NBER WORKING PAPER SERIES

THE ECONOMICS OF SAVING

Mervyn King

Working Paper No. 1247

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 1983

This paper was prepared for the Yrjo Jahnsson Symposium held in Finland, June 1983. I am indebted to numerous colleagues for helpful discussions, but in particular I must thank Tony Atkinson and Nick Stern, my collaborators on the SSRC Research Programme "Taxation, Incentives and the Distribution of Income," as well as Peter Diamond, John Flemming, Roger Gordon, Mordecai Kurz, Jim Poterba, Agnar Sandmo, Bob Shiller and Larry Summers. The research reported here is part of the NBER's research program in Taxation and Financial Markets and Monetary Economics and project in Macroeconomics. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper #1247
December 1983

The Economics of Saving

ABSTRACT

This paper analyses recent contributions to the theory of household saving and examines empirical evidence on the subject. It focuses on (a) the derivation and estimation of first-order conditions for a consumer's optimum life-cycle consumption plan, (b) the conditions under which such conditions may be used to derive an aggregate consumption function, (c) the relationship between constraints in labor and cyclical markets and the notion of a 'representative consumer' in macroeconomic models, and (d) the extent to which existing empirical evidence lends support to a life-cycle model of consumer behavior. Further empirical tests are proposed.

Mervyn King
National Bureau of Economic Research
1050 Massachusetts Avenue
Cambridge, MA 02138
(617) 868-3900

SAVINGS BEHAVIOUR

M.A. King

1. Introduction

The determinants of saving, whether individual or collective, have a major influence on the future welfare of households. There is a substantial literature both on the positive and the normative theories of intertemporal allocation of resources. Enquiries into the nature and causes of the wealth of nations go back a long way.

It is perhaps surprising that given this tradition there remain still some fundamental disagreements about the determinants of household savings and over the circumstances in which government intervention is desirable. An impression frequently conveyed in the literature, for example, is that the empirical status of the life cycle model of savings behaviour is in doubt. Many studies have claimed to reject the model and a typical summary of the literature is that 'for all its elegance and rationality, the life cycle model has not tested out very well... Nor have efforts to test the life cycle model with cross-section micro data worked out very well'. (Courant et al 1982). If there is doubt about the relevance of the life cycle model, then there is equally sharp disagreement about the appropriate policy response to private savings decisions. Few empirical magnitudes can have attracted so much attention and controversy as the impact of social security on the level of private savings and the interest elasticity of savings. Furthermore, the choice of the tax base has returned to the forefront of policy discussion, and this issue centres critically around the optimal tax treatment of savings and its relationship to the current tax system.

There are, therefore, a number of important policy problems for which an understanding of individual savings behaviour is crucial. I shall not attempt to survey the policy debate but will focus instead on the underlying determinants of household consumption and savings decisions. Recent discussions of savings and taxation (both optimal and actual) include Atkinson and Sandmo(1980) Auerbach and Kotlikoff (1981), King (1980),King and Fullerton (forthcoming) Kotlikoff (1983, Sandmo (forthcoming) and Summers (1982). The relationship between pensions and savings is analysed in Aaron (1982), Blinder (1982) and Dicks-Mireaux and King (forthcoming).

I shall argue below that the major failing of existing empirical studies is that, by and large, they have failed to pose the correct questions. There are two reasons for this. First, there is a tendency to regard the life cycle model as an all or nothing model which must explain the behaviour of the entire population or of none of it. It seems that in fact there is evidence that the life cycle model, taken in its broadest sense, is consistent with the observed behaviour of a majority of households, but that there exists a minority of households for which the model appears to be inadequate. The important question is, then, what are the characteristics of those households which appear to deviate from the life cycle pattern? The second reason is that empirical studies test a joint hypothesis of the life cycle model and a particular parameterisation of preferences. If the model fails then the correct inference is that the joint hypothesis can be rejected. One cannot infer from this that the life cycle model is inappropriate. This raises a deeper question of exactly what it is about the life cycle model which is distinctive and which separates it from alternative explanations of savings behaviour. It is immediately clear that one of the problems in testing the life cycle model is the absence of a coherent alternative model with which it can be compared. An alternative rigorous model is required from which to develop empirical predictions in order that we may be able to distinguish one model from the other. I shall assume below that the distinctive characteristics of the life cycle model are that individual behaviour is "forward-looking", that individuals optimise over their expected remaining lifetime, and that the cost of acquiring information about the future is not so prohibitive that only current observations are used to determine current consumption.

The empirical questions which I shall ask will be the following. First, can the observed data be better explained by a model of 'myopic' or 'irrational' behaviour than by the life cycle model? Secondly, is there any evidence that individual behaviour is influenced more by constraints (over and above the lifetime budget constraint) than by preferences? This is a crucial matter for policy purposes, because the main implication of the life cycle model is that individuals can separate their consumption profile from their income profile. Hence government policies directed at individual welfare, which is normally assumed to be defined over consumption, can be aimed at influencing lifetime income. This gives great flexibility about the particular stage in the life cycle at which help can be directed. If, in contrast, individuals face constraints such that the separation of consumption from income can no longer be maintained, then the precise timing of support for individuals becomes important. These issues clearly lie at the heart of the debate over unemployment insurance and social security.

The plan of the paper is as follows. In Section 2 we discuss the life cycle model in both deterministic and stochastic settings. We shall try to derive the empirical predictions implied by the model. Section 3 discusses alternative models, and examines constraints in both labour and capital markets. Section 4 discusses the evidence from aggregate time-series data and microdata for individual households. Bequests are the subject of Section 5, and the conclusions of the survey are summarised in Section 6.

2. The Theory of Household Savings Behaviour

The "life cycle model" of household savings behaviour is normally associated with the names of Modigliani (Modigliani and Brumberg 1954, Modigliani and Ando 1957) and Friedman (1957). But the analysis of the intertemporal allocation of consumption by utility-maximising consumers was stated clearly by Fisher (1907), Ramsey (1928), and Hicks (1939). The contribution of Modigliani and Friedman was to translate the abstract notion of an optimal consumption profile into a model which could be estimated econometrically. In so doing they demonstrated the empirical relevance of the principal implication of the life cycle model, namely the separation of the consumption and income profiles. Current consumption was not determined primarily by current income. Hence Friedman's concept of "permanent income" and Modigliani's stress on the role of wealth in the consumption function.

Duesenberry (1948, 1949), also, demonstrated the weakness of the link between consumption and current income by noting that in the Depression, cross-section data from the 1935-6 Study of Consumer Purchases showed that a large number of families reported expenditure in excess of income for the year. By contrast, in the 1941 survey there were far fewer deficiencies. In other words, the consumption function varied over the trade cycle. Duesenberry attributed this to a relative income effect, by which consumption was positively related not only to current income but also to the ratio of current income to its highest past level. At face value this seems less convincing than the life-cycle or permanent income hypothesis. Professional footballers who base their consumption plans on their highest income ever attained are seen in the bankruptcy courts from time to time but most of their playing colleagues are evidently more careful. In fact, Duesenberry argued for a distributed lag of past incomes (but with a ratchet effect) in the consumption function. The main difference between this formulation and the life-cycle model is that the latter is based on a forward-looking view of behaviour whereas the former stresses the backward-looking aspect of behaviour.

These models represented a departure from the Keynesian consumption function and the empirical work which it spawned. To this day, however, the majority of econometric studies of aggregate consumption have taken income as the main exogenous variable. Yet the theoretical basis for this assumption is weak. In part it rests on "The fundamental psychological law, upon which we are entitled to depend with great confidence both a priori from our knowledge of human nature and from the detailed facts of experience, is that men are disposed, as a rule and on the average, to increase their consumption as their income increases but not by as much as the increase in their income" (Keynes 1936, p 96). It may be argued that the assumption that the marginal propensity to consume out of current income is greater than zero and less than unity is not strong, and in defending his proposition Keynes wrote subsequently, "it does seem to me that I have said one of the simplest and most indisputable things imaginable" (letter to Mrs. E. Gilboy, 1 Feb. 1939, Keynes (1973 p.276). Yet the stability of the relationship between consumption and current income is the cornerstone of the "Keynesian" consumption function, although Keynes himself viewed it as a short-run relationship which, because it was denominated in wage units, shifted with the real wage. Its determinants are usually assumed to be exogenous, and in econometric models of aggregate consumption are represented by fixed coefficients. But as Muth (1960) showed the propensity to consume out of current income depends, in a life-cycle setting, on the stochastic process generating incomes. For any given dynamic model of incomes, a life-cycle theory leads to a particular Keynesian consumption function. One must be careful, therefore, when using aggregate data to discriminate among alternative models of the motives for savings. Schumpeter foresaw the difficulty when, in dismissing Keynes' use of the phrase "fundamental psychological law" as an abuse of language he wrote, "Keynes's well-known psychological law about the propensity to consume... avers that both individuals and societies will, if they experience an increase in income, normally increase their expenditure or consumption but by less than the increase in income. Whether this is so or not, it is a statement of statistically observable fact which Keynes raised to the rank of an assumption. Nothing is gained, except a spurious dignity, by calling it a psychological law" (1954, fn. pp 1059-60).

We shall return to the dynamic structure of the aggregate consumption function in section 4. In this section we concentrate on developing the micro-economic implications of the life-cycle model which may be tested using data for individual households. A nonstochastic model of life-cycle behaviour is explored first, and then we analyse the consequences of introducing uncertainty into the model. The role of constraints (borrowing or liquidity constraints, for example) will be an important issue in later discussion.

2.1 A Nonstochastic Life-Cycle Model

To illustrate the main issues we consider first a simple nonstochastic model of an individual making an optimal consumption plan over the life cycle. No attempt will be made to provide a complete survey of all of the literature on life cycle models (a recent bibliography of the literature may be found in Mann, 1981).

The individual will be assumed to have a well-defined preference ordering over consumption and labour supply in each period and over the value of gifts and bequests made. We denote this by

$$U = U(\underline{C}_1, \ldots \underline{C}_T, L_1 \ldots L_T, B) \tag{2.1}$$

where $\underline{C}_t$ is the vector of consumption of different goods in period t

L_t is labour supply in period t (the length of the period is normalised to unity)

B is the present value at the beginning of period one of lifetime expenditure on gifts and bequests

T is the nonstochastic length of life

Two points may be noted about (2.1). First, concern about succeeding generations is represented solely in terms of the appearance of bequests in the utility function. There are several ways of modelling bequests and these will be discussed below in section 5. At this stage we assume only that bequests are set equal to their optimal value, denoted by B*, and then the allocation of consumption over the life cycle is determined conditional on this value. Whatever bequests are at the optimum, consumption will still be allocated optimally across ages.

