

NBER WORKING PAPER SERIES

IS CONSUMPTION GROWTH
CONSISTENT WITH INTERTEMPORAL
OPTIMIZATION? EVIDENCE FROM THE
CONSUMER EXPENDITURE SURVEY

Orazio P. Attanasio
Guglielmo Weber

Working Paper No. 4795

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 1994

An early and very different version of this paper was circulated at the end of 1992 and presented at the NBER Economic Fluctuations Conference in Cambridge, October 1992. We are grateful to John Cochrane, who discussed the paper, for very insightful comments. Several participants also gave us very good suggestions. This paper was completed while the first author was a National Fellow at the Hoover Institution. Attanasio's research was partially financially supported by NSF grant SES-9057511. This paper is part of NBER's research program in Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

IS CONSUMPTION GROWTH
CONSISTENT WITH INTERTEMPORAL
OPTIMIZATION? EVIDENCE FROM THE
CONSUMER EXPENDITURE SURVEY

ABSTRACT

In this paper we show that some of the predictions of models of consumer intertemporal optimization are not inconsistent with the patterns of non-durable expenditure observed in US household-level data. Our results and our approach are new in several respects.

First, we use the only US micro data set which has direct and complete information on household consumption. The microeconomic data sets used in most of the consumption literature so far contained either very limited information on consumption (like the PSID) or none at all, in which case consumption had to be obtained indirectly from income and changes in assets.

Second, we propose a flexible and novel specification of preferences which is easily estimable and allows a general treatment of multiple commodities. We show that a proper treatment of aggregation over commodities can be important, both theoretically and in practice.

Third, we present empirical results that show that it is possible to find a reasonably simple specification of preferences, which controls for the effects of changes in demographics and labor supply behavior over the life cycle and which is not rejected by the available data. On our preferred specification, we obtain sharp estimates of key behavioral parameters (including the elasticity of intertemporal substitution) and no rejections of theoretical restrictions.

Our results contrast sharply with most of the previous evidence, which has typically been interpreted as rejection of the theory. We show that previous rejections can be explained by the simplifying assumptions made to derive empirically tractable equations. We also show that results obtained using food consumption or aggregate data can be extremely misleading.

Orazio P. Attanasio
Department of Economics
Stanford University
Stanford, CA 94305
and NBER

Guglielmo Weber
Dipartimento di Scienze Economiche
Università di Venezia
Ca' Foscari
Venezia, ITALY

1. Introduction

In this paper we show that some of the predictions of models of consumer intertemporal optimization are not inconsistent with the patterns of non-durable expenditure observed in US household-level data. Our results and our approach are new in several respects.

First, we use the only US micro data set which has direct and complete information on household consumption. The microeconomic data sets used in most of the consumption literature so far contained either very limited information on consumption (like the PSID) or none at all, in which case consumption had to be obtained indirectly from income and changes in assets.

Second, we propose a flexible and novel specification of preferences which is easily estimable and allows a general treatment of multiple commodities. We show that a proper treatment of aggregation over commodities can be important, both theoretically and in practice.

Third, we present empirical results that show that it is possible to find a reasonably simple specification of preferences, which controls for the effects of changes in demographics and labor supply behavior over the life cycle and which is not rejected by the available data.

Our results contrast sharply with most of the previous evidence, which has typically been interpreted as rejection of the theory. Several papers in the last 15 years have used microeconomic data to establish the plausibility of the intertemporal optimization model. Hall and Mishkin (1982), for instance, found that consumption is excessively sensitive to lagged labor income. Zeldes (1989) finds excess sensitivity for low wealth households and suggest that liquidity constraints are important for these individuals. ¹ Results based on aggregate time series data, such as those in Flavin (1981) and Campbell and Mankiw (1989), suggest even stronger rejections.

In this paper we interpret previous rejections as evidence against some simplifying assumptions made to derive empirically tractable equations. In particular, we show that:

a) aggregate data are particularly unsuitable to test the theory. Incorrect aggregation can lead to spurious rejections of the theory. Even in the simple case of isoelastic utility aggregation can be troublesome, as the theory requires knowledge of the logarithm of consumption. We use our household-level data to assess the consequences of neglecting the non-linearities implied by the model - and find that theory restrictions can be rejected just because of an incorrect aggregation procedure;

¹ See also Carroll and Summers (1991) who observe that consumption tracks income over the life cycle. The papers by Runkle (1991), and Keane and Runkle (1992) are exceptions. Altonji and Siow (1987) appeal to measurement error to explain the rejections reported in some of the papers cited above.

b) food consumption, which is often taken as a 'proxy' for total non-durable consumption in studies using micro-data,² is unsuitable because preferences are non-separable between food and other non-durables, and food is a necessity. Furthermore, during the sample period the relative price of food is far from constant. We show that the consequences of using food consumption rather than total non-durable consumption are non-negligible: theory restrictions are rejected, and the elasticity of intertemporal substitution is poorly determined when food is used.

We use a time series of cross-sections (the Consumer Expenditure Survey, 1980-90) to construct consistently aggregated cohort-level data - a synthetic panel- where households are grouped by year-of-birth and education. As discussed below, the availability of a relatively long time period is crucial in the estimation of Euler equations. Furthermore, thanks to the detailed information in the data set, we can construct different consumption measures, which include food and other non-durable commodities.

The use of a micro economic data set which contains comprehensive information on consumption, income, leisure and household composition, allows us to address several questions directly. Our parametrization of preferences takes into account explicitly the possibility of non-homotheticity. The scheme we propose involves the estimation of a demand system, which allows us to construct the household-specific price indices that characterize the intertemporal allocation of consumption expenditure. In addition, we allow for the effects of demographic changes and leisure on intertemporal allocation of expenditure. Finally, we can test theory restrictions in various ways, including the popular excess-sensitivity tests.

We show that a proper consideration of aggregation across commodities can be important for the estimation of key behavioral parameters. We also show that several simplifying assumptions commonly used in the literature, such as homotheticity and separability between consumption and leisure are strongly rejected by the data. On our preferred specification, we obtain sharp estimates of key behavioral parameters (including the elasticity of intertemporal substitution) and no rejections of theoretical restrictions.

The paper is organized as follows. Section 2 describes the data: it shows that income and consumption age profiles track each other. This could be due to myopic behavior (as pointed out by Carroll and Summers, 1991), but could also be explained by changes in demographics and leisure over the life cycle. In Section 2, we also show that both the relative price of food and its share in non-durable expenditure have varied considerably over the 1980s. Section 3 illustrates

² An exception is in a paper by Lusardi (1992).

some of the effects of incorrect aggregation across consumers and of using expenditure on food to describe intertemporal allocation of consumption. Section 4 describes in detail our specification of preferences which allows for non-homotheticity and non-separability between goods and leisure. Section 5 discusses econometric issues, with particular reference to problems which arise when the data are averages over relatively small sub-samples of the population. Section 6 presents estimation results: section 6.1 describes parameter estimates of a small demand system, which confirm the importance of non-homotheticity, and section 6.2 presents our estimates of the Euler equation. Section 7 concludes the paper.

2. Consumption behavior: data and descriptive analysis

Until recently, no micro dataset contained complete information on US household consumption. As stressed above, most of the empirical work on the life-cycle model used the PSID, which contains only a measure of food consumption. Since 1980 the Bureau of Labor Statistics has been running the Consumer Expenditure (CEX) Survey on a continuous basis. The Survey is now available for 11 consecutive years which include two recessions and a long expansion, two major tax changes and witnessed substantial movements in relative wages and prices. The CEX is not a full panel (households are interviewed over four consecutive quarters, and then replaced); however, because it is available over a relatively long time period and it contains considerable demographic information, the use of synthetic cohort analysis allows the study of the evolution of consumption over the life cycle.

Therefore the CEX gives, for the first time, the possibility of analyzing individual consumption behavior over the life cycle and over the business cycle. In the next section we present a structural model for non-durable consumption which is not rejected by the data and allows us to estimate some key behavioral parameters. Before doing so, however, we describe the main features of the data used in estimation.

2.1 The CEX Survey

The data we use cover the period from 1980 to 1990. In this subsection we describe the data and the main selection criteria used. The CEX is based on a comprehensive survey run by the Bureau of Labor Statistics, which interviews about 4,500 households every quarter;³ 80 % of these

³ The unit of reference is what the BLS defines 'consumer unit' which consists of 'all members of a particular housing unit or other type of living quarters who are related by blood, marriage,

are then re-interviewed the following quarter, while the remaining 20 % are replaced by a new, random group. Therefore, each household is interviewed at most four times over a period of a year. During the interviews, a number of questions are asked concerning household characteristics (demographics, work status, education, race, etc.) and detailed expenditures over the three months prior to the interview. The sample is representative of the US population.

We exclude from our sample non-urban households, households residing in student housing, households with incomplete income responses. Furthermore, as discussed below, we consider only households headed by individuals born after 1904 and before 1965 and that are at least 19 and no more than 75 years old. These exclusion restrictions leave us with a sample of 146219 interviews. From each interview we consider, in addition to the demographic and labor supply variables discussed below, the expenditure figures for the month preceding the interview. The first two months of the quarter preceding the interview are excluded to avoid the fairly complicated error structure that the timing of the interviews would imply on quarterly data. ⁴

In what follows we consider various components of non-durable expenditure. In particular, for reasons to be discussed below, we look at food (defined as the sum of food at home, food away from home, alcohol and tobacco) and expenditure on other non-durable goods and services, such as services, heating fuel, public and private transport (including gasoline) and personal care, and semi-durables, defined as clothing and footwear. The major exclusions from total consumption expenditure are consumer durables, housing, health and education expenditure.

