CONSUMPTION DEMAND

Orazio P. Attanasio

Working Paper **6466**

CONSUMPTION DEMAND

Orazio P. Attanasio

Consumption Demand

ABSTRACT

Consumption is the largest component of GDP. Since the 1950s, the life cycle and the permanent income models have constituted the main analytical tools to the study of consumption behavior, both at the micro and at the aggregate levels. Since the late 1970s the literature has focused on versions of the model that incorporate the hypothesis of Rational Expectations and a rigorous treatment of uncertainty. In this paper, I survey the most recent contribution and assess where the life cycle model stands. My reading of the evidence and of recent developments leads me to stress two points: (i) the model can only be tested and estimated using a flexible specification of preferences and individual level data; (ii) it is possible to construct versions of the model that are not rejected by the data. One of the main problems of the approach used in the literature to estimate preferences is the lack of a "consumption function." A challenge for future research is to use preference parameter estimates to construct such functions.

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

1. Introduction.

In most developed economies, consumption accounts for about two thirds of GDP. Moreover, it is from consumption that, in all likelihood, utility and welfare are in large part determined. It is therefore natural that macroeconomists have devoted a considerable amount of research effort to its study. In modern macroeconomics, consumption is typically viewed as part of a dynamic decision problem. There is therefore another sense in which an understanding of consumption is central for macroeconomics. Consumption decisions are also saving decisions from which the funds available for capital accumulation and investment arise. Therefore, consumers attitudes to saving, risk bearing and uncertainty are crucial to understand the behaviour of capital markets, the process of investment and growth and development. It is not by chance that modern consumption theory is also used to characterise asset prices equilibrium conditions. The desire consumers might have to smooth fluctuations over time determines the need for particular financial instruments or institutions. Understanding recent trends in consumption and saving is crucial to the study, both positive and normative, of the development of financial markets, of the institutions that provide social safety nets, of the systems through which retirement income is provided and so on.

One of the main themes of this chapter is that consumption decisions cannot be studied in isolation. Exactly because consumption and saving decisions are part of a dynamic optimisation problem, they are determined jointly with a number of other choices, ranging from labour supply to household formation and fertility decisions, to planned bequests. While modelling all aspects of human economic behaviour simultaneously is probably impossible, it is important to recognise that choices are taken simultaneously and to control for the effects that various aspects of the economic environment in which consumers live might have on any particular choice. This is particularly true if one wants to estimate the parameters that characterise individual preferences.

Implicit in this argument is another of the main themes of this chapter: consumption decisions should be modelled within a well specified and coherent optimisation model. Such a model should be flexible and allow for a variety of factors. Indeed, I think it is crucial that the model should be interpreted as an approximation of reality and should allow for a component of behaviour that we are not able to explain. However, such a model is crucial to organise our thinking and our understanding of the data. Without a structural model it is not possible to make any statement about observed behaviour or to evaluate the effect of any proposed change in economic policy.

This, however, is not a call for a blind faith in structural models. Inferences should always be conditional on the particular identification restrictions used and on the particular structural model used. Such models should also be as flexible as possible and incorporate as much information about individual behaviour as is available. It should be recognised, however, that without such models we cannot provide more than a statistical description of the data.

The other main theme of the analysis in this chapter is that to understand aggregate trends it is necessary to conduct, in most situations, a detailed analysis of individual behaviour. In other words, aggregation problems are too important to be ignored. This obviously does not mean that the analysis of aggregate time series data is not useful. Indeed, I start the chapter with a brief summary of the main time series properties of consumption. Estimation of structural models of economic behaviour, however, cannot be performed using aggregate data only.

This chapter is not an exhaustive survey of the literature on consumption: such a literature has grown so much that it would be hard even to list it, let alone summarise all the contribution. What I offer, instead, is a discussion of the current status of our knowledge, with an eye to what I think are the most interesting directions for future research. In the process of doing so, however, I discuss several of the most important and influential contributions. Omissions and exclusions are unavoidable and should not be read as indicating a negative judgement on a particular contribution. At times, I simply chose, among several contributions, those that most suited my arguments and helped me the most to make a given point. Moreover, notwithstanding the length of the chapter, not every sub-fields and interesting topic has been covered. But a line had to be drawn at some point. There are four fields that I did not included in the chapter and over which I have agonised considerably. The first is asset pricing: while much of the theoretical material I present has direct implications for asset prices, I decided to omit a discussion of these implications as there is an entire chapter of this Handbook devoted to these issues. The second is the axiomatisations of behaviour under uncertainty alternative to expected utility. There are several interesting developments, including some that have been used in consumption and asset pricing theory, such as the Kreps-Porteus axiomatisation used by Epstein and Zin (1989, 1991) in some fascinating papers. The third is the consideration of within household allocation of resources. There is some exciting research being developed in this area, but I decided to draw the line of 'macro' at the level of the individual household. Finally, I do not discuss theories of consumption and saving behaviour that do not assume optimising and fully rational behaviour. Again, there is some exciting work in the area of social norms, mental accounting, time varying preferences, herd behaviour and so on. In the end, however, I decided that it would not fit with the rest of the chapter and rather than giving just a nod to this growing part of the literature I decided to leave it out completely.

The chapter is organised as follows. In section 2, I start with a brief description of some stylised facts about consumption. These include both facts derived from aggregate time series data and from household level data. Throughout the section, I use in parallel data from two countries: the US and the UK.

In section 3, I discuss at length what I think is the most important model of consumption behaviour we have, the life cycle model. In that section, I take a wide view of what I mean by the life cycle model: definitely not the simple textbook version according to which the main motivation for saving is the accumulation of resources to provide for retirement. Instead, I favour a flexible version of the model where demographics, labour supply, uncertainty and precautionary saving and possibly bequests play an important role. In other words, I consider the life cycle model as a model in which consumption decisions are determined within an intertemporal optimisation framework. What elements of this model turn out ot be more important is largely an empirical matter. Indeed, even the presence of liquidity constraints, or borrowing restrictions, can and should be incorporated within this framework.

In section 4, I discuss aggregation problems. In particular, I focus on two different kinds of aggregation: that across consumers and that across commodities. The aim of this section is not just to give lip service to the aggregation issues and proceed to sweep them under the carpet. With the development of computing and storage capability and with the availability of increasing large number of micro data sets, it is important to stress that scientific research on consumption behaviour cannot afford to ignore aggregation issues.

In section 5, I consider the empirical evidence on the life cycle model and discuss both evidence from aggregate time series data and evidence from micro data. In this section I also address a number of econometric problems with the analysis of Euler equations for

consumption. In section 6, I take stock on what I think is the status of the life cycle model, given the evidence presented in Section 5.

In section 7, I address the issues of insurance and inequality. In particular, I present some of the tests of the presence of perfect insurance and discuss the little evidence there is on the evolution of consumption inequality and its relationship to earning inequality.

Most of the models considered up to this point assume time separability of preferences. While such a hypothesis is greatly convenient from an analytical point of view, it is easy to think of situations in which it is violated. In section 8, I discuss to forms of time dependence: that induced by the durability of commodities and habit formation. Section 9 concludes the paper.

2. Stylised Facts.

In this section, I document the main stylised facts about consumption behaviour using both aggregate and individual data. I consider two components of consumption expenditure: on non durable and services and on durables. In addition I also consider disposable income. While most of the facts presented here are quite well established, the evidence in this section constitute the background against which one should set the theoretical model considered in the rest of the chapter.

The data used come from two western countries: the United States and the United Kingdom. I have deliberately excluded from the analysis developing or less developed countries as they involve an additional set of issues which are not part of the present discussion. Among the developed countries I have chosen the US and the UK both because data from these two countries have been among the most widely studied and because the two countries have the best micro data on household consumption. For the UK, in particular, the Family Expenditure Survey runs for 25 consecutive years, giving the possibility of performing interesting exercises.

2.1 Aggregate Time Series Data.

In this section, I present some of the time series properties of consumption expenditure and of disposable income. While the models considered in the following sections refer to household behaviour, typically the consumption aggregates considered in the National Account statistics include outlays of a sector that, together with households, includes other entities, such as charities, whose behaviour is unlikely to be determined by utility maximisation. While this issue is certainly important, especially for structural tests of theoretical models of household behaviour, in the analysis that follows I ignore it and, instead of isolating the part of total expenditure to be attributed to households, I present the time series properties of National Account consumption. Seslnick (1994) has recently stressed the importance of these issues.

In Figure 1, I plot household (log) disposable income along with consumption divided into durables and non durables and services for the UK and the US. The series have quarterly frequency and run from 1959:1 to 1996:3 for the US and from 1965:1 to 1996:2 for the UK. The data are at constant prices and are seasonally adjusted. From the Figure, it is evident that non durable consumption is smoother than disposable income. Durable consumption, on the other hand, which over the sample accounts, on average, for 13% of total consumption in the US and around 14% in the UK, is by far the most volatile of the three

time series. This is even more evident in Figure 2 where I plot the annual rate of changes for the three variables. In Table 1, I report the mean and standard deviation of the three variables. These figures confirm and quantify the differences in the variability of the three variables considered.

In Tables 2 and 3, I consider two alternative ways of summarising the time series properties of the three series I analyse for both countries. In Table 2, I report the estimates of the coefficient of an MA(12) model for the same series. The advantage of such an un-parsimonious model is that it avoids the sometimes difficult choice among competing ARMA representations. Furthermore, its impulse response function can be easily read from the estimated coefficients. I also purposely decided to be agnostic about the presence of random walks in the time series consumption or income, even though this has implications for the so called 'excess smoothness' puzzle briefly discussed below. In Table 3, instead, I report the Maximum Likelihood estimates of a parsimonious ARMA model for the first differences of the log of the three variables. While in some cases there were alternative specifications that fitted the data as well as those reported in the table, the latter all pass several diagnostic tests. The Q- statistics reported in the table indicates that the representations chosen capture adequately the dynamic behaviour of the series over the period considered.

The time series properties of the rate of growth of the three variables are remarkably different. Notice, in particular, the fact that both in the UK and in the US, the sum of the MA coefficients for non durable consumption is positive, while that for durables is negative. The time series properties of non durable consumption differ remarkably: in table 2 the sum of the first 12 MA coefficient is much larger in the UK than in the US. Furthermore, while the US data are well represented by an MA(3) (with the first and third lag large and very strongly significant), the UK require an AR(2) model.[1]

The sum of the MA coefficients for disposable income in both countries is quite small in absolute value, but is positive in the US and negative for the UK. As far as a 'parsimonious' specification is concerned, in the US I chose an MA(1) for the first differences, even though its coefficient is not very large and is statistically insignificant. This model was almost indistinguishable from an AR(1) model. In the UK, the best model for disposable income is an ARMA(1,1). The richer dynamics of the UK series is also evident in the pattern of the MA coefficients in Table 2.

The properties of durable consumption are particularly interesting. The fact that the time series properties are inconsistent with a simple model which adds durability to the standard random walk property derived from some version of the permanent income has been noticed by Mankiw (1982). Such a model would imply an MA(1) model for the changes in expenditure with a coefficient that would differ from minus one by an amount equivalent to the depreciation rate. As can be seen from Table 2, the US series' best representation is indeed an MA(1) with a negative coefficient; but that coefficient is far from minus one.[2] Caballero (1990) has interpreted this and the fact that, as reported in Table 3 for both countries, the sum of the 12 MA coefficients is negative and much larger in absolute value,

---

[1] The presence of an MA(3) effect in the non durable series for the US is evident even in the MA(12) representation but it is not very robust. If one truncates the sample to 1990 or dummies out the few quarters corresponding to the 1990-91 recession, θ_3 is estimated non-significantly different from zero both in the MA(12) and in the MA(3) model. The same result is obtained if one excludes services from this series.

[2] For durable consumption in the UK, the best model is an ARMA(2,1), by far the most complex model I fitted to these data.

as an indication of the presence of inertial behaviour that 'slows down' the process of adjustment of durables.

Having characterised the main time series properties of consumption and income, the next step would be the estimation of a multivariate time series model that would stress the correlations among the variables considered at various leads and lags. Indeed, some of the studies I cite below, such as Flavin (1981), do exactly this with the purpose of testing some of the implications of the life cycle- permanent income hypothesis. For the sake of brevity, I omit the characterisation of the multivariate time series process of consumption and other macro variables. One of the reasons for this omission is the belief, discussed below, that aggregation problems make it very difficult to give structural interpretation to this type of results. This does not mean, however, that aggregate time series studies are not useful.

The careful specification of a flexible time series model for consumption and other variables can be quite informative, especially if the dynamic specification allows for the type of dynamic effects implied by the microeconomic behaviour. Several of the studies by David Hendry and his collaborators are in this spirit; one of the most widely cited examples of this literature is the paper by Davidson et al. (1978).

The approach taken in these papers, which received a further motivation by the development of cointegration techniques, is to estimate a stable error correction model which relates consumption to other variables. The statistical model then allows to identify both short run and long run relationships between consumption and its determinants. While the theory can be informative on the choice of the relevant variables and even on the construction of the data series, it does not provide explicit and tight restrictions on the parameters of the model. A good example of a creative and informative use of this type of techniques is Blinder and Deaton (1985). While it is difficult to relate this type of models to structural models and therefore they cannot be directly used for evaluating economic policy, they constitute useful instruments for summarising the main features of the data and, if used carefully, for forecasting. Often the lack of micro economic data makes the use of aggregate time series data a necessity. The only caveat is that these studies cannot be used to identify structural parameters.

2.2 Household consumption expenditure.

In this section, I use two large microeconomic data set to document the main stylised facts about consumption. The two data sets used are the US Consumption Expenditure Survey (CEX) and the UK Family Expenditure Survey (FES). Both data sets are run on a continuous basis to gather information for the construction of the weights for the CPI (RPI in the UK). They have, however, been extensively used by researchers and have now become an essential tool to study household consumption and saving behaviour. The focus of the analysis is going to be the household. No attempt will be made to attribute consumption to the single household members, even though some (limited) information on this does exist.[3]

---

[3] Both data sets contain very detailed information on the expenditure on individual commodities. Some of this information can be used to attribute some items to some household members. For many items, however, such attribution is difficult both in practice and conceptually. Browning (1987) has imputed expenditure on alcohol and tobacco to the adults to check whether predicted changes in household income and composition (such as the arrival of children with consequent -at least temporary- withdrawal from the labour force of the wife) cause changes in consumption.

Most of the descriptive analysis presented below attempts at describing the main features of the life cycle profile for consumption expenditure and some other variables. This approach reflects the fact that the theoretical discussion in the next sections will be focused around the life cycle model.

2.2.1 Nature of the data sets and their comparability with the National Account data.

The FES is now available for 25 consecutive years. Each year around 7,000 households are interviewed and supply information on their consumption patterns as well as their demographic characteristics and several other economic variables such as employment status, income, education and so on. Each household stays in the sample for two weeks, during which it fills a diary in which all expenditure items are reported. At the end of the two week period an interviewer collects the diaries and asks additional information on durables acquired during the previous three months and on all major expenditure items reported in the diary and periodic expenditures such as utilities.

The CEX is available on a continuous and roughly homogeneous basis since 1980. Each year about 7,000 different households are interviewed for 4 subsequent interviews, with quarterly frequency.[4] Each month new households enter the survey to replace those that have completed their cycle of interviews. During each interview the household is asked to report expenditure on about 500 consumption categories during each of the three months preceding the interview.[5] The panel dimension of the CEX is unfortunately very short: because each household is only interviewed four times, seasonal variability is likely to dominate life cycle and business cycle movements. In what follows, I do not exploit the panel dimension of the survey.

There have been several discussions about the quality of survey data and the importance of measurement error and about their ability to reproduce movements in aggregate consumption. Several studies, both in the US and the UK have addressed the issue.[6] It should be stressed that the aggregated individual data and the National Account aggregate should be expected to differ for several reasons. First of all, for many consumption categories, the definitions used in the surveys and in the National Accounts are quite different. Housing, for instance, includes imputed rents in the National Accounts data but does not in the surveys. In the CEX, health expenditure measures only out-of-pocket expenditures, while the National Accounts definition includes all health expenditures regardless of the payee. Furthermore, the populations of reference are quite different. Surveys, for instance, do not include institutionalised individuals, while the National

---

Gokhale, Kotlikoff and Sabelhaus (1996) in their study of saving behaviour have attempted to impute all of consumption to the individual household members.

[4] In total there are data for over 20,000 interviews per year. Each household is in fact interviewed five times. However, the Bureau for Labor Statistics does not release information on the first (contact interview). The Bureau of Labor Statistics also runs a separate survey based on diaries which collects information on food consumption and 'frequently purchased items'.

[5] Unfortunately, the monthly decomposition of the quarterly expenditure is not very reliable. For several commodities and for many households, the quarterly figure is simply divided by three. Given the rotating nature of the sample, the 'quarters' of expenditure do not coincide perfectly. For instance, somebody interviewed in December will report consumption in September, October and November, while somebody interviewed in November will report consumption in August, September and October.

[6] See, for instance, Seslnick (1992) , and Paulin et al. (1988) for comparisons between the aggregate Personal Consumption Expenditure and the CEX in the US and the papers in Banks and Johnson (1997) for comparisons on the FES and the UK National Accounts.

Accounts do. Finally, National Account data are not exempt from measurement error that, for some items, can be quite substantial. Should major difference emerge, it is not obvious that the National Account data should be considered as being closer to the 'truth'.

The issues that arise are different for the two data sets. Overall, the degree of correspondence between the aggregated individual data and the National Account data seems to be higher in the UK. For most consumption components, aggregating the FES data, one obtains about 90% of the corresponding National Accounts figure, while the same ratio is about 65% for the CEX in the 1980s. This is probably due to the use of diaries rather than recall interviews. The latter, perhaps not surprisingly, tend to underestimate consumption. In both surveys, however, because of the consistent methodology used over time, there is no major trend in the ratio of the aggregated individual data to the corresponding National Accounts aggregates.[7] Furthermore, the dynamics of consumption and income growth and of saving in both the aggregated CEX and FES data do not track the corresponding macroeconomic aggregates badly. The data are therefore not only useful to characterise individual behaviour and its shifts over time, but also to make inferences, based on individual behaviour, about possible explanations of the observed macroeconomic trends.

2.2.2 Life cycle profiles

In the second part of the paper, in which I discuss the main theoretical model of consumption behaviour, a substantial amount of attention is devoted to the life cycle model in its several incarnations. In this section, I present life cycle profiles for consumption, its components and various other variables in the US and the UK. In this sense, the life cycle model is the conceptual framework that I use to organise the presentation of the microeconomic data.

As the data sets I use are not panels, to estimate age profiles, I am forced to use grouping techniques. These techniques were first used within life cycle models by Browning, Deaton and Irish (1985).[8] The idea is quite simple. Rather than following the same individual over time, one can follow the average behaviour of a group of individuals as they age. Groups can be defined in different ways, as long as the membership of the group is constant over time.[9] Within the life cycle framework, the natural group to consider is a 'cohort', that is individuals (household heads) born in the same period. Therefore, to compute the life cycle profile of a given variable, say log consumption, one splits the households interviewed in each individual cross section in groups defined on the basis of the household head's year of birth. This involves, for instance, considering all the individuals aged between 20 and 24 in 1980, those aged between 21 and 25 in 1981 and so on to form the first cohort; those aged

---

[7] There are substantial differences in this ratio between the early CEX surveys (1960-61 and 1972-73) and those of the 1980s, probably due to the differences in the methodology employed. In the FES the one commodity for which a (downward) trend in the ratio is apparent is tobacco.

[8] Ghez and Becker (1975) use observations on individual of *different ages* to study life cycle behaviour. However, as they use a single cross section, they do not control for cohort effects as Browning et al. (1985) do. Deaton (1985) and, more recently, Moffitt (1994) have studied some of the econometric problems connected with the use of average cohort techniques. Heckman and Robb (1987), MaCurdy and Mroz (1989) and Attanasio (1994) discuss identification issues.

[9] Group membership should be fixed over time so that the sample is drawn from the same population and the sample mean is a consistent estimator of the mean of the same population. Attanasio and Hoynes (1996) discuss the implications of differential mortality for the use of average cohort techniques. Other possible problems arise, at the beginning of the life cycle, from the possible endogeneity of household formation and, more generally, from migration.

between 25 and 29 in 1980, between 26 and 30 in 1981 and so on to form the second cohort, etc. Having formed these groups in each year in which the survey is available, one can average log consumption and therefore form pseudo panels: the resulting data will have dimension QxT, where Q is the number of groups (cohorts) formed and T is the number of time periods.[10] Even if the individuals used to compute the means in each year are not the same, they belong to the same group (however defined) and one can therefore study the dynamic behaviour of the average variables. Notice that non linear transformations of the variables do not constitute a problem as they can be computed before averaging.

The resulting age profiles will not cover the entire life cycle of a given cohort, unless the available sample period is longer than any of the micro data set commonly used. Each cohort will be observed over a (different) portion of its life cycle.

These techniques can be and have been used both for descriptive analysis and for estimating structural models. Their big advantage is that they allow to study the dynamic behaviour of the variables of interest even in the absence of panel data. Indeed, in many respects, their use might be superior to that of panel data.[11] Furthermore, as non linear transformations of the data can be handled directly when forming the group means, they allow one to solve various aggregation problems that plague the study of structural models with aggregate time series data.

In what follows, I define groups on the basis of the year of birth and educational attainment of the household head. The length of the interval that defines a birth cohort is chosen taking into account the trade-off between cell size and within cell homogeneity. Table 4 contains the definition of the cohorts and the average sample size for both surveys.

We start, in Figure 3 and 4, with the life cycle profile of (log) consumption and disposable income at constant prices for both countries. The units of measurement for income and consumption are chosen so that the two graphs would be roughly in the same scale, enabling to stress the differences in the shape of the age profile. In the Figures, I plot the average cohort (log) consumption at each point in time, against the median age of the household head. Each connected segment represent the behaviour of a cohort, observed as it ages, at different points in time. As each cohort is defined by a five year interval, and both surveys cover a period longer than five years, at most ages we observe more than one cohort, obviously in different years. It might be tempting to attribute the differences between adjacent cohorts observed at the same age, to 'cohort effects'. It should be remembered, however, that these observations refer to different time periods and might therefore be reflecting business cycle effects. The plotted profiles reflect age, time and cohort effects[12] that, without an arbitrary normalisation or additional information from a structural model, cannot be disentangled.

Several considerations are in order. First of all, both consumption and income age profiles present a characteristic 'hump'. They both peak in the mid 40s and decline afterwards.

---

[10] Here I am implicitly assuming that the pseudo panel is a balanced one. This is not always the case as each group might be observed for a different number of time periods. Suppose, for instance, to have data from 1968 to 1994. One might want to follow the cohort born between 1965 and 1970 only from the late 1980s or the early 1990. On the other hand, at some point during the 1980s one might want to drop the cohort born between 1906 and 1910.