Secondly, changes in household composition over the life cycle mean that in practice the model is usually applied to data for households rather than individuals. This raises a number of questions, such as the nature of decision-making within the household and the variation of "needs" over the life-cycle, which we shall discuss only where they are germane to specific aspects of savings behaviour under consideration.[1] In general, however, we shall use "individual" and "household" interchangeably, although of course this ignores a number of interesting questions concerning household formation.

The essence of the life-cycle model is that individuals face only a single constraint, namely the life-time budget constraint which we may write as

$$\sum_{t=1}^{T} \frac{\underline{p}_t{}' C_t}{d_t} + \sum_{t=1}^{T} \frac{w_t(1 - L_t)}{d_t} + B \leq \sum_{t=1}^{T} \frac{w_t}{d_t} + I \qquad (2.2)$$

where $\underline{p}_t$ is a vector of current prices of the different consumption goods in period t

w_t is the wage rate in period t

d_t is the discount factor in period t and equals $\prod_{j=1}^{t} (1+r_j)$

r_t is the nonstochastic interest rate in period t.

I is the present value at the beginning of period one of inheritances and gifts received.

We shall assume that transactions take place at the end of each period. All prices are net of tax.

Equation (2.2) states that the present value of lifetime expenditure on consumption, on leisure, and on bequests, must equal (at most) the present value of an individual's endowment (or "full wealth") of potential earnings and inheritances. The first-order conditions for a maximum of (2.1) subject to the constraint (2.2) are given by

$$\underline{U}_{ct} - \frac{\lambda \underline{p}_t}{d_t} = 0 \qquad t = 1 \ldots T \quad (2.3)$$

$$U_{Lt} + \frac{\lambda w_t}{d_t} = 0 \qquad t = 1 \ldots T \quad (2.4)$$

$$B = B^* \qquad (2.5)$$

where λ is the multiplier associated with the lifetime budget constraint, $\{\underline{U}_{ct}, U_{Lt}\}$ is the vector of partial derivatives of U with respect to $\underline{C}_t$ and L_t respectively. There are two ways in which we may use these first-order conditions for empirical work. The first is to invert the system of equations (2.3) - (2.5) and solve for the optimal consumption and labour supply vector in terms of the exogenous prices and endowments.

$$\underline{C}_t = \underline{c}^t \, (w_1 \ldots w_T, \ \underline{p}_1 \ldots \underline{p}_T, \ r_1 \ldots r_T, \ I) \qquad (2.6)$$
$$L_t = \ell^t \, (w_1 \ldots w_T, \ \underline{p}_1 \ldots \underline{p}_T, r_1 \ldots r_T, \ I) \qquad (2.7)$$

The functions $\underline{C}$ and ℓ have period superscripts which reminds us that the period here represents not clock time but age. Note also that there will in general be differences in age-specific wage rates among cohorts reflecting productivity growth. The difficulty in estimating equations of the form (2.6) and (2.7) is evident. We require information on prices and endowments in each period of the life cycle, and even with large panel data sets we do not have sufficient observations. Typically panel data sets contain fewer than ten observations on each individual. This problem can be overcome to some extent by restricting the form in which preferences are specified. A common restriction is to assume that preferences are additively separable over time which implies that (2.1) becomes

$$U = \sum_{t=1}^{T} u_t(\underline{C}_t, L_t) \tag{2.8}$$

It is easy to see that this specification means that the system of first-order conditions for current consumption and labour supply may be inverted to yield a demand system which, conditional upon the value of the multiplier λ, depends only upon current prices and interest rates. In other words, the influence of exogenous variables for periods other than the current period comes entirely through their effect on the multiplier λ, the value of which is independent of age. If the form of u_t is independent of age but individuals have a nonzero rate of time preference we have

$$u_t = \frac{u(\underline{C}_t, L_t)}{s_t} \tag{2.9}$$

$$s_t = \prod_{j=1}^{t}(1+\rho_t) \tag{2.10}$$

where ρ_t is the rate of time preference in period t.

In this special case the functional form of the demand system is independent of age and the age profiles of consumption and labour supply depend upon the age profiles of real wages, interest rates, and rates of time preference. If we treat consumption as a single composite good (which assumption we shall maintain for the sake of convenience) then consumption and labour supply at age t may be written as

$$C_t = f(\lambda\tilde{\rho}_t, \lambda\tilde{w}_t) \tag{2.11}$$

$$L_t = g(\lambda\tilde{\rho}_t, \lambda\tilde{w}_t) \tag{2.12}$$

where

$\tilde{\rho}_t = p_t \left(\frac{s_t}{d_t}\right)$, the effective price of consumption

$\tilde{w}_t = w_t \left(\frac{s_t}{d_t}\right)$, the effective wage rate

In this two-commodity framework, with consumption and leisure as substitutes, both consumption and labour supply will be higher in periods with higher "effective" real wages. Normalising the price of consumption to unity, we may interpret w as the real wage and r as the real interest rate.[2] In the particular case where preferences are not only additively separable over time but also the intratemporal preference ordering u is weakly separable between consumption and leisure, the age profile of consumption depends only upon the "effective" price of consumption. When the real interest rate equals the rate of time preference this effective price is constant and the age profile of consumption is flat. In general, however, consumption depends upon the profile of wage rates.

It is important to bear this in mind when evaluating evidence relevant to the life-cycle model. Using US data, Thurow (1969) found that both income and consumption rose with age, peaking in the age range 45-54. This might suggest a much higher correlation between consumption and current income than would be implied by a life-cycle model. But, as Heckman (1974) pointed out, the correlation could be explained by the age profile of wage rates. This can be seen from equations (2.11) and (2.12). Consider the special case in which the interest rate equals the rate of time preference. Normalising the price of consumption to unity we have

$$C_t = f(\lambda, \lambda w_t) \tag{2.13}$$

$$L_t = g(\lambda, \lambda w_t) \tag{2.14}$$

Conditional upon the value of λ, these equations may be regarded as compensated demand functions, where compensation is defined as maintaining constant the value of λ. The partial derivatives of (2.13) and (2.14) with respect to the price of consumption and the wage rate may be thought of as substitution effects. Differentiating (2.3) and (2.4) with respect to own prices, it can be seen that because the direct utility function is quasi-concave in consumption and leisure then the compensated (in terms of λ) own substitution effects are negative. This means that the partial derivatives of (2.13) and (2.14) with respect to the real wage (the only price in the two-commodity case) are positive.

Since λ is independent of age it is clear that both consumption and labour supply reach a maximum at the age at which the wage rate reaches its maximum. Consumption, earnings and the wage rate will all be positively correlated, and will reach a peak at the same age. Typical profiles of wage rates reach a maximum at an age of about fifty, and so this model can provide an explanation of the empirical findings of Thurow.

In the more general case in which the interest rate is not equal to the rate of time preference, consumption, earnings and the wage rate peak at different ages. Their time paths may be explored by differentiating the "constant marginal utility of lifetime wealth" demand functions (2.11) and (2.12). An empirical analysis of labour supply in this framework may be found in MaCurdy (1981). We shall return to the general case later when we shall use empirical evidence on age-earnings profiles and labour supply functions to generate profiles for asset holdings over the life-cycle. We shall argue that under most plausible assumptions savings will generally be low early in working life, will rise as earnings increase with age, and then will flatten out becoming negative in retirement. This pattern implies a "hump-shaped" profile for asset holdings over the life-cycle. Although unsurprising, it suggests that failure to observe a hump-shaped age profile for assets might constitute direct evidence against life-cycle savings behaviour. We assess such evidence in section 3.1. Given the difficulty of estimating the life-cycle model directly, either in general form (equations 2.6 and 2.7) because of lack of data on prices over the life-cycle or in the special case of additive separability (equations 2.11 and 2.12) because of the endogeneity of λ, a simple direct test using data on wealth holdings has considerable appeal.

The difficulty of estimating the life-cycle model is illustrated by the specifications employed in the early studies by Friedman and Modigliani. Friedman (1957) assumed that (a) consumption can be regarded as a composite good, (b) labour supply in each period is exogenous, and (c) preferences defined over consumption at different ages and bequests are homothetic. With these assumptions the first-order conditions (2.3) imply that we may write consumption as

$$C_t = k_t . W \qquad (2.15)$$

where W is lifetime wealth (the present value of labour earnings and inheritances), and k_t depends upon the interest rate and household characteristics but is otherwise independent of economic variables. If we define "permanent income" as the annuity value of lifetime wealth, then consumption is proportional to permanent income. The problem is that permanent income is generally unobservable. Friedman used a distributed lag function of current and past incomes with geometrically declining weights. But there is no reason to suppose that this will measure the annuity value of W, and, in particular, given an age profile for earnings, the appropriate weights will be a function of age. More fundamentally, however, the life-cycle model is essentially forward-looking and the use of distributed lags of past incomes, however general a lag structure we permit, is backward-looking. To incorporate this basic insight into an empirical specification requires an alternative approach. This leads us to the second way in which we may use the first-order conditions (2.3)-(2.5).