In the sequel we use two income variables: total after tax family income and total before tax 'labor' income. 'Labor' income is defined as total family income minus capital income. Unfortunately it is not possible to reconstruct the after tax labor income. ⁵

In this paper we stress the importance of controlling for changes in the demographic structure of the household for a proper modelling of consumption behavior. We also stress the possibility that consumption and leisure might not be separable in the utility function. Therefore, even though it is not necessary to model labor supply explicitly, it is necessary to control for it. Fortunately, the adoption, or some other legal arrangement, such as foster children. Consumer unit determination for unrelated persons is based on financial independence. To be considered financially independent, at least two of the three major expense categories (food, housing, and other living expenses) have to be provided by the respondent'.

⁴ Such a complication would arise because there are households interviewed in each month and therefore their quarterly consumption would refer to overlapping periods.

⁵ Income data are collected at the first and last interview, and refer to the previous 12 months. Labor income is also computed at second or third interview if a member of the household reports changing her employment.

CEX contains a wealth of information on household characteristics and labor supply behavior. The demographic variables we examined in the empirical analysis are family size, number of children by age groups, number of adult members, number of persons older than 64, gender of household head, the educational attainment of the reference person, and the marital status of the household head.

The labor supply variables we have considered are various employment dummies, the number of earners, the number of hours worked by the wife and the dummies for part time and full time female employment.

The demand system in section 6 is estimated for different groups, formed on the basis of the educational attainment of the household head. The first group is that of high school dropouts, the second that of high school graduates, the third that of households with some college education and the fourth that of college graduates. Considering all years together the four groups account for 22.7, 29, 22 and 26.3 per cent of the total sample.

We devote a considerable amount of attention to the construction of appropriate price indices. Besides the parameters of the demand system discussed below, to construct such indices we need the price of the commodities we model (food and other non-durables). These are constructed from the components of the CPI published monthly by the Bureau for Labor Statistics. These prices are region-specific. The price indices for the composite commodities (food, other non-durables) that form our demand system are then constructed as expenditure-share weighted averages of the elementary price indices. These prices are, therefore, household-specific.

To describe the intertemporal allocation of consumption we need also a nominal interest rate. We choose the return on Municipal Bonds as it is tax exempt, therefore avoiding us the difficulty of measuring individual marginal tax rates. We have also experimented with the 3-month treasury bill. Both interest rate series are taken from the *Economic Report of the President*.

2.2 Life cycle and time series variation

In this subsection we describe the main features of the variables that are relevant for the more formal analysis of section 6. In particular, we characterize the life-cycle profile of non-durable consumption, household income, family composition and female labor supply. In addition, we examine the time series variability of prices and average consumption over the sample period under study. This analysis is useful not only as a descriptive tool, but also as an indirect check on the quality of the data set, which is relatively new.

The main problem with using the CEX Survey for the analysis of a dynamic model such as the life cycle one, is that each household is not observed over a long period of time. However, the continuity of the Survey allow us to follow the average consumption behavior of homogeneous groups over time, as they age. This is the main idea behind the synthetic cohort analysis which is used both in this section and in the more structural analysis in section 6.⁶

We divide the households in the sample in cohorts, defined on the basis of the year of birth of the household head.⁷

We then average the variables of interest over all the households belonging to a given cohort observed in a given year. If there are N cohorts observed for T years, this procedure gives use NT observations. In table 1 we report the cohort definition, their age in 1980 and the average cell size. The same cohort definitions are used in the section 5 to construct quarterly times series.

The advantage of grouping by the year of birth rather than by age in studying life cycle behavior is obvious. One follows over time a group of individuals born in the same period and therefore coming of age *at the same time*. Estimating age profiles by pooling together several cross sections and grouping by age is potentially very misleading in the presence of cohort effects.

In Figure 1 we plot average cohort log income and log non-durable consumption against age. Each connecting segment represents mean log income or consumption of a given cohort. Because cohorts are defined by a 5 year interval and the sample covers eleven years, each cohort overlaps with an adjacent cohort at 6 ages.

Both income and consumption are hump-shaped peaking before retirement. Furthermore, consumption is considerably less variable than income: the standard deviation of the residuals obtained regressing first log consumption and then log income on a polynomial in age and cohort dummies is 0.04 and 0.06 respectively. This difference could be explained by smoothing behavior or by greater measurement error in income than in consumption.

The fact that both consumption and income present a pronounced hump (and the fact that differences in the shape of income profiles among occupation groups were reflected by similar differences in consumption profiles), is interpreted by Carroll and Summers (1991) as evidence against the life cycle model. Carroll and Summers, however, ignore family composition: in Figure 2, we plot the age profile of family size. As can be seen, the age profile for family size is also hump

⁶ Details are given in section 5.

⁷ The household head is defined as the person who owns or signs the rental contract of the home where the consumer unit lives. For married couples, however, we define the husband as the household head and therefore use his age to establish to which cohort the household belongs to.

Table 1
Cohort definition

cohort	Year of birth	Age in 1980	Average cell size	Used in estimation
1	1964-1960	16-20	-	no
2	1959-1955	21-25	461	yes
3	1954-1950	26-30	460	yes
4	1949-1945	31-35	426	yes
5	1944-1940	36-40	321	yes
6	1939-1935	41-45	258	yes
7	1934-1930	46-50	241	yes
8	1929-1925	51-55	255	yes
9	1924-1920	56-60	272	yes
10	1919-1915	61-65	-	no
11	1914-1910	66-70	-	no
12	1909-1905	71-75	-	no




Figure 1



Figure 2

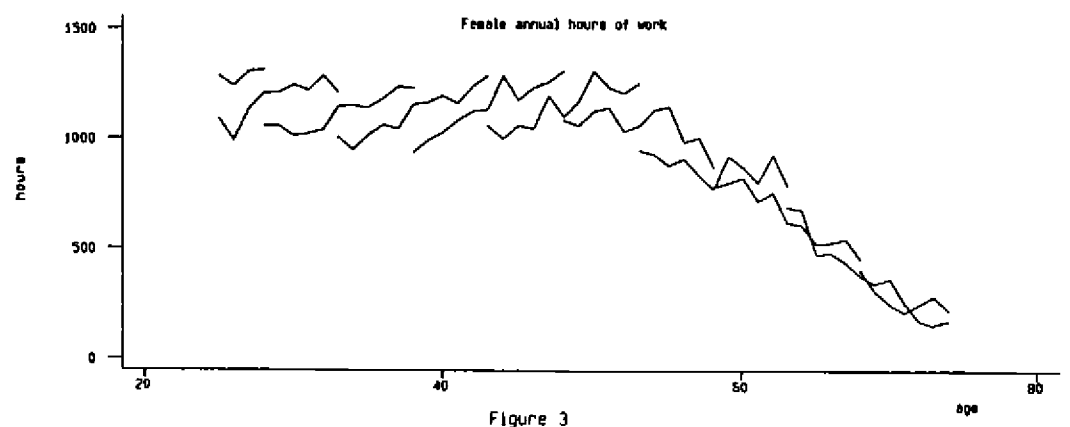


Figure 3

shaped. ⁸

In figure 3, we plot the life-cycle profile for female average annual hours of work. The average is conditional on the presence of the wife, but not on positive hours. Several features are worth noticing. First, female labor supply exhibits a substantial amount of variability both at life cycle and business cycle frequencies. Strong cohort effects are also apparent. Second, there is no apparent dip in female hours corresponding to fertility ages. This feature differentiates this profile from similar ones for other countries (like the UK) or other time periods.

Several influences affect consumption and are partly detectable in Figure 1. We can think of life cycle effects, cohort effects and business cycle effects. In figure 4 we try to remove life cycle and cohort effects to isolate time effects. Of course this decomposition is somewhat artificial as the three kinds of effects are not identifiable. By regressing average cohort consumption on cohort dummies and a 5th degree age polynomial and considering the residuals of such a regression as time effects, we interpret all trends in consumption as deriving from a combination of cohort and age effects. In figure 4 these residuals are plotted together with de-trended aggregate non-durable consumption. Two features are noticeable. First, there is a substantial amount of synchronization across cohorts. The 1981-1982 recession is particularly visible. After that the average residuals seem to rise fairly steadily until the end of the sample. Second, average residuals follow reasonably well aggregate consumption. The correlation coefficient between the average residuals and aggregate de-trended consumption is around 0.4. ⁹


2.3 Movements in consumption shares and relative prices

In what follows, we give considerable emphasis to non-homotheticity and aggregation across commodities. As Gorman (1959) has shown, the conditions under which one can aggregate across commodities and consider a single price index to characterize the allocation of consumption over time, are very stringent. If some goods are necessities and other goods are luxuries, and their relative prices change, then at least *two* price indexes are needed to deflate nominal expenditure. Failing to model the features of preferences which prevent the consideration of a single price index introduces an omitted variable problem which can induce serious biases.

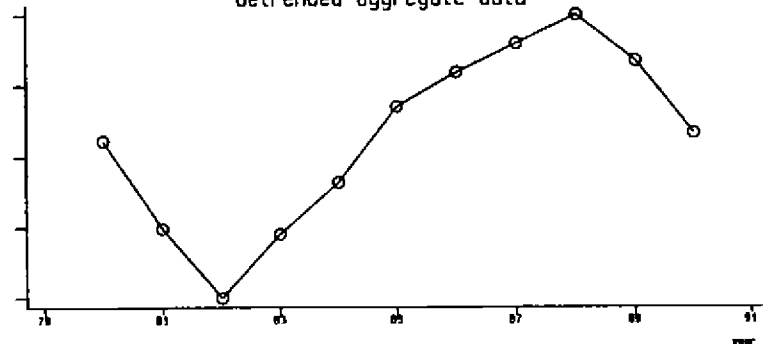
⁸ There are also remarkable differences in the shape of family size across education and occupation groups. It is therefore possible that differences in consumption across education groups could be explained by differences in demographics. See the discussion in Attanasio et al. (1994).

⁹ The presence of measurement error and of definitional differences between CEX and national accounts consumption should also be kept in mind.