[11] Time series of cross sections are probably less affected by non random attrition than panel data. Furthermore, in many situation, averaging across the individuals belonging to a group can eliminate measurement error and purely idiosyncratic factors which are not necessarily of interest. As most grouping techniques, average cohort analysis has an Instrumental Variable interpretation.

[12] As well as measurement error and small sample variability.

The picture seems, at first glance, to contradict the implications of the life cycle model as stressed in the typical textbook picture which draws a 'hump shaped' income profile and a flat consumption profile. For total disposable income, the decline around retirement age is faster in the UK than in the US, but approximately of the same magnitude. This probably reflects the more synchronised retirement of British individuals. The consumption profiles, however, present some strong differences. The most notable is the fact that UK consumption declines much more at retirement than US consumption. Total consumption at age 70 is roughly 35% of the peak in the UK and above 50% in the US. I discuss the decline of consumption at retirement below.

In the UK consumption profile, the consumption boom of the late 1980s, followed by the bust of the early 1990s, is quite apparent. Notice, in particular, the fact that the aggregate consumption boom is accounted for mainly by the youngest cohorts. I have discussed elsewhere how to interpret that episode. It is worth stressing, however, that the analysis of the cross sectional variability of consumption can be useful to shed some light on the nature of episodes that the analysis of the time series data cannot explain. Information about which groups in the populations where mainly responsible for a determinate episode can be informative about alternative hypotheses. [13]

It is not obvious how to assess the time series volatility of (log) consumption and income. The main reason for this is that a large part of the variation of consumption over the life cycle is very predictable and can be explained by age and cohort effects. Furthermore, given the limited size of our samples, the year to year variation in the average cohort data reflects both genuine time series variation and the measurement error induced by sample variation. As Deaton (1985) has stressed, some information about the size of the measurement error can be gained using the within cell variability of the variables used. Using this information, one might correct for that part of variability accounted for by sampling variation and attempt to isolate the genuine time variation. In an attempt to isolate this component, I run a regression of log consumption and income on a fifth order polynomial in age and cohort dummies and consider the deviations of the observed profiles from such a profile. The standard deviation of the changes in these deviations, corrected for that part which can be attributed to sampling error, is my measure of time variability. [14] These estimates of volatility for (log) income and consumption are reported in Table 5 along with those for the other variables considered. The first column refers to the US, while the second and third columns are computed using the UK data. The former includes the whole sample, while the latter truncates it to 1986 to remove the effect of the consumption 'boom and bust' of the last part of the sample.

As in the case of aggregate time series, total consumption appears less volatile than disposable income, both in the UK and in the US. In particular, the standard deviation of changes in total disposable income at the cohort level is above 3% in both countries. That of total consumption is between 0.6% and 0.95% less.

It may be argued that the differences in the consumption profiles for the two countries are due to the differences in the definitions used in the two surveys. For this reason, I next

---

[13] See Attanasio and Weber (1994). Groups do not need to be formed on the basis of age. In Attanasio and Banks (1997) that analysis is extended considering not only the variability across cohorts but also across regions.

[14] The sample mean $\bar{x}$ is distributed around the population mean as a random variable with variance given by σ^2 / N, where N is the cell size and σ^2 is the within cell variance. The latter can be estimated from the available micro data. These estimates can be used to correct our estimates of volatility.

focus on a narrower definition of consumption which excludes a number of items which might be recorded in different fashion in the two countries. In particular, I consider (log) expenditure on non-durables and services. This definition excludes from total consumption durables, housing, health and education expenditure. The other advantage of considering consumption of non durables and services, is that I avoid the issue of durability and the more complicated dynamics that is linked to durables. The main features of the two profiles, however, including the larger decline observed in the UK, are largely unaffected. In Table 5, the volatility of non durable consumption is considerably less than that of total consumption, especially in the UK when data up to 1986 are used.

An important possible explanation for the life cycle variation of consumption over the life cycle (and between the two countries considered), is the variation in needs linked to changes in family size and composition. To control for this possibility, I have deflated total household expenditure by the number of adult equivalents in the household. For such a purpose, I use the OECD adult equivalence scale.[15] The most evident result is that the life cycle profile of consumption looks much flatter now. In this sense, we can say that a large proportion of the variability of consumption over the life cycle is accounted for by changes in needs. This result is perhaps not surprising if one considers that the life cycle profile of the number of adult equivalents (or of family size) is also 'hump-shaped'. It may be argued that changes in needs are, to a large extent, predictable. In terms of the measure of volatility in Table 5, it is greatly reduced for the US, while is slightly increased for the UK.

While the profile for non durable consumption per adult equivalent is quite flat in the first part of the life cycle, a marked decline is still noticeable in the last part. It seems that the decline corresponds roughly to the time of retirement. In the UK, where retirement is much more synchronised than in the US, the decline is much more rapid. The fact that per adult equivalent consumption declines with retirement suggests that this might be due to a link between labour market status and consumption. A possibility, for instance, is that some components of consumption are linked to labour market participation. More generally, it is possible that consumption and leisure are non-separable and, therefore, need to be analysed jointly.

Finally, it is of some interest to consider the life cycle profile of expenditure on durables. The life cycle profiles for durables are plotted in Figure 6. Consistently with the findings in aggregate time series data, the life cycle profiles for durable expenditure are much more volatile than those for non durables and services. The measure in Table 5 for durables is 5 times as large as that of total consumption for the US and almost 4 times as large for the UK.[16]

Several variables are likely to be important determinants of, or determined jointly with consumption. I have already stressed the important role which is likely to be played by demographics and retirement behaviour in shaping the life cycle profiles of consumption. Similar considerations can be made for other labour supply variables such as the participation rate of females to the labour market and the total number of hours of work.

---

[15] No adult equivalence scale is perfect. Different alternatives, however, do not make much difference for the point I want to make here. The OECD scale gives weight 1 to the first adult, 0.67 to the following adults and weight 0.43 to each child below 19.

[16] Because durable expenditure can be zero at the individual level I do not compute the average of the log. Therefore, the deviations from the life cycle profiles are not percentage deviations, but are measured in constant dollars. Because of this, in Table 5, in the row corresponding to durables I report the coefficient of variation, rather than the standard deviation. For comparison, I adopt the same procedure for non durable consumption, in the following row.

A characterisation of the life cycle patterns of these variables and their differences between the UK and the US would go beyond the scope of this section.[17] However, it is important to stress that, as I argue in Section 5 and 6 below, one cannot test any model of consumption without controlling for these factors, that, for the most part, can only be analysed using household level data.

In Table 5, I only report the variability of the various components of consumption and of disposable income. As with the aggregate time series data, it would be interesting to characterise the autocorrelation properties of these variables and their covariances. This analysis could be quite informative about the plausibility of alternative structural models.[18]

One of the implications of the textbook version of the life cycle model I discuss in section 3, is that consumption and *current* income should not be related. And yet, comparing Figure 3 and 4, one cannot help noticing the similarity in the shape of the two life cycle profiles. This similarity was interpreted as a failure of the life cycle model by Thurow (1969) and reinterpreted in terms of non separability of consumption and leisure by Heckman (1974).

To pursue this issue, in Figure 7, I plot the life cycle profile of (log) disposable income and non durable consumption for four education groups in the US defined on the basis of the educational attainment of the household head: high school dropouts, high school graduates, some college and college graduates. An interesting feature of this figure is that the differences across groups in the shape of the income profiles are mirrored in differences in the consumption profiles. In particular, notice that both income and consumption profile of better educated individuals present a more pronounced hump; not only are their income and consumption higher, but the profiles are also much steeper in the first part of the life cycle. These differences where interpreted within a life cycle model by Ghez and Becker (1975), but have interpreted as a failure of the model by Carroll and Summers (1991) in an influential paper. An interesting question, addressed below, is whether a version of the life cycle model I discuss could generate these profile and account for the differences across education groups.

3. The Life Cycle Model

In this and in the next few sections, I will sketch what I think is the most important model of intertemporal consumption behaviour: the life cycle- permanent income model. In doing so, I take a fairly wide definition of the model: I consider a very general framework in which consumption (and saving) decisions are taken as a part of an intertemporal decision process. This general definition includes both the initial formulations of the life cycle and permanent income models and more recent and sophisticated developments, such as the precautionary saving model or the bequest motive. While the emphasis given to various aspects of the problem is different in the various incarnations of the general model I will consider, they have in common the hypothesis that consumption decisions are taken by a decision unit that maximises utility over time. The various versions of the model will then differ for their assumptions about optimisation horizon, uncertainty, curvature of the

---

[17] Interest readers can find the life cycle profiles for several variables in Attanasio (1994), Banks and Blundell (1994) and Attanasio and Banks (1997).

[18] MaCurdy (1983) and Abowd and Card (1989) perform analyses of these kinds for earnings and hours of work and use the results to assess the plausibility of different structural model. No similar analysis exists for consumption and/or its components.

utility function, assumptions about separability and so on. Which of these various versions is the most relevant is in part a matter of taste and, above all, an empirical matter.

3.1 The simple textbook model

The main attractiveness of the life cycle - permanent income model, developed during the 1950s in a number of seminal contributions[19], is the fact that consumption decisions are treated as part of an intertemporal allocation problem. The allocation of consumption over time is treated in a fashion similar to the allocation of total expenditure among different commodities in demand analysis. The model recognises, therefore, that intertemporal prices and the total amount of resources available to an individual are bound to be important determinants of consumption. This approach immediately gives the study of consumption solid microfoundations and constitutes a discontinuous jump with respect to the Keynesian consumption function which assumed consumption to be a simple function of current disposable income.

The main difference between the life cycle and the permanent income model in their original formulation lies in the time horizon considered. The life cycle model is, almost by definition, a finite horizon model, while in the permanent income model the horizon is infinite. In both cases, however, consumers decide how much to consume keeping in mind their future prospects. If no uncertainty is introduced in the model, its predictions are quite straightforward: concavity of the utility function implies a desire to smooth consumption over time; the main motivation for saving is to smooth out fluctuations in income; consumption increases with current income only if that increase is a permanent one. In the case of the life cycle model, the explicit consideration of retirement, that is a period in which income declines considerably, generates the main motivation for saving: households accumulate wealth to provide for their consumption during retirement.

An interesting implication of the life cycle model in its simplest incarnation is the way in which aggregate saving is generated. It is quite obvious that in a stationary life cycle economy with no growth aggregate saving is zero: the younger generations will be accumulating wealth, while the older ones will be decumulating it. Aggregate saving, however, can be generated in the presence of growth. If the amount of resources available over the life cycle to younger generations is larger than that available to older ones, it is possible that the amount accumulated at a point in time exceeds the amount that is decumulated. This introduces a relationship between aggregate saving and growth that Modigliani has stressed in several studies. It should be stressed, however, that such a relationship depends on a number of factors including the life cycle profile, the way in which growth is generated and who benefits from it and so on.

The life cycle- permanent income models were developed to provide an answer to several needs. First, by framing consumption decision within an intertemporal problem, immediately introduces dynamics into the picture. This gives the possibility of fitting some of the empirical facts that seemed at odd with the Keynesian consumption function,[20] such as the difference between average and marginal propensity to consume in the short and long run. In addition, the introduction of dynamics is obtained in a theoretically consistent fashion which is appealing to economists. The model gives an obvious

---

[19] Modigliani and Brumberg (1954) contains the first formulation of the life cycle model. The permanent income model was sketched in Milton Friedman's 1957 volume (Friedman, 1957).
[20] Carroll and Kimball (1996) provide quotes from Keynes' General Theory in which he suggested a concave consumption function.

explanation of the smoothness of consumption relative to disposable income linked to some well defined preference parameter (the concavity of the utility function).

Obviously, the first empirical applications of the model were quite different from the studies of the last 20 years, mainly for the much more sophisticated treatment of uncertainty which I discuss below.[21] The model, however, seemed to score a number of empirical successes. I have already mentioned the fact that the model accounts for differences between short run and long run responses of aggregate consumption to disposable income (or other variables). More generally, it was clear that the model was able to generate very rich dynamic patterns for aggregate consumption and its response to disposable income. It could also explain the relationship between consumption and wealth and provide a rationale for the relationship between wealth- income ratios and growth. Indeed, as Modigliani has pointed out, the simplest version of the model is capable to generate an aggregate wealth to income ratio of 5 which is close to what this number is for the US. Furthermore, the model also seemed able to explain some of the regularities observed in cross sectional data. Just to mention one, Friedman showed how the permanent income model can explain the fact that black households seem to save, at each level of income, a larger fraction of their income than white households.[22]

3.2 Quadratic preferences, certainty equivalence and the permanent income model.

One of the problems with the life cycle- permanent income model is that the dynamic problems that consumers are assumed to solve can be quite complex. As a consequence, if one considers an uncertain environment, unless strong assumptions about the nature of uncertainty and preferences are made it is not possible to obtain a closed form solution for consumption. A popular parametrization of the model which can yield an analytical solution is that of intertemporally separable and quadratic preferences. Indeed, a large part of the profession has come to identify the 'permanent income model' with such a parametrization of preferences with the additional assumptions of infinite horizon, constant interest rates and stochastic labour income.

The analysis of the model is greatly simplified by the linearity of the marginal utility of consumption. This, and the fact that the only uncertainty comes from labour income, allows the derivation of an analytical solution for consumption which depends only on the first moment of future labour income. In particular, under the assumptions listed in the previous paragraph, consumption at time t can be expressed as a simple function of 'permanent income':

(1) $$C_t = kY_t^p$$

where k = 1 if the (fixed) interest rate equals the subjective discount factor and permanent income Y_t^p is defined as:

---

[21] Friedman (1957) essentially approximated permanent income with a distributed lag of current income. Modigliani and Ando (1963) stressed the role played by wealth (in addition to disposable income) in aggregate consumption equations. Both the Modigliani and Ando paper and Friedman's book contained interesting discussions of the aggregation problems that were absent, for a long time, from subsequent empirical studies.

[22] This fact is still true: data from the Consumer Expenditure Survey from 1980 to 1992, confirm that the saving rates of household headed by a black are systematically higher, for any interval of income, than those of household headed by a non black.

$$(2) \qquad Y_t^P = \lambda \left(\sum_{j=0}^{\infty} \frac{E_t[y_{t+j}]}{(1+r)^j} + A_t \right)$$

where r the fixed interest rate, $\lambda = \dfrac{r}{1+r}$, A_t is the value of current wealth and y is disposable labour income.

The main attraction of equations (1) and (2) is that they provide a straightforward relationship between the stochastic process that generates income and consumption. These relationship give rise to a number of testable implications that have been studied at length in the literature.

Flavin (1981) and Sargent (1978) were the first studies to exploit the fact that equations (1) and (2), together with the hypothesis that expectations about future labour income are rational, imply cross equation restrictions on the bivariate VAR representation of consumption and disposable income. Flavin (1981), in particular, estimated such a system using US time series data and rejected the restrictions implied by equations (1) and (2). Flavin finds some evidence of excess sensitivity of consumption to income.

Campbell (1987) proposes a slightly different interpretation of Flavin's results. From equations (1) and (2) it is possible to obtain the following expression for saving[23]:

$$(3) \qquad s_t = -\lambda \sum_{j=0}^{\infty} \frac{E_t y_{t+j} - y_t}{(1+r)^j}$$

The restrictions implied by equation (3) are the same as those implied by (1) and (2). The nice thing about equation (3), however, is its interpretation. The fact that consumers smooth consumption over time is reflected in equation (3) in the fact that saving anticipates expected decline in disposable income. It is for this reason that equation (3) has been dubbed as the 'saving for a rainy day' equation. Formally, the implication of (3) can be written as saying that actual saving should equal the best forecast of labour income declines. Consistently with Flavin's findings, Campbell (1987) rejects the implications of the model. Campbell, however, finds that the time series pattern of actual saving is not far from that implied by the model. He claims that excess sensitivity of consumption to income within this framework 'is more naturally interpreted as insufficient variability of saving than as a correlation between changes in consumption and lagged changes in income' (p.1272).

Related to the tests of excess sensitivity discussed above, and using the same framework, are those papers discussing the issue of 'excess smoothness' of consumption. Campbell and Deaton (1989) were the first to stress that, because equations (1) and (2) can be used to derive the relationship between changes in consumption and innovations to the process generating income, the relationship between the volatility of consumption (or permanent income) and that of current income depends on the stochastic properties of the process generating the latter. In particular, if labour income is difference stationary (rather than trend stationary), permanent income, and therefore consumption, will be more volatile than current income. Intuitively, this result follows the fact that if labour income is not stationary, current innovations are persistent and will therefore imply a permanent

---

[23] s on the left-hand side of equation (9) coincides with saving (i.e. income minus consumption), only when k in equation (7) is equal to 1. Otherwise, $s = y - c / k$.

revision to permanent income. Therefore the observation that consumption growth is less volatile than current disposable income growth contradicts the permanent income hypothesis.[24] This result is ironic as one of the original motivations for the development of the permanent income model was, indeed, the observation that consumption is smoother than income. The most problematic issue with this branch of the literature is the well known difficulty in distinguishing between trend stationary and difference stationary models.[25]

The version of the model with quadratic preferences has also been used to introduce further refinements to the model. Goodfriend (1992) and Pischke (1995), for instance, consider the implications of the lack of complete information on contemporaneous aggregate variables. Pischke, in particular, explains the excess sensitivity results typically obtained with aggregate data with this type of phenomena.[26]

In a recent paper, Blundell and Preston (1996) use the assumption of quadratic preferences to devise a clever way of decomposing transitory and permanent components of income shocks. The idea is quite simple: under the permanent income model, consumption should react to innovations to permanent income and not to transitory income. Consumption variability can therefore be used to isolate that part of the observed volatility of income which is to be attributed to permanent shocks.

3.3 The Euler equation approach.

Without the assumption of quadratic utility (and of uncertainty confined to the exogenous income process), one is left with the problem that it is not possible to derive an analytical solution for the *level* of consumption.

The most important theoretical development since the development of the life cycle-permanent income model is the rigorous treatment of uncertainty introduced in the late 1970s, after the rational expectations revolution in macroeconomics. In a famous paper, Hall (1978) used what is the main implication of the intertemporal optimisation problem faced by a generic consumer to derive empirically testable restrictions that have been at the centre of much of the empirical analysis in the last 20 years. The idea is simple and elegant: in a situation in which consumers maximise expected utility under uncertainty, they act so to keep the expected (discounted) marginal utility of consumption constant.

This condition is equivalent to the equalisation of the marginal rate of substitution to relative prices in consumer demand. The beauty of the approach lies in the way in which the difficulties associated with the presence of uncertainty are circumvented. The effect of future variables on consumption at a given point in time is summarised by the multiplier associated to the budget constraint: the marginal utility of wealth. This object is eliminated by considering the equations for two different periods and considering the optimal pattern for the evolution of the multipliers.

It is now time to introduce a bit of notation and formalise what said so far. Suppose a consumer maximises expected life time utility subject to an intertemporal budget constraint. She consumes a homogeneous consumption good C, receives labour income y

---

[24] For a clear discussions of these issues see chapters 3 and 4 in Deaton (1992).

[25] See, for instance, Christiano and Eichenbaum (1990)

[26] Deaton (1992) also discusses the possibility that the information set used by individual agents differs from that available to the econometrician.

and has the possibility of investing in N different assets A^i that pay a rate of return R_t^i at the end of period t. Both rates of returns and labour income are uncertain. This setup is formalised in the following equation:

$$\text{Max } E_t \sum_{j=0}^{T-t} \beta^j U(C_{t+j}, z_{t+j}, v_{t+j})$$

(4)

$$\text{s.t. } \sum_{i=1}^{N} A_{t+j+1}^i = \sum_{i=1}^{N} A_{t+j}^i (1 + R_{t+j}^i) + y_{t+j} - C_{t+j}$$

where I allow the instantaneous utility function U to depend on a vector of observable variables z, and an unobservable variable v. The operator E_t denotes expectations conditional on the information available at time t. I omit an index for the individual for notational simplicity. Implicit in equation (4) are a number of simplifying assumptions of various nature that will be relaxed in the following sections. It is useful to list some of them along with ways in which they can be rationalised.

(i) Equation (4) assumes that utility is separable over time. This is a strong assumption and rules out at least two important phenomena: habit formation and durable commodities. The marginal utility of consumption in any given time period does not depend on consumption expenditure in any other period. We also assume that expenditure coincides with consumption. One obvious possibility to rationalise this model without excluding the existence of durables is to assume that the instantaneous utility function is additively separable in non durables and in the services provided by durables. In this case, a term for expenditure on durables should be added to the intertemporal budget constraint.

(ii) Utility is derived from an homogeneous consumption good. The conditions under which intertemporal choices can be summarised by a single price index are seldom discussed. The two situations that are treatable are the absence of changes in relative prices, so that one can construct a Hicks composite commodity, or that preferences take the Gorman polar form. These issues are discussed below.

(iii) Labour income is exogenous. No labour supply choices are considered. One can reconcile a situation in which labour supply is endogenous with the model discussed above, assuming that the instantaneous utility function is additively separable in leisure and consumption. In this case, one should modify only the budget constraint of the problem (4) above.

(iv) The duration of life is certain. This assumption is easily relaxed to assume an uncertain life time. Davies (1981) has shown that this equivalent to assuming a discount factor β that varies with age as a consequence of a varying probability of survival. Utility at future ages is discounted not only because it accrues in the future but also because its accrual is uncertain. I will not discuss this issue any further, except when I discuss some of the issues relevant for the analysis of consumption based on numerical methods.

(v) The rate of return on assets does not depend on the net position on that asset or on the total level of wealth held by the consumer. The model, however, can easily accommodate a situation in which several assets are subject to various kinds of constraints, as long as there is at least one asset in which is possible to borrow and lend at the same rate.[27]

(vi) For simplicity, I have not considered explicitly the presence of inflation. Obviously the presence of (uncertain) absolute price changes is simply accommodated in the model above by changing appropriately the budget constraint and the definition of interest rates.

Under these assumptions it is possible to derive an extremely useful first order condition for the intertemporal maximisation problem described above. If we denote with λ_t the multiplier associated to the intertemporal budget constraint at time t, it can be shown that two of the first order conditions for the problem in (4) are:

(5) $$\frac{\partial U(C_t, z_t, v_t)}{\partial C_t} = \lambda_t$$

(6) $$\lambda_t = E_t[\lambda_{t+1} \beta (1 + R_{t+1}^i)] \quad , \quad i = 1, \ldots, m$$

Equation (6) holds for the m (m$\leq$N) assets for which is possible to borrow and lend at the same rate and for which the consumer is not at a corner.

Equation (5) states that the marginal utility of consumption is equal to the marginal utility of wealth at time t. The latter term summarises all the information about the future. Equation (6) is the Euler equation that corresponds to the problem in (4) and states that the intertemporal maximisation problem (4) implies that the discounted value of the marginal utility of wealth is kept constant over time. In other words, the marginal utility of wealth is a martingale.[28]

If we substitute equation (5) into equation (6) we obtain:

(7) $$U_{c_t} = E_t[U_{c_{t+1}} \beta (1 + R_{t+1}^i)] \quad , \quad i = 1, \ldots, m$$

where U_c is the derivative of the instantaneous utility function with respect to consumption. It is equation (7) that Hall (1978) used to derive his famous 'random walk' property of consumption. These can be derived for the *level* of consumption if utility is assumed to be quadratic, or, under some distributional assumptions, for its log if utility is isoelastic.