If we continue to regard consumption as a single composite good and normalise its price to unity, then by taking the ratio of the first-order conditions at successive ages we have

$$\frac{U_{c,t}}{U_{c,t-1}} = \frac{1}{1+r_t} \qquad t=2..T \qquad (2.16)$$

The interest rate here is that paid on assets held at the beginning of period t when consumption and savings decisions in period t-1 have been made. Equation (2.16) is a simple yet rather powerful result. The marginal (lifetime) utility of consumption declines with age at the rate of interest. In other words, for a given interest rate the marginal rate of substitution between consumption in any two neighboring periods is the same. For a particular representation of preference (2.16) can be solved to generate an age profile for consumption and hence savings. Since many of the unobservable characteristics of an individual, such as future earnings, affect current consumption through the shadow price of lifetime wealth, elimination of the multiplier λ produces a more tractable specification. In terms of econometric terminology, λ represents a "fixed effect" for the individual and by taking first differences (or ratios) its influence on the specification may be removed. Since the unobservable fixed effect is usually correlated with the observable independent variables, estimation of equations which include fixed effects ("permanent income", for example) are likely to yield biassed parameter estimates.[3]

To illustrate this approach consider the intertemporal linear expenditure system with lifetime utility given by

$$U = \prod_{t=1}^{T} (C_t - \bar{c})^{\alpha_t} (\bar{I} - L_t)^{\beta_t} \tag{2.17}$$

Substituting (2.17) into (2.16) yields

$$C_t - \bar{c} = \frac{\alpha_t}{\alpha_{t-1}} \cdot (1 + r_t)(C_{t-1} - \bar{c}) \tag{2.18}$$

If we further assume that the α_t coefficients decline at a constant rate with age, then we obtain the following linear regression model for individual consumption

$$C_t = \bar{c} + a_1(1 + r_t) + a_2(1 + r_t)C_{t-1} \tag{2.19}$$

The parameters a_1 and a_2 are independent of age and hence the model could be estimated on aggregate data as well as panel data for individual households. In fact the longer span of most time series for aggregate consumption means that the use of aggregate data might have some advantages because it would enable the variation of real interest rates over time to be exploited.[4]

In the light of the attention paid to the problem of unobservable variables in the early literature on estimation of the life-cycle consumption function, it is perhaps surprising that the approach based on (2.16) was not pursued by such authors. Treatment of the multiplier λ as a fixed effect has been used by Heckman and MaCurdy (1980) and MaCurdy (1981) to analyse labour supply and labour force participation. But it was not until the analysis of <u>stochastic</u> models of life-cycle behaviour (Hall 1978, Sargent 1978), which we discuss in section 2.2, that use of a condition equivalent to (2.16) appeared. This was unfortunate because the use of (2.16) has nothing as such to do with a stochastic model, and the understandable focus in recent work on the incorporation of rational expectations into the life-cycle model has obscured two points which are crucial if the results of estimating the model are to be interpreted as a test of the life-cycle model. The first is that although (2.16) is a general result, empirical tests employing this framework have been embedded in models which assume that preferences are additively separable over time. Without this assumption it may be difficult to solve (2.16) to find the implied difference equation for consumption, and in general the coefficients of the difference equation would be age-dependent. The resulting nonlinear model would not be easy to estimate and, if the model were fitted to aggregate time-series data, there would be insufficient degrees of freedom to identify the model.

The second point is that even if we assume intertemporal separability, there is no reason to impose the further assumption that preferences for consumption and leisure within a given period are separable.[5] If intratemporal separability is assumed, then (2.9) and (2.16) lead to a simple first-order difference equation for the level (or some transformation of the level)of consumption, and this is illustrated in the example given above of the linear expenditure system (see equations 2.17 to 2.19). Even here, however, the coefficients will not be time-invariant if the real interest rate changes over time. With a more general representation of preferences the time series process describing the behaviour of consumption will depend upon the processes governing the evolution of wage and interest rates. The reduced form for consumption will not be a simple first-order difference equation.

The main conclusion from this discussion is that empirical tests of the life-cycle model based on (2.16), although attractive in principle because of the elimination of the unobservable multiplier or fixed effect, are unlikely to be conclusive because they are tests of the joint hypothesis that (a) individuals are life-cycle savers, and (b) preferences may be parameterised in a particular way. In a regression of consumption on its own lagged value, evidence that additional lagged values of either consumption or other variables are statistically significant does not in itself constitute evidence against the life-cycle model. It may equally be interpreted as a rejection of the assumption that preferences are both inter- and intra-temporally separable. Since I am aware of no strong independent evidence for the two separability assumptions, and there is evidence against separability of consumption and leisure (Deaton and Muellbauer, 1980, Chapter 5), it is probably more appropriate to regard (2.16) as a vehicle for research within the life-cycle framework, and not a test of the framework as such.

To test the life-cycle hypothesis itself we shall examine evidence on the amounts of wealth owned by individuals at different ages. We shall compare the observed levels of wealth with the levels which would be predicted by a life-cycle model. Although there is little evidence on the response of savings to wage and interest rates, there have been many studies of the response of labour supply. Using the "stylised facts" about labour supply behaviour we shall construct an indirect utility function describing preferences between consumption and leisure, and compute the implied age profile for wealth. To do this we shall ignore some of the factors which lead to diversity in observed wealth-age profiles such as expenses on child rearing and timing of marriage and birth of children. The labour supply function we shall use is given by

$$L_t = a - b \exp(-\gamma w_t) - cy_t \qquad (2.20)$$

where the price of consumption has been normalised to unity, and y_t denotes "full" expenditure in period t on consumption and leisure (in a static model this is equivalent to Becker's "full income"). This form implicitly assumes intertemporal, though not intratemporal, separability

$$y_t = C_t + w_t(1 - L_t) \qquad (2.21)$$

The choice of the nonlinear labour supply function (2.20) was motivated by the desire to produce a model consistent with the available empirical evidence on life-cycle profiles of hours of work and wage rates, and capable of generating endogenous retirement behaviour. The specification of (2.20) allows this, as shown below. The indirect utility function corresponding to (2.20) may be derived as follows. Using the implicit function theorem and Roy's identity for the demand for leisure we have (omitting age subscripts)[6]

$$\frac{dy}{dw} = 1 - L = (1 - a) + be^{-\gamma w} + cy \qquad (2.22)$$

Integrating we obtain

$$ye^{-cw} = -\left(\frac{1-a}{c}\right) e^{-cw} - \frac{b}{c + \gamma} e^{-(c+\gamma)w} + k \qquad (2.23)$$

where k is the constant of integration which we shall take as the level of indirect utility v. Hence

$$v(w, y) = e^{-cw}\left\{y + \frac{1 - a}{c} + \frac{b}{c + \gamma} e^{-\gamma w}\right\} \qquad (2.24)$$

This indirect utility function is continuous, increasing in income and decreasing in the wage rate (providing $L \leq 1$ i.e. $a - cy \leq 1$).[7] It defines instantaneous utility and with intertemporal separability lifetime utility is defined by[8]

$$U = \sum_{t=1}^{T} \frac{\ln v(w_t, y_t)}{s_t} \qquad (2.25)$$

The individual maximizes U subject to the lifetime budget constraint

$$\sum_{t=1}^{T} \frac{y_t}{d_t} = W \tag{2.26}$$

where W is "full wealth". We shall ignore bequests and inheritances. If the shadow price of full wealth is λ then the evolution of total expenditure (on goods and leisure) over time is determined by the condition

$$\frac{1}{v} \frac{\partial v}{\partial y} = \lambda \left(\frac{s_t}{d_t}\right) \qquad \forall t \tag{2.27}$$

This gives

$$y_t = \frac{1}{\lambda} \left(\frac{d_t}{s_t}\right) - \left(\frac{1-a}{c}\right) - \frac{be^{-\gamma w_t}}{c+\gamma} \tag{2.28}$$

The value of the multiplier λ may be obtained by substituting (2.28) into (2.26).

To justify the specification of (2.20) and (2.25), and also to parameterise the model, we examine evidence on the age profile of wages, earnings, and hours of work. The main exogenous influence on the optimal life-cycle plan is the age profile of wage rates. Maintaining the assumption of perfect certainty we assume that adult life begins at age 15 and ends at 75. From entry into the labour market at age 15 until age 60 wage rates are assumed to be given by

$$\ln w_t = \alpha_o + \alpha_1 (t-15) + \alpha_2 (t-15)^2 \qquad 15 \leq t \leq 60 \tag{2.29}$$

The starting wage is normalized to unity ($\alpha_o = 0$). Wage rates typically rise, reach a maximum in middle age, and then decline. Using data for males from the 1966 Survey of Economic Opportunity, Heckman (1974) found that the maximum wage was reached in the age bracket of 48-50, and a similar result was found by Hurd (1971) using the same data. We assume that the maximum is at age 48 and that the ratio of peak wage rates to the initial wage rate is, in real terms, 1.5. These assumptions define values for α_1 and α_2. After the age of 60, the potential wage rate (net of tax) is assumed to decline rapidly (at a rate μ) toward some minimum level, w_m, according to

$$w_t = w_m + (w_{60} - w_m)e^{-\mu(t-60)} \quad t > 60 \qquad (2.30)$$

This produces a pattern of wage rates which, together with the labour supply function (2.20), imply a sharply decreasing labour supply after the age of 60 which is equivalent to endogenous "retirement", where retirement is defined as a marked reduction in hours worked rather than a constraint that hours worked must be zero. The parameters of the labour supply function are chosen to accord with the following stylized facts. First, hours worked do not vary significantly over the life-cycle in the pre-retirement phase even though wages do vary. Secondly, hours worked peak before the age at which wage rates attain a maximum. This implies that the rate of time preference is less than the rate of interest. For the simulation we take the real interest rate to be 2 per cent per annum, and the rate of time preference to be 1 per cent. Using SEO data for males Heckman (1974) found that hours of work reached a maximum in the age range 39-44, compared with a maximum for wage rates in the age bracket 48-50. In the simulation reported below the wage rate is assumed to reach a maximum at age 48, and the model predicts a peak for labour supply at age 40, with a maximum in earnings at age 45. Empirical studies of earnings have found age profiles for male earnings which peak in the age range 45-50 (Mincer (1976), Psacharopoulos and Layard (1979), Thurow (1969), King and Dicks-Mireaux (1982)). Not all of these studies adjusted for the cohort effect on earnings-age profiles and so their findings are not strictly comparable. Panel data show less evidence of a peak in earnings prior to retirement (US (1976)).

The third stylised fact is that after some conventional "retirement" age, taken to be 60 in our simulation, both potential earnings and hours of work decline quite rapidly. We do not impose the constraint that after some given age labour supply and hence earnings are zero. Instead, the model generates endogenous withdrawal from full-time employment as the wage rate falls and allows for the aged to supply a small number of hours worked after the normal retirement age. This is perhaps more realistic than the assumption that hours of work fall to zero, although I make no attempt to justify the wage equation (2.30) on the basis of statistical evidence. The value of the minimum wage rate in retirement was assumed in the simulation to be 0.1 which may be compared with the maximum pre-retirement wage of 1.5 and μ was taken to be 0.5. In addition to any earnings the individual is assumed to receive social security or a private pension after the age of 60 equal to fifty per cent of the net of tax wage at age 60. The pension remains constant in real terms until death.