Figure 4
cohort data



detrended aggregate data



Business cycle effects

Figure 5


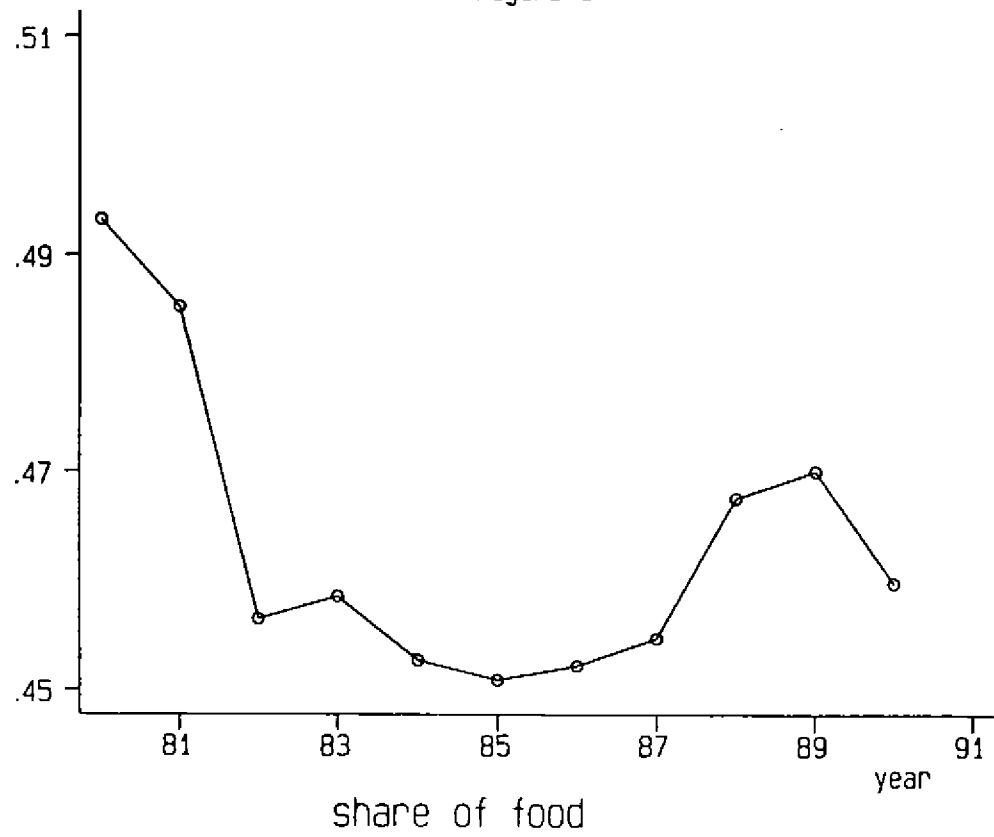


Figure 6



In figure 5 we plot the relative price of food and other non-durables over the 1980s. As can be seen there is a substantial amount of variability. In particular, around 1986 the relative price of food and other non-durables exhibits a substantial shock. This is due to the decline in oil products prices which are included in non-durables.

Corresponding to these changes in relative prices one can observe changes in consumption shares. In Figure 6 we plot the average share of food in non-durable expenditure over the 1980s. The graphs exhibit a substantial drop in 1982 followed by a more gentle decline up until 1985, when the share of food starts to increase again. Part of the dramatic drop of 1982 is probably explained by data problems.¹⁰ For this reason, in what follows, we report results for both the whole sample and a subsample which excludes the first two years.

In the first part of section 6 we model the relationship between relative prices and food share when we estimate a simple demand system. This is the first block of our empirical strategy to model non-durable consumption expenditure. The necessity of such a step is evident in Figure 5.

3. What can we learn from aggregate data and from food consumption?

As stressed in the introduction, most of the empirical work on the life cycle model of consumption has used either aggregate data or micro data which contained only limited information on consumption, namely expenditure on food. The use of the CES puts us in a vantage position, in that we can replicate the results of other researchers and address the issues of aggregation across consumers and commodities.

3.1 Aggregation across consumers

Nobody would doubt that actual consumers are heterogenous. The issue, however, is to establish to what extent heterogeneity affects inferences on the intertemporal optimization model based on aggregate data. Attanasio and Weber (1993), using a long time series of UK cross-sections, have shown that aggregation bias can explain the rejection of the overidentifying restrictions implied by

¹⁰ The BLS runs a separate sample, based on diaries, rather than interviews, to collect information on expenditure on food and other frequently purchased items. In 1980 and 1981 (but not in later years) average food consumption was substantially higher in the interview survey than in the more reliable diary survey. A direct comparison with the share of food which can be computed from NIPA figures is not feasible because of definitional differences in some of the items included in non-durable consumption. The aggregate food share shows a constant and gentle decline which slows down around 1985.

the model. The Euler equation derived from isoelastic preferences implies a non-linear relationship between consumption and the interest rate at the individual level. It is easy to show that the use of aggregate data to estimate and test such a relationship is equivalent to taking the log of the mean rather than the mean of the log. The difference between these two quantities is an index of inequality which, as shown by Attanasio and Weber (1993), has strong cyclical properties and therefore might cause the rejection of overidentifying restrictions.¹¹ A similar exercise can be performed using the CEX data. One of the advantages of working with a time series of cross-sections is that one can control the aggregation process directly: we can aggregate any known non-linear function of the individual data.

A typical Euler equation derived from intertemporally separable and isoelastic preferences is the following:

$$(1) \quad \Delta \log(c_{t+1}^h) = \text{constant} + \sigma \log(1 + r_{t+1}) + \epsilon_{t+1}^h$$

where c^h is individual consumption, r is the real interest rate, σ is the elasticity of intertemporal substitution and ϵ_t is a residual uncorrelated with all the information available as of time $t-1$.

Equation (1) can be aggregated across consumers to obtain a similar equation for aggregated data.

$$(1') \quad \frac{1}{H} \sum_h \Delta \log(c_{t+1}^h) = \text{constant} + \sigma \log(1 + r_{t+1}) + \frac{1}{H} \sum_h \epsilon_{t+1}^h$$

In the absence of individual data researchers have estimated :

$$(1'') \quad \Delta \log \frac{1}{H} \sum_h (c_{t+1}^h) = \text{constant} + \sigma \log(1 + r_{t+1}) + v_{t+1}$$

In table 2 we present estimates of equations (1') and (1''). We did not search for a satisfactory specification: the exercise is only meant to show the effects of incorrect aggregation.

The left-hand side column in table 2 reports estimates for the correctly aggregated non-durable consumption measure. The right-hand side column reports estimates for the incorrectly aggregated

¹¹ The use of a representative consumer can be justified, theoretically, assuming the existence of perfect insurance markets. Attanasio and Davis (1994) show overwhelming evidence against this hypothesis. Furthermore, even in the presence of complete markets, the non-linearity issue would still be relevant, unless we impose special preferences on the individual agents. In a recent paper, Mace (1991) reports results obtained using the CEX which could be interpreted in favor of the perfect insurance hypothesis. However, it is very easy to show that they could be caused by the presence of measurement error.

Table 2

Estimates based on aggregate CEX data (weighted)

	geometric mean	arithmetic mean
$\Delta \log(\text{farmsize})$	0.942 (0.168)	0.668 (0.207)
r	0.565 (0.230)	0.244 (0.303)
Sargan Criterion	11.40 (11)	12.39 (11)
Breusch-Godfrey Criterion	6.64 (4)	27.34 (4)

Notes: MA(1)-consistent standard errors in parentheses. Instruments used are second, third and fourth lags of consumption growth, income growth, inflation and the real nominal interest rate, plus the following exogenous explanatory variables: S1-S4 and $\Delta \log(\text{farmsize})$. The Sargan criterion is a χ^2 test of the overidentifying restriction (in this case with 11 degrees of freedom) the Breusch-Godfrey criterion is a χ^2 test for serial correlation of order 2 to 5 (with 4 degrees of freedom).

model (where we have taken the logarithm of the arithmetic mean, as is normally done on aggregate time series data).¹² Four seasonal dummies (S1-S4) are introduced to take into account seasonality in preferences. Changes in family size are also used as an explanatory variable. The real interest rate is the final explanatory variable, and is treated as endogenous in estimation. We allow for an MA(1) error term reflecting measurement error and/or time aggregation, by correcting all standard errors and choosing instruments lagged 2 or more quarters.¹³

Two differences between the two columns are noticeable. First, the point estimates of the parameters are different. In particular, the coefficient on the interest rate, which is usually interpreted as the elasticity of intertemporal substitution, is less than half in size when estimated using incorrectly aggregated data. Second, and more importantly, even though the Sargan test of overidentifying restrictions does not signal any violation for either equation, the Breusch-Godfrey test for correlation of second to fourth order strongly rejects the null for the equation estimated on incorrectly aggregated data.¹⁴

Higher order serial correlation is inconsistent with theory predictions, and would - given a parsimonious instrument list - lead to the type of rejections of the overidentifying restrictions often reported in studies using aggregate data (e.g.: Hansen and Singleton, 1982).

As we argue in the following section, there are reasons to believe that equation (1') is misspecified. The differences obtained estimating the two equations, however, are an indication of the importance of aggregation effects.

3.2 Food consumption

Most of the US literature on the life cycle model based on micro data has used the PSID, which contains only information on food consumption.¹⁵ This is a very serious limitation. The assumption necessary to justify the use of food to study the intertemporal allocation of consumption is that utility is separable between food and other consumption goods. All available studies of demand systems strongly reject such hypothesis. Furthermore, if the aim of the exercise is to estimate behavioral parameters, such as the elasticity of intertemporal substitution, it is ques-

¹² We have used CEX-provided population weights to enhance comparability with published National Accounts data, but similar results obtain when unweighted averages are taken instead.

¹³ The GMM estimator used to obtain the estimates in table 2 is the same as that used for the subsequent tables and is discussed in section 5 and in the appendix.

¹⁴ These results are similar to those reported for UK micro data in Attanasio and Weber (1993). Excess sensitivity tests fail to find a significant coefficient on income growth in either specification.

