

NBER WORKING PAPER SERIES

**HUMPS AND BUMPS IN LIFETIME
CONSUMPTION**

**Orazio P. Attanasio
James Banks
Costas Meghir
Guglielmo Weber**

Working Paper 5350

**NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 1995**

We are grateful to Richard Blundell, Angus Deaton and Andrew Samwick for comments on an earlier version of this paper, 'Dynamic Consumption and Saving Behaviour in the US and the UK' which was presented at the Sardinia conference on dynamic models of household behavior in June 1993. This paper was developed while the first author was a National Fellow at the Hoover Institution and the last author was at Università di Venezia. This study is part of the research program of the ESRC Centre for the Microeconomic Analysis of Fiscal Policy at IFS. Financial support from CNR (Italy) and the European Union (Human Capital Mobility program) is gratefully acknowledged. The usual disclaimer applies. 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.

© 1995 by Orazio P. Attanasio, James Banks, Costas Meghir and Guglielmo Weber. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

HUMPS AND BUMPS IN LIFETIME
CONSUMPTION

ABSTRACT

In this paper we argue that once one departs from the simple classroom example, or 'stripped down life-cycle model,' the empirical model for consumption growth can be made flexible enough to fit the main features of the data. More specifically, we show that allowing demographics to affect household preferences and relaxing the assumption of certainty equivalence can generate hump-shaped consumption profiles over age that are very similar to those observed in household-level data sources, without appealing to alternative explanations (such as liquidity constraints, myopia or mental accounting). The hump-shape is partly attributable to precautionary savings, and partly due to demographics; the tracking (whereby consumption jumps with income) is instead due to the permanent nature of the income shocks.

We use US household-level data to estimate preference parameters and income profiles, and then simulate consumption profiles for different education groups. Our simulated profiles show that the key features observed in the data can be closely matched in simulation. We also show that neglecting uncertainty produces consumption profiles that are 'too flat,' whereas neglecting demographics generates consumption profiles that peak 'too late.'

Orazio P. Attanasio
Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE
UNITED KINGDOM
and NBER

James Banks
Institute for Fiscal Studies
7 Ridgmount Street
London WC1E 7AE
UNITED KINGDOM

Costas Meghir
University College London
Gower Street
London WC1D 6BT
UNITED KINGDOM

Guglielmo Weber
Dipartimento de Scienze Economiche
Università di Padova
Via del Santo, 28
35123 Padova
ITALY

1. Introduction

The life cycle model is a coherent theory of consumption behavior that is theoretically attractive primarily because it treats consumption choices as deriving from an intertemporal maximization problem. The determination of consumption in different time periods can be thought of as analogous to the determination of quantities of different goods in a demand system. The standard model suggests that households smooth their consumption across periods of high and low income to keep expected marginal utility constant. However, the observation that, in many datasets, consumption tracks income as individuals age has been used as evidence that the life-cycle model cannot explain observed consumption behaviour (see Carroll and Summers, 1991).

In this paper we argue that once one departs from the simple classroom example or 'stripped down' life-cycle model¹ the empirical model for consumption growth can be made flexible enough to fit the main features of the data. More specifically, we show that allowing demographics to affect household preferences and relaxing the assumption of certainty equivalence can generate hump-shaped consumption profiles over age that are very similar to those observed in household level data sources, without appealing to alternative explanations (such as liquidity constraints, myopia or mental accounting).

With flexible preferences it is very difficult to reject the life cycle model using traditional statistical tests. It is therefore important to establish the plausibility of estimated preference parameters. While the first order conditions from the consumer's optimization problem have been used to estimate structural parameters for the consumption growth equation, the model has not been used to describe saving and consumption levels. The main reason for this is the difficulty in deriving a closed form solution that relates consumption to the state variables of the problem, except in very simple cases. In addition, we are not aware of any systematic attempt at quantifying in a precise way the combined effects of

¹see Deaton, 1992 for a survey of theoretical and empirical evidence.

demographics, uncertainty and other variables on consumption profiles.

In recent papers Hubbard, Skinner and Zeldes (1995) and Carroll (1994) show that life-cycle consumption profiles can be hump-shaped because of income uncertainty. Our paper adds to the picture the effects of demographics and a non-stationary income process: even though we treat demographic profiles as known and exogenously given, demographic variables affect the solution to the dynamic optimization problem because of their direct influence on the marginal utility of consumption, which is then reflected in perceived income uncertainty. For example: the same shock to income affects households in different ways according to their size and composition. Demographics and persistent income shocks generate both humps and bumps in life-cycle consumption profiles.

The preferences employed in this paper are such that it is not possible to derive the closed form solution for consumption as a function of the state variables. It is therefore necessary to simulate the model to compute consumption paths. The simulation technique we use is an extension of the techniques proposed by Deaton (1991), Marcet and Singleton (1991) and Judd (1992) insofar as we develop an algorithm that parameterises the effects of second and subsequent state variables in such a way as to keep computation time to a minimum. The resulting simulations allow us to construct life cycle and business cycle paths of consumption in the US, isolate the effects of specific variables and also perform counterfactual experiments.

The structure of the paper is as follows. In the next section we present the micro-evidence that needs to be explained — in this case from a US dataset. Section three describes the simulation methodology and section four presents simulation results which are compared to the profiles described earlier. Section five concludes.