To produce a model capable of simulating life-cycle behaviour described by these three stylised facts we choose the following values for the parameters of the indirect utility function [9]

$a = 1.0$

$b = 0.9$

$c = 0.1$

$\gamma = 0.65$

The life-cycle pattern of labour supply generated by the model is shown in figure 1. To calibrate the number of hours worked per week I have assumed that ten hours per day are devoted to sleep and other necessary activities leaving 98 hours a week available to allocate between work and leisure. Hours worked rise early in life as the wage rate increases, remain roughly constant for the bulk of working life (reaching a maximum at age 40), decline slowly as "retirement" approaches, and, finally, fall sharply after the age of 60. It should be emphasised that this pattern is not imposed but is the optimal repsonse to the assumed age profile of wage rates.

The pattern of asset accumulation over the life cycle is given by the equation

FIGURE 1 LABOUR SUPPLY OVER THE LIFE-CYCLE

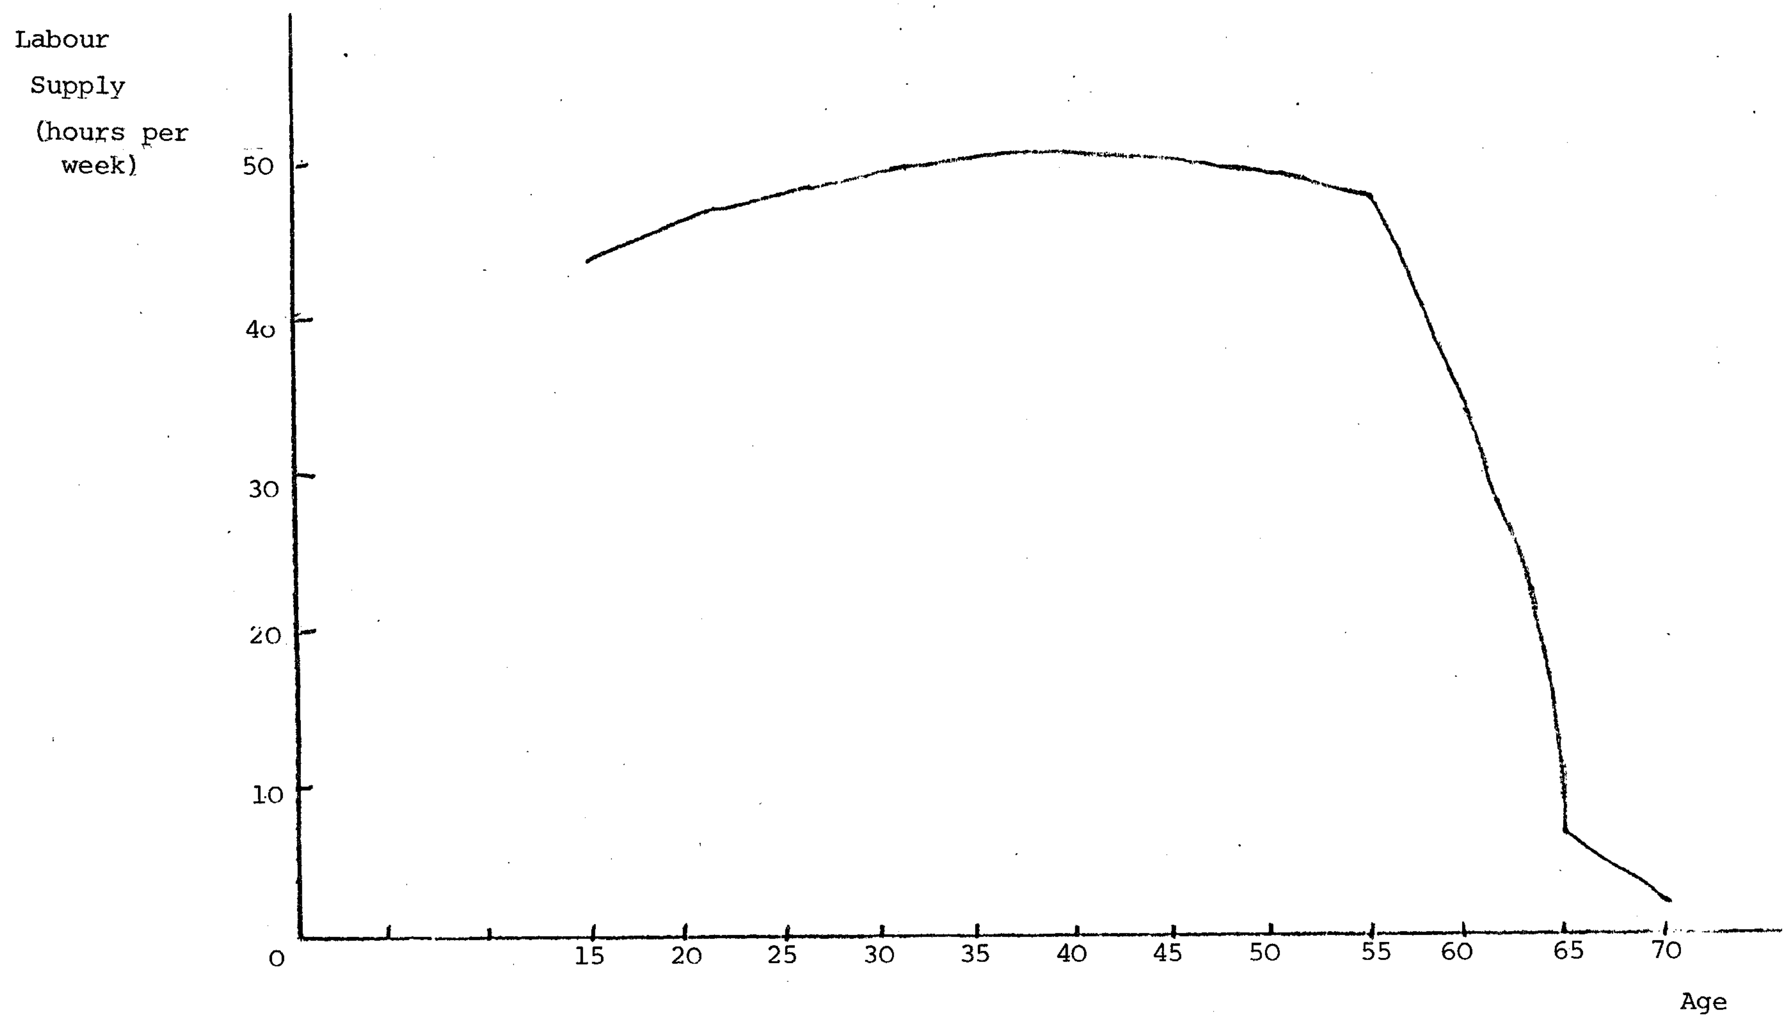

$$A_t = (1+r_t)A_{t-1} + w_t L_t + x_t - C_t \quad t=1...T \quad (2.31)$$

where A_t is the level of assets excluding pension wealth held at the end of period t (note the distinction between assets, A_t, and full wealth, W), and x_t is receipts of social security and pensions in period t. If, as before, we ignore inheritances, the initial condition is $A_o = 0$. The age profile of asset holdings is shown by the solid line in figure 2. It measures the ratio of assets to permanent income where the latter measures the constant annual annuity value of full wealth. It shows the "hump-shaped"pattern of the life cycle model. Holdings of assets rise during working life and are consumed in retirement. Asset holdings reach a maximum before retirement, and in figure 2 this occurs at age 51.

The simulated age profile shown in figure 2 may be compared with actual age profiles estimated fron household data. To do this we require a survey of asset-holdings of a random sample of the population drawn from all ages. Unfortunately many of the available surveys (such as the National Longitudinal Survey of Mature Men used by Diamond (1977) and Diamond and Hausman (1982), and the Retirement History Survey used by Blinder et al (1980)) cover only part of the life-cycle, such as the period immediately prior to and after retirement. But the 1977 Survey of Consumer Finances in Canada collected data on assets and debts for a sample of 12,734 families of all ages. An analysis of these data in relation to the life-cycle model may be found in King