Notice that I have allowed instantaneous utility to depend on a number of observable variables (the vector z) for which I have not specified any property. Such variables could in principle be choice variables determined in the same intertemporal maximisation

---

[27] This is the condition under which the first order condition derived below holds. If the rate of return on a given asset changes with the net position in that asset in a continuous and differentiable fashion, that is if the intertemporal budget constraint is concave and does not present kinks, the first order condition derived below can be easily modified. More complicated is the situation in which there are discontinuities and kinks for all assets at some level of net worth (for instance zero).

[28] MaCurdy (1981) uses this framework to construct his 'λ-constant' labour supply function.

problem. Indeed, the variables z constitute the vehicle through which I introduce a number of factors such as the effect of demographics or of labour supply. As long as we control for their effect on the marginal utility of consumption (and treat them as endogenous at the estimation level), I do not need to model them explicitly.

It should be stressed that equation (7) is not a consumption function, but only an equilibrium relationship that can be used (and has been used) to estimate structural (behavioural) parameters and/or to test some of the implications of the model. The big advantage of equation (7) is the elimination of the term that represents the marginal utility of wealth and therefore the necessity of explicitly modelling the way in which the distribution of future variables influences consumption choices. The price of this simplification, however, is not a small one: we loose the ability to say anything about the levels of consumption. While it is true that given the level of current consumption, we can use equation (7) to forecast the expected level of future consumption, we do not know how consumption reacts to unexpected changes in the economic environment. These include changes to taxes or any other policy instruments that might affect consumption decisions. The elimination of the unobservable marginal utility of wealth is similar to the elimination by quasi differencing of fixed effects in econometrics. The problem here is that the 'fixed effect' that is differenced out is one of great importance.

Hall (1978) used quadratic utility to derive his famous 'random walk' result. After his paper, several studies developed and used the Euler equation approach with different assumptions about preferences. Hansen and Singleton (1982), for example, used isoelastic preferences to derive an expression whose parameters they estimated by GMM on aggregate time series data for consumption and several rates of return. In particular, equation (7) with isoelastic preferences (and neglecting the z and v variables) implies:

$$(8) \qquad E_t[\frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}}(1+R_{t+1}^i)\beta] = 1$$

The assumption of rational expectations and the fact that expectations in equation (8) are conditional to the information available at time t can be used to find valid instruments to identify the structural parameters of equation (8), which, once again, can be considered for several rates of returns. If we denote with q the dimension of the vector of parameters (in equation (8), q=2), with m the assets for which (8) holds and with k the number of instruments considered, equation (8) yields mk-q overidentifying restrictions that can be used to test the model. The results obtained by Hansen and Singleton on aggregate data indicate that the model is strongly rejected whenever several returns are considered simultaneously. On the other hand, when one asset is considered in isolation, the overidentifying restrictions are not violated, but the preference parameters are estimated at somewhat implausible values.

Hansen and Singleton (1983) considered a log-linear version of equation (8). If one assumes that the rate of returns and consumption growth are joint log-normal, from equation (8) one can derive the following expression.

$$(9) \qquad \Delta \log(C_{t+1}) = \frac{1}{\gamma}[\log(\beta) + \gamma^2 \, var_t(\Delta \log(C_{t+1})) + var_t(\log(1+R_{t+1}^i)) +$$

$$2Cov(\Delta \log(C_{t+1}), \log(1+R_{t+1}^i))] + \frac{1}{\gamma}\log(1+R_{t+1}^i) + \varepsilon_{t+1}$$

where ε_{t+1} is an expectational error uncorrelated with all the information available at time t. The advantages of equation (9) are that is (log) linear and some of its coefficients, as the one on the interest rate, have a natural and interesting interpretation. When the utility function in equation (4) is isoelastic, its curvature parameter γ plays a double role. On the one hand it is equal to the coefficient of relative risk aversion and therefore summarises consumer's attitude towards risk. On the other, its reciprocal, is equal to the elasticity of intertemporal substitution and therefore measures as consumption growth changes when the relative price of present and future consumption changes.[29]

The linearity of equation (9) allowed a number of researchers to estimate it by linear instrumental variables methods. Valid instruments include, as before, any information available to consumers at time t-1. Hall (1988), however, noted that if the frequency with which consumption decisions are taken is higher than the frequency of observations, the residuals of equations will not be uncorrelated over time, as implied by the assumption of rational expectations, but have, under certain assumptions, an MA(1) structure with a positive coefficient. If this is the case, valid instruments will be any variable dated t-2 or earlier. The choice of lagged two instruments has become the common practice for the studies who estimate versions of equation (9) using aggregate time series data.[30]

Most empirical applications that used a version of equation (9), including Hansen and Singleton (1983), typically assume that the conditional second moments that appear on its right hand side are constant over time so that they can be estimated as an intercept term. I discuss the role of the conditional variances in section 5. For the time being it will suffice to notice that assumption of constant second moments can be relaxed to the weaker one that the innovations to the second moments in equation (9) are uncorrelated with the instruments used in estimating the equation.

Having found the instruments that allow the (over) identification of the parameters of equation (9), it is possible to estimate its parameters and test the model. Alternatively, it is possible to test for the significance of additional variables that might be of particular interest and that, according to the simplest versions of the model, should not appear in the equation. A test which has received considerable attention is the addition of the expected value of current labour income growth. One of the main examples of this approach, is a series of papers by Campbell and Mankiw (1989, 1991) who have estimated on aggregate time series data an equation like the following:

$$(10) \qquad \Delta \log(C_{t+1}) = \alpha + \sigma \, r_{t+1} + \lambda \, \Delta \log(y_{t+1}) + \varepsilon_{t+1}$$

where y is labour income and r an interest rate. Both of the variables on the right hand side are instrumented. According to the model the parameter λ should be zero and the residual term should be orthogonal to the information available at time t.

Equations (8) and (9) illustrate one of the main advantages of the Euler equation approach. Even if it is not possible to obtain a closed form solution for consumption, it is possible to consider equilibrium relationships that can be used to estimate structural parameters. While these, as I discuss below, are not sufficient to answer many important policy

---

[29] For a discussion of the interpretation of such a parameter and the links between intertemporal substitution and risk aversion see Hall (1988), Attanasio and Weber (1989), Epstein and Zin (1991, 1993).

[30] Carroll et al. (1994) argue for the explicit consideration of the MA(1) process that would allow the use of the orthogonality conditions with instruments dated t-1 and therefore yield more precise estimates and more powerful tests.

questions, they constitute a basic ingredient of any answer. Furthermore, the orthogonality restrictions implied by the equations (and the assumption of rational expectations) can be used to test the validity of the model. Finally, the model can encompass a number of features that make it quite realistic without loosing its empirical tractability. In particular, it is possible to consider the effects that variables such as durables, labour supply, children and so on have on the marginal utility of non durable consumption without having to model explicitly these variables. Having said this, however, one should also stress that the empirical implementation of the Euler equation is not without problems. First of all, data requirements can be quite formidable. Furthermore, a number of subtle econometric problems needs to be considered. I discuss these problems and some of the available empirical evidence in Section 5.

3.4 Precautionary motives for saving.

The assumption of quadratic preferences, which makes it possible to derive a closed form solution for consumption, is not very appealing. In addition to various shortcomings of quadratic preferences, many would find certainty equivalence, i.e. what makes the model easy to handle, questionable. It is therefore not surprising that a large literature has developed in the attempt to go beyond quadratic preferences.

The first paper to consider explicitly the effects of non-linear marginal utility is Dreze and Modigliani (1972). While the Dreze and Modigliani contribution focuses on a two period problem, it contains many of the insights of the precautionary saving literature, such as the fact that the importance of the precautionary motive for saving depends on the third derivative of the utility function.

Kimball (1990) discusses within a rigorous framework the conditions under which one can expect to observe precautionary savings. In particular, he proves that precautionary saving will occur when the utility function exhibits 'prudence', that is when the 3rd derivative of the utility function is positive. Under a CRRA utility function this will always be the case. To see this it is sufficient to consider the equation (6) above, that is the log-linearization of the Euler equation under the assumption of log-normally distributed random variables.[31] Notice that, keeping the other variables constant, an increase in the conditional variance of consumption increases, at time t, the expected rate of growth of consumption between t and t+1. This can be achieved by decreasing consumption at t, that is increasing saving at t. It is this that is usually referred to as precautionary motive for saving. As Browning and Lusardi (1996) stress, the fact that the effect that the conditional variance has on the rate of growth of consumption is scaled by the coefficient of relative risk aversion is an artefact of the CRRA utility function in which a single parameter controls both risk aversion and prudence (as well as the elasticity of intertemporal substitution). More generally, Carroll and Kimball (1996) have proved the concavity of the consumption function under precautionary saving.

It should be noticed that the variance of consumption is not an exogenous variable. While it is likely to depend on the overall uncertainty facing the consumer, it is not obvious how it reacts to the arrival of new information or how it is related to uncertainty in income and interest rates. To evaluate this relationship it would be necessary to solve for the level of

---

[31] Alternatively, an equation similar to (6) can be obtained by a Taylor expansion of the marginal rate of substitution, as in Dynan (1994). Banks, Blundell and Brugiavini (1997) use an approximation of the Euler equation developed by Blundell and Stoker (1996).

consumption and establish how consumption at t+1 reacts to changes in the economic environment which, under CRRA preferences, is not possible.

Caballero (1990, 1991) using the exponential utility function as a parametrization of within period utility obtains, with some additional assumptions, a closed form solution for consumption. In an interesting example, he expresses consumption as a function of permanent income (as in the case of certainty equivalence), minus a term which summarises the effect of the precautionary motive. He then goes on to evaluate the effect that precautionary saving is likely to have in reality.

While the results obtained with the exponential utility give useful insights about the potential importance of the precautionary motive, such a parametrization of the utility function is not exempt from criticism. It is therefore important to evaluate the importance of precautionary saving using more general functional forms for the utility function. Unfortunately, when one uses different utility functions is not possible to obtain a closed form solution for consumption. It is therefore necessary to use numerical methods to obtain solutions and additional assumptions about the nature of the problem faced by consumers. This is the approach taken in a number of studies, such as those of Skinner (1988), Zeldes (1989b), Deaton (1991), Carroll (1994) and Hubbard, Skinner and Zeldes (1993, 1995). A disturbing feature of these studies is that the quantitative importance of the precautionary motive depends crucially on the properties of the distribution of the income process in the left tail. Bounding away the income process from zero (or from arbitrarily small realisations) greatly reduces the precautionary motive. While the reason for this is clear (consumers will want to insure themselves against disastrous events), the realism of such a mechanism is questionable. It might be worth investigating how the precautionary motive would change in the presence of insurance mechanisms other than self insurance (such as a safety net supplied either by society or by family and relations).

Carroll (1994, 1997a) has strongly advocated the precautionary saving motive (or 'buffer stock saving') as an explanation of most empirical puzzles in consumption, including the tracking of expected consumption and income. I discuss the empirical findings on the precautionary motive and more generally on the Euler equation in Section 5. From a theoretical point of view, however, it should be stressed that while the precautionary motive is likely to be important for impatient consumers and especially in the early part of the life cycle, it is also clear that if retirement savings become, at some point during the life cycle, important, they can be used to buffer unexpected fluctuations to income. This effect, however, can be limited if most retirement wealth is held in the form of claims to future benefit or in other illiquid assets (such as special retirement accounts, housing, social security etc.). The relevance of the precautionary saving motive is ultimately an empirical matter, as it depends on preference parameters, on the features of the individual income process and on the availability of safety nets and insurance mechanisms.

An attractive possibility to test the relevance of the precautionary motive for saving is to relate observed saving behaviour to perceived uncertainty. Jappelli, Guiso and Terlizzese (1993) is the only paper that uses direct observations on the perceived variance of the income process to assess the importance of precautionary savings. They use a unique data set (the Bank of Italy Survey of Household Income and Wealth) in which individuals are asked not only about their income expectations, but about the complete probability distribution of future income.[32] This allows the authors to compute individual variances, which they then relate to individual savings, finding some mild evidence of precautionary savings. One obvious problem with this approach is the possibility that the individuals

---

[32] Individuals are asked to divide 100 probability points over several intervals of income growth.

who have selected themselves into riskier occupation are less risk averse and, maybe, less 'prudent'. A similar argument was used by Skinner (1988) to justify his finding that self-employed individuals seem to save less than the average.

More recently, Jappelli, Guiso and Terlizzese (1996) have studied the implications of the precautionary motive for portfolio composition. They show that, in the presence of different and unrelated sources of risk, individuals with higher and undiversifiable earning uncertainty will tend to invest in relatively 'safer' portfolios. Jappelli et al. (1996) use the same data used in their previous study and find that this implication seems to be borne out.

3.5 Borrowing restrictions

Related to the issue of precautionary saving is that of the presence of borrowing constraints. The standard model sketched above assumes that individuals can borrow against future labour income to finance current consumption at the same rate at which they can lend. Of course, if this is not the case, the basic model has to be amended in that the maximisation problem in equation (4) has to take into account the additional constraint. It has now become customary to interpret evidence of 'excess sensitivity' of consumption growth to labour income as an indication of 'liquidity constraints', by which it is usually meant the presence of some imperfection in financial markets that prevents people from borrowing. However, there is no reason to believe that liquidity constrained individuals consume their disposable income. Only when the constraints are actually binding will this occur. Therefore, excess sensitivity tests and, more generally, Euler equations are not the best way to identify the presence of liquidity constraints. Furthermore, as we discuss below, there are several reasons why predicted income and consumption might be related that have nothing to do with the presence of liquidity constraints.

With these considerations, I do not want to dismiss the possibility of liquidity constraints as unimportant or unrealistic.[33] Especially for some groups in the populations, they might be quite relevant and have an important effect on aggregate consumption. Early contributions to the literature on the policy implications of the permanent income - life cycle models, such as the papers by Flemming (1973) and Tobin and Dolde (1975), were quite aware of the importance of liquidity constraints. It is crucial, however, that the presence of liquidity constraints is incorporated in the optimising framework sketched above.

The definition of liquidity constraints I use in what follows appeal to partial equilibrium considerations: the interest rate schedule is taken as given. General equilibrium considerations, however, even though they are rarely made, can be quite important. In a world of identical consumers, the possibility of smoothing consumption over time will be constrained by the technology available to transfer resources over time. In equilibrium, interest rates and asset prices will adjust depending on the demand for borrowing (saving) and on the available technology. In such a situation, nobody is liquidity constrained at the current interest rates. With heterogeneous consumers, it is possible that some will want to borrow and others will be saving. The equilibrium interest rates will then reflect these

---

[33] As Hayashi (1996) stresses, most specifications of preference implicitly imply a form of borrowing restrictions. If the marginal utility of consumption is infinite a zero consumption, consumers will not want to borrow more than the present discounted value of the minimum realisation of income, even though the probability of this event is very low. This is because they want to avoid the possibility of zero consumption even with a very small probability.

factors. While the focus of most of the considerations in this section is on the partial equilibrium effects, it is worth keeping in mind that aggregate fluctuations, for instance a recession, might have effects on asset prices that reduce the demand for loans relative to what would be observed under constant interest rates.

The first step in the integration of borrowing restrictions in the model above must be their exact definition. There are several possibilities. The first and most general alternative is to allow the interest rate paid on assets to depend on the net asset position. As a simple example of this alternative, consider the possibility of a difference between borrowing and lending rate. Such a wedge induces a kink in the intertemporal budget constraint. One can then consider Euler equations for individuals who are net borrower and net savers: the interest rate relevant for the two groups will be different, but the Euler equation still holds as an equality. For the individuals that will cluster at zero net assets, however, the Euler equation (with either interest rate) holds as an inequality.

It is also possible that the interest rate varies continuously with the quantity borrowed or saved. If that is the case, at any point in which the function is differentiable, it is still possible to write down an Euler equation which, relative to equation (8), contains an additional term referring to the derivative of the interest rate with respect to the asset position.[34]

$$(11) \qquad E_t[\frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}}((1+R_{t+1}^i)+\frac{\partial R_{t+1}^i}{\partial A_t^i})\beta] = 1$$

At those points of the intertemporal budget constraint where the interest rate changes discontinuously (such as zero), equation (11) will be replaced by an inequality.

An alternative to the consideration of interest rates varying with the amount borrowed (or saved) is the assumption that individuals face a limit to the amount they can borrow (which can be zero). Obviously this can be interpreted as a case of the previous situation, with the borrowing rate being infinite at the limit. Even in this case, however, there are several alternatives. It is possible, for instance, that the limit an individual can borrow is not fixed but a function of some variables which, in turn, can be endogenous.[35] Finally, it is possible to consider the existence of collateralizable loans.

When liquidity constraints take the form of a limit to borrowing, it is still possible to write the Euler equation for consumption, as long as the constraint is not binding. When it is binding, instead, the Euler equation will hold as an inequality, or as an equality with the addition of a slack variable (a Kuhn-Tucker multiplier). Equation (8) becomes:

$$(12) \qquad E_t[\frac{C_{t+1}^{-\gamma}}{C_t^{-\gamma}}(1+R_{t+1}^i)\beta] = 1+\mu_t$$

where μ_t is an unobservable Kuhn-Tucker multiplier associated to the borrowing restriction.

If instead of considering a homogeneous consumption good, one considers several commodities explicitly, assuming that none of them is durable or can be used as a

[34] A paper which contains this type of analysis is Pissarides (1978).
[35] Alessie et al. (1989) and Weber (199?) consider the possibility that the limit to borrowing is a function of earnings.

collateral, for each of these commodities it is possible to obtain an equation of the following form:

$$(12) \qquad \frac{\partial U()}{\partial q_t^i} = \lambda_t p_t^i + \mu_t$$

where λ_t is the marginal utility of wealth and p_t^i is the relative price of commodity i. Notice that μ_t appears in the first order conditions for all commodities, so that any two of these equations can be used to eliminate it. This implies that the intratemporal first order conditions hold regardless of the presence of liquidity. Meghir and Weber (1996) used this intuition to distinguish between liquidity constraints and intertemporal dependence in preferences. They consider three different non durable commodities and stress that the presence of dynamic effects in the Euler equation can be rationalised with intertemporal non separabilities only if the same dynamic effects are found in the intratemporal first order conditions. On the contrary, if one finds that the intratemporal conditions do not show any sign of dynamic effects while these appear in the Euler equations, one should interpret this evidence as a sign of binding liquidity constraints.

The possibility that some commodities, such as durables, can be used as collateral to relax the severity of borrowing restrictions has been examined in several studies. Because durables can be used as collateral, when liquidity constraints are binding, they become relatively more attractive than non-durables. This fact has implications for the within period allocation of resources between durables and non durable consumption. In the absence of liquidity constraints, this would depend only on the relative price and on their marginal rate of substitution (where the relevant price for durables would be their user cost). In the presence of liquidity constraints, however, an additional term has to be added to the first order condition to reflect the fact that it is possible to use durables to borrow against future resources. Therefore, as Chah, Ramey and Starr, (1995) Brugiavini and Weber (1994) and Alessie, Devereux and Weber (1997) have noted, the possibility of using durables as collateral and the presence of liquidity constraints distorts the intratemporal allocation of resources between durables and non durables.

At several points in the discussion so far I have stressed that even when liquidity constraints are present, if they are not binding, the Euler equation (8) will hold. Therefore, such an equation and the mis-specification tests conducted on it (such as tests of 'excess sensitivity') are likely to be a poor tool to identify the presence of borrowing restrictions. In addition to the power considerations just made, in what follows I also stress that evidence on excess sensitivity can often be interpreted as evidence of non separability between consumption and leisure.

The fact that the Euler equation (8) holds whenever the borrowing restrictions are not binding does not mean that these constraints have no effect on the level of consumption. Indeed, as discussed clearly by Hayashi (1987), the presence of potential liquidity constraints has the same effect as that of a shortening of the horizon relevant for current choices. Alternatively, one can interpret the presence of liquidity constraints (when they are not binding) as an increase in the discount factor.

From a policy perspective, liquidity constraints are important because of their effect on the level of consumption, rather than on its changes. In other words, what matters is how consumption reacts to unexpected changes in the economic environment (including policy changes). Euler equations are not informative about this.

Deaton (1991) provides one of the first analysis of the effect of liquidity constraints on the level of consumption. By solving the Euler equation numerically, he shows that the behaviour induced by the presence of liquidity constraints is similar to that associated with a precautionary motive for saving. As in the precautionary motive for saving, people will accumulate a buffer stock to avoid the possibility of needing a loan that they cannot obtain. Furthermore, the liquidity constraint will be binding only occasionally as during most periods, individuals avoid, by their optimal behaviour, to find themselves constrained.

For liquidity constraints (or precautionary saving) to be relevant for individual behaviour, it is necessary that the households concerned want to borrow, that is they face an increasing income path and are impatient enough to want to bring resources from the future to the present. Even in such a situation, however, it is not necessary that the Euler equation restrictions are violated. If one rules out the possibility of dying in debt and considers finite lives and a marginal utility of consumption that goes to infinity when consumption goes to zero, it is possible that 'liquidity constraints' are generated endogenously by the model. In particular, to avoid the possibility of having zero resources (and therefore zero consumption) in the last period of life, individuals will not want to borrow any amount in excess of what they can repay with probability one. Notice that in such a situation, consumption in any two periods satisfies the Euler equation.[36]

From a theoretical point of view, the considerations above indicate that a profitable research strategy is one which aims at characterising the response of consumption to various news when liquidity constraints are relevant (regardless of whether they are binding). Deaton (1991) constitutes a first important step in this direction. On a more specific level, the analysis in Hubbard, Skinner and Zeldes (1993, 1995) constitutes another good example of how a consistent theoretical model incorporating borrowing restrictions can be used and be informative about important policy issues.[37]

Several other interesting problems remain to be modelled and understood. If labour supply choices are endogenous, the presence of liquidity constraints might induce female labour force participation. These effects might be strengthened if the ability to borrow is linked to earnings.[38] All these issues are examples of the need to be able to use the model to make statements about consumption levels that I discuss again below.

Identifying the presence and the relevance of borrowing restrictions is not easy. One possibility is to use direct questions on the matter. Jappelli (1990) used this strategy and analysed the answers to some questions contained in the Survey of Consumer Finances in the US. The households interviewed in 1983 were asked whether they were denied credit or whether they did not applied because they felt they would be denied. The main problem with this research strategy is that very few household surveys contain questions similar to those analysed by Jappelli.