2. 'Life Cycle' Behavior in the United States

We draw the data used in this paper from the major US survey of household income and expenditure — Consumer Expenditure Survey (CEX). Data are available on a continuous basis throughout the 1980's and early 1990's, and contain detailed information on income, expenditure and household characteristics. The sampling frame used is representative of the US population and thus can be aggregated, at least in principle, to estimate aggregate consumption. Unfortunately the quality of the income data is worse than that of consumption but the figures are not totally unreasonable. The correlation coefficient between aggregate consumption growth and the corresponding aggregate from the micro data is 0.71.

It should be stressed that the definition of consumption used in the survey is very different to that used in the National Account Statistics so that a perfect match is impossible. For instance, National Accounts consumption includes imputed rents for homeowners while the CEX consumption does not. National Accounts includes total health expenditure regardless of the proportion covered by insurance while the CEX includes only out-of-pocket health expenditures. The correlation coefficient between rates of growth of aggregated micro income and National Account income is as low as 0.39 for the US (see Attanasio, 1994).

One of the main points of this paper is that the life cycle model with a flexible and realistic preference specification is able to generate the life cycle behavior we observe in the data. It is, therefore, important to present the main features of the life cycle patterns of consumption and other related variables that are found in this and many other micro data sources, and that require explanation.

2.1. Life-time profiles for consumption, income and demographic indicators

Although the CEX is a panel data set each household is observed only 4 times (over a period of one year). Therefore, to study life cycle behavior we

Table 2.1: CEX Cohort Definition

Cohort Number	Year of Birth	Age in 1980	Cellsize
1	1955–59	21–25	735
2	1950–54	26–30	658
3	1945–49	31–35	562
4	1940–44	36–40	420
5	1935–39	41–45	342
6	1930–34	46–50	317
7	1925–29	51–55	323
8	1920–24	56–60	331


divide our sample into more or less homogeneous groups that can be followed over time. This procedure allows us to study the dynamics of group averages for the relevant variables. One obvious grouping criterion when studying life cycle behavior is the year of birth of the household head. Forming groups (typically called ‘cohorts’) on the basis of such a variable allows us to study the life cycle behavior of individuals that come of age at the same time². Of course, other variables beside birth year can be used to form groups or sub-groups. In what follows we use the educational attainment of the household head, but one could use region of residence, race and so on. The tradeoff is between within group heterogeneity and cell size. Given the large amount of unobserved household heterogeneity, it is essential to work with cells of reasonable size.

We start by presenting life cycle profiles for several variables ignoring any variable other than year of birth to form cohorts. These are the data used in the estimation of the Euler equations in the next section. The cohorts are defined by five year bands, their definitions being given in Table 2.1.


In Figure 2.1, we present the life cycle profiles of log non-durable consumption and three other variables which are going to be relevant in our analysis: log disposable household income, log family size and wife’s annual hours of work

²It has been pointed out in several papers that, in the presence of strong generational effects, interpreting cross-sectional age profiles as life cycle profiles would be extremely misleading.


Figure 2.1




US- Log consumption life cycle profile



US- Log income life cycle profile



US- Log family size age profile



wife's annual hours of work


In the graph, each variable is plotted against age and the data points of each cohort are joined by a continuous line. Because a cohort is defined by a five year interval and the period covers more than five years the profile for each cohort will overlap with that of the adjacent cohort³. The averages are taken quarterly and are used in the estimation of the Euler equation discussed below.

Notice that, in the top two panels both consumption and income profiles are hump shaped. This observation, and the fact that the shape of the two profiles seems to covary across occupation and education groups, has been interpreted as evidence against the life cycle model. In the bottom two panels we plot log family size and female annual hours of work (not conditional on participation). Notice that family size is also hump shaped. Two features are noticeable about female labor supply. First, there is substantial variability of female labor supply over the life cycle and strong cohort effects can be identified. Second, there is no visible dip in average female labor supply corresponding to child bearing ages.


In Figures 2.2 to 2.5, we present the life cycle profile for the same variables for the four education groups considered in the simulations in the following section: high school dropouts, high school graduates, some college and college graduates⁴. As stressed in the introduction, the fact that differences in income profiles are mirrored in differences in consumption profiles has been interpreted as evidence

³The idea of drawing age profiles for cohorts of individuals born in the same year was developed in Browning, Deaton and Irish (1985) and is consistent with the life-cycle theory. At any point in time individuals of different age also differ in their year of birth. Cross-sectional age profiles reflect both genuine age effects, and differences across year-of-birth cohorts, or cohort effects. With at least two points in time, cohort effects can be controlled for, by taking separate averages over individuals of a certain age according to their year-of-birth. These averages will differ if individuals born later have access to higher life-time resources, say, but may also differ because the sampling period is not the same (take individuals 30 years of age: those sampled in 1980 would be born in 1950; those sampled in 1985 would be born in 1955. If we find a systematic difference between the two groups of 30-year old we may attribute it to cohort effects or to the different sampling period). The latter is a business cycle effect. We interpret the vertical distances between age-profiles as pure cohort effects. This way we neglect the business cycle interpretation, which may be relevant when the vertical distance varies across cohorts in the same way over time.

⁴These cohort data are computed using annual averages to maintain an adequate number of observations in each cell once we split by education group.