Figure 2 Assets over the Life-cycle

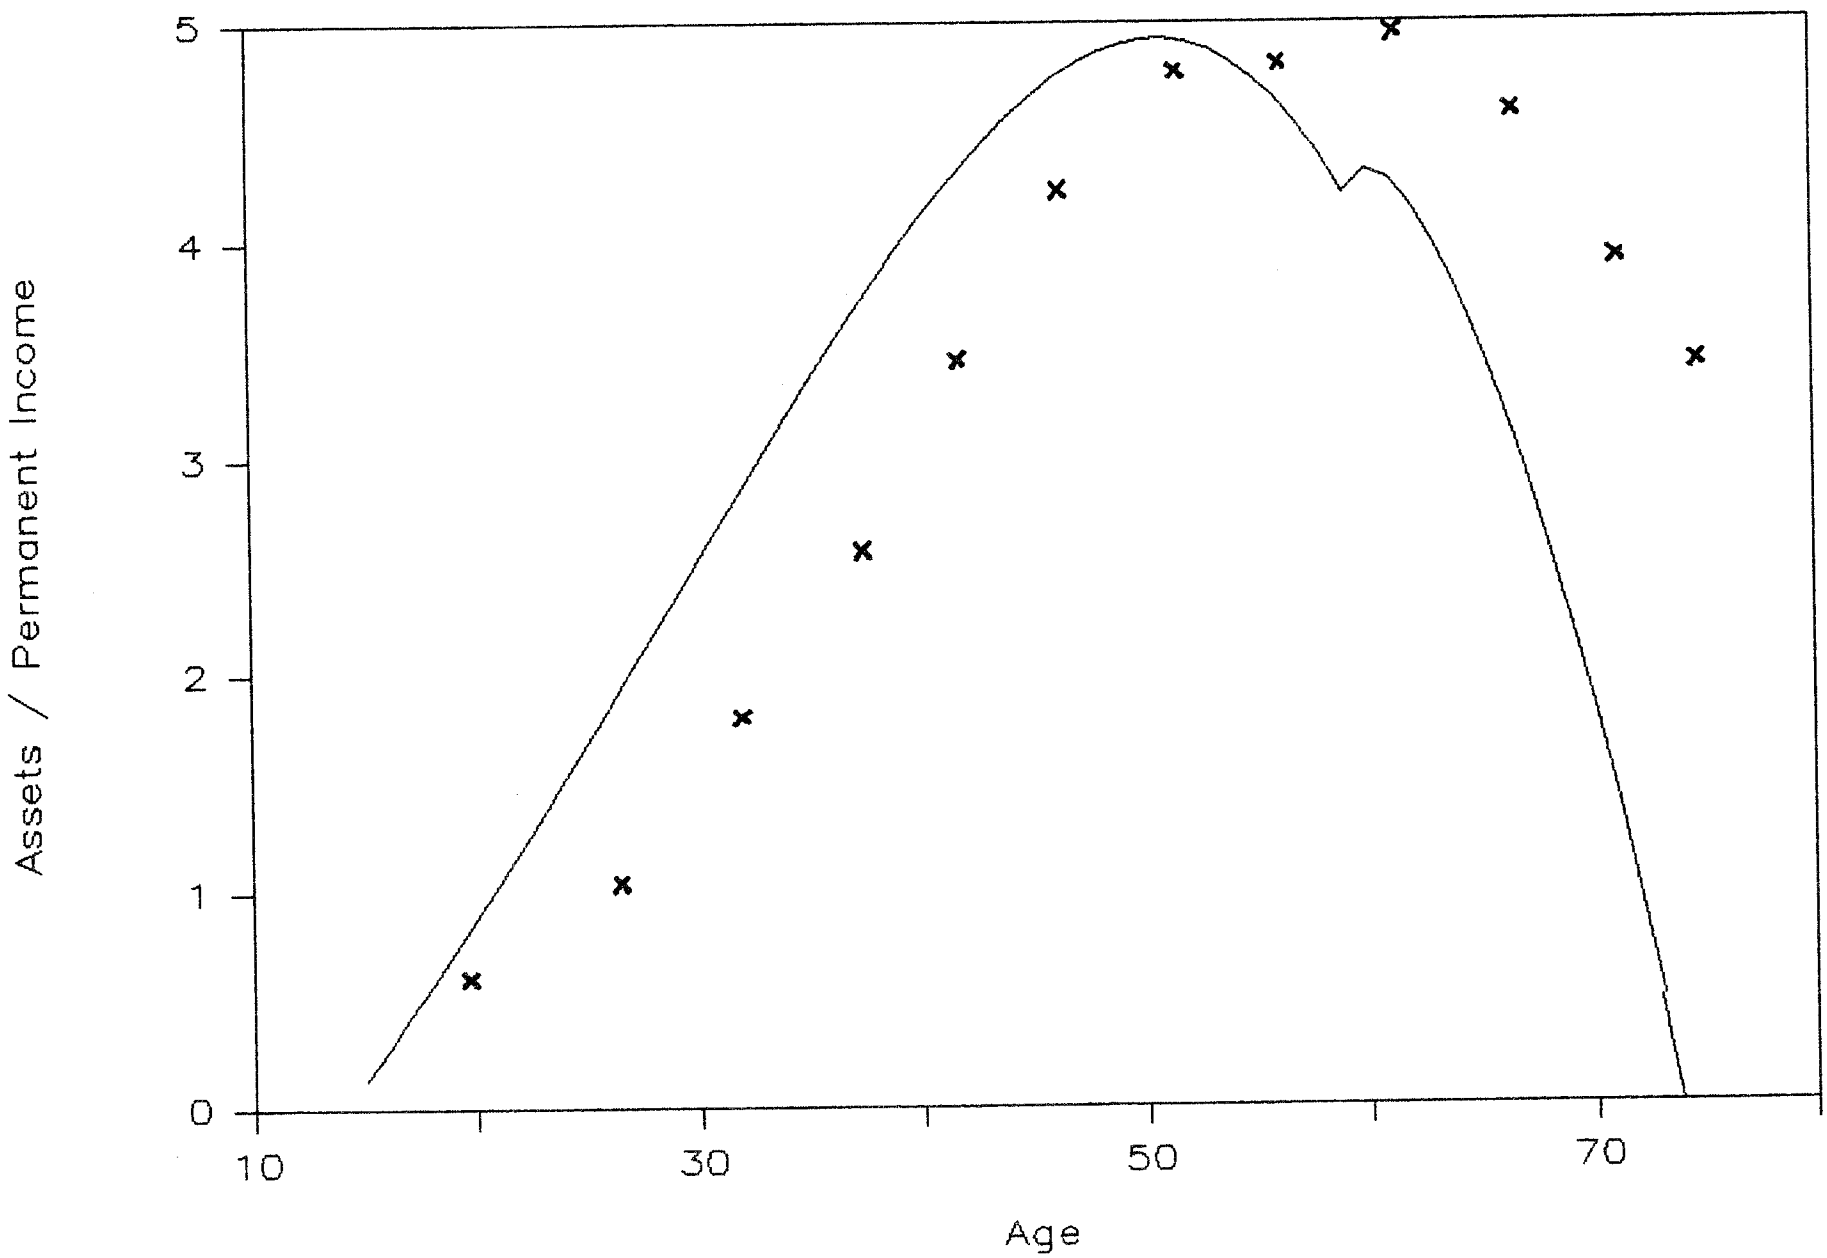

—— Simulated

x x Observations from King and Dicks-Mireaux (1982) and Dicks-Mireaux (1982). These authors constructed an estimate of "permanent income" for each household in the sample, and the average values of the ratio of asset-holdings (excluding pension wealth) to permanent income for each five-year age group are plotted on figure 2.

The "hump-shaped" pattern is clearly evident for the actual and simulated age profiles of asset-holdings. At first glance, therefore, the data are consistent with the hypothesis that on average household behaviour can be explained by the life-cycle model. The major respect in which the actual deviates from the simulated profile is the low rate at which wealth is run down after retirement. These observed low rates of decumulation are inconsistent with a life-cycle model in which the date of death is known with certainty. This is hardly surprising. Benjamin Franklin's well-known observation that "in this world nothing can be said to be certain, except death and taxes" omits some crucial considerations of timing. If we assume that (a) the only uncertainty is that about date of death and (b) preferences are intertemporally separable with instantaneous utility given by (2.9), then we may follow Yaari (1965) and Tobin (1967) and write expected utility over remaining lifespan as

$$U = \sum_{t=\tau}^{M} p(t,\tau) \frac{u(C_t, L_t)}{s_t} \qquad (2.32)$$

where τ is current age

M is maximum possible age

$p(t,\tau)$ is probability of survival until age t given that one has survived to age τ.

The first-order conditions for the optimal consumption and labour supply plan may be written as (eliminating the multiplier on the budget constraint)

$$\frac{u_{c,t}}{u_{c,t-1}} = \frac{1}{\Pi_t}\left(\frac{1+\rho_t}{1+r_t}\right) \qquad (2.33)$$

where Π_t is the probability of surviving from age t-1 to age t (unity minus the mortality rate). Hence we may regard uncertainty about date of death as equivalent to an age-specific increase in the rate of time preference. The impact of this on consumption and asset-holdings is ambiguous. On the one hand, greater uncertainty about length of life leads to a desire for certain consumption in the present, and, on the other hand it results in a desire to save for a potentially longer future life. The outcome will depend upon the household's aversion to risk, and its rate of time preference. For very high rates of risk aversion the wish to avoid the risk of a low level of utility can clearly produce an optimal savings plan characterised by low rates of decumulation of wealth after retirement. In the limiting case in which the household maximises the minimum value of (discounted) instantaneous utility, then "full expenditure" in each period, y_t $(t = \tau \ldots M)$, will be equal to the annuity value of full wealth over the maximum lifespan. Hence, even though survival to great age may be highly improbable, the household will retain sufficient wealth to enable it to maintain its standard of living up to the maximum possible age. It is immediately apparent that the need for such savings is greatly diminished by the existence of insurance against date of death. Purchases of annuities, especially indexed annuities, would enable households to decumulate wealth more rapidly. The absence or small size of markets in such annuities has been used to explain low rates of decumulation. Davies (1981) used data on life expectancy at different ages to simulate a life-cycle model with exogenous labour supply assuming constant relative risk aversion. For degrees of risk aversion in excess of three the predicted rates of decumulation were low, and comparable with the actual rates shown in figure 2(see King and Dicks-Mireaux op.cit. for further discussion of this point).

The problem with this explanation is that the results are sensitive to the existence of pensions and social security. Pensions provide an alternative to annuities and, indeed, one of the arguments for publicly provided pensions is that they compensate for the market's failure to provide indexed annuities (Diamond, 1977). With indexed pensions or social security the need to maintain asset-holdings against the eventuality of longevity is much reduced. This suggests a way of testing the hypothesis that the life cycle model can explain observed low rates of decumulation. For a given level of assets, the rate of decumulation in the later stages of life should be positively correlated with the ratio of the pension to permanent income. The level of assets itself would normally be negatively correlated with pension and social security wealth. These two aspects of the impacts of pensions on asset-holdings, the "insurance" effect and the "wealth" effect respectively, must be disentangled if the hypothesis that observed low rates of decumulation can be explained by the life-cycle model is to be tested empirically. Unfortunately, existing studies do not separate these two effects (for example, neither Diamond and Hausman (1982) nor King and Dicks-Mireaux (1982) test the hypothesis). For the moment, therefore, the question remains open.

Another natural question is how sensitive is the hump-shaped pattern for asset holdings to changes in the assumptions made to generate the simulated profile. We shall comment briefly on four assumptions. First, bequests were omitted from the model. Detailed discussion of this topic is deferred until section 5, but if target bequests are equal to some fraction of permanent income and there are no inter vivos gifts then terminal wealth tends to this value rather than zero. This does not, however, explain low rates of decumulation. The main effect is to raise the peak value of asset-holdings and shift the profile upwards. The second assumption is the rate at which real wages grow over the life-cycle. In the simulation the ratio of peak wages to starting wages was taken to be 1.5. With a value of 3.0 the maximum value of asset holdings is little changed, and the main consequence is that households go into debt in the early years of the life-cycle. Thirdly, a higher pension (with unchanged taxes in working life) lowers the level of household saving. If the pension is 75 per cent, instead of 50 per cent, of final earnings then the ratio of assets to permanent income reaches a maximum of 2.9 at age 45 rather than 4.9 at age 51.