The presence of (binding) liquidity constraints has some direct implications for the demand for loans. Its analysis constitute therefore an interesting possibility for the identification of borrowing restriction and for evaluating their importance. If consumers are subject to binding borrowing restrictions they will be at a kink (or in a relatively steep section) of an intertemporal budget constraint. Their demand for loans, therefore, will not be much affected by changes in the interest rate. On the contrary, if consumers are not at a corner

---

[36] This point was also noted in Hayashi (1987).

[37] Jappelli and Pagano (1994) characterise the link between liquidity constraints, saving and growth.

[38] See O'Brien and Hawley (1986).

(or in a relatively flat part of the budget constraint) their demand for loans will be elastic to the interest rate. The effects of maturity, however, should be opposite. A consumer who is not affected by borrowing restrictions will be indifferent to changes in maturity. On the other hand, a constrained consumer, for whom an increase in maturity will effectively increase his ability to borrow by reducing the size of the payments to maturity, will increase its demand for loan. Juster and Shay (1964) were the first to use this intuition to analyse semi-experimental data. They asked a sample of consumers whether they would finance the hypothetical purchase of an automobile when faced with different packages of interest rates and maturity. The packages of interest rate and maturity were randomised to the individuals in the sample so to enable the estimation of the interest rate and maturity elasticity of the demand for loans. In a recent paper, Attanasio and Goldberg (1996) have performed a similar exercise but on a sample of households who actually purchased (and occasionally financed) automobiles. The main advantage of the Attanasio and Goldberg exercise is that their data refer to actual choices. The main disadvantage relative to the work of Juster and Shay (1964) is that for the households that decided not to finance their car purchases neither interest rates nor maturity are observed. This poses a number of econometric problems that are similar to those that labour economists deal with when analysing participation choices and labour supply.

3.6 Taking into account demographics, labour supply and unobserved heterogeneity.

In the previous sections, for the sake of expositional simplicity, I have neglected the vector of observable variables z and the unobserved component v in equation (4). Their exclusion, however, especially if one wants to estimate or test the model, would be unrealistic. It is quite obvious, for instance, that family size and composition affect the marginal utility of a given amount of consumption expenditure. It is also likely that labour supply behaviour, and in particular female labour supply, will also affect the utility derived from expenditure. Working often implies bearing a certain number of costs that range from transport, to eating out, to clothing. Furthermore, if both spouses work, a number of services that would be produced at home by the partner not participating in the labour force would have to be purchased on the market and would be counted as consumption. More generally, it is plausible that consumption and leisure are not separable in the utility function and that consumption, saving and labour supply choices are taken simultaneously.

Given these considerations, it is important, especially if one wants to bring the model to the data, to consider a specification that allows for these factors. While it is true, as shown in Section 2, that consumption life cycle profiles often mirror income profiles, it is not obvious that one should interpret this as evidence of the empirical failure of the model without considering explicitly the possibility of important demographic effects and that consumption and leisure are not separable. These are not new arguments and I discuss them at length in Section 5, where I interpret the available evidence. The point of this section, however, is to stress that the consideration of these factor is not only important, but also relatively simple. If one does not want to model explicitly these variables, controlling for them does not jeopardise the empirical tractability of the model illustrated by equations (8) and (9).

I re-write a slight modification of equation (4) for convenience.

(4')
$$\text{Max} \, E_t \sum_{j=0}^{T-t} \beta^j U(C_{t+j}^{nd}, z_{t+j}, v_{t+j})$$

where the superscript nd indicates that I am now modelling explicitly only non durable consumption. As before, the vector of observable variables z indicates variables that are relevant for the intertemporal optimisation problem. The variable v is unobservable to the econometrician. The latter variable is sometimes referred to as 'taste shift'. More generally, it represents all the unobservable factors that affect consumption choices and that we do not model or control for. The variables included in the vector z may range from demographic variables, to labour supply variable, to other components of consumption (such as services from durables). Notice that these variables may be either exogenous to the choice problem (age is a good example), or chosen simultaneously with non durable consumption (such as labour supply). Even if one does not wish to model these latter variables explicitly, it is possible to identify, using equation (4'), *conditional* preferences and the associated parameters under very mild conditions.[39]

If the z variables in equation (4') enter the utility function in an additive separable fashion there is no need to consider them. If, on the other hand, they affect the marginal utility of non durable consumption, they will enter the corresponding Euler equation. As an example, let us consider the following parametrization:

$$(13) \qquad U(C_t^{nd}, z_t, v_t) = \frac{(C_t^{nd})^{1-\gamma}}{1-\gamma} \exp(\theta' z_t + v_t)$$

A possible interpretation of equation (13) is that the discount factor varies with variables Z and v. Such an interpretation is particularly attractive for demographic variables. In this case, the vector of parameters θ implicitly represents an equivalence scale.

Under the parametrization used in equation (13), the Euler equation for non durable consumption will be:

$$(14) \qquad E_t \left[\left(\frac{C_{t+1}^{nd}}{C_t^{nd}} \right)^{-\gamma} (1 + R_{t+1}^i) \beta e^{\theta'(z_{t+1} - z_t) + v_{t+1} - v_t} \right] = 1$$

It should be stressed that equation (14) holds even if some of the variables included in z are endogenous choice variable, regardless of the way in which they are determined. Modelling labour supply or durable consumption, for example, can be quite difficult because of corner solutions, intertemporal (non)-separability and transaction costs. These problems do not need to be tackled explicitly and equation (14) can be considered as an equilibrium relationship that holds at the optimal values of the endogenous z, regardless of how these are obtained.

Equation (14) can be linearised in the same fashion as equation (8) to obtain a log-linear expression similar to equation (9). Neglecting the second moments, that I incorporate into the constant of the equation, we have:

$$(15) \qquad \Delta \log(C_{t+1}^{nd}) = constant + \sigma R_{t+1}^i + \theta' \Delta Z_{t+1} + \Delta v_{t+1} + \varepsilon_{t+1}$$

Notice that because of the presence of the term representing unobserved heterogeneity and taste shocks, the residual of equation has now two components, one which represents an expectational error and, by the assumption of rational expectations, is likely to be

---

[39] See the discussion in Browning and Meghir (1991) on demand systems *conditional* on labour supply.

uncorrelated over time, and another that has an MA(1) structure if taste shocks are i.i.d. over time.

If the unobserved heterogeneity term is term v has a time invariant component which is individual specific, it will be eliminated in the first differences. On the other hand, it is possible that such a term is persistent (rather than white noise) but not fixed and/or that the 'pure' discount factor β is individual specific. In the former case the term Δv in equation (15) would have a structure more complex than a simple MA(1), while in the latter, equation (15) would have a fixed effect. In section 5, I discuss these econometric problems.

The particular parametrization considered in equation (13) is just an example. More complex structures may be and have been considered. It is possible, for instance, to allow some of the variables in the vector z to affect the curvature of the utility function as well as the rate at which utility is discounted. Browning, Blundell and Meghir (1994) and Attanasio and Browning (1995) have used this approach and found significant deviations of estimated differences from the simple isoelastic case. The issue is particularly important in that preferences of this kind can allow for systematic differences in the elasticity of intertemporal substitution across consumers.

Attanasio and Browning (1995) also stress the fact that to estimate an Euler equation and its parameters, it is not necessary to specify explicitly the within period utility function. It is possible and analytically convenient to start from a flexible specification for the marginal utility of consumption and, if needed, obtain the corresponding utility function by integration.[40]

3.7 Bequest motives

According to the simplest version of the life cycle model sketched at the beginning of the section, the main motivation for individual saving is to provide resources during the last part of the life cycle, when, following retirement, income is low. In such a situation aggregate wealth can be generated by and associated with productivity growth, if the generations that save (the young) are relatively wealthier than those that dis-save (the old). If the bequest motive is operative, on the other hand, the mechanism through which aggregate wealth is accumulated is quite different. A lively exchange on whether a large or small proportion of aggregate wealth in the US is accounted for by bequeathed wealth or retirement saving wealth developed in the 1980s between Kotlikoff and Summers (1981, 1989) on one side and Modigliani (1989) on the other. The result of that debate remained somewhat ambiguous, as the answer seems to depend mainly on whether one considers the interests earned on bequeathed wealth as originating from bequests or not.

A bequest motive can be simply added to the basic model (1) by considering a term which is a function of the bequest left. As it is obvious such a term does not affect the Euler equation for consumption in subsequent periods. It will, however, affect the level of consumption (and saving).

One of the implications of the simplest version of the life cycle model is that wealth is decumulated in the last part of the life cycle. While the rate at which wealth is decumulated depends on the parameters of the model and in particular about beliefs about longevity, the result that wealth should decline seems to be robust. The evidence on this

---

[40] If one follows such a procedure, one should check the integrability conditions.

point is mixed, in that several studies do not find strong evidence of decumulation of wealth by the elderly.[41]

When bequests motives are operative, from a theoretical point of view, the wealth age profile can take different shapes in the last part of the life cycle. In an important paper Hurd (1989) has characterised several of these profiles. Hurd also showed that for several realistic sets of parameters, the wealth age profile under bequests is also declining in the last part of the life cycle. From an empirical point of view, Hurd (1989) stresses the importance of conditioning on the labour force status of the individuals in the sample. In particular, whether individuals are retired or not seems to be crucial for the decumulation of their wealth.

4. Aggregation issues

The models considered in Section 3 refer to the dynamic optimisation problem faced by an individual consumer (or household). The early contributors to the life cycle - permanent income model were quite aware of the aggregation issues involved with the empirical implementation of the theory of intertemporal optimisation. These issues, however, were largely ignored by the literature of the late seventies and early eighties which focused mainly on the rigorous introduction of uncertainty in the model. Aggregation problems, however, cannot be ignored when the model is tested and when is estimated to evaluate structural parameters. In this section, I consider briefly two problems: the aggregation across consumers and that across commodities.

4.1 Aggregation across consumers

The Euler equations for (non durable) consumption that can be derived from the problem in equation (1) are, for most specifications of preferences, non linear. As they refer to individual households, their aggregation is problematic. A number papers have shown that the dynamics of aggregate consumption implied by the Euler equation for individual consumption cannot be described simply by the first moments of the cross- sectional distribution of consumption. Attanasio and Weber(1993), Blundell, Pashardes and Weber (1993), Blundell, Browning and Meghir (1994), have shown that higher order moments can play an important role. Attanasio and Weber (1993), in particular, have stressed that the use of aggregate data to estimate an Euler equation is equivalent to omit high moments of the cross sectional distribution of consumption and might cause systematic biases in the estimation of structural parameters and lead to rejections of the over-identifying restrictions.[42] These contributions will be discussed when I evaluate the empirical evidence, at this point I only wanted to stress that exact aggregation is in general impossible when considering the Euler equation for consumption. This implies that aggregate data cannot be used to estimate structural preference parameters and/or to test the model.

Individual panel data on consumption might be very difficult to obtain. Partly because of this, many of the papers I discuss in Section 5, use the average cohort data that I have used in Section 2 to describe the individual data. The use of grouped data allows one to use time series of repeated cross section to study dynamic models. In this sense, average (or synthetic) cohort data are particularly useful to study life cycle models. It is important to

---

[41] Jappelli and Modigliani (1996) have forcefully argued that pension benefits should be considered as decumulation of pension wealth.
[42] Blundell, Pashardes and Weber (1993) make a similar point for a demand system. Similar results were obtained in an asset pricing framework by Constantinides and Duffie (1996).

stress, however, that synthetic cohort data are *aggregate* data; the difference relative to National Accounts data is that the aggregation process is controlled directly. As a cohort ages, the researcher can follow the evolution of the variables of interest, that range from consumption to income, to family size and composition, to labour supply behaviour. The presence of non linearities in the relevant equations (as long as they are linear in the parameters) does not constitute a problem, as one can compute the relevant non linear transformations *before* averaging the data. Unlike in the aggregate data, therefore, time series of the average of any non linear transformations of the variables of interest is readily available.

In addition to the standard aggregation problems caused by the non linearity of the Euler equations, there is another sense in which aggregation across consumers can be problematic. As discussed above, it is likely that demographic variables are likely to be important for the determination of consumption. If one does not want to make arbitrary assumptions about the effects that these demographic variables have on the utility function, it might be necessary to consider individual data explicitly.[43] In the aggregate, demographic variables move quite slowly so that, even neglecting non- linearities, their effect cannot be estimated precisely from those data. On the contrary, differences across cohorts, observed over different parts of their life cycle, can be profitably be exploited to identify these effects.

Finally, because repeated cross sections are much more common than long panels, one can estimate the structural parameters using a relatively long time period. The importance of a 'large T' in getting consistent estimates is discussed in section 5.

The use of average cohort data is not without problems. Probably the most important is that of the presence of measurement error and sample variation. Particular care has to be devoted to evaluating the quality of the data and to use the appropriate econometric techniques.

4.2 Aggregation across commodities

Another important aspect which is often neglected in the literature on the Euler equation is the aggregation of expenditure on different commodities. In principle, one should consider simultaneously the allocation of total expenditure across time periods and the allocation within each period across different commodities. If one had detailed enough data on expenditure on individual commodities and was willing to specify the form of the direct utility function, the problem could be addressed in a straightforward manner. One could consider the Euler equation defined in terms of the marginal utility of each individual commodity.

The issue, however, is to determine under what conditions one can consider the within period utility function defined in terms of total consumption and assume that the allocation of expenditure over time can be determined with the help of a single price index. One obvious and not particularly interesting answer is when relative price do not change so that it is possible to construct a Hicks composite commodity. Gorman (1959) was the first to provide a more interesting answer in that he derived the conditions that the utility function has to satisfy so that it is possible to consider a single commodity.[44] In particular,

---

[43] Estimation of the effects of demographic variables can be interpreted as the estimation of adult equivalent scales and is therefore important for a variety of reasons.

[44] See Deaton and Muellbauer (1980) for a clear discussion of these issues.

it is necessary that preferences take the generalised Gorman Polar form, that is, if X is the Kx1 vector of commodities and **p** the vector of corresponding prices, the indirect utility function has to take the form:

(16) $$V(X,\mathbf{p}) = F(X / b(\mathbf{p}), a(\mathbf{p}))$$

where b(p) and a(p) are functions homogeneous of degree 1 and degree 0 respectively and depend on within period preferences.

From the specification in equation (16) one can derive an Euler equation for total consumption expenditure: two-stage budgeting can be used to separate the intertemporal from the within period allocation. To implement equation (16) it is necessary, however, to estimate the price functions a(p) and b(p), which requires the specification and estimation of a within period demand system. In an important paper, Blundell, Browning and Meghir (1994) estimate a demand system and use the resulting price indexes to estimate the Euler equation for total consumption expenditure. A remarkable result they obtain is that a Stone price index (which does not require the knowledge of preference parameters) constitutes a good approximation to the price indexes in equation (16). This result is important because points to a simple way to implement Euler equations for non durable consumption without estimating the entire demand system. The results in Blundell, Browning and Meghir are generalised to a demand system with quadratic Engel curves by Banks, Blundell and Preston (1994). Attanasio and Weber (1995) with a different parametrization of the indirect utility function find a statistically significant role for the zero-homogeneous price index a(p) on US data.

If one is willing to specify a utility function defined over several commodities, one can derive an Euler equation for each of these commodities. Under the assumption of additive separability, each of these Euler equation depends only on expenditure on that commodity, nominal interest rates, changes in the commodity- specific price index and on the relevant controls. On the other hand, when additive separability does not hold, the individual commodity Euler equation depends also on the consumption of other commodities. This is an important point in that creates problems for the use of data sources which contain information only on some components of total expenditure, such as the US PSID which gives only information on food consumption. Attanasio and Weber (1995) show that the consideration of food consumption in isolation can yield very misleading results.

5. Econometric Issues and Empirical Evidence

I now move to consider the empirical evidence on the life cycle- permanent income model. In the process of doing so, I also discuss some of the econometric problems relevant for the analysis. As I want to have a rigorous treatment of uncertainty, I only consider the contributions that followed Hall (1978) paper. I start with a discussion of studies based on aggregate data and then discuss some based on individual data. As I focus mainly on studies based on Euler equations, I do not discuss a large empirical literature, which has estimated flexible error correction models for aggregate consumption that I have mentioned in Section 2.

5.1 Aggregate time series studies

Because of the aggregation problems discussed in Section 4.1, it is my opinion that aggregate time series data cannot be used to estimate structural preference parameters and

test the life-cycle model. Indeed, some papers, such as Attanasio and Weber (1993, 1995), have shown that aggregation issues can easily explain some of the rejections of the model found in the literature. Nonetheless, the papers that have tested the life cycle - permanent income model using aggregate time series data are many and almost impossible to list. Many of the most influential papers in the literature on consumption, such as those by Hall (1978, 1988), Hansen and Singleton (1982,1983), Flavin (1981), have been useful conceptual exercises that have brought to the attention of the profession many important issues.

Among the most influential and most widely cited studies of aggregate time series papers are those of Campbell and Mankiw (1989, 1991). In these papers the authors consider a simple version of the Euler equation (5) and estimate it on aggregate data with the addition of income growth on the right hand side. The justification they use for such a specification is that there is a fraction λ of consumers behaving according to the permanent income-life cycle model while the remaining follow a 'rule-of-thumb' which consists in setting their consumption equal to their disposable income. Several scholars have since interpreted the coefficient $(1-\lambda)$ estimated by Campbell and Mankiw as the fraction of consumers that are 'liquidity constrained'. Campbell and Mankiw estimate their equation by instrumental variables[45] and obtain a coefficient $(1-\lambda)$ of 0.4.

There is no reason to believe that the 'excess sensitivity' of consumption to income, as summarised in such a coefficient, represents the fraction of consumers that behave according to Campbell and Mankiw's rule of thumb. First of all, one would have to assume that, if the consumers behaving according to such a rule were the same, the fraction of income (and consumption) accruing to them was constant over time. Second, one would have to believe that the non linearities in the Euler equation have no effect and could not generate such a result. Finally, to exclude that income and consumption are unrelated under the life cycle model, one would have to believe that consumption and leisure are separable in the utility function.

More recently, in the spirit of the Campbell and Mankinw studies, several papers have studied the time series properties of aggregate consumption and have tried to interpret its relation to a number of other variables in terms of the life cycle- permanent income model or as alternative deviations from it. Carroll et al. (1994), for instance, relate consumption to the index of consumer confidence and provide some interpretation of the high predictive power that such a variable has. Ludvigson (1997), on the other hand, relates consumption to consumer credit and finds that 'excess sensitivity' of consumption to such a variable is even more marked than that to labour income.

In the past, some of the papers that used aggregate time series data have also tried to incorporate in the model the possibility that consumption and leisure are not separable in the utility function. Bean (1986), Eichenbaum, Hansen and Singleton (1988), Mankiw, Rotemberg and Summers (1985) are examples of these attempts. The presence of corner solutions in labour supply, such as those observed in the case a spouse does not participate to the labour force or in the case of retirement, make the aggregation issues even more complicated.

---

[45] Because they worry about time aggregation, Campbell and Mankiw use instruments dated t-2 and earlier.

5.2 Micro data: some econometric problems

The analysis of household data on consumption presents a variety of problems that range from the availability and reliability of individual consumption data, to some subtle econometric problems. Of the many problems that is worth stressing three in particular are relevant for the discussion at hand. The first relates to the way in which consistent estimates are achieved from the orthogonality conditions implied by an Euler equation for consumption. The second concerns instead estimation and inference with average cohort data instead of genuine panels. The last point is about the presence of conditional second moments in the (log) linearised Euler equation

5.2.1 Consistency of estimators derived from Euler equations.

The first issue was first pointed out by Chamberlain (1984) and subsequently discussed by Hayashi (1987), Altug and Miller (1990) and Deaton (1992) among others. The Euler equation for consumption for a generic individual i, can be written as follows:

(17) $$E_t[h(x_{t+1}^i, x_t^i, \theta] = 0$$

where θ is a vector of parameters and x a vector of observable and unobservable variables. Equation (17) states that a function of parameters and variables is orthogonal to information available at time t. If the unobservables are i.i.d. innovations to stationary processes, we can, without loss of generality, restrict the vector x to be made of observable variables. In such a situation one can identify a vector of instruments w_t that, if its size is greater than the dimension of the parameter vector, can be used to identify it.

The problem arises if one tries to exploit the cross sectional dimension rather than the time dimension to construct the sample equivalent of equation (17). Equation (17) implicitly defines expectational errors, which by rational expectations are orthogonal with lagged information, of which, typically, the instruments are part. There is no reason to believe, however, that expectational errors average to zero in a cross section, at a given point in time. Analogously, there is no reason to believe that the cross sectional covariance of h in equation (17) with the vector of instrument w equals zero. The addition of time dummies, therefore, is not a solution for the problem at hand: each orthogonality condition one considers would imply the inclusion of a vector of dummy variables.

The only condition under which one can use the cross sectional equivalent of equation (17) under which introducing a vector of time dummies can solve the problem is when the expectational errors at a point in time are known to be exactly the same across individuals. But this is equivalent to assume complete markets. This is the option chosen, for instance, by Altug and Miller (1990) and by Atkenson and Ogaki (1997).

If one is not willing to assume complete markets, the only alternative is to have a sample covering a long time horizon. This is necessary so that expectation errors average out to zero. Notice that the 'large T' refers to the total length of the sample period over which individuals (or groups of individuals) are observed and not to the length over which a single individual is observed. [46]

---

[46] A problem similar to the one just discussed is the issue of the presence of individual fixed effects in the Euler equation for consumption. These could arise, for instance, if the discount factors contain an individual specific component. In this case, the use of weakly exogenous instrument

The necessity of a 'large T' to obtain consistent estimates of the parameters of an Euler equation follows from the nature of the residuals of such an equation which incorporate expectational errors. The same argument does not apply to the residuals of intratemporal (within period) first order conditions. Indeed, if expectational errors are the only component of the residuals, such equations should have a perfect fit. It is for this reason that the consideration of unobserved heterogeneity is crucial in the specification of preferences (see equation 13 above).

As mentioned in section 3.6, the presence of unobserved heterogeneity has implications on the nature of the residuals of the Euler equation. The specification of preferences in equation (13) implies that in equation (9') unobserved heterogeneity enters in first differences, so that, depending on the time series properties of v_t^i, the properties of the Euler equation residuals will be different. If the unobserved heterogeneity term v_t^i has a time invariant component which is individual specific, it is differenced. If v_t^i has a unit root, its differences would be white noise. If the deviation of v_t^i from a constant component are white noise, the residuals of the Euler equation are an MA(1) process. This has implications for the choice of instruments and the computation of the standard errors.

The situation is more complicated if the deviations of v_t^i from its fixed component are persistent (rather than white noise) and/or if the 'pure' discount factor β is individual specific. In the former case the term Δv in equation (9') would have a structure more complex than a simple MA(1), while in the latter, equation (9') would have a fixed effect. In such a situation the choice of instruments is not trivial, in that lagged individual variables are, in principle, correlated with the residuals of the Euler equation.

The problem might be less severe if one uses average cohort data: grouping the individuals belonging to a given cohort averages out the individual fixed effects and leaves only the cohort specific ones. Here the availability of a long time period is once again crucial: it is possible to estimate cohort specific fixed effects (say in discount factors) by introducing cohort dummies in the Euler equation.