Log non-durable expenditure
Figure 2.2

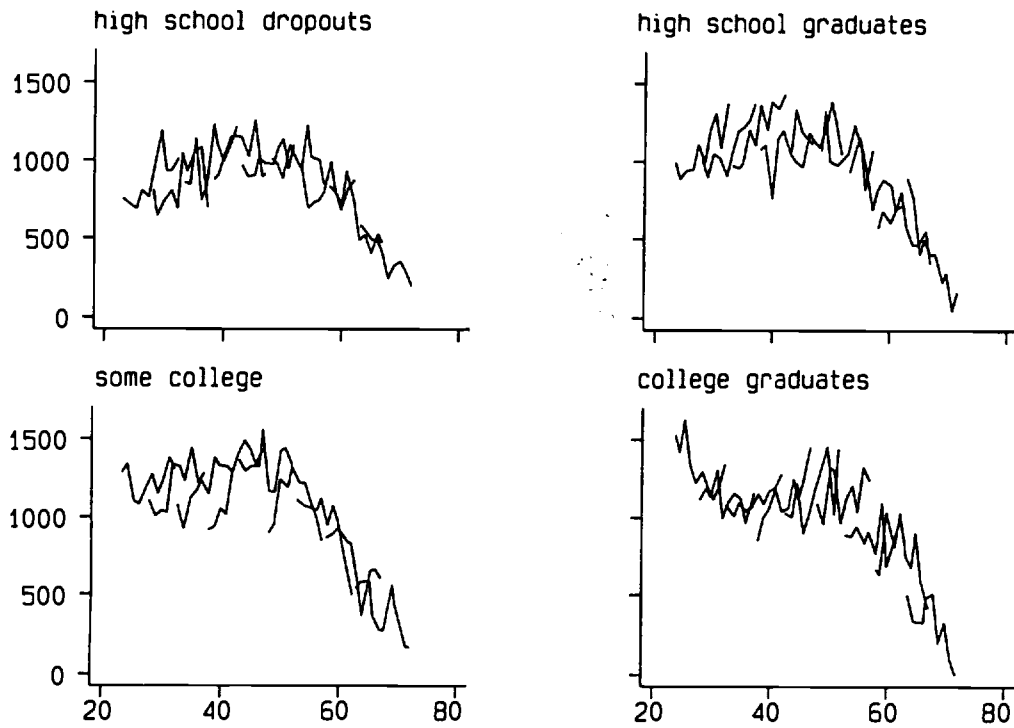


Log after-tax income

Figure 2.3




Log family size
Figure 2.4



Wife's annual hours of work

Figure 2.5



Log 'net' income

Figure 2.6

against the life cycle model. It is therefore interesting to notice that the life cycle pattern of other variables such as family size and female labor supply is also considerably different across education groups. Notice, for instance, that the life cycle profile of log family size is lower and steeper for higher educated individuals. Furthermore, the peak occurs considerably later in the life cycle than for lower educated individuals. On the other hand, college graduates, is the only group for which there is a visible decline in female labor supply corresponding to child bearing ages.

In Figure 2.6 we present age profiles for a new variable: disposable income net of ‘precommitted expenditure’, or net income. This variable is the relevant income measure for the simulations we discuss in the what follows, and is defined as the difference between household disposable income and the sum of expenditure on durable goods, education and health. This last component, health expenditure, is of particular relevance for elderly households, and explains why the net income profile keeps falling well into retirement age. Relatively small cell size is responsible for some of the bumps among the relatively old, but we shall only take a smoothed out version of these profiles in our application.

2.2. Preferences and estimation of the life-cycle model

The maximization problem we consider is the following:

$$\text{Max} \sum_{j=1}^T \beta^j U(C_{t+j}, Z_{t+j}, W_{t+j}, v_{t+j}) \quad (2.1)$$

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

where C , y , A and R are consumption, income, assets and interest rate respectively. The definition of consumption is discussed below. Z is a vector of observable endogenous variables (such as the labor supply of various household members), W is a vector of observable variables that are considered exogenous

for the determination of consumption but that can affect the marginal utility of consumption; v represent unobservable preference shocks. As usual in this kind of study, we will be limited to preferences which imply the presence of additive shocks for the first order conditions we consider.

We specify the within period utility function in the following way:

$$U(C_t, Z_t, W_t, v_t) = \frac{1}{1-\gamma} C_t^{1-\gamma} \exp(\theta'_1 W_t + \theta'_2 Z_t + v_t) \quad (2.2)$$

where γ is the reciprocal of the elasticity of intertemporal substitution. We could, in addition, make γ a function of a number of observable demographic variables as in Blundell, Browning and Meghir (1994) for the UK⁵.

One possible interpretation of the role of the W and Z variables in equation (2.2) is as determinants of the discount factors. The level of utility achieved by a given amount of consumption expenditure depends upon both family composition and other demographic and labor supply variables. Rather than parametrizing adult equivalent schemes and the way in which labor supply choices affect utility in a precise way, we prefer the relatively flexible specification in (2.2). This is equivalent to letting the data determine the way in which demographic variable enter the utility function and to modeling preferences conditional on labor supply behavior (see Browning and Meghir (1992)).

Implicit in the formulation used above is the assumption that the components of expenditure excluded from C influence utility in a separable way.

From equations (2.1) and (2.2) it is possible to derive an Euler equation for consumption which is linear in the parameters and in the residuals. This procedure is standard in the literature; examples of applications of these techniques are, *inter alia*, the papers by Altug and Miller (1991) and Attanasio and Weber (1993). The individual Euler equation can be then aggregated for a particular

⁵We experimented with allowing demographic variation in the estimation of both γ and θ but found insufficient information in the data to identify separate effects. Although we did not test formally where to include demographic variation we use the 'taste shifters', θ , because of their intuitive appeal.

group (such as a cohort) and estimated on average group data. Therefore, to estimate preference parameters it is not necessary to have panel data; a time series of cross sections is sufficient. The advantage of having micro data consists in the ability to control the aggregation process directly. Consistency is obtained letting T , the number of time periods, go to infinity: there is no need to assume the existence of complete markets, as done by Altug and Miller (1991) who achieve consistency keeping T fixed and increasing the number of households N .