The fourth assumption made was that the rate of time preference is less than the real interest rate. Some households may, however, exhibit much higher rates of time preference. It seems implausible to suppose that this could be true of a majority of households, because it would then be difficult to account for the sustained growth rate in aggregate per capita consumption in the absence of a very high rate of technical progress. But it may not be unrelated to the findings of both Diamond and Hausman (1982) and King and Dicks-Mireaux (1982) that around 20 per cent of the households in their samples had very low or negative levels of assets. We discuss this point below but it is worth pointing out that in the simulation reported above merely setting the rate of time preference to 2 per cent (holding all other parameters constant) is sufficient to ensure that the optimal plan leaves the household in debt in every period of life except for the initial and terminal values (this result depends on the pension to final earnings ratio). With this rate of time preference (equal to the assumed real interest rate) consumption, labour supply, the wage rate and earnings, all reach a peak at the same age (see the discussion following 2.14). Consumption is approximately constant over the life-cycle and households borrow against future earnings to finance this level of consumption in the early part of life, running down the level of debt as they become older but paying off the last instalment of debt only in the final period. In practice, it may be difficult to borrow against future labour earnings in which case households may consume all of their income, at least in the first half of their working life. From this example it is clear that observed values for asset-holdings which are either low or negative may be perfectly consistent with life-cycle behaviour. But for this to be the case, households with small or negative values for asset-holdings should exhibit different patterns of labour supply behaviour from those of households with low rates of time preference. Labour supply should peak at a later age for households with lower ratios of assets to permanent income. If this correlation is not present then the ability of the life-cycle model to explain low asset-holdings is much diminished. To my knowledge this test has not been carried out and the demands on data are substantial. One would require panel data covering both labour supply and asset-holdings.

The preceding discussion has ignored the complications which result from variations in the timing of household formation, entry into the labour force, imperfections in the capital market, marriage, and the birth of offspring, all of which will influence the age profile of asset holdings. In any empirical study these factors must be modelled explicitly or left to increase the magnitude of the residual variance. Their omission here is not intended as a reflection on their quantitative significance. Rather, the aim has been to focus on the empirical implications of forward-looking life-cycle behaviour.

Two main conclusions emerge from the above. The first concerns the nature of the empirical tests which could be regarded as constituting evidence against the life-cyle model. There are many policy issues for which we need to know whether households are forward rather than backward-looking, or whether they are likely to face constraints additional to the lifetime budget constraint. In other words, is the life-cycle model in its broadest sense, a good description of household behaviour? As with static demand analysis tests of the model are in fact tests of the joint hypothesis of life-cycle behaviour on the one hand and a particular parameterisation of preferences on the other. Most empirical findings which claim to reject the life-cycle model could, often more plausibly, be said to reject the assumptions made about preferences. In particular, even if we are prepared to accept the assumption of inter-temporal separability (which is crucial to empirical tests of a stochastic life-cyle model as we shall see below), the assumption of intratemporal separability is very strong and requires verification.

The two major respects in which observed savings behaviour might appear to deviate from the life-cycle model are (i) low rates of decumulation among the aged, and (ii) low or negative asset holdings for a sizeable minority of the population. But, as explained above, there are circumstances in which these phenomena are consistent with life-cycle behaviour and we proposed two empirical tests which would throw light on the issue. The first was to estimate the insurance and the wealth effects of pensions separately. The second was to compare the labour supply behaviour of small (or negative) asset holders with that of wealthier individuals. To carry out these tests will require detailed and high-quality microdata on individual households. Aggregate data are inadequate for the task. As in the case of static demand analysis, almost any set of aggregate observations can be consistent with an underlying life-cycle model.

The second conclusion relates to the treatment of unobservable variables in the consumption function, a problem of long standing concern. We have seen that one way to avoid this problem is to take the ratio of successive first-order conditions as in equations (2.16) and (2.33). In econometric terminology, taking first differences eliminates the fixed effect associated with each observation. The idea behind this is that instead of solving for the complete consumption profile as a function of all of the exogenous variables, we simply exploit the Euler equation for the household's optimal control problem. In other words, we abandon the econometric specifications of Friedman and Modigliani, and revert to Frank Ramsey. This approach has played a central role in recent work on consumption decisions under uncertainty in which the first-order conditions are expressed in terms of the stochastic Euler equation. We now turn, therefore, to an analysis of the life-cycle model under uncertainty.

2.2 A Stochastic Life-Cycle Model

We have already discussed the impact of uncertainty over the length of life on the optimal savings plan. There is a conflict between the desire to save more for a possibly longer future life and the desire to enjoy current consumption which is certain (Yaari 1965, Champernowne 1969, Levhari and Mirman 1977 and Davies 1981). More important perhaps is the uncertainty surrounding future earnings and rates of return, (as well as future changes in household consumption). Among those who have investigated the former are Hall (1978), Sargent (1978) and Eden and Pakes (1981); Sandmo (1970) and Levhari and Mirman (1977) analyse uncertainty about interest rates.

It is clear that once we admit uncertainty into the picture the market structure assumed is of critical importance. With a complete (but finite) set of Arrow-Debreu contingent commodity markets the optimal savings plan can be analysed in the framework of the certainty model of section 2.1 expanded to include all relevant contingent commodities. In practice not all, or even many, of these markets operate, and we might consider the opposite extreme case in which in any period only spot markets exist. The absence of futures markets means that although the current real wage is nonstochastic, future wage rates and asset returns are uncertain. The economy follows a sequence of temporary equilibria, and each period individuals receive new information which in general leads them to revise their savings plans. In such a model expectations are obviously crucial, and an assumption about how expectations are formed and modified over time is required to close the model. This assumption replaces that of a complete set of Arrow-Debreu markets.

Consider an individual who has reached age t. For simplicity we maintain the assumption of a composite consumption good, and for the moment assume a fixed length of life. If we adopt the axioms required for expected utility maximisation then the individual's objective function is

$$U_t = E_t U(\bar{C}_1, \ldots \bar{C}_{t-1}, C_t, \ldots C_T; \bar{L}_1, \ldots \bar{L}_{t-1}, L_t, \ldots L_T; B) \tag{2.34}$$

where the expectations operator E_t uses all available information at age t.

This replaces the nonstochastic objective function (2.1), and reflects the fact that past levels of consumption and labour supply are given. The budget constraint is still given by (2.2) but (a) past levels of consumption and earnings may be replaced by the scalar value of current assets which captures all that is relevant about previous periods, (b) the constraint is defined over stochastic future wage and interest rates.

In the absence of futures markets the only decision variables are current consumption and labour supply. The first-order conditions are

$$E_t U_{ct} = \lambda_t p_t \tag{2.35}$$

$$E_t U_{Lt} = \lambda_t w_t \tag{2.36}$$

λ_t is the multiplier associated with the budget constraint defined over the remaining lifespan and is the shadow price of current assets. Even though the current real wage is nonstochastic it is not possible in general to express labour supply conditional on current "full expenditure" as a deterministic function of parameters and observable variables. Expectations, as represented by the joint distribution of future wage and interest rates, will influence labour supply not only through a wealth effect on current full expenditure but also through their impact on marginal rates of substitution.

As before we may eliminate the multiplier by writing down the stochastic Euler equations (normalizing the price of consumption to unity).

$$E_t\, U_{ct} = E_t\{U_{c,t+j} \cdot d_j\} \qquad j = 1 \ldots T\text{-}t \quad (2.37)$$

$$E_t\, U_{Lt} = E_t\{U_{L,t+j} \cdot \frac{w_t}{w_{t+j}} \cdot d_j\} \quad j = 1 \ldots T\text{-}t \quad (2.38)$$

where the real stochastic discount factor is

$$d_j = \prod_{k=1}^{j}(1 + r_{t+k})$$

These equations may be derived from a straightforward variational argument (note that this ignores the possibility of corner solutions). The expected utility lost by giving up one unit of consumption now must be equal to the expected utility gained by consuming the proceeds of the extra saving at any future date. Equally, the expected utility gained by taking an extra unit of leisure now, holding consumption at each date constant, must be equal to the expected utility lost at any future date from the sacrifice in leisure required to make up the foregone earnings. In addition to (2.37) and (2.38) there is a third equation describing the rate at which current consumption will be substituted for current leisure which is derived from the ratio of (2.35) to (2.36).

As with the nonstochastic Euler equation (2.16), the first-order conditions are of limited use unless the utility function is restricted. In general, it is impossible to solve for the explicit stochastic time path of consumption and asset holdings, but it is clear that the past (lagged values of consumption and labour supply) can affect current decisions. An obvious restriction to examine is intertemporal separability where U is given by equation (2.9). In this case the left-hand side of equations (2.37) and (2.38) is nonstochastic enabling the conditions to be written as

$$E_t \left\{ \frac{u_{c,t+j} \cdot d_j}{u_{ct} \cdot s_j} - 1 \right\} = 0 \qquad j = 1, \ldots T-t \quad (2.39)$$

$$E_t \left\{ \frac{u_{L,t+j}}{u_{Lt}} \cdot \frac{w_t}{w_{t+j}} \cdot \frac{d_j}{s_j} - 1 \right\} = 0 \quad j = 1 \ldots T-t \quad (2.40)$$

where s_j, the time preference factor, is defined analogously to d_j.

These conditions must hold for all individuals, and since the expectation equals zero it is clearly independent of the information set available to the individual. Hence, as pointed out by Grossman and Shiller (1981), any specification based on (2.39) is robust with respect to any asymmetries in access to information. For additive separability over time (2.35) and (2.36) may be written in the conventional nonstochastic form.

$$\frac{u_{ct}}{u_{Lt}} = \frac{1}{w_t} \qquad (2.41)$$

In special cases it may be possible to derive explicitly the decision rule which relates consumption and labour supply to the stochastic processes generating wage and interest rates (Hansen and Sargent 1980, 1981). But for more general assumptions about preferences this is not possible, and empirical work must be based on the estimation of the stochastic Euler equations (2.39) and (2.40). These equations state that the expectation of a given nonlinear function, conditional on information available at time t, is zero. The information set at time t includes current and past values of variables such as wage rates, and, whatever values these variables may take, the expectation of the nonlinear function is zero. Hence the value of the function is orthogonal to any variable that appears in the information set.

This suggests that the appropriate method for estimating the stochastic Euler equations is a nonlinear instrumental variables procedure and this was proposed by Hansen and Singleton (1982). The method minimises the deviations of the sample values of the orthogonality conditions from zero, and the metric in which these deviations are measured is chosen to produce the smallest asymptotic covariance matrix. The resulting estimates for the parameters of the direct or indirect utility function are consistent and, under relatively weak conditions on the underlying stochastic processes, are also asymptotically normal (Hansen and Singleton, op.cit). Since it is possible to imagine a large number of variables entering the information set at time t, there are likely to be more orthogonality conditions than parameters to be estimated. In this sense the model is overidentified. The overidentifying restrictions can be tested using the distance of the sample values of the orthogonality conditions from zero as a test statistic.