The necessity of a 'large T' to obtain consistent estimates of the parameters of an Euler equation follows from the nature of the residuals of such an equation which incorporate expectational errors. The same argument does not apply to the residuals of intratemporal (within period) first order conditions. Indeed, if expectational errors are the only component of the residuals, such equations should have a perfect fit. It is for this reason that the consideration of unobserved heterogeneity is crucial in the specification of preferences (see equation 13 above).

5.2.2 Average cohort techniques

The lack of true panel data has often forced researchers to use the synthetic cohort data, which I have used in Section 2 for descriptive purposes, to estimate and test dynamic models, such as the Euler equation for consumption. The use of these techniques, pioneered by Browning, Deaton and Irish (1985) and discussed by Deaton (1985) and more recently by Moffitt (1994), has several advantages over real panels.[47] Non random attrition,

---

would be invalid. For a discussion of the issues related to this problem see Keane and Runkle (1994) and the comments to that paper in the Journal of Business and Economic Statistics.

[47] Average cohort techniques can be interpreted as an example of an estimator using complementary data sources of the kind discussed in Arellano and Meghir (1992).

for instance, is much less of an issue. On the other hand, one should be careful about the econometric problems induced by these data. In particular, because the size of the cell one uses to compute averages of the variables of interest is less than infinite, one should consider explicitly the presence of sampling error in the data one uses. Furthermore, because the averages are computed on the levels of variables and one is typically interested in the first differences of variables, the presence of sampling error induces an MA(1) structure on the residuals of the Euler equation estimated on synthetic cohort data. The presence of MA(1) residuals obviously has implications for the choice of the instruments and the computation of standard errors. The fact that synthetic cohort data are typically constructed out of independent repeated cross section, however, suggests simple valid instruments: by lagging the instruments an extra period should guarantee consistency of the IV estimator. [48]

The structure of the errors is even more complicated when one considers more cohorts simultaneously. In this case, one has to keep into account the possibility that the residuals of different cohorts in the same time period are correlated. The variance covariance matrix of the residuals is therefore quite complex. Its computation is, however, important to make correct inferences. [49]

Besides adjusting the estimated standard errors for the presence of an MA(1) error and the correlation among the expectational errors of different cohorts, it is also possible to use a GLS type of scheme to improve the efficiency of the estimator. One should be careful, however, in filtering the data so to avoid the inconsistency caused by the correlation of lagged expectational errors with the instruments. [50]

5.2.3 Conditional second (and higher) moments

In equation (9'), as in equation (10), I have neglected the presence of conditional second (and, if the log-normality assumption does not hold, higher) moments that follows from the log-linearization of the Euler equation in (6). In theory, this is not quite correct in that these variables are likely to vary as new information is made available to the individual consumer. This is particularly true for the conditional variance of consumption which is endogenously determined by the model (consumption is a choice variable). Incorporating these variables in the constant is equivalent to assuming that innovations to the conditional second (or higher) moments of consumption and interest rates are uncorrelated with the variables typically used as instruments.

Assessing the plausibility of this assumption is very difficult, especially because the answer is likely to depend on the time series properties (and in particular on the heteroscedasticity) of the determinants of consumption levels (income, wages and so on). In a recent paper, Ludvigson and Paxson (1997) have used numerical methods to compare the consumption function obtained from a log-linearised Euler equation to the 'true' one. Their method, however, does not provide estimates for the bias introduced in Euler equation estimation. Carroll (1997b) also uses numerical techniques to simulate log-linearised Euler equations, but focuses on the cross sectional rather than time series dimension, which, given the considerations in 5.2.1 is the relevant one.

---

[48] The Consumer Expenditure Survey used, for instance, in Attanasio and Weber (1995) is a rotating panel so that the construction of valid instruments is more complicated.
[49] A technical problem arises from the fact that is not easy to guarantee that the estimated variance covariance matrix is positive definite in finite samples.
[50] See Hayashi and Sims (1983).

An interesting and novel approach to this problem is the one used by Banks, Blundell and Brugiavini (1997) in a recent paper. These authors use an approximation to the consumption function derived by Blundell and Stoker (1995) which implies weighting the conditional second moments by the ratio of income to wealth at a given point in time. Banks et al. (1997) also try to identify separately the effect of aggregate and cohort specific shocks. The main findings of the exercise are two. First, conditional second moments are significant in the Euler equation; and second, the estimate of some important structural parameters, such as the elasticity of intertemporal substitution, is not much affected by the introduction of the conditional second moments.

Yet another alternative to the log-linearization procedure is the one proposed in Attanasio and Browning (1995) who start with a flexible functional form for the (log) marginal utility of consumption so to avoid the necessity of approximating an Euler equation. Should one be interested in the utility function that generates such a function, one can obtain it by integration. Not only does this method avoid dealing with the issue of linearization, but gives the possibility of estimating much more flexible functional forms than the isoelastic one. This is not to say that precautionary motives are unimportant, but that their relevance is going to depend on the curvature of the marginal utility of consumption and is, in the end, an empirical matter.

The alternative of estimating a non - linearised Euler equation (such as equation 8) is particularly unappealing because it would imply assuming away measurement error and would, in any case, make the use of average cohort techniques much harder, if not impossible.

5.3 Micro data: some evidence.

In this section, I do not discuss in detail all the papers that have estimated and tested Euler equations for consumption using micro data. Rather, I summarise some of the main contributions with an eye to the overall evaluation of the evidence which I give in the next sub-section.

One of the first papers to consider the implications of the permanent income - life cycle models with micro data is Hall and Mishkin (1982) in which the authors used PSID data to test and reject the implication that consumption changes were uncorrelated with lagged values of current income. Hall and Mishkin (1982) evidence was later criticised by Altonji and Siow (1987) because did not allow for measurement error.

Hall and Mishkin (1982) focused on the permanent income hypothesis and did not consider explicitly the Euler equation that one can get from the consumer intertemporal optimisation problem. One of the first, and probably the most influential article to take such an equation to the data, was the paper by Zeldes (1989a), who implemented ideas very similar to those in Runkle (1991). In both articles, the Euler equation is fitted to observations on food consumption from the PSID. In both articles, the authors consider explicitly the possibility of liquidity constraints and the possibility that these constraints affect different group in the population differently. Effectively, Zeldes (1989a) and Runkle (1991) estimated versions of equation (12) for different groups of the populations that had, on the basis of an observable variable, different probabilities of having $\mu_t = 0$.

The results they get, however, are quite different. Zeldes (1989a), in particular, splits the sample according to the wealth held and finds that the rate of growth of consumption is related with the lagged level of income for the low wealth sample. The same result does

not hold for the high wealth sample. Zeldes interprets this results as evidence of binding liquidity constraints for a large fraction of the population. Runkle (1991), who uses a different extract of the PSID, and different econometric techniques, obtains different results.[51]

Since Zeldes (1989) and Runkle (1991), many papers, including Hayashi (1985), Garcia, Lusardi and Ng (1997), Lusardi (1996), Mankiw and Zeldes (1994), Shea (1995) have used PSID data, either alone or in conjunction with other data, to estimate and test various versions of the model. Shea (1995), in particular, identifies a subsample of PSID households for whom he can track their union wage contracts and therefore construct a good measure of their expected wage growth. As many others, Shea (1995) tests (and marginally rejects) the hypothesis that consumption growth is not related to wage growth. Shea (1995), however, also notices that liquidity constraints imply an asymmetry between households who expect a decline and a rise in wages, because the former should be saving rather than borrowing. As he fails to identify these effects, his results cast doubts about the plausibility of liquidity constraints as an explanation of the excess sensitivity of consumption growth to income or wage growth.

One of the main problems with the PSID is that the measure of consumption that is included in the survey refers only to food consumption. I have stressed above the theoretical problems with such a measure. Attanasio and Weber (1995) show how the use of such a measure of consumption can lead to misleading results. In that paper, my co-author and I use average cohort data constructed from the CEX for the period 1980-1992 to show that some of the excess sensitivity results obtained on micro data can be accounted for by the non separability of food and other consumption.

The other two issues on which Attanasio and Weber (1995) focuses are the effects of aggregation over consumers and over commodities. As far as the aggregation over commodities is concerned, Attanasio and Weber (1995) find significant effects of the zero-homogeneous price index in equation (14) that, however, are not quantitatively important. This evidence is consistent with that for the UK FES reported in Blundell, Browning and Meghir (1994) and Banks, Blundell and Preston (1994).

About the aggregation over individuals, Attanasio and Weber (1995) report evidence which shows that the effect of the non-linearities in the Euler equation can be quite important. This is consistent with the evidence reported in Attanasio and Weber (1993) on FES data. In the latter paper, my co-author and I show that aggregating the individual data so to obtain an aggregate conceptually similar to the National Accounts statistics (i.e. taking the log of the arithmetic mean of consumption) one obtains results that are quite similar to those obtained with aggregate times series data. In particular, it is possible to obtain excess sensitivity to predicted income, a low estimated elasticity of intertemporal substitution and rejection of the overidentifying restrictions. The results obtained aggregating the data in the proper way (i.e. taking the average of log consumption), however, are much more consistent with the theoretical model. In particular, there is no rejection of the overidentifying restrictions. The difference between the log of the average and the average of the log is a measure of inequality which varies over the business cycle and is likely to be correlated with the instruments used in estimating the Euler equation.

---

[51] Keane and Runkle (1994) have recently re-estimated Zeldes equations on his data but using different econometric techniques. Jappelli, Pischke and Souleles (1996) match SCF data, which contain information on self reported 'liquidity constrained' status with PSID data.

The last group of papers I have cited (Attanasio and Weber (1993, 1995), Browning, Blundell and Meghir (1994), Banks, Blundell and Preston (1994), Attanasio and Browning (1995)) all stress the importance of controlling for demographic factors and for labour supply effects. These papers, which use average cohort data for the UK and the US (the data presented in Section 2) show that once one controls for the influence that demographics and labour supply might have on the marginal utility of consumption, there is no evidence of excess sensitivity of consumption to income or rejection of the overidentifying restrictions. Female labour force participation and family size seem to be particularly important in this respect.[52]

If one believes that the estimation of the Euler equations discussed in this section yields consistent estimates, the parameter on the real interest rate can be interpreted as the elasticity of intertemporal substitution. Such a parameter is of some importance for a number of policy issues and its size has been discussed at length. Hall (1988) in particular, claims that the use of the correct instruments delivers very low estimates of this elasticity. However, the evidence that emerges from the micro studies which use an isoelastic specification of preferences (such as Attanasio and Weber (1993, 1995), and Blundell, Browning and Meghir (1994)), is that, both in the UK and in the US, the elasticity of intertemporal substitution of consumption (EIS) is just below 1.[53]

From the discussion above, it is clear that in addition to the EIS, a number of other parameters, measuring the effect of leisure, that of demographics or other variables affecting the marginal utility of consumption are likely to be extremely important in determining the level of consumption. Their interpretation, however, is not easy without a solution for the level of consumption.

6. Where does the life cycle model stand?

It is now time to take stock on the empirical relevance of the models discussed so far. My reading of the evidence briefly presented in Section 5 is that the life cycle model, enriched to account for the effect of demographic and labour supply variables, is not rejected by the available data. Alternatively, and perhaps more accurately, one could say that the life cycle model can be made complex enough not to be inconsistent with the available data. This view is not widely accepted in the profession so that it deserves some discussion.

Before discussing these issues, however, a number of caveats are in order. First, the model has been fitted with success only to households in the middle of their life cycle. While these account for a large fraction of aggregate consumption, additional work is needed to understand the behaviour of young and elderly households.

The behaviour of retirees in particular, can be quite difficult to model. If leisure and consumption are non separable in the utility function, a radical change in labour supply could be linked to a change in consumption. Furthermore, a number of other important factors, ranging from family size, to health status, to the probability of death, changes dramatically in the years after retirement. In section 2, we have seen that both in the US and especially in the UK, consumption drops substantially at retirement. This could be

---

[52] Demographic variables might be capturing the effect of the conditional second (and higher) moments ignored in the log linearization procedure. Lusardi (1996) criticises the use of income data in the CEX as they might be affected by measurement error.

[53] As mentioned above, the use of aggregate rather than individual data and aggregation problems can explain part of the differences in the size of the estimates of the elasticity of intertemporal substitution.

related to insufficient savings and a misperception about the amount of (public and private) pension benefits. Alternatively, the drop in consumption could be explained within the optimisation framework of the life cycle model if considered together with all the changes mentioned above. This important topic is studied in a recent paper by Banks, Blundell and Tanner (1998) who find that, even ignoring health status and mortality issues, two thirds of the drop in consumption observed around retirement can be rationalised by an optimisation model.

Second, while the endogeneity of labour supply choices can be controlled for in the empirical analysis, the joint modelling of consumption and of labour supply is extremely valuable and is still at the beginning.[54] Similarly, expenditure on durables, which is an important component of total expenditure, needs to be modelled in a different way, as discussed in section 8.

Third, the fact that a particular specification of preferences fits the data reasonably well and that one can obtain consistent estimates of the structural parameters, does not mean that borrowing restrictions and liquidity constraints are unimportant. The Euler equation can be a very poor tool for identifying the presence of such phenomena, as the equation holds whenever the constraints are not binding. For the notion of liquidity constraints to be useful, however, it is necessary to model the behaviour of constrained individual explicitly and determine what the aggregate implications of these constraints might be.

A literal interpretation of the papers mentioned in the last part of Section 5 is that a flexible version of the life cycle model, which allows for demographic and labour supply effects in the utility function, cannot be rejected by the data. Neither excess sensitivity tests, nor tests of overidentifying restrictions reject the null. The strength of a result that does not reject the null, however, obviously depends on the power of the test used.

The power of the tests of overidentifying restrictions or of the excess sensitivity tests that fail to reject within a sufficiently complex version of the model might be questioned. If an Euler equation is saturated with demographic and labour supply variables, it might be hard to measure precisely the coefficient on expected income or, more generally, reject overidentifying restrictions. This problem might be particularly serious if some of these variables, such as labour supply, are correlated with expected income. In this sense, this criticism is equivalent to state the difficulty in distinguishing hypotheses about preferences from hypotheses on budget constraints.

A possible response to this kind of criticism is to check whether the preferences parameters of the model that is not rejected by the data are sensible and, most importantly, whether the model is able to explain facts other than those that have been used to fit it. An example will make this argument clear.

In Attanasio et al. (1996), my co-authors and I consider a relatively simple version of the life cycle model where the within period utility function includes demographic variables and whose parameters are assumed to be the same across education groups. Such a specification of the model is estimated on US average cohort data and it is not rejected by

---

[54] Surprisingly few studies analyse consumption and labour supply choices jointly. An almost exhaustive list includes MaCurdy (1983), Browning, Deaton and Irish (1985), Hotz et al. (1988), Altug and Miller (1990), Blundell, Meghir and Neves (1993), Attanasio and MaCurdy (1996). Of these only the last two papers consider female and male labour supply jointly.

the data. The elasticity of intertemporal substitution is estimated to be just below 1. Given some hypotheses about the stochastic process that generates income, interest rates and demographic variables, and given an assumption on terminal conditions, one can solve numerically for the consumption function at each age. We do so and then simulate the model for income and demographic profiles calibrated on different education groups.

With this technique we are able to reproduce some of the main feature observed in the data, that is, that the consumption profiles are steeper for the groups that have a steeped income profiles. This feature of the data has been interpreted as a failure of the life cycle model by Carroll and Summers (1991). Our exercise, however, shows that a flexible version of the life cycle model incorporating some realistic features, such as the effect of family composition on utility, can explain it. The result is remarkable because there is nothing, at the estimation level, that fits the particular feature of the data (differences in consumption age profiles by education groups) that the estimated preferences (assumed to be the same across groups) and the differences in demographic profiles are able to explain. In this sense the exercise constitute a genuine 'out-of-sample' verification of the ability of the model to fit the data.

As the utility function is of the CRRA type, the model also incorporates the precautionary motive linked to income uncertainty which Carroll (1997a) has recently advocated as the most likely explanation of the relationship between the shape of consumption profiles and that of income profiles. The simulations show, however, that this effect is quantitatively less important than that generated by differences in demographic profiles. Indeed, most of the 'tracking' of consumption and income profiles is generated by corresponding similarities in demographics. This argument is also consistent with the evidence presented in Section 2 where consumption 'per adult equivalent' profiles are quite flat over the life cycle. [55]

The simulation techniques used in Attanasio et al. (1996) are similar to those used by Deaton (1991) and by Hubbard, Skinner and Zeldes (1994), among others, to obtain a closed form solution for consumption. These techniques are the only way, if one is not willing to assume quadratic utility or CARA utility (as in Caballero (1990))[56], to obtain a consumption function in the presence of uncertainty. They are quite expensive as they are numerically intensive for any problem which goes beyond the simplest assumptions and require the researcher to take a stand on any detail of the optimisation problem (from the terminal condition to the nature of the income process). The payoff one can obtain can, however, be quite large.

The Euler equation approach, first used by Hall (1978) in the consumption literature, was a major innovation. It allowed the rigorous consideration of uncertainty in a complicated dynamic decision problem while preserving the empirical tractability of the model. The beauty of the approach consists in the fact that exploits the main implication of the life cycle - permanent income model, that is that an optimising consumer keeps the discounted expected value of the marginal utility of consumption constant, to eliminate the unobservable 'fixed effect', the marginal utility of wealth, which incorporates the influence of expectations and all the variables that are relevant for the consumer optimisation problem. The price that one pays by differencing out the marginal utility of wealth is, however, non negligible. While it is true that one can use the Euler equation to estimate

---

[55] It can be argued that the demographics in the Euler equation for consumption are picking up the effect of the omitted conditional higher moments.
[56] Both in the case of quadratic and CARA utility a number of additional assumptions about interest rates are necessary.

structural parameters, one looses the ability of saying anything about the *level* of consumption. This implies that without additional information and structure, it is not possible to answer a number of extremely important questions. In other words, the Euler equation for consumption *is not* a consumption function and therefore it does not tell us anything about how consumption reacts to news about the economic environment in which economic agents operate.

The symptom of this unsatisfactory state of affairs are apparent in the literature. If one reads, for instance, the literature on saving, which in recent years has dealt with a number of interesting issues that range from the decline in personal saving rates in the US to the effects of tax incentives on personal and national saving, one rarely finds a systematic use of Euler equation estimates or even references to the Euler equation literature. Indeed, the use of structural models of consumption behaviour is quite rare. On the other hand, the same literature suffers from a number of identification problems that arise from the inability or unwillingness to put more structure onto the descriptive analysis. A good example is the debate on the effectiveness of fiscal incentives to retirement saving.

It seems that, 40 years after the publication of Friedman's book, we still have not been able to construct a 'consumption function' which incorporates the main features of the dynamic optimisation model without loosing analytical and empirical tractability. The development of the numerical solution techniques referred to above is one possible avenue. The main problem with these techniques, however, is that they require a complete specification of the economic environment (including terminal conditions) in which the agent lives. Therefore, to make the problem treatable, even with the use of powerful computation methods, very strong assumptions are needed. Furthermore, we cannot use the conditioning arguments which simplify the estimation of an Euler equation. While at the estimation level variables such as durables and labour supply can be rigorously treated as endogenous without modelling them explicitly, at the level of numerical solutions, they cannot be ignored.

7. Insurance and Inequality

One of the main implications of the life cycle model sketched in Section 3, is the fact that households attempt to smooth consumption over time. If one sees an individual household in isolation, this can be achieved only through the accumulation of buffer stock saving or through borrowing. However, if one considers the interaction of many households, one realises that if individual shocks are not perfectly correlated, there is scope for the diversification of idiosyncratic risk and for welfare improving contracts that allow the individual households to share part of the risk they face. The essence of a risk sharing situation, as formalised, for instance, by Townsend (1987), is well described by one of the characters in *Shipping News*, the novel by Anne Proulx set in Newfoundland. Describing life in one of the small islands off the coast, he declares "...it was never easy....on Gaze Island, but they had the cows and a bit of hay, and the berries, the fish and the potato patches, and they'd get their flour and bacon in the fall from the merchant over at Killick-Claw, and if it was hard times, they shared, they helped their neighbor. No they didn't have any money, the sea was dangerous and men were lost, but it was a satisfying life in a way people today do not understand. There was a joinery of lives all worked together, smooth in places, or lumpy, but joined." (pp. 168-9).

Social and economic interactions in modern western societies are probably much more complex than those that prevailed on Gaze Island. It is however an interesting question to establish to what extent institutions of various kinds are used to smooth consumption

across households and to diversify idiosyncratic risk. In other words, it might be interesting to establish the extent to which implicit or explicit contracts, family networks, social safety nets and so on can approximate the intertemporal allocation that would prevail under complete contingent markets of the kind described in an Arrow- Debreu equilibrium. Furthermore, it might be interesting to establish what are the welfare costs implied by the lack of complete contingent markets, given the nature of the observed idiosyncratic shocks.

From a theoretical point of view, the basic proposition can be derived if one considers the allocation problem faced by a central planner who maximises, given a set of Pareto weights, the average of individual utilities given the resource constraint. This framework is particularly useful because it can be used to control easily for a number of factors and in particular for the presence of multiple commodities, non separable leisure and so on. The basic implication is that the rate of change of the marginal utility of consumption (which might depend on leisure and other commodities) is equalised across individuals.

Townsend (1994) was one of the first to consider the empirical implications of complete market, that is that under certain conditions "... individual consumptions are determined by aggregate consumption, no matter what the date and history of shocks, and so individuals' consumption will move together"(p.540).[57]   While Townsend (1994) considered data from India, other studies, such as Mace (1991), Cochrane (1991) and Hayashi et al. (1996) performed similar exercises using US data, that is either the CEX or the PSID.

Both Cohcrane (1991) and Hayashi et al. (1996) reject the implications of the risk sharing hypothesis. Cochrane (1991) finds that food consumption changes in the PSID are related to changes in health and employment status. Hayashi et al. (1996) use the same data but control for the possibility that the relationship between changes in wages and consumption is generated by the non-separability of leisure and consumption in the utility function. By following spin-offs in the PSID, they also test (and reject) the hypothesis that risk is shared among families.

Attanasio and Davies (1996) have matched data from the Current Population Survey and the CEX to test the hypothesis that movements in relative wages are reflected into movements in relative consumption or, to be precise, in marginal utilities of consumption. Attanasio and Davis consider education and year of birth cohorts. Most of the variability in relative wages comes from the shifts in the wage distribution across education groups that occurred during the 1980s. They find that such movements were mirrored in movements in the consumption distribution, indicating a 'spectacular failure' of the perfect insurance hypothesis.[58]   The analysis of Attanasio and Davis is particularly damning for the perfect insurance paradigm because it focuses on economy wide, well observed shifts, and therefore the relationship between consumption and income cannot be explained with models which consider private information such as those in Phelan and Townsend (1993).