The residuals of the Euler equation for consumption consist of three components. The first is an expectational error which arises from the difference between the expected and the actual marginal utility of wealth. The second derives from the presence of unobservable taste shifts in the utility function (v_t), while the third arises from the fact that the cell size is finite and reflects the difference between sample and population means. The first component is presumably white noise and should be uncorrelated with all the information available at time t . Because we take averages of *levels* and subsequently compute first differences the third component is an MA(1) with a coefficient of minus unity. Note that if the size of the year cohort cells is large enough, this component becomes negligible. The stochastic properties of the second component are not obvious. We assume that this residual is uncorrelated with the (group) instruments we use. The presence of the MA(1) component forces the use of instruments dated $t-2$ and earlier.

The labor supply variables on which we condition are treated as endogenous and therefore instrumented. The instrument set includes three groups of variables: first, deterministic variables such as seasonals and a polynomial in age, second, the second to fourth lag of consumption and income growth, the interest rate, and the inflation rate; the second and third lag of the demographic variables used in the equation. The standard errors are computed using standard formulae robust to the presence of heteroscedasticity.

2.3. Estimation of structural parameters

In view of the main task of this paper which is to simulate a life cycle model, the specification of the Euler equation is very parsimonious. Relative to other studies which have analyzed the same data, the specification presented below differs in two respects. First, the demographic and labor supply variables that affect the discount rate are very few. They are, however, able to capture the main features of the data and eliminate the evidence of excess sensitivity.

Second, we do not consider the possibility of varying the elasticity of intertemporal substitution. The presence of variables affecting the elasticity of intertemporal substitution would imply the presence, in the log-linearized Euler equation that we estimate, of log-consumption terms interacted with the variables we consider to be the determinants of the elasticity. The consideration of interaction terms that do not exhibit much variability (at least in our US data) would introduce a substantial amount of collinearity in the equation making the overall estimates rather imprecise.

The Euler equation is estimated on quarterly average cohort data with seasonal preference shifts. The estimation is performed on several cohorts simultaneously allowing the identification of life cycle effects which could not be measured on aggregated data. Details about the econometric problems relevant in this context can be found in the papers cited above.

The results are reported in Table 2.2. After some specification search we allow the discount factor to depend on the log of family size and the log of the number of annual leisure hours of the spouse, computed as 5000 minus the number of work hours. The equation is estimated using data from the 1982-1990 CEX Survey. Similar results were obtained adding the 1980 and 1981 data. In estimation, we use 8 cohorts, each defined by a 5 years band. The synthetic panel is unbalanced in that some cohorts enter the sample later than 1980 and others leave it earlier than 1990 because in estimation we do not consider individuals younger than 25 or older than 60.

Table 2.2: CEX Euler equation estimates

Variable	Coefficient	Standard Error
S1	0.0463	0.0111
S2	0.0473	0.0114
S3	0.0632	0.0115
Constant	-0.0449	0.0116
log(family size)	0.7935	0.3744
log (spouse's leisure)	0.0902	0.0725
Real interest rate	0.7324	0.4218


Sargan test of over-identifying restrictions = 11.74 (17 df)

The coefficient on the interest rate, which measures the elasticity of intertemporal substitution, is estimated at 0.73 with a fairly large standard error (0.42). The coefficient on log family size is quite significant while that for female labour supply is estimated without much precision. The Sargan criterion for overidentifying restrictions does not reject the null hypothesis.

We have considered additional demographics and labor supply variables as possible determinants of the discount factor. The parameters were not estimated with much precision and the main results did not change. Therefore, we prefer to report the more parsimonious specification.


We have also investigated, without much success, the possibility of estimating the Euler equation separately for different education groups. The estimates turned out to be extremely noisy, probably because the relatively small size of the cells increases the variance of the measurement error component of the residuals considerably.

Given the estimated parameters, one can construct an estimate of the discount factor in equation (2.2) with the v_t term set to its mean value of zero — $\exp(\theta_1'W_t + \theta_2'Z_t)$. Even though we assume that households in different education groups have the same preference parameters, the discount factors can differ substantially because of differences in the life-cycle patterns of demographics and



discount factors

Figure 2.7



Smoothed discount factors
Figure 2.8

labor supply variables.

In Figure 2.7 we plot the unsmoothed quarterly profiles for the discount factors of the four education groups. Even though the small cell size makes these pictures quite noisy, substantial differences among the four groups are evident.

In Figure 2.8, we plot the annual smoothed profiles used to solve the model. These profiles are computed, for each education group, as follows. First, we aggregate the quarterly data to obtain annual series. We then smooth the components of the discount factor for each education group by regressing the average cohort data for that component on a 5th order polynomial in age and cohort specific intercepts. We finally take, for each component, the smoothed profile for an arbitrarily chosen cohort and construct the discount factors using the estimates from table 2.1. Notice that because what matters for the solution is the rate of change in discount factors over age, the choice of the cohort is immaterial.

The discount factors are extremely important in determining the shape of the life-cycle profile of consumption. Notice that these average profiles are extremely different across education groups and, in particular, are much steeper for more highly educated individuals. There is nothing in the estimation of the Euler equation, in which we imposed the same coefficient on the education groups, which guarantees these differences in the profiles. They are generated by differences in demographics and labor supply behavior.