An alternative explanation of the estimation procedure is as follows. We may rewrite the equations as

$$E_t\, f^c(C_t, L_t, C_{t+j}, L_{t+j}, \frac{d_j}{s_j}) = 0 \tag{2.41}$$

$$E_t f^L(C_t, L_t, C_{t+j}, L_{t+j}, w_t, w_{t+j}, \frac{d_j}{s_j}) = 0 \tag{2.42}$$

The actual values of f^c and f^L may deviate from zero because new information accrues after time t. If individuals have expectations such that their revisions are uncorrelated with their original plans (for which rational expectations is a sufficient condition) then we may express (2.41) and (2.42) in terms of the following regression model with C_{t+j} and L_{t+j} as the dependent variables.

$$f^c(C_t, L_t, C_{t+j}, L_{t+j}, \frac{d_j}{s_j}) = \varepsilon^c_{j,t+1} \tag{2.43}$$

$$f^L(C_t, L_t, C_{t+j}, L_{t+j}, w_t, w_{t+j}, \frac{d_j}{s_j}) = \varepsilon^L_{j,t+1} \tag{2.44}$$

The error terms have zero mean but the covariance matrix is a function of the underlying stochastic processes generating wage and interest rates. This implies cross-equation restrictions and suggests that the distribution of the error terms should be consistent with the observed joint distribution of wage and interest rates. But if we regard the elements of the covariance matrix as free parameters, the model to be estimated consists of a pair of nonlinear simultaneous equations for each value of j. The inclusion of future wage and interest rates in the model necessitates the use of instrumental variables for consistent estimates.

This model provides a framework for the investigation of the role of intertemporal substitution in consumption and labour supply. Recently, much attention has been devoted to the question of whether such substitution can explain aggregate fluctuations in employment and consumption. The magnitude of the substitution effects is relevant to major policy issues such as the nature of unemployment, the design of the tax system (income versus expenditure as the tax base), and the effects of public borrowing. It is important to note, therefore, that existing empirical studies have examined only very special cases of the model given by equations (2.43) and (2.44). First, only estimates for the case j = 1 have been reported. For forecasting purposes higher values of j would be relevant to a comparison of alternative models. Secondly, except for the study by Mankiw, Rotemberg and Summers (1982), it has been assumed that preferences are weakly separable between consumption and leisure. This implies that the model for consumption is

$$f^{c}\left(C_{t+1}, C_{t}, \frac{1+r_{t+1}}{1+p_{t+1}}\right) = \varepsilon_{t+1} \qquad (2.45)$$

In other words either the level, or some transform of the level, of consumption follows a first-order difference equation, and if real interest rates (and the rate of time preference) are constant the parameters of this equation are time-invariant. This is the "Hall consumption function" (Hall, 1978). From (2.39) we may write this as

$$u_{c,t+1} = \frac{1+\rho}{1+r} \cdot u_{ct} + \eta_{t+1} \tag{2.46}$$

where $E(\eta_t) = 0$, but if the error term is introduced as in (2.39) then there is no reason to suppose that it is homoscedastic. If the stochastic shocks are small then we can take a linear approximation to the marginal utility of consumption, and consumption is related only to its own previous value. In fact consumption follows a random walk with drift. Given the assumption of weak separability between consumption and leisure and the validity of the linear approximation, then no lagged variables other than last period's consumption itself will contain any explanatory power in a regression equation for consumption. This is the analogue to the orthogonality condition discussed above. Unfortunately, it has been widely stated in the literature that a test of the life-cycle model is the significance or otherwise of additional lagged variables in a regression model for consumption. This is false. If any relevant explanatory variables are omitted from the model then additional lagged values of consumption or income may appear to be significant. Two examples are important. First, consider the case when interest rates vary over time. If interest rates are omitted from the estimation of (2.46), as in Hall's original paper, then the reduced form for consumption will not follow a random walk with drift. Suppose that the real interest rate itself follows a simple autoregressive process. Then consumption would follow a higher-order autoregressive process, and the significance of additional lagged values of consumption would indicate nothing about the consistency of observed behaviour with a life-cycle model.

More interesting, perhaps, is the fact that most of the studies of consumption assume that preferences for consumption and leisure are separable. To analyse the case when preferences are not separable it is helpful to use the indirect rather than the direct utility function. In the stochastic case the objective function (2.25) becomes

$$U_t = E_t \sum_{j=t}^{T} \frac{\ln v(w_j, y_j)}{s_j} \tag{2.47}$$

A simple variational argument leads to the following Euler equation for full expenditure y_t

$$E_t \left\{ \frac{1}{v_{t+1}} \frac{\partial v_{t+1}}{\partial y_{t+1}} \frac{(1+r_{t+1})}{(1+\rho_{t+1})} - \frac{1}{v_t} \frac{\partial v_t}{\partial y_t} \right\} = 0 \tag{2.48}$$

which may be written as

$$E_t \, f^y \left\{ y_{t+1}, \; y_t, \; w_{t+1}, \; w_t, \; \frac{1+r_{t+1}}{1+\rho_{t+1}} \right\} = 0 \tag{2.49}$$

Within each period labour supply is given by a non stochastic application of Roy's identity and consumption then follows from the within-period budget constraint (2.21). The attraction of this dual approach is that it exploits the two-stage budgeting structure implied by the assumption that preferences are intertemporally separable. At the first stage full expenditure is given by the stochastic process implied by (2.48). Then in any period, given the level of full expenditure, consumption and labour supply are determined exactly as in the nonstochastic model. Commodity demands, conditional on the level of full expenditure, are given by a conventional static model, and the behaviour of full expenditure is governed by the single stochastic Euler equation. Given a functional form for the indirect utility function (for example 2.24 above) this system can be estimated by instrumental variables procedures, and may well be easier to implement than instrumental variables estimation of the equations derived from the direct utility function.[11]

For the preferences described by (2.24) used in the simulation of section 2.1, full expenditure follows the path

$$y_{t+1} = \frac{1-a}{c}\left(\frac{r_{t+1} - \rho_{t+1}}{1+\rho_{t+1}}\right) + \left(\frac{1+r_{t+1}}{1+\rho_{t+1}}\right) y_t + \frac{b}{c+\gamma}\left(\frac{1+r_{t+1}}{1+\rho_{t+1}} e^{-\gamma w_t} - e^{-\gamma w_{t+1}}\right) + \eta^y_{t+1} \qquad (2.50)$$

The time path of consumption is given by

$$c_{t+1} = \frac{1-a}{c}\left(\frac{r_{t+1}-\rho_{t+1}}{1+\rho_{t+1}}\right) + \left(\frac{1+r_{t+1}}{1+\rho_{t+1}} \cdot \frac{1-cw_{t+1}}{1-cw_t}\right) c_t + h(w_{t+1}, w_t) + \hat{\eta}_{t+1} \qquad (2.51)$$

where the function h is given by

$$h(w_{t+1}, w_t) = - w_{t+1}\left\{(1-a)\ \frac{1+r_{t+1}}{1+\rho_{t+1}} + be^{-\gamma w_{t+1}}\right\} + \frac{b}{c+\gamma}$$

$$(1-cw_{t+1})\left[\frac{1+r_{t+1}}{1+\rho_{t+1}}\ e^{-\gamma w_t} - e^{-\gamma w_{t+1}}\right]$$

$$+ \frac{1+r_{t+1}}{1+\rho_{t+1}} \cdot \frac{1-cw_{t+1}}{1-cw_t}\ (1-a+be^{-\gamma w_t})w_t \tag{2.52}$$

It is clear that consumption does not follow a simple random walk with drift. If we compare (2.51) with the Hall consumption function (Hall, 1978)

$$C_{t+1} = a_o + a_1 C_t + \mu_{t+1} \tag{2.53}$$

we may observe two obvious differences. The first is that in (2.51) the coefficient on lagged consumption is not time invariant and will depend on changes in real wage and interest rates. The second is that the linear specification (2.53) omits the h function defined over current and lagged wage rates. Suppose that the specification (2.51) were correct. Estimates of the Hall consumption function would then produce evidence of (a) parameter instability in the coefficient of lagged consumption, (b) the existence of additional variables with significant explanatory power in the regression, and (c) serially correlated residuals generated by a stochastic process the order of which would depend upon the process generating wage and interest rates. The extra variables would be correlated with the h function and hence obvious candidates would be current and lagged labour earnings.

Alternatively, by substituting out from the process generating wage and interest rates we could obtain a reduced form for consumption in terms of its own lagged values which would involve a nontrivial higher order lag distribution.

All of this raises the issue of the objectives of an econometric investigation into savings and consumption behaviour. If the aim is simply to obtain the best forecasting equation for purposes of projecting developments in the economy then there is no compelling theoretical reason to restrict the specifications over which we search to models such as (2.53). A time-series analysis using unrestricted lags would be the appropriate search procedure. But if we are interested in predicting the response of consumption to changes in policy variables (for example, changes in the taxation of capital income or the introduction of permanent or temporary taxes on expenditure) then the use of the reduced form equation for consumption is unhelpful (for reasons known as the Lucas critique). Even if we can model the effect of the policy changes on the parameters of the Euler equation (as in (2.51) for example) the resulting model simply tells us the rate at which consumption changes once we are in the new regime. It does not tell us how consumption will change as we move from one regime to another. To examine this we need to solve for the optimum decision rule which relates consumption and labour supply to the joint distribution of wage and interest rates. As mentioned above, only in special cases will it be possible to find an explicit form for the decision rule, and the comparative statics of stochastic models of this kind is an area requiring further research.