¹⁵ A notable exception is Lusardi (1992) who used the CEX.

tionable whether those obtained using food consumption are indicative of the substitutability over time of total consumption. Finally, the isoelastic specification often assumed in empirical studies is particularly ill-suited for food, which is a necessity.

In Table 3, we report the results obtained estimating equation (1) with consumption defined as expenditure on food. While the econometric technique used is slightly different,¹⁶ the specification of the equation is similar to that estimated by several authors, such as Zeldes (1989), Runkle (1991) and Keane and Runkle (1992). In column (1) the rate of growth in food consumption at the cohort level is related to the real interest rate (obtained subtracting the rate of growth in food prices from the nominal rate), and to other control variables such as seasonal dummies and the log of family size. Instruments used include second to the fourth lags of interest and inflation rates, second to the fourth lags of income, food and other non-durable consumption growth as well as a polynomial in age and lagged demographics. The results are not greatly affected by the introduction of other demographic variables or by a change in the instruments.

The estimate of the coefficient on the interest rate is negative and statistically different from zero. The Sargan criterion does not indicate a rejection of the overidentifying restrictions.

In the second column we add to the specification the rate of growth in labor income. The coefficient of this variable is estimated at 0.18 with a standard error of 0.07. Its introduction does not affect considerably the remaining coefficients or the test of overidentifying restrictions.

In the third column, we add to the specification in the second column the rate of growth of consumption of other non-durable commodities. The coefficient on this additional variable is estimated at 0.17 with a s.e. of 0.10, which makes it marginally different from zero. The point estimate of the coefficient on income is greatly reduced and it is no longer statistically different from zero. The coefficient on the interest rate is still negative and (marginally) different from zero.

In the final column we remove the rate of growth of income. The coefficient on other non-durables is now strongly significant. The coefficient on the interest rate, however, is still negative, but not statistically different from zero.¹⁷

We interpret the evidence in Table 3 as indicating that the non-separability between food and other non-durables is a potentially important problem.

¹⁶ The techniques used to obtain the results in table (3) are analogous to those used for the subsequent tables and are discussed in detail in section 5 and in the appendix.

¹⁷ As we pointed above, the food figures for 1980 and 1981 might be of particularly low quality. Because of this we re-estimated the columns in Table 3 over the period 1982:3 to 1990:4. The results, available upon request, are very similar to those in Table 3.

Table 3

Estimates based on food consumption

	(1)	(2)	(3)	(4)
$\Delta \log$ <i>fam. size</i>	0.459 (0.142)	0.271 (0.165)	0.292 (0.158)	0.383 (0.160)
<i>r</i>	-0.733 (0.253)	-0.862 (0.277)	-0.551 (0.363)	-0.369 (0.323)
$\Delta \log$ <i>labor income</i>	- -	0.177 (0.067)	0.101 (0.089)	- -
$\Delta \log$ <i>other non - durab.-</i>	- -	- -	0.170 (0.104)	0.189 (0.088)
Sargan Crit. p-value	12.1 (0.88)	11.6 (0.86)	10.2 (0.89)	10.2 (0.95)
Number of observations	288	288	288	288

Notes: MA(1)-consistent standard errors in parentheses. Instruments used are second, third and fourth lags of food and other durable consumption growth, income growth, inflation and the nominal interest rate, the second lag of family size growth plus the following exogenous variables: S1-S4 age and age squared. The Sargan criterion is a χ^2 test of the overidentifying restriction.

The estimates of the elasticity of food consumption growth to the real interest rate presented here differ from those available in the literature.¹⁸ These differences could be due to a variety of factors including, for instance, the use of a different survey in which the timing of food expenditure is better determined. An important difference is in the econometric methodology. The use of synthetic panels (rather than a short panel at the individual level) affords important gains in the time dimension of the sample. Given that the error term has an expectational component, this is a matter of great importance (consistency issues are further discussed in section 5).

4. The allocation of consumer expenditure over time with multiple commodities

Food is a necessity and it is unlikely, as we have shown in the previous section, to be separable from other non-durables. Therefore food expenditure is inadequate to study the intertemporal allocation of consumption. On the other hand, even when data on total non-durable expenditure are available, it is not obvious that its intertemporal allocation could be described, in the presence of large changes in relative prices, by a single price index. It might be necessary to model the intratemporal and the intertemporal allocation of consumption simultaneously.

In this respect several modelling strategies are available. One could formulate within period utility as a function of several commodities and consider the Euler equation for each of them. The main problem here is to find a flexible direct utility function which nests the isoelastic case (and for which integrability conditions can be imposed by setting data independent restrictions on the parameter space). Alternatively, one could consider flexible specifications for marginal utilities and estimate the Euler equations from those. In this case, however, quasi-concavity of the implied utility function might be hard to impose. Finally, one can choose to work with an indirect utility function which captures both non-homotheticity and non-separability.

We adopt this last strategy. While most of the theoretical results used in the analysis are well known at least since Gorman (1959), the specification of preferences we present is relatively new. We introduce an empirically tractable way to consider the intertemporal allocation problem when within period utility depends on several commodities. Our specification relaxes the assumption of homotheticity and therefore does not allow the characterization of intertemporal allocation by means of a single price index. The homothetic case, however, is nested as a special case and can be obtained with simple restrictions on the parameters we estimate.

Our specification of preferences is similar to that used by Blundell, Browning and Meghir

¹⁸ See, for instance, Zeldes (1989), Runkle (1991), Keane and Runkle (1992).

(1994) in a recent paper, but presents the key advantage of producing an Euler equation for consumption which nests equation (1) when preferences are homothetic.

We proceed in two steps. We first estimate a flexible demand system which satisfies integrability conditions. The results obtained in this first step are then used to construct the price indexes necessary to characterize intertemporal allocation. The parameters that enter the Euler equation can then be estimated.

4.1 Multiple Commodities and the Euler Equation

Let us first consider the representative agent problem as formulated in most macro economic papers:

$$(2a) \quad \text{Max } E_t \sum_{j=0}^{T-t} U(c_{t+j})(1+\delta)^{-j}$$

$$(2b) \quad \text{subject to } A_{t+1} = (1+r_{t+1})A_t + y_t - p_t c_t$$

where A_t are the assets at the beginning of period t , c_t is expenditure on an homogeneous and non-durable consumption good, y_t is income in period t , and r_{t+1} is the nominal interest rate between periods t and $t+1$.

Equations (2a) and (2b) give rise to the standard Euler equation for consumption. In the literature, the instantaneous utility function is often parametrized as the following CRRA utility function which gives constant elasticity of intertemporal substitution equal to σ .

$$(2c) \quad U(c_t) = \frac{1}{1-\frac{1}{\sigma}} c_t^{1-\frac{1}{\sigma}}$$

Equation (1) in the previous section can be obtained log-linearizing the first order condition for the maximization problem (2a)-(2c).¹⁹ When one considers several commodities, one can think of c_t as total expenditure deflated by an appropriate price index. However, as Gorman (1959) proved,

¹⁹ The constant of equation (1) includes the log of the discount rate β and various terms reflecting second and higher moments of the conditional distribution of ϵ_{t+1} . If ϵ_{t+1} is conditionally log-normal, the constant will include only the variances of consumption growth and the interest rate as well as their covariance. An implicit assumption which is usually made is that changes in these conditional moments are uncorrelated with the instruments used.

only under very stringent conditions the intertemporal optimization problem can be determined on the basis of a single price index.

One can interpret (2a) as the utility index of a consumer who breaks her optimization problems in two steps: in the first step, she decides how much total expenditure X to allocate to each time period. In the second step, she allocates X to different goods (q_1, \dots, q_N), according to their relative prices and to X itself (some goods will be luxuries, some necessities). Suppose the second step produces demand equations of the Almost Ideal type (see Deaton and Muellbauer (1980)):

$$(3) \quad \frac{q_{it} p_{it}}{X_t} = \alpha_i + \sum_j \gamma_{ij} \ln(p_{jt}) + \beta_i [\ln(X_t) - \ln a(p_t)]$$

where the p 's are individual prices (and p is the corresponding price vector), and α , γ and β 's are preference parameters (which will be functions of demographic characteristics, employment, etc., as in Blundell, Pashardes and Weber (1993)). If all the β 's are zero, preferences are homothetic, and the indirect utility function for period- t consumption is:

$$(4) \quad V = F\left[\frac{X_t}{a(p_t)}\right] = F[c_t]$$

where $F[\cdot]$ is a monotonic transformation (which cannot be identified from the demand system alone and determines the intertemporal allocation), and $\ln(a(p_t)) = a + \sum \alpha_i \ln(p_{it}) + .5 \sum_i \sum_j \gamma_{ij} \ln(p_{it}) \ln(p_{jt})$. Equation (4) implicitly defines c_t as nominal expenditure X_t deflated by $a(p_t)$.

On the other hand, if the β 's are non-zero, violating homotheticity, the indirect utility function becomes:

$$(5) \quad V^* = F\left[\left(\frac{X_t}{a(p_t)}\right)^{\frac{1}{b(p_t)}}\right]$$

where $b(p_t) = \Pi_i p_{it}^{\beta_i}$, i.e a zero-degree homogeneous price index (adding-up implies $\sum \beta_i = 0$). This second price index takes into account the different impact price changes have on utility according to the type of good they refer to ($\beta_i > 0$ for luxury goods, $\beta_i < 0$ for necessities).

Neglecting the existence of this second price index may lead to spurious inferences. Equation (1) will suffer from omitted variable bias, because no account is taken of changes in $b(p_t)$ over time. This problem is particularly severe if, over the sample period analyzed there are large

changes in relative prices across luxuries and necessities (for instance, if luxuries become relatively more expensive in booms, and cheaper during recessions).