3. Simulation of the life-cycle model

In this section, we show how a flexible version of the life cycle model can be used to characterize and quantify the differences in consumption and saving behavior across countries and over time. In Section 3.1, we outline the model and the estimation technique used. In section 3.2 we discuss the solution and simulation techniques.

3.1. Solution and Simulation Techniques

Having estimated the structural parameters in (2.1) and (2.2) we first solve the model by a backward recursion method and then simulate it. To implement this methodology, however, we need to be specific about the processes which generate income, interest rates and demographics and make a number of stringent assumptions which we list below without justification. We assume that:

- Households lives are finite and life length is known in advance. We also assume that households leave no bequest.
- The life-cycle profile of demographic variables is deterministic and known;
- Net income (Y) until retirement is made of two components: a deterministic one (T) which is calibrated on the basis of available data and a stochastic one (W).

$$Y_t = T_t W_t \quad (3.1)$$

The stochastic component is a random walk with MA(1) innovations.

$$W_t = W_{t-1} v_t \quad (3.2)$$

$$v_t = e_t e_{t-1}^\theta \quad (3.3)$$

This specification is roughly in accordance with the available evidence on earnings found in micro data (see MaCurdy (1983) and Abowd and Card (1989));

- The retirement age is exogenous. However, even after retirement net income uncertainty persists, because of unpredictable changes in health which affect committed expenditure. The deterministic component of net income is given by a smooth approximation of observed net income (shown in Figure 2.6);

- The interest rate can be fixed or stochastic; however, it does not exhibit persistence. Our methodology allows the introduction of a persistent interest rate in a straightforward manner.

Of crucial importance to our analysis is that durable expenditure as well as other components of consumption such as health and personal education expenditure are considered exogenous. For this reason we subtract these items from disposable income, and use net income instead. Therefore, the measure of income we use both to find solutions and in simulation is given by total labor income (inclusive of transfers) minus what we call 'committed expenditure'. Alternatively one could assume that the utility function in (2.1) is defined in terms of total consumption (which differs from total expenditure) and use total income. The problem with this latter approach is that the consumption figures obtained out of the simulations are not directly comparable to any immediately observable variable.

The approach we take to solve the model is to construct an approximation of an age dependent consumption function which satisfies both the first-order condition derived from (2.1) and an intertemporal budget constraint. At each age one takes the consumption function at the following age as given. The model, therefore, needs to be solved backwards starting from a consumption function for the last period of life. The method we use is an extension of that used by Deaton (1991).

It is convenient to rewrite the intertemporal budget constraint in terms of what Deaton defines as 'cash-in-hand' ($X_t = A_t + Y_t$):

$$X_{t+1} = (1 + R_{t+1})(X_t - C_t) + Y_{t+1} \quad (3.4)$$

The nature of the problem implies that the consumption function at each age is a function of two state variables: the current value of cash-in-hand (to income ratio) and the current shock to income. This latter variable is relevant only

because it gives information about future income. In the absence of autocorrelation, one state variable would suffice. When this is the case, Deaton's method consists in choosing a grid of points for the state variable and solving for the consumption function at each of these points by standard non linear methods. The consumption function at points others than those on the grid is obtained by interpolation.

Consumption in the last period T is given by the no bequest condition. Consumption in T-1 is a function of X_{T-1} and has to satisfy the Euler equation between T-1 and T. Omitting the demographic variables and the deterministic component of income for notational simplicity, we have that:

$$C_{T-1}^{-\gamma} = F(X_{T-1}, T-1)^{-\gamma} = E_t[C_T^{-\gamma}\beta(1+R_T)] \quad (3.5)$$

where, in turn, C_T is given by:

$$C_T = X_T = (1+R_T)(X_{T-1} - F(X_{T-1}, T-1)) + y_T \quad (3.6)$$

Because the income process is non-stationary, the maximization problem faced by the consumer needs to be re-written so to obtain an Euler equation in terms of stationary variables. Given the homotheticity of the utility function and the linearity of the budget constraint, this can be achieved if all the variables are expressed in terms of ratios to current income. Because the rate of income growth is a stationary stochastic process this gives an Euler equation which involves only stationary variables.

We denote the variables expressed as ratios to current income with smaller cases and re-write equations (3.5) and (3.6) as follows:

$$c_{T-1}^{-\gamma} = F(x_{T-1}, T-1)^{-\gamma} = E_t[(g_T c_T)^{-\gamma}\beta(1+R_T)] \quad (3.7)$$

$$g_t c_T = (1+R_T)(x_{T-1} - F(x_{T-1}, T-1)) + g_T y_T \quad (3.8)$$

where $g_T = y_T/y_{T-1}$.

Substituting (3.8) into (3.7) yields an equation for $F(x)$ that can be solved numerically for each of a number of values of x . The integrals in equation (3.7) are computed using Gauss-Hermite quadrature formulae.

Having solved for the consumption function in T-1 for each point on a grid of values for x_{T-1} , we can go back one period and solve for consumption at T-2. The consumption function at time T-2 is a function of x_{T-2} . Now, however, consumption in the next period is not given by the first identity in (3.6) but by the consumption function at T-1. Finding the solution for C_{T-2} involves evaluating the consumption function at $T-1$ at points which are not necessarily on the grid used to find the solution at $T-1$. These values of the consumption function are obtained by interpolation.