Attanasio and Davis (1996) also try to evaluate the welfare cost of the lack of an institutional framework that allows for the diversification of idiosyncratic risk. The evaluation of such a cost obviously depends on preference parameters and on the nature of the shocks considered. Attanasio and Davis (1996) under a set of plausible assumptions and

---

[57] Townsend (1994) credits Diamond (1967) and Wilson (1968) with the first versions of this proposition.
[58] As Hayashi et al. (1996), Attanasio and Davis (1996) use lead as well as lags of wages as instruments. Both studies consider short *and* long run shocks.

a coefficient of relative risk aversion of two, evaluate the cost of the lack of insurance to 'group specific shocks' at 2.5% of consumption. This estimate is considerably larger than the cost of business cycle fluctuations evaluated by Lucas (1987) and is more in line with the results reported by Imrohoroglu (1989). It should also be stressed that Attanasio and Davis (1996) estimate completely ignores 'within groups' shocks and focuses mainly on relative low (rather than business cycle) frequencies.

Whilst the analysis of Attanasio and Davis (1996) can be framed in terms of a test of the perfect insurance hypothesis, it can also be interpreted as documenting the evolution in inequality in consumption and wages in the US during the 1980s. The increase in inequality in wages in the US has been analysed in many studies by labour economists. On the other hand, the evolution of inequality in consumption has not received much attention until quite recently. Little is known about the evolution of consumption inequality over the business cycle and over different stages of the process of economic growth. In addition to Attanasio and Davis (1996), Cutler and Katz (1991) analyse the evolution of inequality in consumption during the 1980s using CEX data. Goodman, Johnson and Webb (1997) provide an exhaustive analysis of the recent trends in both income and expenditure inequality in the UK.

The lack of evidence on consumption inequality is somewhat disconcerting, especially if one compares it to the amount of evidence on wage and income inequality. From a theoretical point of view, it is not completely obvious which of the two measures of inequality is more interesting. Blundell and Preston (1997) discuss the advantages and drawbacks of both: consumption inequality is more likely to reflect the cross sectional variability of permanent income, while income inequality is more affected by temporary shocks. The relationship between the two, however, depends on a number of factors that range from the specification of preferences to the nature of the income shocks, to the institutions and instruments available to households to (self) insure against idiosyncratic shocks.

More generally, a theoretically consistent analysis of consumption inequality cannot avoid the discussion of the implications of different preference specifications and market institutions. Suppose, for instance, that individual households are prevented from borrowing. If labour supply behaviour is considered exogenous, these households will engage in precautionary saving to avoid the effects of extremely negative shocks to income. On the other hand, if one models labour supply choices and in particular female participation in the labour market, it is quite possible that the role of precautionary saving is taken by labour supply. That is to say, it is possible that households, rather than consuming less goods, decide to consume less leisure.

Another important element, which has been only recently analysed, is the availability of means tested safety nets. Hubbard, Skinner and Zeldes (1994) consider the effect that liquidity constraints and means tested safety nets have on the pattern of consumption using numerical techniques to solve the consumption function.

Browning and Crossley (1997) have recently considered another potentially important channel that some households might use to smooth the marginal utility of consumption over time: what they call small durables. The idea is intuitive: individuals that have no access to savings or loans to smooth out fluctuations in income, might postpone the replacement of small durables. The features that characterise the 'small durables' in Browning and Crossley's problem are the durability and the irreversibility (due to the lack of a second hand market) of such items, as well as the fact that scrappage decisions are not only determined by the physical depreciation but also by the economic situation faced by

the agent. Browning and Crossley (1997) present both some theoretical results and some empirical evidence on a sample of Canadian unemployed supporting the importance of this type of smoothing behaviour. The evidence I presented in Table 2.5, about the relative variability of durable expenditure for groups of households with different levels of education is consistent with the idea that poorer individuals might use durables to smooth out fluctuations in income.

Browning and Crossley (1997) focus on items that cannot be used as collateral (such as clothes or small appliances) to obtain loans. It would be interesting to consider the implications of the same ideas for the replacement of items for which reasonably efficient second hand markets exists and which are collateralizable such as automobiles and, to a certain extent, 'white goods'. Greespan and Cohen (1997) have stressed the important role that scrappage of old cars plays in forecasting expenditure on automobiles which, in turn, is a very important indicator of the status of the business cycle.

A few studies, recently, have started to study the evolution of inequality over the life cycle. The idea is quite simple: if consumption follows a random walk, the cross sectional distribution of individual consumption should 'fan out' over the life cycle. That is, the cross sectional variance of individual consumption should be increasing, on average, over long periods of time. This idea was first exploited by Deaton and Paxson (1994) who analysed average cohort profiles for the variance of (log) non durable consumption in three different countries, the US, the UK and Taiwan.

8. Intertemporal non- separability

In the discussion so far, I have assumed that preferences are separable over time. While this is an extremely convenient assumption, it is easy to think of situations in which it is violated. In this section I consider the implications of the fact that current expenditure might have lasting effects. After sketching the general problem, I discuss in detail two particular examples of time non-separability of some importance: that of durability and of habit formation. As is clear from the discussion, the complication implied by time dependence of preferences is only one of the things that make the treatment of these phenomena extremely hard.

The general problems can be stated in reasonably simple terms if one states the problem in terms of a flow of services derived from a 'stock' and assumes that there are no costs in changing the 'stock'. The simple problem in section 3, can therefore be re-written as:

$$\text{Max} \sum_{j=0}^{T-t} E_t U(S_{t+j}) \beta^j \; ; \text{subject to}$$

(18)
$$A_{t+j+1} = (1 + r_{t+j}) A_{t+j} + y_{t+j} - s_{t+j}$$

$$S_{t+j} = \sum_{k=0}^{\infty} \alpha_k s_{t-k}; \quad \alpha_1 = 1.$$

where the S_t is the 'stock' from which utility is derived, s_t is expenditure and the coefficients α_k define the type of time dependency. As we see below, such a model can easily incorporate durability as well as habit formation. It is also straightforward to

consider the case in which S is a vector whose components have different degrees of durability.[59]

A rational individual, in choosing the level of expenditure s will take into account the effect that this has on the level of the 'stock' in the current as well as in the future period. It is straightforward to verify that the first order condition for such a problem is going to be:

$$(19) \quad E_t m_t = E_t \beta (1 + r_{t+1}) m_{t+1}$$

where $m_t = \sum_{k=0}^{T-t} \beta^k U'(S_{t+k}) \alpha_k$.

Notice that if the α_k's are different from zero, m_t is not a variable known at time t as it depends on the future marginal utility of the stock S. In general, therefore, it will not be possible to express equation (19) in terms of the rate of growth of m_t as is usually done (see, for instance, equation 8). Notice also that the marginal utility of expenditure m_t depends on a potentially very large number of terms. To make an equation such as (19) operational it will be necessary to simplify it by some algebraic manipulation whose nature depends, once again, on the pattern of the coefficients α_k.

The two models of non separability I consider below (durables and habits) differ on the nature of the time dependence. In the first case we have substitutability, while in the latter we have complementarity of expenditure over time.

8.1 Durables

As illustrated in section 2, expenditure on durables, both at the aggregate level and at the cohort level, is by far the most volatile component of consumption. In addition, the dynamic of durable expenditure seems much more complex of that of the other components of consumption. A simple model of durable consumption can be obtained from the system in (18) if it is assumed that $\alpha_k = (1 - \delta)^k$. In this case S can be interpreted as the value of the stock of durables, δ the depreciation rate and the third equation in (18) becomes:

$$(20) \quad S_t = (1 - \delta) S_{t-1} + s_t$$

Assuming quadratic utility, Mankiw (1982) generalised Hall's random walk model to durable expenditure. In particular, he showed that under these circumstances the model implies that changes in expenditure should follow an MA(1) process with coefficient equal to (1-δ). This implication is strongly rejected by the aggregate data, as shown, for instance, in section 2. Dunn and Singleton (1986) have used a more attractive preference specification and considered the Euler equation for durable consumption and its implications for asset pricing. The most comprehensive treatment of durability and of the Euler equations associated with a variety of preferences is found in Eichenbaum and Hansen (1990) who use a Gorman- Lancaster technology that converts expenditure into stocks and stocks into services over which utility is defined. The aggregation issues I discussed in section 4 are obviously relevant for durables as well. The aggregation problems

---

[59] See, for instance, Eichnbaum and Hansen (1990).

are actually even more complicated because of the possibility of corner solutions, i.e. households that decide not to own a durable (a car for instance), that can be safely ruled out for non-durables under the assumption that the marginal utility goes to infinity at zero or low levels of consumption.[60]

A possibility for the deviations of the dynamics of aggregate durable expenditure from that implied by the simple model is that of the existence of adjustment costs. It seems natural to assume that adjusting the stock of durables involves costs, motivated both by the existence of imperfections in the second hand market and by 'pure' adjustment costs (such as search costs, taxes and similar).

The literature first studied convex (typically quadratic) costs of adjustment. A typical example is Bernanke (1984, 1985), who estimated models with quadratic costs both with aggregate and individual data. Eichenbaum and Hansen (1990) also incorporate the possibility of adjustment costs in their technology. Such attempts, however, have turned out to be unsuccessful. The main reason for this is that convex adjustment costs predict a smooth adjustment towards an equilibrium. To avoid increasing costs households will adjust their stock of durables often and by small amounts. Even casual observation suggests that this is not a very accurate description of household behaviour.

The next step is then to consider non convex costs of adjustment; that is either fix or proportional costs. The characterisation of optimal behaviour under this type of costs is obviously much harder because for many households in many periods the optimal policy involves no adjustment. The first paper to characterise the optimal adjustment policy under non convex costs was Grossman and Laroque (1990) who studied a problem of a consumer deriving utility from a single durable and solving a dynamic optimisation problem involving the choice of the optimal size for the durable and the optimal portfolio investment for her financial wealth. Grossman and Laroque (1990) proved that the optimal strategy involves an (S,s) rule of the kind considered in the optimal inventory literature of the 1950s. Specifically, they proved that the value function associated to the consumer problem is a function of a single state variable, the ratio of the value of the durable to financial wealth, and that the durable is adjusted to a 'target' level when the state variable crosses a lower or an upper bound. When the state variable is within the two bounds the optimal policy is not to adjust the durable. In addition, Grossman and Laroque (1990) characterised the size of the inaction band as a function of the parameters of the problem, as well as the optimal portfolio strategy.

One of the problems of the Grossman-Laroque model is that a solution can be obtained only if the problem can be expressed as a function of a single state variable. This precludes, for instance, the consideration of labour income. More recently, Eberly (1994) and Beaulieu (1993) have extended the Grossman- Laroque model in various directions. Eberly (1994), has showed that changes in the durable stock between periods in which an adjustment is performed obey to an Euler equation. Beaulieu (1993) has reformulated the Grossman -Laroque model and expressed the (S,s) rule in terms of the ratio of durables to non durables.

The attractiveness of the (S,s) model of adjustment lies in its implications that, in most periods, consumers do not adjust their stock of durables and when they do adjust, they usually make substantial adjustments. These features (large and infrequent adjustments),

---

[60] Bernanke (1984) and Hayashi (1985) are among the few studies that have used individual data to analyse durable expenditure.

that can be generated by the *(S,s)* model, are the exact contrary of those of the quadratic adjustment model which yields small and frequent adjustments.

The consideration of large and infrequent adjustments poses an entire new category of aggregation problems. Caplin (1985) was the first to consider them. Caplin and Spulber (1987), in the context of a model of price adjustment, provided (stringent) conditions under which lumpy individual adjustment could result in a smooth aggregate adjustment. Bertola and Caballero (1990) analysed the implications of the *(S,s)* type of adjustment for the dynamics of aggregate durable expenditure and stressed that the assumptions needed to generate the type of neutrality results obtained by Caplin and Spulber are indeed very fragile. *(S,s)* models can generate very rich and complicate dynamics that seem not to be inconsistent with the observed autocorrelation of durable expenditure. Caballero (1990, 1993), in particular, has stressed how *(S,s)* models can generate patterns of MA coefficients that are similar to those I reported in section 2. Caballero and Engel (1991) have studied more generally the aggregate properties of *(S,s)* economies. Caballero (1993) has tried to estimate the parameters of the *(S,s)* model by fitting the highly non linear model resulting from the aggregation of a simple *(S,s)* rule to aggregate time series data.

The estimation of *(S,s)* models with individual data is not an easy task for several reasons. First of all, the data requirements can be quite formidable: information on the value of the stock of durables before and after the adjustment is necessary. It is also desirable to follow households over some time to bound the range of inaction by the households that are observed not to adjust. In addition, the numerical problems in estimation can be quite formidable.

The first attempt at estimating such a model with individual data was by Lam (1991) who estimated by maximum likelihood an *(S,s)* rule defined in terms of the ratio of the stock of automobiles over permanent income. Eberly (1994), also estimated the width of the *(S,s)* rule for the subset of consumers who were observed to adjust the stock of their automobiles in her sample.[61] Beaulieu (1992) also uses individual observations but follows a different approach. He considers a single cross section and computes the ergodic cross sectional distribution to which the economy would converge in the absence of aggregate shocks. The parameters of this distribution will depend on the parameters of the *(S,s)* rule. He can then estimate them from the parameters of the observed cross sectional distribution which is assumed to coincide with the ergodic distribution.

Like Lam (1991), I have estimated the parameters of the *(S,s)* rule by maximum likelihood on individual data in Attanasio (1995). However, I formulate the *(S,s)* rule in terms of the ratio of the stock of automobiles to non-durable consumption, allow for observed and unobserved heterogeneity in both the target and the band width and consider a more flexible stochastic specification. The consideration of unobserved heterogeneity in both the target level and the band width makes the model equivalent to a recent formulation of lumpy adjustment by Caballero and Engle (1992) who consider, rather than a strict *(S,s)* rule, a 'hazard function' which gives the probability of adjusting as a function of the difference between the current stock and its target level.

While substantial progress in the understanding of the behaviour of models with lumpy adjustment has been made (not only in consumption but also in investment and labour demand), a lot of research is still necessary to establish the empirical relevance of these

---

[61] Eberly (1994) also splits the sample between consumers who are likely to be liquidity constrained and those that are liquid. In doing so she uses both the procedure used by Zeldes (1989a) and a switching regression model.

models and to quantify their parameters. The aggregate dynamics of an *(S,s)* system depends crucially on the properties of several stochastic processes about which we still have little or no information. These include the process by which the state variable changes when no adjustment occurs, the process by which the target level changes, the degree of heterogeneity in target levels and band width, the degree of persistence of individual shocks to band width (cost of adjustment) and the correlations among these variables.

8.2 Habit formation

A form of intertemporal persistence of preferences alternative to durability is that of habits. As said above, habits can be obtained using a specific patterns for the coefficients α in equation (18). For instance, Heaton (1995) and Constantinides (1990) use the following specification:

$$(21) \quad S_t = c_t - \alpha(1-\theta)\sum_{j=0}^{\infty}\theta^j c_{t-1-j} \; ; \quad ; \quad 0 \le \alpha < 1, \quad 0 \le \theta < 1$$

The term $(1-\theta)\sum_{j=0}^{\infty}\theta^j c_{t-1-j}$ is referred to as the 'stock of habits', which depreciates according to the parameter θ. A particularly simple specification is obtained when $\theta = 0$, in which case the stock of habits is simply given by last period consumption. As with the cases considered above, the marginal utility of time t expenditure is a function of both present and future variables, so that it is now known at time t. The habit formation model has a long history, dating back at least to Gorman (1967), Pollack (1970), Houthakker and Taylor (1970)[62]. Spynnewin (1981) presents an ingenious reparametrization of the optimisation problem which, by an appropriate redefinition of prices and quantities, allows the maximisation of an intertemporally additive function. However, it is only in the 1980s that habits models are coupled with the hypothesis of rational expectations and preferences incorporating temporal persistence are used to derive Euler equations analogous to equation (19) above. Examples of this practice include Hansen, Eichenbaum and Singleton (1988) (for a model with habit forming consumption and leisure), Hansen and Eichenbaum (1990), Novales (1990), Constantinides (1990) while Hotz, Kydland and Sedlacek (1988) and Kennan (1988) consider habit formation in leisure. Browning (1991) uses a dual approach to derive the equivalent of 'Frisch' consumption functions in the presence of intertemporal non separability.[63]

More recently, Heaton (1993, 1995) has considered the interplay between time aggregation and habit formation. His argument is that at high frequency consumption seems to exhibit substitutability, while at lower frequencies, there is evidence of complementarities. Heaton argues that the local substitutability of consumption could be explained by time aggregation. In general, he presents preferences that are flexible enough to accommodate the 'slow' formation of habits.

Most of the work on habit models so far has been done on aggregate time series data. One of the reasons for this is the fact that very few panel data contain information on consumption. The CEX, which I used in section 2, contains only up to four quarterly

---

[62] Browning (1991), however, has citations from Marshall and Haavelmo.

[63] Browning's simple structure can encompass both substitutability and complementarities over time. The main problem with his approach is that he can only introduce uncertainty by considering points expectations about the future.

observations per households. The average cohort analysis that is used in the study of dynamic models of consumption, cannot be used in the analysis of habit formation, because these models involve the covariance of subsequent consumption observations *for the same household.* In other words, we cannot aggregate the product $C_t^h C_{t-k}^h$ over the household index h, unless we have observations for times t and t-k for the same households.

The only paper that studies time dependence in preferences using micro data is, to the best of my knowledge, Meghir and Weber (1996), that I discussed in the section on liquidity constraints. Another novelty of that paper is that it considers different components of consumption and allows for different levels of persistence depending on the commodity considered. The interesting evidence is that, when durable commodities are controlled for, there is no evidence of persistence in the 3 equations demand system that Meghir and Weber (1996) consider.

9. Conclusions

Rather than summarising the various sections that compose the chapter I prefer to conclude the chapter comparing the status of our knowledge and understanding of consumption behaviour to that of twenty years ago. It is fair to say that substantial progress has been made. In the last twenty years we have learned how to deal with uncertainty in a rigorous fashion and have recognised the importance that this may have for consumption and saving decision. Indeed an entire branch of the literature, that on precautionary saving, has developed to deal with these issues.

While the emphasis given to the proper treatment of uncertainty and the lack of appropriate data sources has meant that many studies have focused on aggregate data, it has now become clear that it is very hard, if not impossible, to estimate preference parameters from aggregate time series data. Aggregation issues are important in many dimensions. In addition to the standard aggregation problems created by the non linearity of the relevant theoretical relationships, there are other ways in which aggregation issues become important. Corner solutions and participation decisions (in labour and financial markets), inertial behaviour and transaction costs, all add a new dimension of complexity and make the cross sectional distribution of the variables under study relevant for the dynamics of the aggregate.

The Euler equation has been the main instrument to analyse consumption both in micro and macro data, to estimate preference parameters and to test the overidentifying restrictions implied by the consumers' optimisation problem. In the process we have learned a lot about the econometric problems of estimating Euler equations. In particular, we have learned what are the identifying assumptions that one needs to make to get consistent estimates from the available data.

The Euler equation is a powerful tool in that it allows the consideration of complex and flexible preference specifications without loosing the empirical tractability. As long as the variable over which one is optimising can be adjusted without cost and is not at a corner, one can derive an Euler equation which, even if it includes the values of other endogenous variables, delivers orthogonality restrictions that can be used to estimate preference parameters. The empirical research has used Euler equations for non durable consumption and its components that have become, in the attempt to fit the observable data, increasingly complicated. This level of complexity is probably unavoidable if one is serious about taking the model to the data.

As I stressed several times during the chapter, the price one pays in dealing with Euler equations is not negligible: one looses the ability of saying anything about the *level* of consumption. The empirical tractability of the Euler equation is obtained differencing out the marginal utility of wealth, and therefore one of the main determinants of consumption *levels*. The challenge for future research is to construct a consumption function that incorporates the insights of the Euler equation and yet allows us to say something about the levels of consumption and about how consumption reacts to changes in the economic environment. Such a consumption function is necessary to make predictions about future consumption, about saving behaviour, and about the effects of alternative policy measures.

One possibility that is being explored is the use of numerical techniques. They have proved to be useful both to improve our understanding of the dynamic optimization problems typically postulated and to characterise the implications of different preference structures and economic environments on consumption and saving behaviour. They can also be used to validate, indirectly, the preference specifications implied by the Euler equation estimates that best fit the consumption growth data. Their main limitation, however, is an important one: they can only be used to analyse extremely simplified models and they require the full characterisation of the agents' economic environment.

Several other areas of research are important and exciting. I will just mention two. Our understanding of consumption smoothing and of the evolution of inequality is still at the beginning. The development of this understanding and of the importance that different institutional frameworks and financial instruments have for these issues is an extremely important research agenda both from a positive and from a normative point of view. Related to this is also the issue of the time series properties of individual consumption: unlike for earnings and hours of work, next to nothing is known about this. Our understanding of durable expenditure is still very limited and yet it is an extremely important issue, both because durables are the most volatile component of consumption and because they have a number of features (the fact that they can be used as collateral, transaction costs, the durability itself) that make them difficult to model.

References

Abel, A.B. (1990): "Asset Prices Under Habit Formation and Catching Up with the Joneses", *American Economic Review Papers and Proceedings*,80, 38-42.

Abel, Andrew and Kotlikoff, Laurence J. (1992) "Intergenerational Altruism and the Effectiveness of Fiscal Policy: New Tests Based on Cohort Data, " forthcoming in Toshiaki Tachibanaki, ed., Savings and Bequests, University of Michigan Press.

Abowd, J. and D. Card (1989): "On the Covariance Structure of Earnings and Hours Changes, " *Econometrica*, 57, 411-45.

Alessie, R., Melenberg, B. and G. Weber (1989): "Consumption, Leisure and Earnings Related Liquidity Constraints: A Note, " Economics Letters, 27, 101-4.

Alessie, R., Devereux, M. and G. Weber (1997): "Intertemporal Consumption, Durables and Liquidity Constraints: A Cohort Analysis," *European Economics Review*, 41, 37-60.

Altonji, Joseph G., Hayashi, Fumio and Kotlikoff, Laurence J. "Is the Extended Family Altruistically Linked?" *American Economic Review*, 82, no .5 (December 1992), 1177-1198

Altonji, J.G. and A. Siow, (1987): "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data, " *Quarterly Journal of Economics* , 102, 293-328.

Arellano, M. and C. Meghir (1992): "Female Labour Supply and On The Job Search. An Empirical Model Estimated Using Complementary Data Sets" *Review of Economics Studies*, 59, 537-57.

Atkeson, A. G. and R.E. Lucas,(1992): "On Efficient Distribution with Private Information", *Review of Economic Studies*, 92, 427-53.

Atkeson, A. G. and M.Ogaki, (1996): "Wealth-Varying Intertemporal Elasticities of Substitution: Evidence from Panel and Aggregate Data", *Journal of Monetary Economics*, 38, 507-34.