If we substitute equation (2a) with the following,

$$(6) \quad U = \sum_t \frac{1}{1 - \frac{1}{\sigma}} \left[\left(\frac{X_t}{a(p_t)} \right)^{\frac{1}{1+\beta}} \right]^{1-\frac{1}{\sigma}} (1+\delta)^{-t}$$

the Euler equation (1) becomes:

$$(7) \quad \Delta \left[\frac{1 + \sigma(b(p_{t+1}) - 1)}{b(p_{t+1})} \ln(c_{t+1}) \right] = \psi + \sigma [r_{t+1}^* - \Delta \ln(b(p_{t+1}))] + c_{t+1}$$

This expression looks daunting, but neatly simplifies to equation (1) when all the β 's are zero (in which case $b(p) = 1$).²⁰

When the β 's are not zero, the second price index is subtracted from the standard definition of the real interest rate, and it affects the coefficient on consumption growth. The equation is easy to estimate, particularly if the β 's are known in advance: this is the case if we estimate Engel curves separately in each year. However, σ is no longer the elasticity of intertemporal substitution: its multi-good definition is (Browning, (1987)) $EIS = (-) \frac{V^*}{X V^*_{xx}}$ where an x subscript denotes the partial derivative with respect to X . This implies:

$$EIS = - \frac{\sigma b(p)}{1 + \sigma(b(p) - 1)}$$

We can rewrite equation (7) as :

$$(7') \quad \Delta \frac{\ln(c_{t+1})}{b(p_{t+1})} = \psi + \sigma [r_{t+1}^* - \Delta \ln(b(p_{t+1})) - \Delta \left(\ln(c_{t+1}) - \frac{\ln(c_{t+1})}{b(p_{t+1})} \right)]$$

²⁰ Blundell, Browning and Meghir (1994) estimate an Euler equation similar to (7), but with further non-linearities. Their analysis is different in two respects. On the one hand, their demand system is more general than ours as it is consistent with a wider pattern of Engel curves (and nests the Almost Ideal case from which we start). On the other hand, their monotonic transformation in V^* takes as argument a non-linear function in X . This second feature implies that even in the homothetic case their Euler equation does not simplify to equation (1). In the Almost Ideal case, for instance, their Euler equation involves taking the logarithm of $\log(\text{consumption})$ - a particularly unappealing data transformation, which rules out very low, but theory-consistent consumption levels.

which is the specification we use in our empirical work. Here, the dependent variable is the first difference in the log real consumption divided by $b(p)$ (a number close to 1). On the right-hand side, the real interest rate (obtained subtracting from the nominal rate the rate of inflation in $a(p)$) is further deflated by the growth rate of the second price index, $b(p)$, and by a correction term, which is the difference between consumption growth as normally defined and the dependent variable.

4.2 Tuning up the model

Before even trying to fit the model described above to the data, a number of simple modifications are necessary. These allow the model to explain some obvious, and yet quantitatively important, features of the data.

(i) Household consumption exhibits large seasonal fluctuations which are clearly inconsistent with a simple minded version of the life-cycle model. A simple way of introducing seasonal fluctuations is to have the utility function depending on seasonal shifts, so that a given amount of expenditure gives different levels of utility in different quarters. This framework has been used (among others) by Miron (1986) on aggregate data and by Attanasio and Weber (1993) on micro data. In the log-linearized Euler equation used in this paper, this specification implies the use of seasonal intercepts in a regression of the rate of growth of consumption on the interest rate.

(ii) The utility derived from a given amount of expenditure obviously depends, among other things, on family composition. In general, it is quite difficult to model properly intra-household decisions. Fertility choices are probably endogenous and should be modeled simultaneously with consumption and labor supply behavior. However, it is quite easy to introduce some simple corrections to make the model more realistic. We assume that utility is shifted by a number of demographic variables such as the number of children of various ages, the number of adults, etc. Such a framework allows for fairly flexible adult equivalent schemes. The instantaneous utility function for a generic household h that will be used in the empirical application below is the following:

$$U_t^h = U(C_t^h)\phi(Z_t^h, \theta)$$

where C_t^h is total family expenditure and $\phi(Z_t^h, \theta)$ is a function of various demographic variables. If the function ϕ is given by $\phi(Z_t^h, \theta) = \exp(\theta' Z_t^h)$, the term $\theta \Delta Z_t^h$ will enter the Euler equation for consumption. Changes in ϕ are equivalent to a time-varying discount δ . For this reason we

shall refer to ϕ as 'the discount factor'.²¹

Demographic variables are also likely to affect the demand system. As a consequence they will also have an indirect effect on the Euler equation through the price indexes $a(p)$ and $b(p)$.

(iii) An implicit and potentially controversial assumption often used in papers that estimate Euler equations for consumption is that of separability of the utility function between consumption and leisure. While in this paper we do not model explicitly labor supply behavior, we argue that non-separability between consumption and leisure can be taken into account in a very simple way. Common sense tells us that the level of utility obtained from a given amount of expenditure depends on labor force participation variables: when a member of the household works, he or she will have to bear a number of job-related expenses that will be reflected in total consumption expenditure. These job-related expenses will affect both the intertemporal allocation of expenditure (normally increasing expenditure when leisure is low) and the intratemporal allocation (by making leisure-intensive commodities relatively more expensive). For this reason, we allow for leisure effects in the demand system - thus producing price deflators that depend on leisure, but also capture direct intertemporal leisure effects by introducing leisure-related variables in the Euler equation.

In practice, we introduced as determinants of the marginal utility of consumption a number of labor supply variables, reflecting both employment status and hours. This strategy allows us to avoid the formal modeling of labor supply, with the complications arising from corner solutions and institutional constraints. This obviously only identifies conditional preferences (see Browning and Meghir (1991)).

5. Econometric Issues

As we said in Section 2, the CEX survey is a rotating panel. Rather than employing the (short) panel dimension of the survey, which is probably dominated by seasonal factors, we decided to construct synthetic panels. We define cohort by the year of birth of the household head. Our technique is equivalent to using the interaction of time and group dummies as instruments. We impose an age limit (23-60) and a cell size limit (150) on the cohorts we consider, so that our synthetic panel is not balanced.²²

²¹ An analogous parametrization applies to the non-homothetic case, where C is defined as $(\frac{X_{t+1}}{a(p_{t+1})})^{\frac{1}{\sigma(p_{t+1})}}$.

²² See Deaton (1985), Browning, Deaton and Irish (1985) and Moffitt (1993). We also experimented with groups formed by year of birth and educational attainment. The education groups were the same as those used in the estimation of the demand system. The reason we only report

There are several advantages in the use of synthetic panel techniques to estimate an Euler equation. First, averaging over individuals belonging to a group should eliminate additive idiosyncratic measurement error.

Second, it is known that if the panel dimension is short, the introduction of household specific fixed effects gives inconsistent estimates, unless the instrument used are strictly exogenous (see Runkle (1991)). Taking cohort averages over long time periods can get round this problem: because we observe groups for the whole sample period (at least potentially), the relevant dimension is the total length of the period covered by the survey, not how long each household stays in it. A similar argument can be made about the presence of aggregate shocks. If we think our sample is long enough so that expectational errors are averaged out, the use of more or less standard IV techniques gives us consistent estimates. This solves the small T problem discussed by Chamberlain (1984) and Hayashi (1987).²³

Third, we do not need to worry about attrition as much as we should if we were using a long panel. In this respect, as stressed by Moffitt(1993), the use of a time series of cross section has some advantages relative to the use of panel data.

Fourth, even though we work with aggregate data, in that we sum over the individuals belonging to a certain group, we can control the aggregation process directly. In this respect, as we saw in section 3, we can perform interesting exercises to evaluate the extent of aggregation biases.

The use of synthetic panels, however, does not by itself solve some important econometric problems. First, it does not help deal with non-additive measurement error. We should stress that any form of non-additive measurement error in the variables of which we take averages, induces inconsistent estimates. Second, it does not eliminate concerns about non-random attrition. In particular, we have to rely on the assumption that the population from which the sample is drawn is homogeneous over time. This assumption might be violated if, for instance, there is a relationship between mortality and wealth. If this is the case each cohort would become progressively 'richer' as it ages and therefore we would overestimate the rate of growth of consumption for older cohorts.

Averaging over cells of relatively small size induces measurement error in the levels, which

results based on birth year cohorts is that by crossing cohorts and education groups we are left with very small cells giving rise to extremely noisy data. The results were extremely imprecise.

²³ Of course this argument impinges on the assumption that our sample is large enough so that expectational errors are averaged out. While 44 quarterly observations might not be a very large number, it should be stressed that our sample period includes two recessions and a long period of moderate growth.

in turn implies an MA(1) structure in the first differences.²⁴ The reason for this is obvious. If the sample for a given quarter t includes a very rich household, this will induce a positive measurement error in the consumption growth at time t followed by a negative measurement error a time $t + 1$. The error of equation (1) is therefore going to be made of two components. A white noise component which reflects expectational errors and an MA(1) component with a coefficient of -1. The sum of a white noise and an MA(1) is an MA(1). The size of the coefficient of this MA(1) depends on the relative magnitude of the variances of the expectational error and of the measurement error.²⁵

If the only error to equation (1) and its extensions was an expectational one, instruments dated $t - 1$ and earlier would be valid ones. Our data, however, do indicate negative first order autocorrelation, thus suggesting that measurement error is an important issue. Because of this, the instruments dated $t - 1$ are invalid, but instruments lagged 2 and more periods yield consistent estimates.

The panel dimension of the CEX implies that temporally adjacent cells do not include completely different households. For instance, households at their first interview in time period t , appear also at time $t + 1$, $t + 2$ and $t + 3$. On the other hand those at the 4th and last interview at time t , also appear at time $t - 1$, $t - 2$ and $t - 3$. If we ignore the rotating nature of the panel and use all the households in the construction of the relevant variables and of the instruments, we get inconsistent estimates in the presence of household specific fixed effects. On the other hand, using only one interview per household involves the loss of a substantial amount of information.