Having solved for T-2, we can repeat the process for T-3, and so on, back to the first period. We found that to compute a sensible solution it is important to use a wide and wisely chosen grid for cash-in-hand. It is also advisable to move the grid as one goes backward in finding a solution for consumption. It is important that at each point in time $t-1$, consumption at t is evaluated at points inside the grid used to find the solution at t .

In the case of two state variables one could work, in principle, with a two dimensional grid. While this extension is conceptually trivial, its computational cost can be daunting. If we have N points on the grid, and have to solve a non linear equation for each of them, we would have N^2 points on a bidimensional grid and would also face a more complicated interpolation problem.

Therefore, rather than following this procedure we devised an alternative method that proved to be efficient and reliable. We still consider a grid of values for cash-in-hand and we assume that, for each point of the grid, the consumption function is given by a polynomial in the second state variable, which is the (stationary) innovation to income. We then evaluate the Euler equation at $q+1$ points for the second state variable, where q is the degree of the polynomial

considered. This gives us $q + 1$ equations in $q + 1$ unknowns (the coefficient of the polynomial) which can be solved by standard non linear methods. The advantage of this method is that one has to solve for $N * (q + 1)$ rather than N^2 coefficients. The method can be trivially extended to cases in which there are more than two state variables. Having solved for the consumption function at each point on the grid for the first state variable, one can interpolate each of the coefficients of the polynomial on the second asset variable to evaluate the consumption functions at points not on the grid.

The choice of cash-in-hand as the first state variable (for which we consider a grid) and of the shock to income as the second is not arbitrary: income shocks are completely exogenous and their range is fully determined by the assumptions on the stochastic process e_t . The only role played by the shock to income is to provide information about future income. Cash-in-hand, instead, is an endogenous variable which summarizes all the information about the level of consumption. We, therefore, want to have greater flexibility in the way the consumption function depends on cash in hand than on the shocks to income.

The assumption that the utility function is isoelastic implies that optimal consumption is always bounded away from zero. This means that the consumption function, F , should be positive and such that C_T is positive with probability one. The latter part means that consumers will not want to borrow more than what they will be almost certain to be able to repay by time $T - 1$, so that they can afford a positive level of consumption in the last period. It turns out that the solution algorithm is more efficient if one supplies it with this kind of information by choosing an appropriate functional form. We parametrize F as:

$$F(X_i, y, t) = \frac{\phi_{it}(y)}{1 + \phi_{it}(y)} (X_i + \bar{y}_{min}(y, t)) \quad (3.9)$$

where:

$$\phi_{it}(y) = \exp(\alpha_{it} + \sum_{j=1}^p \beta_{ijt} y^j) \quad (3.10)$$

and $y_{min}(y, t)$ reflects the minimum present discounted value of income in the remaining periods of life under the hypothesis that income and interest rate realizations are consistently the worst the consumer can expect *conditional* on income at time t .

Having found the solution for consumption in terms of the state variables, we can simulate the model and study the aggregate behavior of consumption implied by our preferences and by the pattern of our forcing variables. We can simulate random draws from the income (and interest process) and, using the consumption function, study the reaction of both micro and macro consumption. A number of issues, however, need to be resolved.

As we stressed above, demographic variables play a crucial role in the determination of consumption. To simulate aggregate consumption we can either solve and simulate the model for an average path of demographics or solve and simulate it for many different patterns and aggregate. The latter alternative is obviously more satisfactory, but is more complicated. In this paper we present only simulations for average patterns and a few experiments to study the effects of changes in various demographics on the level of consumption.

Finally, not all parameters can be estimated. For instance, the pure discount rate cannot be identified from the log linearized Euler equation we estimate in the next subsection. Therefore, we fix that parameter at 'plausible' values. The variability and persistence of income (which are crucial parameters to the determination of the consumption function) are extremely difficult to estimate. The main problem lies in the difficulty in distinguishing between uncertainty as perceived by the individual and as measured by the econometrician. In the end, we decided to calibrate these parameters using the results obtained by other authors and to check how our results change when we vary them.


4. Simulation Results

In this section we focus on two aspects of the simulation results — the relative shapes of the consumption profiles across education groups and the relative importance of precautionary saving versus demographic variables. The first set of simulation results aims at establishing whether the life-cycle model, with the preferences specified and estimated above, is able to generate the kind of life cycle profiles for different education groups reported in Section 2. As we mention above, we solve and simulate the model for the average discount factors for each education group. We, therefore, neglect the aggregation across households with different demographic and labor supply profiles.⁶

The first step in the simulation process is to solve the model to get policy functions for the estimated parameters. As stressed above not all parameters can be estimated. Initially, we set the 'pure' discount rate equal to 0.03 (i.e. we set the discount factor β in (2.1) to 0.97). The discount factors used in the solution in this first exercise are those plotted in Figure 2.8. The process for deterministic income minus 'committed consumption' is calibrated using the average cohort data. We regressed this variable on a 5-th order polynomial in age and cohort specific intercepts and used the profile for an arbitrarily chosen cohort. The profiles used for the four education groups are plotted in Figure 4.1.

In Figure 4.2, we show a typical consumption profile simulated for a household where the head is a high-school graduate, and the average profile for 1000 such households. Details of this simulation are discussed below, but a few interesting features are worth noticing. First, individual consumption profiles display both 'humps' and 'bumps'. The humps arise because of both precautionary saving

⁶Solving and simulating the problem for individual (rather than average) demographic profiles and then aggregating the corresponding consumption paths would be more appealing and would allow us to evaluate the extent to which the consumption function aggregates and what the effects on consumption of marginally changing demographic and labor supply patterns are. However, reproducing the cross-sectional variability of demographic patterns, while leaving the average group profiles unchanged, is not easy and is left as a topic for future research.