Attanasio, O.P. (1994) "Personal Saving in the US", in J. Poterba (ed). *International Comparisons of Household Savings*, NBER, Chicago University Press.

Attanasio, O.P. (1995a): "Consumer Durables and Inertial Behavior: Estimation and Aggregation of *(S,s)* Rules", NBER Working Paper No. 5282.

Attanasio, O.P. (1995b): "The Intertemporal Allocation of Consumption", *Cernegie Rochester Conference on Public Policy*, Supplement to the *Journal of Monetary Economics*.

Attanasio, O.P. and J. Banks (1997): "A Tale of Two Countries", *Economic Policy*, *forthcoming*.

Attanasio, O.P., Banks, J., Meghir, C. and G. Weber (1996): "Humps and Bumps in Lifetime Consumption," NBER WP No. 5350.

Attanasio, O. P. and M. Browning, (1995): "Consumption over the Life Cycle and over the Business Cycle," *American Economic Review*, 85, 1118-1137.

Attanasio, O.P. and S.J. Davis (1996): "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, 104, 1227-1262.

Attanasio, O.P. and P.K. Goldberg (1997): "On The Relevance of Borrowing Restrictions. Evidence from Car Loans in the US", In progress.

Attanasio, O.P. and Hoynes, H. W. (1995): "Differential Mortality and Wealth Accumulation", NBER WP No. 5126.

Attanasio, O. and T.E. MaCurdy (1997): "Interactions in Household Labor Supply and their Implications for the Impact of EITC" Mimeo, Stanford University.

Attanasio, O. and G. Weber, (1993): "Consumption Growth, the Interest Rate and Aggregation," *Review of Economics Studies*, 60, 631-49.

Attanasio, O. and G. Weber, (1994): "The UK Consumption Boom of the Late 1980s : Aggregate Implications of Microeconomic Evidence ", *Economic Journal*, 104, 1269-1302.

Attanasio, O.P. and G. Weber, (1995): "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103, 1121-1157.

Avery, R.B. (1981): "Estimating Credit Constraints by Switching Regressions, " in C.F. Manski and D. McFadden (eds.) *Structural Analysis of Discrete Data with Econometric Applications*, MIT Press, Cambridge, Mass.

Banks, J., Blundell, R., and A. Brugiavini (1997): "Risk Pooling and Consumption Growth", University College London, Mimeo.

Banks, J., Blundell, R. and I. Preston (1994): "Life-Cycle Expenditure Allocations and the Consumption Cost of Children" *European Economic Review*, 38, 1234-78.

Banks, J., Blundell, R. and S. Tanner (1998): "Is There a Retirement Saving Puzzle?", *American Economic Review*, forthcoming.

Banks, J. and P.G. Johnson (1997) (eds): *How Reliable is the Family Expenditure Survey? Trends in Income and Expenditures over Time*, The Institute for Fiscal Studies, London.

Barro, R. (1974): "Are Government Bonds Net Wealth?, " *Journal of Political Economy*, 81, 1095-117.

Bean, C.R., (1986): "The Estimation of `Surprise' Models and the `Surprise' Consumption Function," *Review of Economics Studies*, 53, 497-516.

Beaulieu, J.J. (1992): "Aggregate Movements in Durable Goods with Fixed Costs of Adjustment", Miemo, MIT.

Beaulieu, J.J. (1993): "Optimal Durable and Non-Durable Consumption with Transaction Costs", Federal Reserve Board Washington, Finance and Economics Discussion Series, 93-12.

Bernanke, B. (1984): " Permanent Income, Liquidity, and Expenditure on Automobiles: Evidence From Panel Data", *Quarterly Journal of Economics*, 99, 587-614.

Bernanke, B. (1985): "Adjustment Costs, Durables, and Aggregate Consumption", *Journal of Monetary Economics*, 15, 41-68.

Bertola, G. and R.J. Caballero (1990): "Kinked Adjustment Costs and Aggregate Dynamics", in O.J.Blanchard and S. Fischer (eds) *NBER Macroeconomics Annual*, Vol 5.

Blinder A. and A. Deaton (1985): " The Time-Series Consumption Revisited, " Brookings Papers on Economic Activity, 465-521.

Blundell, R., Browning, M. and C. Meghir, (1994): "Consumer Demand and the Lifetime Allocation of Consumption," *Review of Economics Studies* , 61, 57-80.

Blundell, R. , Pashardes, P. and G. Weber (1993): "What Do We Learn about Consumer Demand Patterns from Micro Data?", *American Economic Review*, 83, 570-597.

Blundell, R. and I. Preston (1996): "Consumption Inequality and Income Uncertainty", UCL Discussion Paper No. 96-07

Blundell, R. and T. Stoker (1995): "Consumption and the Timing of Income Risk". Mimeo.

Browning, M. (1987): "Eating, Drinking, Smoking, and Testing the Lifecycle Hypothesis", *Quarterly Journal of Economics;*102, 329-45..

Browning, M. (1991): "A Simple Nonadditive Preference Structure for Models of Household Behavior over Time" , *Journal of Political Economy*, 99, 607-637.

Browning, M. and A. Lusardi, (1996): "Household Saving: Micro Theories and Micro Facts", *Journal of Economic Literature*, 34, 1797-1855.

Browning M. and T. Crossley (1997) : "Socks and Stocks", McMaster University, Mimeo.

Browning, M., Deaton, A. and M. Irish, (1985): "A Profitable Approach to Labor Supply and Commodity Demand over the Life Cycle ", *Econometrica*, 53, 503-44.

Browning, M. and C. Meghir (1991): "The Effect of Male and Female Labour Supply on Commodity Demands, " *Econometrica*, 59, 925-51.

Brugiavini, A. and G. Weber (1992): "Durable and Non Durable Consumption: Evidence from Italian Household Data, " in A.Ando, L. Guiso and I. Visco (eds) : *Saving and the Accumulation of Wealth*, 305-329, Cambridge University Press.

Caballero, R.J. (1990a): "Consumption Puzzles and Precautionary Savings", *Journal of Monetary Economics*, 25, 113-136.

Caballero, R.J. (1990b): "Expenditure on Durable Goods: A Case for Slow Adjustment", *Quarterly Journal of Economics* , 105, 727-43.

Caballero, R.J. (1991): "Earnings Uncertainty and Aggregate Wealth Accumulation" , *American Economic Review*, 81, 859-871.

Caballero, R.J. (1993): "Durable Goods: An Explanation for Their Slow Adjustment", *Journal of Political Economy*, 101, 351-384.

Caballero, R.J. and E. Engel (1991): "Dynamic (S,s) Economies", *Econometrica* , 59, 1659-86.

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 A.S. Deaton (1989): "Why Is Consumption So Smooth?" *Review of Economic Studies*, 56, 357-74.

Campbell, J.Y. and N.G. Mankiw (1989): ."Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence, " in O.J. Blanchard and S. Fischer (eds.) *NBER Macroeconomics Annual* 1989 , MIT Press, Cambridge, Mass., 185-216.

Campbell, J.Y. and N.G. Mankiw (1991): "The Response of Consumption to Income: A Cross-Country Investigation", *European Economic Review*, 35, 715-21.

Caplin, A.S. (1985): "The Variability of Aggregate Demand with (S,s) Inventory Policies", *Econometrica*, 53, 1395-1409.

Caplin, A.S. and D. Spulber (1987): "Menu Costs and The Neutrality of Money", *Quarterly Journal of Economics*, 102, 703-726.

Carroll, C. D. (1994): "How Does Future Income Affect Current Consumption?", *Quarterly Journal of Economics*;109,111-47.

Carroll, C. (1997a): "Buffer Stock Saving and the Life Cycle -Permanent Income Hypothesis, " *Quarterly Journal of Economics*, 112, 1-55.

Carroll, C. (1997b): "Death to the log-linearised Euler equation", Mimeo.

Carroll, C.D., J.C. Fuhrer and D. W. Wilcox, (1994): "Does Consumer Sentiment Forecast Household Spending? If So Why?", *American Economic Review*, 84, 1397-1408.

Carroll, C. and M. Kimball (1996): "On the Concavity of the Consumption Function", *Econometrica*, 64, 981-992.

Carroll C. and L.H. Summers (1991): "Consumption Growth Parallels Income Growth: Some New Evidence, " in B.D. Bernheim and J.B. Shoven (eds.), *National Saving and Economic Performance*, Chicago University Press for NBER, Chicago, 305-43.

Chah, E.Y., Ramey, V.A. and R.M. Starr (1995) : "Liquidity Constraints and Intertemporal Consumer Optimization: Theory and Evidence from Durable Goods, " *Journal of Money Credit and Banking*, 27, 272-287.

Chamberlain, G. (1984): "Panel Data, " in Z. Griliches and M.D. Intriligator (eds.), *Handbook of Econometrics* , Vol. II, North Holland, Amsterdam, 1247-318.

Christiano, L. and M. Eichenbaum (1990): "Unit Roots in Real GNP: Do We Know and Do We Care?", *Carnegie Rochester Conference Series on Public Policy*, 32, 7-61.

Cochrane, J. H. (1991): "A Simple Test of Consumption Insurance. " *Journal of Political Economy*, 99, 957-976.

Constantinides, G. M. (1982) "Intertemporal Asset Pricing with Heterogeneous Consumers and without Demand Aggregation," *Journal of Business*, 55 (April), 253-267.

Constantinides, G.M. (1986): "Capital Market Equilibrium with Transaction Costs", *Journal of Political Economy*, 94, 842-862.

Constantinides, G.M. (1990): "Habit Formation: A Resolution of the Equity Premium Puzzle", *Journal of Political Economy*, 98, 519-543.

Constantinides, G.M. and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers", *Journal of Political Economy*, 104, 219-240.

Cutler, D.M. and Katz, L.F. (1991): "Macroeconomic Performance and the Disadvantaged," *Brookings Papers on Economic Activity*, 1991:1 , 1-74.

Davidson, J., Hendry, D., Srba, F. and S. Yeo (1978): "Econometric Modelling of the Aggregate Time-Series Relationship Between Consumers' Expenditure and Income in the UK", *Economic Journal*, 88, 661-92.

Davies, J. B. (1981): "Uncertain Lifetime, Consumption, and Dissaving in Retirement" *Journal of Political Economy*; 89, 561-77.

Deaton, A. (1985): "Panel Data from Time Series of Cross Sections," *Journal of Econometrics*, 30, 109-26.

Deaton, A. (1991): "Saving and Liquidity Constraints," "*Econometrica*, 59, 1221-48.

Deaton, A. (1992): *Understanding Consumption* , Oxford University Press, Oxford.

Deaton, A. and C. Paxson (1994): "Intertemporal Choice and Inequality", *Journal of Political Economy*, .

Diamond, P. (1967): "The role of a Stock Market in a General Equilibrium Model with Technological Uncertainty", *American Economic Review*, 57, 759-776.

Dreze, J. and F. Modigliani (1972) : "Consumption Decisions under Uncercainty", *Journal of Economics Theory*, 5, 308-335.

Dunn, K. B. and K. J. Singleton, (1986): "Modeling the Term Structure of Interest Rates under Non-separable Utility and Durability of Goods", *Journal of Financial Economics*, 17, 27-55.

Dynan, K.E. (1993): "How Prudent Are Consumers?", *Journal of Political Economy*, 101, 1104-1113.

Eberly, J. (1994): "Adjustment of Counsumers' Durables Stocks: Evidence from Automobile Purchases", *Journal of Political Economy*, 102, 403-436.

Eichenbaum, M. and Hansen, L.P. (1990) "Estimating Models with Intertemporal Substitution Using Aggregate Time Series Data", *Journal of Business & Economic Statistics*, 8, 53-69.

Eichenbaum, M. , Hansen, L.P. and K.J. Singleton, (1988) "A Time Series Analysis of Representative Agent Models of Consumption and Leisure under Uncertainty", *The Quarterly Journal of Economics*, 103, 51-78.

Epstein,L.G. and S.E. Zin, (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework", *Econometrica*, 57, 937-69.

Epstein, L.G. and S.E. Zin, (1991): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis", *Journal of Political Economy*, 99, 263-86.

Flavin, M. (1981): "The Adjustment of Consumption to Changing Expectations About Future Income," *Journal of Political Economy*, 89, 974-1009.

Flemming, J.S. (1973): "The Consumption Function when Capital Markets are Imperfect, The Permanent Income Hypothesis Reconsidered, " Oxford Economic Papers , 25, 160-72.

Friedman, M. (1957): A *Theory of the Consumption Function*, Princeton University Press, Princeton.

Garcia, R., Lusardi, A. and S. Ng (1997): "Excess Sensitivity and Asymmetries in Consumption", *Journal of Money, Credit and Banking, forthcoming*.

Ghez, G. and G.S. Becker (1975): *The Allocation of Time and Goods over the Life Cycle*, NBER, Columbia University Press.

Gokhale, J., Kotlikoff, L. J. and J. Sabelhaus,(1996) "Understanding the Postwar Decline in U. S. Saving: A Cohort Analysis", *Brookings Papers on Economic Activity*, 1996(1), 315-90.

Goodfriend, M. (1992): "Information Aggregation Bias," *American Economic Review*, 82, 508-519.

Goodman, A., Johnson, P. and S. Webb (1997): *Inequality in the UK*, London.

Gorman, W.M. (1959): "Separable Utility and Aggregation", *Econometrica*, 27, 469-481.

Gorman, W.M. (1967): "Tastes, Habits and Choices", *International Economic Review*, 8, 218-22.

Grossman S.J. and G. Laroque (1990) "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods", *Econometrica*, 58, 25-51.

Greenspan A. and D. Cohen (1996): "Motor Vehicle Stocks, Scrappage and Sales" Federal Reserve Board Dinance and Economics Discussion Paper, No 96-40.

Guiso, L., Jappelli, T. and D. Terlizzese (1991): "Why is Italy's Saving So High?" in Ando, A., Guiso, L. and I. Visco (eds): *Saving and the accumulation of wealth*, Cambridge University Press.

Guiso, L., Jappelli, T. and D. Terlizzese (1992): "Earnings Uncertainty and Precautionary Savings", *Journal of Monetary Economics*, 30, 307-37.

Guiso, L., Jappelli, T. and D. Terlizzese (1996): "Income Risk, Borrowing Constraints and Portfolio Choice", *American Economic Review*, 86, 158-72.

Hall, R.E., (1978): "The Stochastic Implications of the Life Cycle- Permanent Income Hypothesis: Theory and Evidence, "*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: Estimate from Panel Data on Households, "*Econometrica*, 50, 461-81.

Hansen, G. and A. Imrohoroglu (1992): "The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard, " *Journal of Political Economy*, 100, 118-42.

Hansen, L.P. and R. Jagannathan (1991): "Implications of Security Market Data for Models of Dynamic Economics", *Journal of Political Economy*, 99, 225-262.

Hansen, L. P. and K.J. Singleton (1982): "Generalized instrumental variables estimation of nonlinear rational expectation models, "*Econometrica*, 50, 1269-86.

Hansen, L.P. and K.J. Singleton (1983): "Stochastic Consumption, Risk Aversion and the Temporal Behavior of Asset Returns", *Journal of Political Economy*, 91, 249-265.

Hayashi, F. (1982): "The Permanent Income Hypothesis: Estimation and Testing by Instrumental Variables, "*Journal of Political Economy*, 90, 895-916.

Hayashi, F. (1985a): "The Effects of Liquidity Constraints on Consumption: A Cross Section Analysis," *Quarterly Journal of Economics*, 100, 183-206.

Hayashi, F. (1985b): "The Permanent Income Hypothesis and Consumption Durability: Analysis Based On Japanese Panel Data," *Quarterly Journal of Economics*, 100, 1083-1116.

Hayashi, F.(1987): "Tests for Liquidity Constraints: A Critical Survey and Some New Observations, " in T.F. Bewley (ed.), *Advances in Econometrics: Fifth World Congress*, Vol II, 91-120.

Hayashi, F. (1996): "A Review of Empirical Studies of Household Saving" , Mimeo, Columbia University.

Hayashi, F., J. Altonji and L. Kotlikoff, (1996): "Risk-Sharing Between and Within Families", *Econometrica*, 64, 261-294.

Hayashi, F. and C. Sims (1983): "Nearly Efficient Estimation of Time Series Models with Predetermined but not Exogenous Instruments", *Econometrica*, 51, 783-792.

Heaton, J. (1993): "The Interaction between Time-Nonseparable Preferences and Time Aggregation", *Econometrica*, 61, 353-386.

Heaton, J. (1995): "An Empirical Investigation of Asset Pricing with Temporally Dependent Preferences", *Econometrica*, 63, 681-718.

Heaton, J. and D. Lucas (1996): "Evaluating the Effects of Incomplete Markets for Risk Sharing and Asset Prices", *Journal of Political Economy*, 104, 443-487.

Heckman, J.J. (1974): "Life-cycle Consumption and *labour* Supply: An Exploration of the Relationship Between Income and Consumption Over the Life Cycle, " *American Economic Review*, 64, 188-94.

Heckman, J.J. and T. MaCurdy (1988): "Empirical Tests of Labor Market Equilibrium: An Evaluation", Carnegie Rochester Conference Series, 28, 231-58.

Heckmann, J.J. and R. Robb (1987): "Using Longitudinal Dataa to Estimate Age Period and Cohort Effects in Earnings Equations", in W.M.Mason and S.E. Fienberg (eds): *Cohort Analysis in Social Research*, Springer and Verlag.

Hotz, V.J., Kydland, F.E. and G.J. Sedlacek (1988): "Intertemporal Preferences and Labor Supply", *Econometrica*, 56, 335-360.

Houthakker, H.S. and L.D. Taylor (1970): *Consumer Demand in the United States: Analyses and Projections*, 2d ed. Cambridge (MA); Harvard Univ. Press.

Hubbard, G. and K. Judd (1986): "Liquidity Constraints, Fiscal Policy, and Consumption, " Brookings Papers on Economic Activity, 1-50.

Hubbard, G., Skinner, J. and S. Zeldes, (1993): "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving ", NBER WP. No. 4516.

Hubbard, G., Skinner, J. and S. Zeldes, (1995): "Precautionary Saving and Social Insurance", *Journal of Political Economy*, 103, 360-399.

Hurd, M. D. (1989): "Mortality Risk and Bequests" *Econometrica*;57, 779-813.

Imrohoroglu, A. (1989): "Cost of Business Cycle with Indivisibilities and Liquidity Constraints", *Journal of Political Economy*, 100, 118-42.

Imrohoroglu, A., Imrohoroglu, S. and D.H. Joines (1993): "A Numerical Algorithm for Solving Models with Incomplete Markets," *The International Journal of Supercomputer Applications*, 7, 212-30.

Jackman R. and J. Sutton (1982): "Imperfect Capital Markets and the Monetarist Black Box: Liquidity Constraints and the Assymetric Effects of Interest Rate Policy," *Economic Journal*, 92, 108-28.

Jappelli, T. (1990): "Who is Credit Constrained in the US Economy?" *Quarterly Journal of Economics*, 105, 219-34.

Jappelli, T. and F. Modigliani (1997): "Is The Age - Saving Profile Consistent with The Life Cycle Hypothesis?", Mimeo.

Jappelli, T. and M. Pagano (1989): "Consumption and Capital Market Imperfections: An International Comparison, " *American Economic Review*, 79, 1088-1105.

Jappelli, T. and M. Pagano (1994): "Saving, Growth and Liquidity Constraints" *Quarterly Journal of Economics*, 109, 83-110.

Jappelli, T., J.-S. Pischke and N. Souleles, (1997): "Testing for Liquidity Constraints Using Complementary Data Sources", *Review of Economics and Statistics, forthcoming.*

Juster, F.T. and R.P. Shay (1964): "Consumer Sensitivity to Finance Rates: An Empirical and Analytical Investigation," NBER Occasional Paper No. 88.

Keane, M. and D. Runkle (1992): "On the Estimation of Panel-Data Models with Serial Correlation when Instruments are not Strictly Exogenous," *Journal of Business and Economic Statistics,* 10, 1-9.

Kennan, J. (1988): "An Econometric Analysis of Fluctuations in Aggregate Labor Supply and Demand", *Econometrica,* 56, 317-334.

Kimball, M. S. (1990): "Precautionary Saving in the Small and in the Large" *Econometrica;* 58, 53-73.

Lam, P.S. (1991): "Permanent Income, Liquidity and Adjustment of Automobile Stocks: Evidence from Panel Data", *Quarterly Journal of Economics,* 106, 203-230.

Lucas, R. E. Jr. (1978): "Asset Prices in an Exchange Economy" *Econometrica,* 46, 1429-45.

Lucas, R. E. Jr. (1987): *Models of Business Cycles,* New York: Blackwell.

Ludvigson, S. (1996): "Consumption and Credit: A Model of Time-varying Liquidity Constraints", Mimeo, Federal Reserve Bank of New York.

Ludvigson, S. and C. Paxson (1997): "Approximation Bias in Linearized Euler Equations", Mimeo, Princeton University.

Lusardi, A. (1996): "Permanent Income, Current Income and Consumption: Evidence from Two Panel Data Sets, "*Journal of Business Economics and Statistics,* 14, 81-90.

Mace, B.J. (1991): "Full Insurance in the Presence of Aggregate Uncertainty, "*Journal of Political Economy,* 99, 928-956.

MaCurdy , T.E. (1981): "An Intertemporal Model of Portfolio Choice and Human Capital Accumulation Under Uncertainty with Extentions Incorporating Taxes, Consumer Durables, Imperfections in Capital Markets and Non Separable Preferences ", Hoover Institution, WP E-81-18.

MaCurdy, T.E. (1982): "The Use of Time Series Processes to Model the Error Structure of Earnings in Longitudinal Data Analysis," *Journal of Econometrics,* 18, 83-114.

MaCurdy, T.E. (1983): "A Simple Scheme for Estimating an Intertemporal Model of Labor Supply and Consumption in the Presence of Taxes and Uncertainty," *International Economic Review,* 24, 265-89.

MaCurdy, T.E. and T. Mroz (1989): "Measuring Macroeconomic Shifts in Wages From Cohort Specifications", Mimeo Stanford University.

Maki, D. (1994a): "Consumer Debt and the Tax Reform Act of 1986", Mimeo, Stanford University.

Mankiw, N.G. (1982): "Hall's Consumption Hypothesis and Durable Goods", *Journal of Monetary Economics*, 10, 417-25.

Mankiw, N.G., Rotemberg, J.J., and L.H. Summers (1985): "Intertemporal Substitution in Macroeconomics", *Quarterly Journal of Economics*, 100, 225-53.

Mankiw, N.G. and S. Zeldes (1994): "The Consumption of Stock Owners and the Consumption of Non Stock Owners", *Journal of Financial Economics*.

Meghir, C. and G. Weber (1993): "Intertemporal Non-Separability or Borrowing Restrictions? A disaggregated analysis using a US Panel, "Mimeo, IFS.

Mehra, R. and E.C. Prescott (1985): "The Equity Premiuim Puzzle", *Journal of Monetary Economics*, 15, 145-161.