To get around this problem we use all observations in the construction of the variables that enter our regression and select subsamples on the basis of the interview number in the construction of instruments. Namely, we use only households at the fourth interview in the construction of lag 2 instruments, households at the fourth and third interview for lag 3 instruments and we exclude households at the first interview for lag 4 instruments. This scheme guarantees that there is no overlap between the households used in the construction of the instruments and those used in the construction of the variables that enter the estimated equation.

²⁴ The presence of measurement error induced by small cell size is relevant for all the household-specific variables considered in the equation, even for those, such as family composition, that could be conceivably be considered as exogenous. The only exception is age, that we define as median cohort age, and is therefore unaffected by sampling variability.

²⁵ Hall (1988) suggests that if the planning horizon is shorter than the frequency of the observed data, the Euler equation has MA(1) errors. The sum of two independent MA(1) processes is again an MA(1) process.

The presence of MA(1) residuals for each cohort is not the only problem with the error structure of equation (1). Because we estimate it for N cohorts simultaneously, the expectational errors for a given time period for different cohorts are likely to be correlated. We allow for contemporaneous correlation among the residuals of different cohorts. We also allow for the presence of arbitrary heteroscedasticity which is likely to arise because of differences in cell sizes.

The complicated error structure of equation (1) estimated for several cohorts simultaneously has to be taken into account in the construction of an efficient estimator and in the estimation of its standard errors.²⁶ Details are provided in the appendix, where we describe in detail the GMM estimator we use.

A final issue is the way we treat instruments. In principle, we could 'stack' the instruments for each cohort (effectively imposing the same reduced form for all cohorts), or we could have different first stage regressions for each cohort. Given the limited number of observations, we decided to 'stack' the instruments.²⁷

6. Results.

This section describes the results obtained estimating our preferred specification. First, we present the estimation of a simple demand system. This shows that we need at least two price indexes to describe intertemporal allocation over time. This follows from the fact that food consumption is a necessity and from the fact that over the sample period the relative price of food and other non-durable changed dramatically.

The estimation of the demand system allows us to construct the price indexes which are necessary to determine intertemporal allocation. We show the time series behavior of these price indexes and discuss the implications of their omission.

Second, we present the results obtained estimating equation (7') and contrast them to those obtained estimating equation (1).

6.1 Demand system

²⁶ A GLS type transformation can generate inconsistent estimates if involves filtering the system backward (see Hayashi and Sims (1983)). Deaton (1985) shows that the cross sectional second moments can be used to improve the efficiency of the standard Instrumental Variable estimator, by giving less weight to the instruments more affected by sectional variability. Unfortunately, as noted by Fuller (1987), there is no guarantee that the resulting projection matrix be positive definite in finite samples. We do not make use of these corrections because of their finite sample unreliability.

²⁷ For a discussion of these issues, see Attanasio and Browning (1991a).

We split total non-durable expenditure on food and all other non-durable goods and services. This split, while arbitrary, has the advantage of grouping (potential) necessities separately from (potential) luxuries, and of defining commodities for which zero expenditures are not reported.

The two-commodity demand system we estimate is of the Almost Ideal type (see equation (3) above): the dependent variables are the budget shares of food and of other non-durables, the explanatory variables are their prices and total non-durable expenditure (the budget), deflated by a linearly homogenous price index, $a(p)$. Because the budget shares add up to one an equation is redundant and we estimate and report just the food equation without loss of generality.

In principle, all preference parameters (α , γ and β 's) could vary across households. We restrict their variation in the following ways: we assume γ 's and β 's to be constant within educational groups, and the α 's to depend on a few demographic and labor market variables. We therefore estimate separate budget share equations for food for each of four groups formed on the basis of the education attainment of the household head: high school dropouts, high school graduates, college dropouts and college graduates. For each educational group the α parameter (which affects both the intercept and the $a(p)$ index deflating nominal expenditure) is allowed to depend on some deterministic variables (seasonal dummies and a zero-one indicator for the 1980-1 wave of interviews), some demographic indicators (age of the head, single adult dummy, total number of family members, number of children) and a few zero-one labour market variables (head unemployed, second adults works full time, second adult works part time, log of female leisure). All of these variables affect significantly the budget share of food for at least one educational group, but none of the non-deterministic variables plays a key role in determining the β parameter, i.e. the degree of non-homotheticity.

In order to avoid the measurement error problems (discussed in Blundell, Pashardes and Weber, 1993) which plague household-level data, we have estimated consistently aggregated budget share equations at the cohort level. This is equivalent to treating all explanatory variables as endogenous, and using year-quarter-cohort dummies as instruments. We concentrated on 10 year-of-birth cohorts and therefore have 430 observations for each equation. In estimation we impose all theory restrictions (homogeneity is never rejected; symmetry is marginally rejected in two of the four equations).

In Table 4, we report parameter estimates for β and descriptive statistics of some key elasticities by educational group. Two things are worth noticing. First, the estimates confirm, not surprisingly, that food is a necessity. Second, the budget elasticities of food expenditure decline

Table 4

Elasticities derived from the demand system

	β (s.e)	budget elasticity at the mean	uncompensated price elast. at the mean
High School dropouts	-.064 (.016)	0.877 (0.03)	-.577 (0.06)
High School Graduates	-.074 (.016)	0.841 (0.035)	-.513 (0.06)
College Dropouts	-.099 (.011)	0.781 (0.025)	-.580 (0.06)
College Graduates	-.166 (.011)	0.616 (0.027)	-.438 (0.08)

Note: Standard errors in parentheses. The standard errors of the elasticities at the mean are computed using the delta method.

monotonically with educational attainment, ranging from 0.88 for high school dropouts to 0.61 for college graduates. Price elasticity is lowest for college graduates and (in absolute value) considerably less than unity.

6.2 Euler equations

In Table 5, we report the results obtained estimating the Euler equation (7). They incorporate the estimates of the demand system discussed above through the price indexes and are obtained by the GMM techniques discussed in section 5 applied to average cohort data.²⁸

In the first 3, columns we present estimates of three different specifications for the entire sample period and 8 cohorts. Because, as stressed above, the 1980 and 1981 survey might be of lower quality we re-estimate the three specifications on a shorter sample period. We report these results in columns 4 to 6.

To the basic specification (7'), we add several variables which are meant to capture the effects of changing family composition and labor supply on the discount factor $\phi(\cdot)$. After trying several specifications we settled on the one reported. The main conclusions we draw are robust to the inclusion of additional demographic variables or to the exclusion of some of the less significant ones.

In the specification reported the discount factor $\phi(\cdot)$ is assumed to depend on seasonal dummies, on the log of family size, on the number of children between the ages of 0 and 15, on a dummy which equals unity if the wife works full time, on the log of annual hours of leisure enjoyed by the wife (computed as 5000 minus the number of hours of work) and a dummy for single individuals. Other variables that were considered during the specification search include the number of children of various ages, a dummy for part time working wife, and the number of earners.²⁹ We also considered variables such as the number of vehicles and dummies for home-ownership.

All variables in the equation, with the exception of seasonals, are instrumented. There are two reasons to use this procedure. First, some of the variables considered are choice variables determined simultaneously with consumption. Second, all of them are subject, given the size of our sample, to measurement error. As argued above the presence of measurement error and therefore of MA(1) residuals makes lag-one instrument invalid. The instruments used were second,

²⁸ No attempt is made to correct the standard errors for the use of generated regressors.

²⁹ The cohort average of a dummy variable (such as that for working wives) measure the proportion of households, within a particular quarter-cohort cell, which satisfy a particular condition.

Table 5
Euler equation for total consumption expenditure
Using results from the demand system

	(1)	(2)	(3)	(4)	(5)	(6)
r	0.559	0.421	0.271	0.668	0.589	0.424
0.43	(0.224)	(0.250)	(0.189)	(0.189)	(0.201)	(0.156)
$\Delta l f m s z$	1.351	1.213	-	1.619	1.563	-
0.17	(0.357)	(0.375)	-	(0.340)	(0.356)	-
$\Delta c h i l d r e n$	-0.498	-0.451	-	-0.487	-0.431	-
0.18	(0.156)	(0.167)	-	(0.155)	(0.163)	-
$\Delta w w$	-0.816	-0.857	-	-1.469	-0.924	-
0.12	(0.442)	(0.435)	-	(0.967)	(0.462)	-
$\Delta l n w l$	-1.437	-1.222	-	-0.696	-1.314	-
0.09	(0.835)	(0.797)	-	(0.751)	(0.902)	-
$\Delta s i n g l e$	-1.214	-1.038	-	-1.211	-1.712	-
0.12	(0.712)	(0.683)	-	(0.828)	(0.773)	-
$\Delta l y$	-	0.121	0.306	0.107	0.090	0.247
0.25	-	(0.090)	(0.068)	(0.155)	(0.068)	(0.048)
Sargan crit.	17.18	18.29	25.70	18.40	19.31	30.29
(p-value)	(0.64)	(0.50)	(0.36)	(0.56)	(0.44)	(0.18)
Est. period	81:3-90:4	81:3-90:4	81:3-90:4	82:3-90:4	82:3-90:4	82:3-90:4
Cohorts	1-8	1-8	1-8	1-8	1-8	1-8
Obs.	288	288	288	256	256	256

Notes: Asymptotic standard errors in parentheses. All specifications include a constant and three seasonal dummies. The instruments set is the same across columns and includes the second to fourth lag of consumption growth, inflation, nominal interest rates and labor income growth, the second and third lag of all the other variables in column (1), the second and third lag of the number of earners, three seasonal dummies, age, age squared and a constant. The numbers under the variable names are R^2 of the first step regression on the 81:3-90:4 sample.

Legend: r = real interest rate, $l f m s z$ = log of family size, $w w$ = dummy for wife working full time, $l n w l$ = log of wife's annual hours of leisure, $s i n g l e$ = dummy for single consumers, $l y$ = log of labor income, $c h i l d r e n$ = household members between 0 and 15.

third and fourth lag of interest rates, consumption and income growth and inflation, the second and third lag of all the variables considered in the discount factor, nad of the number of earners, age, age squared, a constant and three seasonal dummies. The reasons behind this choice of instruments are discussed below.