Deterministic 'net' income

Figure 4.1

Figure 4.2a: One path only, Education = 2





Figure 4.2b: Average paths, Education = 2, N = 1000



and demographic effects; we assess the relative importance of these factors below. The 'bumps', on the other hand, are attributable to the information about future income contained in the current income shock: by taking a high-memory income process we can simulate a high degree of income tracking in consumption. As we do not assume income shocks are common across households these bumps average out in the aggregate profiles.

Before presenting results for other education groups we describe the details of our simulation procedure. In the first set of simulations below, we set the value of the persistence parameter, θ in (3.3), at -0.2 and consider values for the standard deviation of log income innovations of 0.025 and 0.05. These are considerably lower than those reported in some of the best known studies of the time series properties of microeconomic income processes (see MaCurdy (1982) and Abowd and Card (1989)). However, both these studies use models for earnings while the relevant concept for us is disposable non-asset income (minus committed consumption) and they do not consider different education groups separately. Furthermore, it is not obvious how to distinguish that part of earnings variability which is due to measurement error and unobserved heterogeneity from that which reflects genuine uncertainty as perceived by the household⁷. Initially, the mean of the interest rate and the pure rate of time preferences are both fixed at 0.03. The standard deviation of the interest rate (which is assumed to be i.i.d. log normal) is set equal to zero since the results obtained in the cases with interest rate uncertainty are remarkably similar.

Given all these elements, it is possible to apply the techniques described in Section 3 to solve the model. We assume that when income innovations are persistent, the polynomial in equation (3.10) is of order two. We solve for thirty values of the cash-in-hand to income ratio, ranging between -1.75 to 7 in the penultimate period. We check that both during the solution and in the

⁷The issue of how to compute conditional variances for use in a consumption growth equation from repeated cross-sectional data on income is addressed in Banks, Blundell and Brugiavini (1995)

Table 4.1: Actual and simulated consumption profiles, $\delta_0 = r = 0.03$
Age of Peak and ratio of peak consumption to that at age 25

Education	Actual Data		Simulation 1		Simulation 2		Simulation 3	
	Age	Ratio	Age	Ratio	Age	Ratio	Age	Ratio
High school dropouts	37	1.27	36	1.07	40	1.15	53	2.22
High school graduates	41	1.33	36	1.17	39	1.23	59	2.84
Some college	45	1.52	40	1.34	44	1.44	56	3.17
College graduates	50	2.05	39	1.43	44	1.56	73	3.51

Note to Table

Simulation 1 has no uncertainty in income profiles.

Simulation 2 has 2.5% income uncertainty (with -20% persistence)

Simulation 3 has 5% income uncertainty (with -20% persistence)

simulations the cash-in-hand to income ratio stays within these extremes.

Having solved the model for each education group, we simulate the model for 1000 draws of income and average the consumption profiles. The seeds for the random number generator are constant across education groups. In Table 4.1 we report, for each education group, the age at which the consumption peaks and the ratio between the peak and age 25 consumption for the actual and simulated data.

Consumption profiles of better educated individuals are higher and steeper relative to those for less highly educated individuals both in actual and in simulated data. However, in the simulation without uncertainty (simulation 1) the peak age for earnings increases only slightly with education. In actual data, the peak age goes from 37 for the less educated to 50 for college graduates. When we allow 2.5% income uncertainty (simulation 2) we find higher ratios of peak consumption to age 25 consumption, closer to the values observed in the CEX data. Increasing uncertainty even further — to 5% in simulation 3 — pushes these effects to unreasonable levels. Therefore we take simulation 2 as our first benchmark case.

This comparison of simulated to actual peak ages in this table suggests that the pure rate of time preference⁸ could well be different across education groups (more highly educated individuals may well be less impatient than less educated

⁸In a model with time-varying demographic variables, this is the discount rate only if demographics stay constant.

Table 4.2: Actual and simulated consumption profiles (δ_0 variable)
Age of Peak and ratio of peak consumption to that at age 25

Education Group	Actual Data			Simulation	
	Age	Ratio	δ_0	Age	Ratio
High School Dropouts	37	1.27	.040	38	1.06
High School Graduates	41	1.33	.025	41	1.29
Some College	45	1.52	.025	45	1.52
College Graduates	50	2.05	.010	48	1.90

individuals). Therefore in Table 4.2 we show results for the case where the pure discount rate is chosen to vary across education groups in order to match peak ages as closely as possible. Given our chosen degree of uncertainty and the variability in demographic characteristics across education groups, we must allow the pure discount rate to fall from 4% a year for the least educated to 2.5% for High School graduates and those who have some college education and to 1% for college graduates. When we do this, we get a close match in peak ages, and a reasonable match in the ratios of peak consumption to consumption at 25 (with some downward bias for the least educated). These values will be used in the counterfactual simulations of the next section.

These simulations show that the empirical evidence presented by Carroll and Summers (1991) is fully consistent with the life cycle model if patience and education correlate in the way one might expect (highly educated people are more patient, i.e. more willing to postpone their consumption). We should stress that this result is obtained by using preference parameters estimated using the restrictions implied by an Euler equation. There is nothing in the process of estimation which yields this result. The Euler equation is estimated exploiting high frequency variation and the preference parameters are constrained to be the same across education groups.