Miles D. (1992): "Housing Markets, Consumption, and Financial Liberalisation in the Major Economies," *European Economic Review*, 36, 1093-1127.

Modigliani, F. and A. Ando (1963): "The Life Cycle Hypothesis of Saving: Aggregate Implications and Tests", *American Economic Review*, 53, 55-84.

Modigliani, F. and R. Brumberg (1954): "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data, " in K.K. Kurihara (ed.), *Post-Keynesian Economics*, Rutgers University Press, New Brunswick, N.J., 128-97.

Moffitt, R. (1993) : "Identification and Estimation of Dynamic Models with a Time Series of Repeated Cross-Sections" , *Journal of Econometrics*, 59, 99-124.

Muellbauer, J. (1988): "Habits, Rationality and Myopia in the Life Cycle Consumption Function," *Annales d'Economie et de Statistique*, 9, 47-70.

Muellbauer, J. (1996): 'Consumer Expenditure', in T Jenkinson (ed.), *Readings in Macroeconomics*. Oxford: Oxford University Press, 1996.

Muellbauer, J. and R. Lattimore (1994) "The Consumption Function: A Theoretical and Empirical Overview" in H Pesaran and M Wickens (eds.) *Handbook of Applied Econometrics*. Oxford: Blackwell, 1995.

Muellbauer, J. and A. Murphy (1990): "Is the U.K. Balance of Payments Sustainable?" *Economic Policy*, 5, 347-95.

Novales, A. (1990): "Solving Nonlinear Rational Expectations Models: A Stochastic Equilibrium Model of Interest Rates", *Econometrica*, 58, 93-111.

O'Brien A.M. and C.B Hawley (1986): "The Labor Force Participation Behavior of Married Women Under Conditions of Constraints on Borrowing," *The Journal of Human Resources*, 21, 267-78.

Paulin, G., Boyle, M., Brauch, R. and R. Cage (1990): "Comparison of 1987-1988 CE Intewgrated Survey and PCE Annual Estimates", BLS Mimeo.

Pencavel, J. (1986): "Labour Supply of Men: A Survey" in O.C. Ashenfelter and R. Layard (eds.), *Handbook of Labor Economics*, Vol I, Elsevier Science Publishers, Amsterdam, 1-102.

Pischke, J.-S. (1995): "Individual Income, Incomplete Information, and Aggregate Consumption" , *Econometrica,* 63, No. 4 (July 1995), 805-40.

Pissarides, C. (1978): "Liquidity Considerations in the Theory of Consumption, " Quarterly Journal of Economics, 92, 279-296.

Pollak, R.A. (1970): "Habit Formation and Dynamic Demand Functions", *Journal of Political Economy,* 78, 745-63.

Rosenzweig , M. and K. Wolpin (1993): "Credit Market Constraints, Consumption Smoothing and Accumulation of Durable Production Assets in Low Income Countries: Investment in Bullocks in India", *Journal oif Political Economy,* 101, 223-244.

Runkle, D. (1991): "Liquidity Constraints and the Permanent Income Hypothesis, " Journal of Monetary Economics, 27, 73-98.

Sargent, T.J. (1978): "Rational Expectations, Econometric Exogeneity and Consumption", *Journal of Political Economy,* 86, 673-700.

Scheinkman, J. and A. Weiss (1986): "Borrowing Constraints and Aggregate Economic Activity, " *Econometrica,* 54, 23-45.

Seslnick, D. (1992) : "Aggregate Consumption and Saving in the Postwar United States" , *Review of Economics and Statistics, 585-97,*

Seslnick, D. (1994) : "Are our Data Relevant to the Theory? The Case of Aggregate Consumption' , Mimeo.

Shea, J. (1995): "Union Contracts and the Life-Cycle / Permanent- Income Hypothesis", *American Economic Review,* 85, 186-200 .

Skinner, J. (1988): "Risky Income. Life Cycle Consumption and Precautionary Savings", *Journal of Monetary Economics,* 22, 237-255.

Spynnewin, F. (1981) :"Rational Habit Formation", *European Economic Review,* 15, 91-109.

Stiglitz, J. and A. Weiss (1981): "Credit Rationing in Markets with Imperfect Information, " *American Economic Review,* 71, 393-410.

Stoker, T. (1993): "Empirical Approaches to the Problem of Aggregation over Individuals ", *Journal of Economic Literature* , 31, 1827-74.

Thurow, L. (1969): "The Optimum Lifetime Distribution of Consumption Expenditures, " *American Economic Review,* 59, 371-96.

Tobin , J. (1980): Asset Accumulation and Economic Activity , Basil Blackwell, Oxford.

Tobin, J. and W. C. Dolde (1971): "Wealth, Liquidity and Consumption, " in Reserve Bank of Boston, Monetary Conference Series No. 5.

Townsend, R.M. (1994): "Risk and Insurance in Village India", *Econometrica,* 62, 539-592.

Weber, G. (1993): "Earnings Related Borrowing Restrictions: Empirical Evidence from a Pseudo Panel for the UK", *Annales d' Economie et de Statistique,* 29, 157-173.

Wilson, R. (1968): "The Thoery of Syndicates", *Econometrica*, 36, 119-132.

Zeldes, S. (1989): "Consumption and Liquidity Constraints: An Empirical Analysis, " *Journal of Political Economy*, 97, 305-46.

Zeldes, S. (1989b): "Optimal Consumption with Stochastic Income", *Quarterly Journal of Economics*, 104, 275-298.

Table 1
Mean and Standard deviations
(annual rates of growth)

	US		UK	
	Mean	St.Dev.	Mean	St.Dev.
Disposable Income	0.032	0.025	0.026	0.026
Nondurable consumption	0.023	0.018	0.017	0.021
Durable Expenditure	0.048	0.069	0.043	0.112

Table 2
MA(12) representation

	US			UK		
	Disposable Income	Non-Durable Consumption	Durable Consumption	Disposable Income	Non-Durable Consumption	Durable Consumption
θ_1	-0.30 (.091)	0.41 (.096)	-0.092 (.088)	-0.10 (.094)	0.005 (.094)	-0.29 (.094)
θ_2	0.15 (.094)	0.18 (.103)	-0.035 (.089)	0.12 (.095)	0.20 (.093)	-0.14 (.096)
θ_3	0.092 (.094)	0.43 (.104)	0.063 (.089)	-0.06 (.092)	0.004 (.093)	0.24 (.097)
θ_4	-0.092 (.088)	0.12 (.110)	0.084 (.086)	-0.18 (.088)	0.28 (.092)	-0.45 (.099)
θ_5	-0.15 (.087)	-0.057 (.108)	-0.16 (.085)	-0.19 (.089)	0.19 (.093)	0.15 (.106)
θ_6	0.11 (.088)	.100 (.108)	0.15 (.077)	0.22 (.088)	0.19 (.094)	0.05 (.107)
θ_7	-0.13 (.087)	0.11 (.107)	-0.45 (.077)	0.21 (.088)	0.09 (.094)	-0.07 (.106)
θ_8	-0.17 (.088)	-0.20 (.107)	-0.021 (.085)	0.14 (.087)	0.22 (.092)	-0.18 (.104)
θ_9	0.38 (.088)	0.05 (.109)	-0.23 (.085)	-0.14 (.086)	-0.11 (.090)	-0.08 (.098)
θ_{10}	0.20 (.095)	-0.03 (.100)	-0.03 (.088)	-0.20 (.087)	0.23 (.091)	0.02 (.095)
θ_{11}	-0.06 (.096)	0.05 (.099)	0.005 (.087)	0.05 (.088)	0.18 (.091)	-0.20 (.094)
θ_{12}	-0.27 (.091)	0.08 (.092)	-0.23 (.086)	-0.05 (.086)	0.03 (.093)	-0.02 (.089)
$\sum_{i=1}^{12} \theta_i$	-0.25	1.23	-0.95	-0.18	1.51	-0.97

Table 3
ARMA representation

	US			UK		
	Disposable Income	Non-Durable	Durable Consumption	Disposable Income	Non-Durable Consumption	Durable Consumption
ψ_1	–	–	-0.103 (.089)	-0.77 (.293)	0.015 (.087)	-1.09 (.098)
ψ_2	–	–	–	–	0.28 (.087)	-0.45 (.082)
θ_1	-0.19 (.088)	0.38 (.083)	–	0.684 (.339)	–	0.85 (.077)
θ_2	–	0.18 (.088)	–	–	–	–
θ_3	–	0.39 (.082)	–	0.684 (.339)	–	0.85 (.077)
Q-stat (p-value)	13.40 (0.10)	7.35 (0.28)	10.09 (0.26)	11.79 (0.11)	11.49 (0.12)	7.22 (0.30)

Note: Sample 1965:3 - 1996:3 (125 observations). Standard errors in brackets.

		US		UK	
		Table 4			
		Cohort definition and cell size			
cohort	year of birth	average cell size	years in the sample	average cell size	years in the sample
1	1895-99	-	-	338	68-77
2	1900-04	-	-	459	68-82
3	1905-09	-	-	526	68-87
4	1910-14	232	80-92	560	68-92
5	1915-19	390	80-92	519	68-92
6	1920-24	333	80-92	653	68-92
7	1925-29	325	80-92	572	68-92
8	1930-34	317	80-92	546	68-92
9	1935-39	345	80-92	562	68-92
10	1940-44	420	80-92	594	68-92
11	1945-49	566	80-92	652	68-92
12	1950-54	657	80-92	547	71-92
13	1955-59	734	80-92	508	76-92
14	1960-64	-	-	463	81-92
15	1965-69	-	-	334	86-92

Table 5
Variability of Consumption and Income

	United States (CEX)	United Kingdom (FES) age<81	United Kingdom (FES) 10 cohorts, year<86
	Standard Error (%)	Standard Error (%)	Standard Error (%)
Total Consumption	2.94	2.46	2.65
Total Consumption per adult equivalent	2.39	2.62	2.64
Non-durable Consumption	2.60	2.30	1.88
Non-durable cons. per adult equivalent	1.95	2.49	2.05
Durable Consumption	15.79	9.54	8.54
Non-durable Consumption (from levels)	2.58	2.31	1.86
Income	3.68	3.05	3.60

Table 6:
Variability of Consumption and Income by educational group

	High school dropouts	High school graduates	More than high school
Total consumption	2.88	2.74	2.88
Non-durable Consumption	2.40	2.93	2.25
Durable Consumption	22.85	16.58	9.66
Income	5.98	6.53	5.17

US - disposable income and consumption

UK - disposable income and consumption

Figure 1

US - disposable income and consumption rates of growth

UK - disposable income and consumption rates of growth
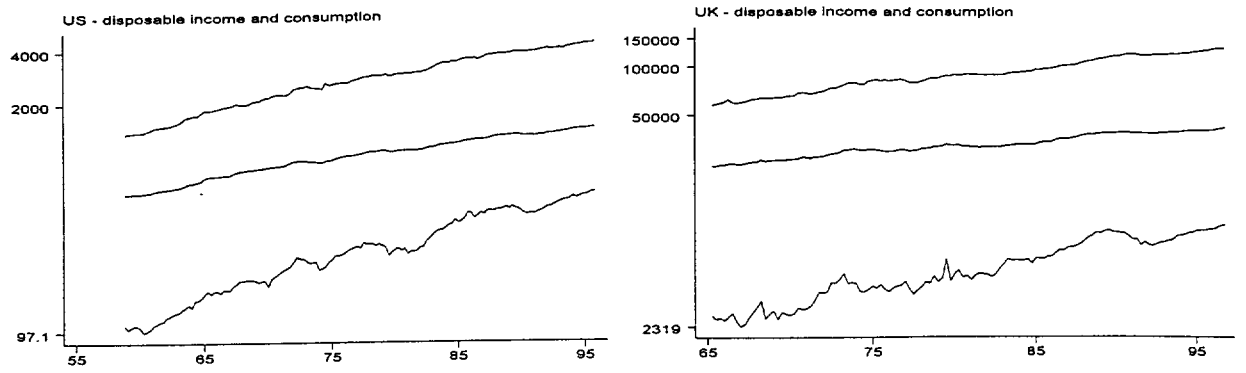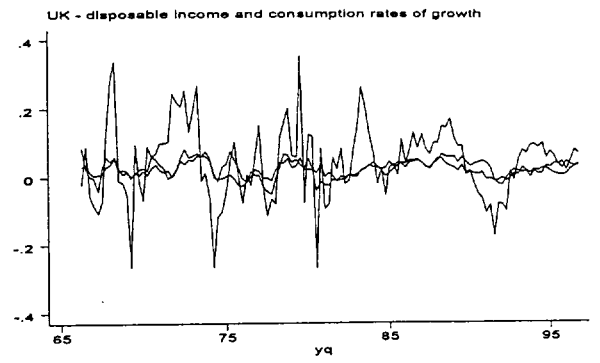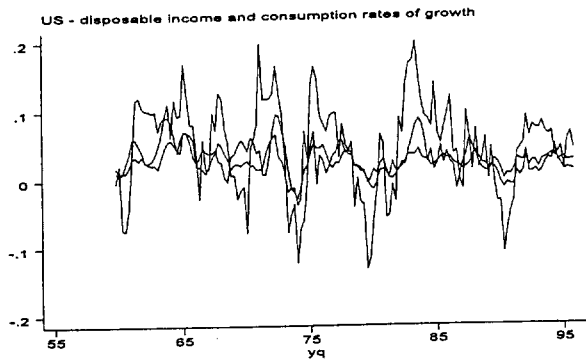
Figure 2

Figure 3

Figure 4

US - Non-durable consumption
per household and per adult equivalent

UK - Non-durable consumption
per household and per adult equivalent
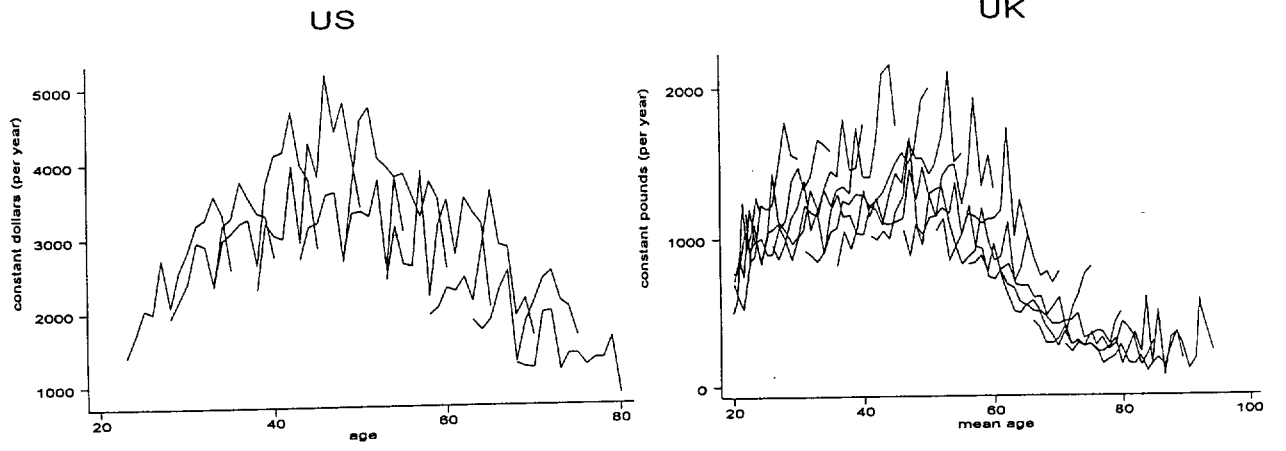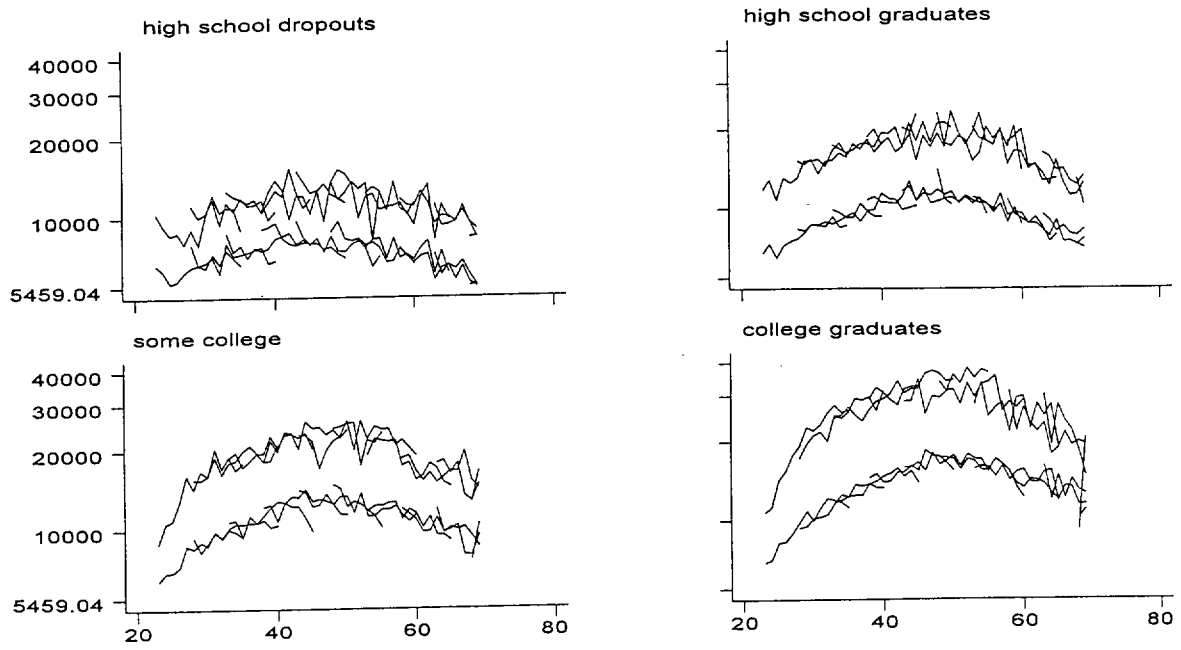
Figure 5

Durables

US

UK
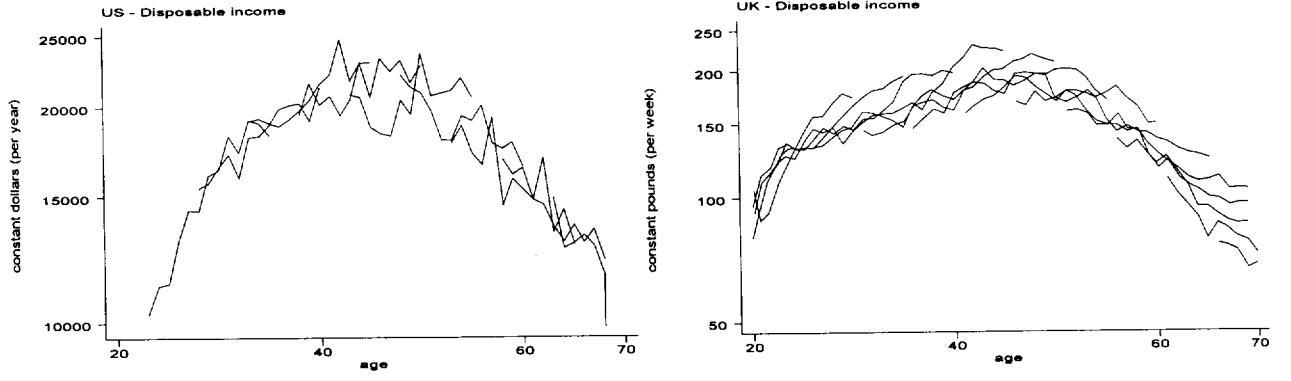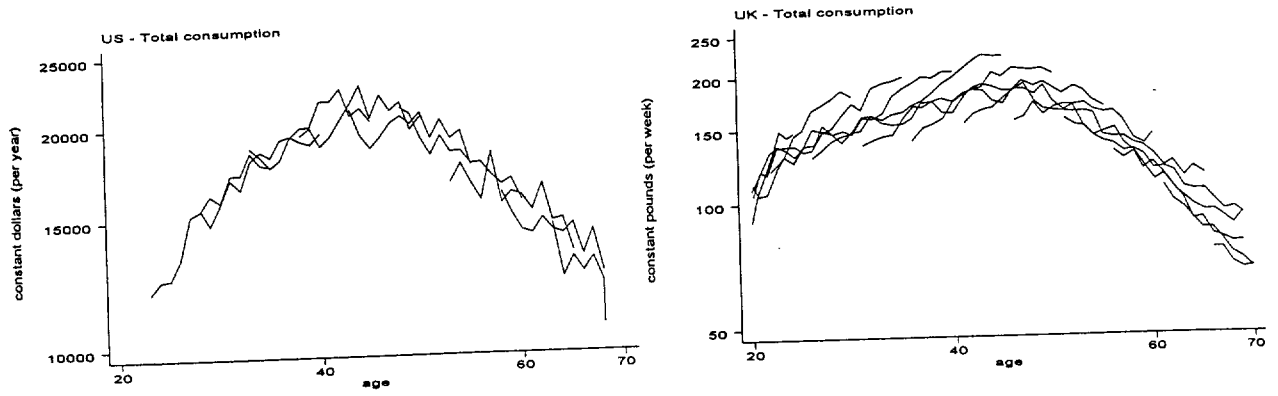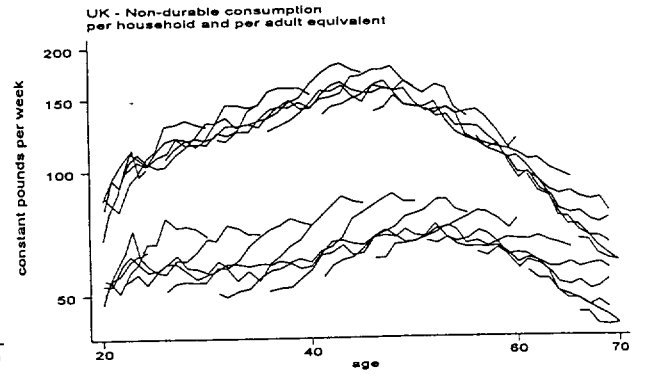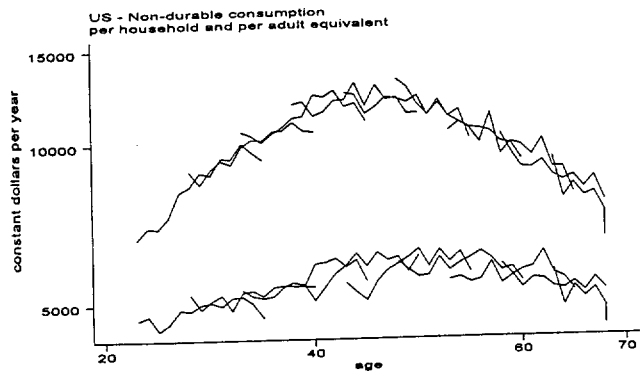
Figure 6

log income and non durable consumption by education

Figure7