* 27, pp. 469-481.
- Grossman, S. J. and Shiller, R. J. (1982): «Consumption Correlatedness and Risk Measurement in Economies with Non Traded Assets and Heterogeneous Information», *Journal of Financial Economics* 10, pp. 195-210.
- Hall, R. E. (1978): «Stochastic Implications of the Life Cycle - Permanent Income Hypothesis: Theory and Evidence», *Journal of Political Economy* 86, pp. 971-987.
- Hall, R. E. (1988): «Intertemporal Substitution in Consumption», *Journal of Political Economy* 96, pp. 339-357.
- Hall, R. E. and Mishkin, F. S. (1982): «The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households», *Econometrica* 50, pp. 461-481.
- Hansen, L. P. and Singleton, K. J. (1982): «Generalized Instrumental Variables Estimation of Non-Linear Rational Expectations Models», *Econometrica* 50, pp. 461-481.
- Heckman, J. (1974): Life Cycle Consumption and Labor Supply: An Explanation of the Relationship Between Consumption and Income over the Life Cycle», *American Economic Review* 64, pp. 188-199.
- Heckman, J. and Robb (1987): «Using Longitudinal Data to Estimate Age, Period and Cohort Effects in Earnings Equations», in: Mason, W. and Fienberg, S. (eds.): *Cohort Analysis in Social Research*, Springer-Verlag, New York.
- Jackman, R. and Sutton, J. (1982): «Imperfect Capital Markets and the Monetarist Black Box: Liquidity Constraints and the Asymmetric Effects of Interest Rate Policy», *Economic Journal* 92, pp. 108-128.
- Keane, M. and Runkle, D. (1992): «On The Estimation of Panel Data Models with Serial Correlation when Instruments are not Strictly Exogenous», *Journal of Business and Economic Statistics* 109, pp. 83-110.
- MaCurdy, T. E. (1981): «An Empirical Model of Labor Supply in a Life Cycle Setting», *Journal of Political Economy* 89, pp. 1059-1085.
- MaCurdy, T. E. and Mroz, T. (1991): «Measuring Macroeconomic Shifts in Wages from Cohort Specifications», Stanford University, Mimeo.
- Mankiw, N. G.; Rotemberg, J. J. and Summers, L. H. (1985): «Intertemporal Substitution in Macroeconomics», *Quarterly Journal of Economics* 100, pp. 225-251.
- Miron, J. (1986): «Seasonal Fluctuations and the Life Cycle - Permanent Income Model of Consumption», *Journal of Political Economy* 94, pp. 249-265.
- Modigliani, F. and Brumberg, R. (1954): «Utility Analysis and the Consumption Function: An Interpretation of Cross Section Data», in Kurihara, K. K. (ed.): *Post Keynesian Economics*, George Allan and Unwin.
- Modigliani, F. and Ando, A. (1960): «The 'Permanent Income' and the 'Life Cycle' Hypothesis of Saving Behavior: Comparisons and Tests», *Consumption and Saving*, vol. 2, University of Pennsylvania.
- Modigliani, F. and Ando, A. (1963): «The Life Cycle Hypothesis of Saving: Aggregate Implications and Tests», *American Economic Review* 53, pp. 55-84.
- Moffitt, R. (1993): «Identification and Estimation of Dynamic Models with a Time Series of Repeated Cross Sections», *Journal of Econometrics* 59, pp. 73-98.
- Muellbauer, J. (1975): «Aggregation, Income Distribution and Consumer Demand», *Review of Economic Studies* 62, pp. 525-543.

- Runkle, D. (1991): «Liquidity Constraints and the Permanent Income Hypothesis», *Journal of Monetary Economics* 27, pp. 73-98.
- Shorrocks, A. F. (1975): «The age-wealth relationship: a cross-section and cohort analysis», *Review of Economics and Statistics* 57, pp. 155-163.
- Zeldes, S. (1989): «Consumption and Liquidity Constraints: an Empirical Investigation», *Journal of Political Economy* 97, pp. 305-345.

Resumen

En este trabajo analizamos las implicaciones de la teoría del ciclo vital del consumo y consideramos la posibilidad de contrastarla utilizando datos micro y macroeconómicos, concluyendo que la teoría no se puede contrastar con los segundos. Asimismo, proporcionamos alguna evidencia, basada en datos microeconómicos norteamericanos e ingleses, que demuestra que la teoría no es, en principio, incompatible con los datos.

Recepción del original, abril de 1994

Versión final, agosto de 1994