4.1. Counter-factual simulations

One of the advantages of a simulation algorithm is the ability to artificially restrict demographic variation or uncertainty in the construction of counter-factual experiments. In this section we use this technique to assess the degree to which the ‘humps’ observed in the profiles of average consumption are attributable

Table 4.3: Counterfactual simulations, High-school graduates ($\delta_0 = 0.025, r = 0.03$)
Age of Peak and ratio of peak consumption to that at age 25


Scenario	Peak Age	Ratio
Actual data	41	1.33
1. Uncertainty on, demographic effects on	41	1.29
2. Uncertainty on, demographic effects off	62	1.47
3. Uncertainty off, demographic effects on	37	1.21
4. Uncertainty off, demographic effects off	75	1.15

solely to either time-varying demographic effects or uncertainty. We choose the group of high-school graduates to illustrate these effects.

Figure 4.3 shows the results of these counter-factual experiments. Perhaps the most noticeable feature that emerges from Figure 4.3 is the key role played by demographic variables in changing the age profile of consumption: if we compare scenario 2 (no demographics) to the benchmark case (scenario 1) we see that demographics move the peak age back by a very large amount. Comparing scenario 3 (no uncertainty) to the benchmark case we find a sizeable effect of the precautionary savings motive on the actual peak, but a very minor effect on the peak age.

Results of these simulations are presented in a comparable way to those of the previous section in Table 4.3 below. As in the figure, we define four scenarios. The first allows full demographic and uncertainty effects and just reproduces the simulation for this group in table 4.2. In scenario two, however, we do not allow demographic variables to affect intertemporal discount rates (setting θ_2 to zero in (2.2)). The resulting profile is steeper and peaks much later than the benchmark case in the first scenario. On the other hand, when we allow for demographic effects but turn off the uncertainty the profiles are flatter and the peak age is earlier. If we allow neither to affect consumption growth then we get a steadily rising profile corresponding to the ‘stripped down’ life-cycle model with $\delta < r$.

Figure 4.3: Counterfactual simulations



5. Conclusions

It is a noted feature of household-level data that the age -profile of consumption is hump-shaped and tracks the age-profile of income. This is in apparent contrast to the life-cycle model of consumer behaviour, which in its simplest formulation predicts a smooth consumption profile independently of the shape of the income profile.

In this paper we have considered a life-cycle model with uncertainty and time-varying demographic factors. We have shown that this model is quite capable of generating both the hump and the income tracking of the consumption profile: the hump is partly attributable to precautionary savings, and partly due to demographics; the tracking (whereby consumption jumps with income) is instead due to the permanent nature of the income shocks.

We have used US household-level data to estimate preference parameters and income profiles, and simulated consumption profiles for different education groups. We have found that the key features of the data can be closely matched in simulation. We have also shown that neglecting uncertainty produces consumption profiles that are 'too flat', whereas neglecting demographics generates consumption profiles that peak 'too late'.

References

- [1] Abowd and Card (1989), 'On the covariance structure of earnings and hours changes, *Econometrica*, 57, 411-445
- [2] Altonji, J. and A. Siow (1987) : 'Testing the Response of Consumption to Income with (Noisy) Panel Data', *Quarterly Journal of Economics*, 102, 293-328
- [3] Altug, S. and R.A. Miller (1990) : 'Household Choices in Equilibrium', *Econometrica*, 543-70.
- [4] Attanasio, O.P. (1994) : 'Personal Saving in the US", in J.M. Poterba (ed.) *International Comparisons of Household Saving*, University of Chicago Press.

- [5] Attanasio, O.P., Banks, J., Meghir, C. and G. Weber, (1994): 'Dynamic Consumption and Saving Decisions in the US and in the UK', Mimeo.
- [6] Attanasio, O.P., and Weber, G. (1993): 'Consumption Growth, the Interest Rate and Aggregation', *Review of Economic Studies*, 60, 631-649.
- [7] Attanasio, O.P., and Weber, G. (1994): 'Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey', forthcoming *Journal of Political Economy*
- [8] Banks, J.W., R.W. Blundell and A. Brugiavini, (1995), 'Income uncertainty and consumption growth', IFS Working paper 95/13
- [9] Blundell, R.W., M. Browning and C. Meghir (1994): 'A Microeconomic Model of Intertemporal Substitution and Consumer Demand', *Review of Economic Studies*, 61, 57-80.
- [10] 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.
- [11] Browning, M. and C. Meghir (1991): 'Testing for Separability of Commodity Demands from Male and Female Labour Supply', *Econometrica*, 59, 925-952.
- [12] 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.
- [13] Carroll, C.D., (1994) : 'How does future income affect current consumption', *Quarterly Journal of Economics*, 109, 111-148.
- [14] Deaton A., (1991) : 'Saving and liquidity constraints', *Econometrica*, 59, 1221-1248.
- [15] Deaton, A., (1992) : *Understanding Consumption*, Oxford: Clarendon Press.

- [16] Hubbard, R.G., J. Skinner and S.P. Zeldes (1994), 'Precautionary Saving and Social Insurance', *Journal of Political Economy*. 103, 360–399.
- [17] Judd, K.L., (1992), 'Projection methods for solving aggregate growth models', *Journal of Economic Theory*, 58, 410–452
- [18] MaCurdy, T.F., (1982), 'The use of time-series processes to model the error structure of earnings in a longitudinal data analysis' *Journal of Econometrics*, 18, 83–114.
- [19] Marcet, A. and K.J. Singleton, (1991), 'Equilibrium Asset Prices and Savings of Heterogeneous Agents in the presence of incomplete markets and portfolio constraints', Carnegie-Mellon University, mimeo