In the column with the variable names we report the R^2 's of the first step regression of each of the variables on the instruments (for the longer sample). In columns (1) and (4) we report our favourite specification. The coefficient on the interest rate is related to the elasticity of intertemporal substitution (ϵ), as evident from equation 7'. When preferences are homothetic so that $b(p) = 1$, such a coefficient is actually equal to the ϵ . In practice, the estimated parameters of our demand system and the behavior of relative prices, imply a very small variability of $b(p)$ which, in the sample, ranges from 0.99 to 1.01. Therefore, for all practical purposes, we can consider the estimated coefficients on the real interest rate as an estimate of the ϵ . This coefficient is estimated at 0.56 (s.e. = 0.22) for the whole sample and at 0.67 (s.e. = 0.19) for the shorter period. This relatively high estimate of the ϵ is consistent with the result reported for the UK in Attanasio and Weber (1993).

As is evident from the estimates, some of the demographic variables are quite important. Attanasio et al. (1994) plot the life cycle profiles for the discount factors implied by a specification similar to the present one and estimated on the same data set. They show that, when used to solve and simulate a life cycle model, these estimated discount factors are able to generate not only the hump shaped profile which characterizes life cycle consumption, but also the differences across education groups.³⁰

In columns (2) and (5), we add to the specification in columns (1) and (4), the rate of growth in labor income (including transfers).³¹ The finding of a non-zero coefficient on this variable has been interpreted in the literature as evidence of excess sensitivity of consumption to labor income. The coefficient we estimate is relatively small and is not significantly different from zero. This result cannot be explained with the fact that the instruments used are unable to capture the variability in income growth: the R^2 of the first step regression for income growth is approximately 0.25, which is higher than the R^2 for all the demographic and labor supply variables. Indeed, the main motivation for the inclusion of so many instruments and in particular of the lagged values

³⁰ Even though the taste parameters are assumed to be the same across education groups, the discount factors (and therefore the implied consumption profiles) will differ because the forcing variables - demographics and labor supply- will differ across education groups.

³¹ We also tried the rate of growth in total after tax family income obtaining comparable results.

Table 6
Euler equation for total consumption expenditure
Using Stone price index to deflate total non-durable expenditure

	(1)	(2)	(3)	(4)	(5)	(6)
<i>r</i>	0.341 (0.276)	0.149 (0.347)	0.221 (0.263)	0.480 (0.282)	0.331 (0.316)	0.442 (0.218)
$\Delta l f m s z$	1.172 (0.399)	0.948 (0.479)	- -	1.539 (0.383)	1.413 (0.417)	- -
$\Delta c h i l d r e n$	-0.539 (0.169)	-0.453 (0.200)	- -	-0.617 (0.186)	-0.558 (0.192)	- -
$\Delta w w$	-1.551 (0.666)	-1.560 (0.639)	- -	-1.808 (0.665)	-1.826 (0.649)	- -
$\Delta l n w l$	-2.578 (0.835)	-2.486 (1.046)	- -	-3.207 (1.185)	-3.011 (1.144)	- -
$\Delta s i n g l e$	-2.239 (0.912)	-2.157 (0.906)	- -	-2.744 (0.828)	-2.567 (0.987)	- -
$\Delta l y$	- -	0.100 (0.103)	0.268 (0.067)	- -	0.094 (0.089)	0.223 (0.053)
Sargan crit. (p-value)	11.66 (0.92)	12.34 (0.87)	23.32 (0.50)	12.11 (0.91)	13.06 (0.84)	25.99 (0.35)
Est. period	81:3-90:4	81:3-90:4	81:3-90:4	82:3-90:4	82:3-90:4	82:3-90:4
Cohorts	1-8	1-8	1-8	1-8	1-8	1-8
Obs.	288	288	288	256	256	256

of the number of earners, is the attempt to capture a substantial part of income variability. All results in Table 5 are unaffected by a reduction in the number of instruments (even though a few coefficients are estimated with less precision).

In columns (3) and (6) we exclude all demographic and labor supply variables. The coefficient on income is now estimated as almost three times as large as in columns 2 and 5 and it is significantly different from zero. This result confirms that the lack of excess sensitivity in the other columns is not due to the inability of our instrument set to capture income growth variability.

In none of the equations reported the Sargan test of overidentifying restrictions rejects the null. However, given the large number of instruments used, the power of such a test can be low.

We interpret the results in Table 5 as indicating that a flexible specification of the life cycle model which allows for the effect of demographic and labor supply variables and models in a coherent way the aggregation across commodities is not inconsistent with the available evidence on microeconomic behavior.

Having gone through the exercise of estimating a small demand system, a legitimate question to ask is: was it worth the effort? To answer this question, in Table 6 we report the results obtained when the specifications in Table 5 are estimated under the restriction that $b(p) = 1$. This is equivalent to deflate consumption and income growth and the nominal price rate by the inflation in a Stone price index of the different commodities considered. The evidence indicates that the use of the appropriate price index can, indeed, make a difference, especially as far as the estimate of the coefficient on the interest rate is concerned. In particular, we obtain point estimates of the coefficient that determines the elasticity of intertemporal substitution consistently smaller than those reported in Table 5. Furthermore, these estimates are much less precise. The absence of excess sensitivity to income and the failure to reject the overidentifying restrictions, instead, are robust to the use of the more conventional price index used in Table 6.

Overall, a comparison of Tables 5 and 6 indicates that the explicit consideration of aggregation over commodities is important especially to evaluate the magnitude of the curvature of the within period utility function and therefore the degree of substitutability.

7. Conclusions

In this paper, we have shown that the intertemporal model of optimizing behavior for consumption is not inconsistent with US micro data. It is crucial, however, that preferences are modeled so as to take into account changes in family composition and labor supply behavior over

the life cycle. As we recognize in the empirical analysis, such factors are not necessarily exogenous. Modelling fertility and labour supply is beyond the scope of this paper, and our analysis is therefore limited to the estimation of *conditional* preferences.

The introduction of demographic and labor supply variables as determinants of the discount factor may be open to criticism. It could be argued that the large number of parameters we estimate reduces the power of standard tests of overidentifying restrictions or of tests of excess sensitivity of consumption to income. If income is affected by severe measurement error then demographic and labor supply variables may capture the correlation between expected consumption and expected income which would imply a rejection of the model. Furthermore, the sign and size of the coefficients on the demographic and labor supply variables might be difficult to interpret.

The answer to these criticisms requires a further step, which we do not take in this paper. To establish whether the coefficients on demographics (and labor supply) are sensible, one should derive the consumption patterns implied by the estimated parameters, given 'reasonable' demographic profiles as well as income and interest rate patterns. In a recent paper, Attanasio et al. (1994) show that a life-cycle model similar to that presented here is not only able to generate the hump-shaped life-cycle consumption that we observe in reality, but also to explain the differences in the shapes of life-cycle profiles across education groups. This sort of evidence goes some way towards proving that the preference specifications we estimate are reasonable.

We have devoted considerable attention to the issue of aggregation over different commodities. We have shown that ignoring it can affect considerably the estimates of some key behavioral parameters and in particular of the elasticity of intertemporal substitution. This is to be expected if we consider that preferences might not be homothetic (as shown in the analysis of the demand system) and that there have been large changes in relative prices over the 1980s. On the other hand, ignoring the aggregation over commodities does not lead to a rejection of the model.

The results we have obtained contrast sharply with most of the existing literature either in terms of tests of the model (Hall and Mishkin (1982), Zeldes (1989)) or in terms of estimates of behavioral parameters (Runkle (1991), Keane and Runkle (1992)). The contrast is even sharper with papers that use aggregate time series data (Hall (1988), Campbell and Mankiw (1989)). We have shown that these differences can be explained either by the use of a very poor proxy for total consumption or by aggregation problems.

The answer to the question posed in the title of this paper is a qualified yes. We have shown that some of implications of an intertemporal optimization model cannot be easily rejected.

However, we are aware of the limitations of our methodology. As Deaton (1993) has pointed out, tests of the life cycle model based on consumption growth might fail to detect the presence of operative liquidity constraints as they affect the Euler equation only when they are binding. More generally, the Euler equation approach, while enabling us to estimate key behavioral parameters, does not provide much information on the *level* of consumption and, therefore, on saving.

There are several directions in which the present research could and should be developed. We believe more work is necessary to model the relationship between consumption and labor supply. In this paper, we have estimated only conditional preferences. It is important and useful, however, to model household labor supply choices in a rigorous fashion. Furthermore, several aspects of the consumer optimization problem, which were considered only marginally or ignored in this paper, are clearly crucial for a complete understanding of saving and consumption behavior both at the macro and micro levels. An obvious example is expenditure on durables, which is the most volatile component of consumption expenditure over the business cycle. Other important issues are the behavior of the elderly, expenditure on housing, health, education and so on.

Given all these qualifications, however, we claim that the intertemporal optimization model of consumer behavior cannot be easily dismissed and constitutes a useful starting point for the understanding of individual and aggregate consumption expenditure.

Appendix: Estimation and inference

As stated in the text, the construction of an efficient IV estimator has to take into account the presence of MA(1) residuals *for each group* used in estimation and the contemporaneous correlation among the residuals of the different groups.

A further difficulty arises because the variance of the measurement error component of the residuals is a function of cell size. Variable cell size, therefore, may induce a substantial amount of heteroscedasticity. We do not specify a model for heteroscedasticity, but allow for it in the construction of the estimator and in the estimation of its standard errors.

Finally, given the age limitations we impose on our sample, the synthetic panel we use in estimation is unbalanced, in that some cohorts satisfy the age restrictions only after or before a given year. This does not induce selectivity bias (age is exogenous) but requires extra care in computing contemporaneous correlations across cohorts.

We write the estimatable equation as:

$$y = X\beta + u$$

where X denotes the matrix of k explanatory variables, and u is the error term, which contains both an expectational component, and an MA(1) measurement error component. Given this error structure, valid instruments are the exogenous variables in X (seasonal dummies), other deterministic contemporaneous variables (such as age), and second and further lags of the remaining variables. We denote the instruments matrix by Z , and assume its columns are $m \geq k$.

The GMM estimator used in the paper is given by the following expression.

$$\hat{\beta} = P_{zz}^{-1} X'Z(Z'\hat{\Omega}Z)^{-1} Z'y$$

where $P_{zz} = X'Z(Z'\hat{\Omega}Z)^{-1} Z'X$. Its asymptotic variance covariance matrix is given by:

$$V(\hat{\beta}) = P_{zz}^{-1}$$

where $\hat{\Omega}$ is an NT by NT block matrix. Each block on the main diagonal is a T by T matrix which represents the variance covariance matrix of the residuals of one cohort. These matrices have the diagonal and the band surrounding the diagonal different from zero. This reflects the MA(1) structure of the residuals of each cohort induced by measurement error. Both parameters of these matrices are estimated from the residuals of each cohort and are allowed to be different across cohorts. The off-diagonal blocks of $\hat{\Omega}$ represent the correlation of the residuals of different cohorts. We assume that only contemporaneous correlation is possible, so that these matrices are diagonal. Furthermore we assume that this correlation is constant over time.

For this estimator to be feasible we need an estimate of $\hat{\Omega}$. This is obtained from a first round of consistent estimates computed using the formula above with the identity matrix instead of $\hat{\Omega}$. Rather than estimating $\hat{\Omega}$ we use the first round residuals to construct an estimate of $Z'\hat{\Omega}Z$ robust to the presence of heteroscedasticity of unknown form. This estimate is obtained using the following expressions:

$$\widehat{Z'\hat{\Omega}Z} = P_0 + \alpha_1 P_1 + \alpha_2 P_2, \quad 0 < \alpha_i \leq 1 \quad i = 1, 2$$

where

$$P_0 = \frac{1}{N} \sum_{j=1}^N \frac{1}{T_j} \sum_r^{T_j} z_{j,r} z'_{j,r} u_{j,r}^2, \quad \text{and}$$

$$P_1 = \frac{1}{N} \sum_{j=1}^N \frac{1}{T_j} \sum_{t=2}^{T_j} [z_{j,t} z'_{j,t-1} u_{j,t} u_{j,t-1} + z_{j,t-1} z'_{j,t} u_{j,t} u_{j,t-1}]$$

$$P_j = \frac{1}{N} \sum_{j=1}^N \sum_{i=1}^N \frac{1}{T_{ji}} \sum_{t=1}^{T_{ji}} [z_{j,t} z'_{i,t} u_{i,t} u_{j,t} + z_{i,t} z'_{j,t} u_{j,t} u_{i,t}]$$

T_j is the number of periods over which group j is observed, T_{ij} is the number of periods over which groups i and j are both observed, N is the number of groups. α_1 and α_2 are ad-hoc weights used to guarantee that the estimated variance covariance matrix is positive definite in small samples. This procedure (suggested, for instance, by Fuller (1987); see also Newey and West (1987) for a case where α_1 can be chosen optimally), is arbitrary, but does not affect points estimates in a substantial way.

No correction to the variance covariance matrix is made for the fact that some of the regressors we use are generated from estimated parameters (the demand system).

The Sargan (Hansen) test of overidentifying restrictions is given by:

$$\hat{u}' Z (Z' \hat{\Omega} Z)^{-1} Z' \hat{u}$$

which is distributed as a χ^2 with $m - k$ degrees of freedom

References

- Altonji, J. and A. Siow (1987) : "Testing the Response of Consumption to Income with (Noisy) Panel Data", *Quarterly Journal of Economics* , 102, 293-328
- Altug, S. and R.A. Miller (1990) : "Household Choices in Equilibrium" *Econometrica* , 543-70.
- Attanasio, O.P. (1993) : "A Cohort Analysis of Saving Behavior by Us Households", NBER W.P. No. 4454.
- Attanasio, O.P. and M. Browning (1991a): "Consumption over the Life Cycle and over the Business Cycle", NBER W.P. No. 4453.
- Attanasio, O.P. and M. Browning (1991b): "Testing the Life Cycle Model of Consumption: What Can We Learn From Micro and Macro Data?", Stanford University, Mimeo.
- Attanasio, O.P., Banks, J., Meghir, C. and G. Weber, (1994): "Dynamic Consumption and Saving Decisions in the US and in the UK", Mimeo.
- Attanasio, O.P. and S.J. Davis (1994): "Relative Wage Movements and the Distribution of Consumption", Mimeo.
- Attanasio, O.P., and Weber, G. (1993): "Consumption Growth, the Interest Rate and Aggregation", *Review of Economic Studies* , 60, 631-649.
- Blinder, A. and A. Deaton (1985) : "The Time Series Consumption Function Revisited", *Brookings Papers on Economic Activity*, 2, 465-521.
- Blundell, R., Browning, M. and C. Meghir , (1994) : "A Microeconomic Model of Intertemporal Substitution and Consumer Demand" *Review of Economic Studies* , 61, 57-80.
- Blundell, R., Pashardes, P. and G. Weber (1993): "What Do We Learn About Consumer Demand Patterbs from Micro Data?" *American Economic Review* , 83, 570-97.
- Breusch, T.S., and L.G. Godfrey (1981) "A Review of Recent Work on Testing for Autocorrelation in Dynamic Simultaneous Models", in *Macroeconomic Analysis*, ed. D. Currie, A.R. Nobay and D. Peel, 63-105. London: Croom Helm.
- Browning, M. (1987) "Which Demand Elasticities Do We Know and Which Do We Need to Know for Policy Analysis?", McMaster, mimeo.
- Browning, M., Deaton, A. and M. Irish (1985): "A Profitable Approach to Labor Supply and Commodity Demands over the Life Cycle", *Econometrica* , 53, pp.503-44.
- Browning, M., and C. Meghir (1991): "Testing for Separability of Commodity Demands from Male and Female Labour Supply", *Econometrica* , 59 925-952.
- Campbell, J.Y., and Mankiw, N.G. (1989): "Permanent Income, Current Income, and Consumption", *NBER Macroeconomic Annual*, Cambridge.
- Carroll, C.D. and L.H. Summers (1991): "Consumption Growth Parallels Income Growth: Some New Evidence", in: B.D. Bernheim and J.B. Shoven (eds.) : *National Savings and Economic Performance*, Chicago.
- Chamberlain, G. (1984) : "Panel Data", in: *Handbook of Econometrics* , North Holland.
- Cochrane, J. (1991) : "A Simple Test of Consumption Insurance", *Journal of Political Economy* , 99, 957-81.
- Deaton, A., (1985) : " Panel Data from Time Series of Cross Sections", *Journal of Econometrics* , 30, 109-26.
- Deaton, A., (1993) : *Understanding Consumption*, Oxford Economic Press.
- Deaton, A. and J. Muellbauer (1980): "An Almost Ideal Demand System", *American Economic Review* , 70, 312-36.
- Flavin, M. (1981) : "The Adjustment of Consumption to Changing Expectations about Future Income" *Journal of Political Economy* , 89, 974-1009.
- Fuller, W. A. (1987) : *Measurement Error Models* , New York: John Wiley.
- Gorman, W.T. (1959) : "Separable Utility and Aggregation", *Econometrica* , 27, 469-481.

- Hall, R.E., (1988) : " Intertemporal Substitution in Consumption", *Journal of Political Economy* , 96, 339-57.
- Hall, R.E. and F. Mishkin, (1982) : "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households", *Econometrica* , 50, 461-82.
- Hansen, L.P. and K.J. Singleton, (1982) : "Generalized Instrumental Variable Estimation of Non-Linear Rational Expectations Models", *Econometrica* , 50, 1269-86.
- Hayashi, F. (1985) : "The Effects of Liquidity Constraints on Consumption: A Cross Section Analysis", *Quarterly Journal of Economics* , C, 183-206.
- Hayashi, F. (1987) : "Tests for Liquidity Constraints: A Critical Survey and Some New Results", in T.F. Bewley (ed.) *Advances in Econometrics: Fifth World Congress*, vol. 2, 91-120.
- Hayashi, F. and C. Sims (1983): "Nearly Efficient Estimation of Time Series Models with Predetermined but not Exogenous Instruments" , *Econometrica* , 51, 783-92.
- Keane, M.P. and D.E. Runkle, (1991) : "On the Estimation of Panel Data Models With Serial Correlation When Instruments Are Not Strictly Exogenous" , Minneapolis, Mimeo.
- Lusardi, A.M. (1992) : "Permanent Income, Consumption and Precautionary Saving: An Empirical Investigation" , Ph.D. thesis, Princeton University.
- Mace, B. (1991) : "Full Insurance in the Presence of Aggregate Uncertainty ", *Journal of Political Economy* , 99, 928-56.
- Meghir, C. and G. Weber, (1991) : "Intertemporal Non-Separability or Borrowing Restrictions?", Institute for Fiscal Studies, Mimeo.
- Miron, J.A. (1986) : "Seasonal Fluctuations and the Life Cycle - Permanent Income Model of Consumption", *Journal of Political Economy* , 94, 1258-79.
- Moffitt, R. (1993) : "Identification and Estimation of Dynamic Models with a Time Series of Repeated Cross-Sections", *Journal of Econometrics* , 59, 99-124.
- Newey, W.K., and K.D. West (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix", *Econometrica* , 55, 703-8.
- Runkle, D.E. (1991) : "Liquidity Constraints and the Permanent Income Hypothesis: Evidence from Panel Data", *Journal of Monetary Economics* , 27, 73-98.
- Zeldes, S. (1989): "Consumption and Liquidity Constraints: An Empirical Investigation", *Journal of Political Economy* , 97, 305-46.