Finally, we might wish to examine the hypothesis that consumers are forward-looking rational planners unconstrained in current markets. The main insight in the Hall model is that in such a world consumers' plans would change only upon the receipt of new information about the future, and that new information would by definition be unpredictable. Hence changes in lifetime wealth would follow a random walk. It is easy to show that the basic result behind this intuition is that the marginal utility of lifetime wealth follows a random walk. The assumption required is that preferences are additively separable over time. With this assumption then from the objective function (2.47) and the budget constraint defined over remaining lifetime, at any time t

$$\frac{1}{v_t}\frac{\partial v_t}{\partial y_t} = \lambda_t \tag{2.54}$$

where λ_t is the multiplier associated with the budget constraint and is the marginal utility of wealth. Substituting into the Euler equation (2.48) yields

$$E_t\left\{\left(\frac{1+r_{t+1}}{1+\rho_{t+1}}\right)\lambda_{t+1} - \lambda_t\right\} = 0 \tag{2.55}$$

The marginal utility of wealth follows a random walk with drift where the drift parameter depends upon the real interest rate and the rate of time preference. From (2.35) it immediately follows that the marginal utility of consumption follows a random walk with the same drift parameters. But unless preferences between consumption and leisure are separable it does not follow that consumption itself follows a random walk (even as a linear approximation). No amount of sophisticated econometric analysis of the time-series behaviour of consumption will throw light on the relevance of the stochastic life-cycle model. We need to examine either the cross-equation restrictions implied by the model for the joint behaviour of consumption and labour supply (an investigation for which the data requirements are prodigious) or to examine the predictions of alternative models. In the following section we ask, therefore, what are the implications of a model which allows for the possibility that individuals may be constrained in either the capital or labour markets?

To conclude the discussion of stochastic life-cycle models, we may note the following three points. First, uncertainty about length of life may be analysed, as noted in section 2.1, by allowing an age-specific rate of time preference. This means that age will enter into models of consumption using microdata (whether cross-section or panel data), and that demographic variables will be relevant to the aggregate consumption function.

Secondly, in this section we have concentrated on drawing out the implications of the first-order conditions for individual optimization. The results are equations for optimal individual or aggregate consumption and labour supply conditional upon the joint distribution of wage and interest rates. But this is only part of the story. As in static demand analysis, we may aggregate equations (2.39) and (2.40) across individuals and invert the system to obtain a model for equilibrium asset prices. Work along these lines has been carried out by Lucas (1978), Shiller (1981), and Grossman and Shiller (1981). The endogenous distribution of prices is related to the exogenous distribution of shocks to output (originating with technological or other changes). Hence just as income in a Keynesian consumption function may be jointly determined with consumption, so the distribution of wage and interest rates in a stochastic life-cycle model may be jointly determined with aggregate consumption and labour supply. Much more study of these issues is required if we are to construct models of tax incidence, for example, in a stochastic setting.

Thirdly, it is apparent that the extension of the life-cycle model to a stochastic world alters less the specification of the equations to be estimated than the methods of estimation. Given the assumption of inter-temporal separability the dynamic specifications for consumption and labour supply which follow from the result that the marginal utility of wealth follows a random walk are the stochastic analogue of the "constant marginal utility of wealth" demand functions discussed in section 2.1. The stochastic Euler equations for consumption and labour supply can be estimated by a generalised instrumental variables procedure. Tests of the significance of variables other than lagged consumption in a regression equation for consumption tell us more about the appropriate parameter-isation of preferences than about the relevance of the life-cycle model. In particular, if the intratemporal preference ordering is not separable then consumption does not follow a random walk with drift. A much wider range of models needs to be tested before we can be confident of robust results in this area. It is also important not to confuse the reduced form with the underlying structural model. The omission of relevant variables can lead to a complicated distributed lag model for consumption. This is uninformative about the underlying behavioural model and the lag structure may exhibit parameter instability (a common feature of aggregate time-series studies of consumption).

3. Alternative Models

In this section we examine the empirical predictions of models which might be regarded as alternatives to the life-cycle model. To some extent this is semantic. Our definition of the life-cycle model is that individuals are assumed to be "forward-looking" and face a single lifetime budget constraint. In contrast, Kurz (1981) has argued that the essence of the life-cycle model is that individuals plan only over their own lifespan because the inclusion of bequests in a satisfactory manner leads inexorably to a concept of intergenerational equilibrium which cannot be represented as the outcome of an individual optimisation problem of the kind discussed above. In other words, "optimal bequests" are not appropriately analysed in a model of individual family optimization.

We discuss first the consequence of quantity constraints in either the labour or capital markets. The most obvious example of such a constraint is involuntary unemployment. In the absence of constraints labour supply and consumption are given by Roy's identity and the within period budget constraint. Suppose, however, that some individuals are constrained at zero labour supply.[12] Then we may define a virtual wage, w_t^*, and virtual full expenditure y_t^*, such that at the virtual budget constraint the individual would choose to work zero hours and consume the observed consumption level C_t. The virtual wage and full expenditure are determined by

$$1 = - \frac{-\partial v(w_t^*, y_t^*)}{\partial w_t^*} \Bigg/ \frac{\partial v(w^*, y^*_t)}{\partial y_t^*} \tag{3.1}$$

$$y_t - y_t^* = w_t - w_t^* \tag{3.2}$$

The Euler equation in the constrained case is more complicated. Assuming that there are no constraints in the capital market (or, more accurately, that there is at least one asset market in which individuals are unconstrained), the variational argument for the optimal allocation of full expenditure over time yields the following modified version of (2.48)

$$E_t \left\{ \frac{1+r_{t+1}}{1+\rho_{t+1}} \cdot \frac{1}{v_{t+1}} \left(\frac{\partial v_{t+1}}{\partial w_{t+1}^*} . \frac{dw_{t+1}^*}{dy_{t+1}} + \frac{\partial v_{t+1}}{\partial y_{t+1}^*} \cdot \frac{dy_{t+1}^*}{dy_{t+1}} \right) \right.$$

$$\left. - \frac{1}{v_t} \left(\frac{\partial v_t}{\partial w_t^*} \cdot \frac{dw_t^*}{dy_t} + \frac{\partial v_t}{\partial y_t^*} \cdot \frac{dy_t^*}{dy_t} \right) \right\} = 0 \tag{3.3}$$

When there are no constraints (w^*, y^*) is equal to (w, y) and (3.3) reduces to (2.48). Note, however, that if we substitute the constrained version of (2.54) into (3.3) we obtain (2.55). Even in the presence of labour market constraints the marginal utility of wealth, and hence the marginal utility of consumption, follows a random walk. If preferences between consumption and leisure are separable then consumption itself will (approximately) follow a random walk, and unemployment, or any other quantity constraints, will have no place in the consumption function. In general, however, with nonseparable preferences equation (3.3) has some interesting implications for the time-series analysis of consumption.

In order to illustrate the way in which constraints affect the specification of the aggregate consumption function, we shall consider a simple model in which it is assumed that individuals know in period t whether or not they will be constrained in period t+1.

There are four cases to examine:

Case 1 : individual unconstrained in both periods t and t+1

Case 2 : individual constrained in period t but not t+1

Case 3 : individual constrained in period t+1 but not t

Case 4 : individual constrained in both periods t and t+1

Consider the simple case in which the real rate of interest is nonstochastic and equal to the rate of time preference. Assume also that the real wage is expected to remain constant. Then with these assumptions (of unchanged relative prices) both full expenditure and consumption will follow a random walk when individuals are unconstrained (Case 1). But relative virtual prices will also be constant and so if individuals are constrained in both periods, consumption will follow a random walk in Case 4 also.[13] Consumption does not, however, follow a random walk when individuals are constrained in one period but not in the other. In Cases 2 and 3, consumption follows a nonlinear stochastic process in which lagged earnings and wage rates appear in addition to lagged consumption. This suggests that it is changes in the unemployment rate, not the level of unemployment itself, which should enter the consumption function when the relevant constraints are those in the labour market.

To illustrate the empirical specifications which constraints may produce we consider the nonlinear labour supply function used above (equation 2.20). Rather than solve explicitly for the virtual prices we adopt a rather different view of unemployment in which the individual faces not a constraint on the number of hours worked but a very low wage rate. Redundancy leads to a contraction in market opportunities and the potential wage rate falls sharply. The new wage rate depends upon alternative employment opportunities. How much, for example, could a redundant professor of economics earn by cleaning his neighbour's car or windows? I shall assume that the answer is zero. At a zero wage optimal labour supply is very small (see Figure 1). In a formal sense the resulting unemployment is "voluntary" although the individual is not receiving the same wage offer as otherwise identical workers and in this sense it is "involuntary". For this preference ordering the stochastic time path of consumption is given by equation (2.51). When $r_t = \rho_t$ and $w_{t+1} = w_t$ then the h function has the property that

$$h(w, w) = 0 \tag{3.4}$$

This is true for any level of the wage rate. Hence in both the "no unemployment" case <u>and</u> the "continuing unemployment" (zero wage in both periods) case, consumption follows a pure random walk.

$$C_{t+1} = C_t + \eta_{t+1} \tag{3.5}$$

In cases 2 and 3 consumption is given by

$$C_{t+1} = \left(\frac{1 - cw_{t+1}}{1 - cw_t}\right) C_t + h(w_{t+1}, w_t) + \eta_{t+1} \tag{3.6}$$

where (w_{t+1}, w_t) is equal to (w, 0) in Case 2 and (0, w) in Case 3. From (2.52)

$$h(w, 0) = -(1 - cw)h(0,w) \tag{3.7}$$

This implies that the changes in consumption ΔC, conditional on the same level of current consumption, in the two cases are inversely related according to

$$(\Delta C)_2 = -(1-cw)(\Delta C)_3 \tag{3.8}$$

It might be interesting to test this model on panel data for four distinct groups in the population comprising those (i) in continuing employment, (ii) continuing unemployment, (iii) entering unemployment and (iv) becoming re-employed. The model implies strong cross-equation restrictions on the consumption functions for the four groups (Equations (3.5) and (3.8)).

Suppose that a fraction α of consumption in period t is by people who are becoming re-employed and a fraction β is by people entering the state of unemployment. Suppose also that the average wage of those losing jobs is equal to the average wage of those finding jobs. Then the aggregate per capita consumption function is of the form

$$C_{t+1} = \left(1 + cw\left(\frac{\beta}{1-cw} - \alpha\right)\right)C_t + h(w, 0)\left(\alpha - \frac{\beta}{1-cw}\right) + \eta_{t+1} \tag{3.9}$$

If labour supply is income-inelastic ($c = 0$) then[14]

$$C_{t+1} = C_t + (\alpha - \beta)h(w,0) + \eta_{t+1} \tag{3.10}$$

The random walk path of consumption is shifted up or down by a factor proportional to the <u>change</u> in the unemployment rate (measured by $\alpha - \beta$).

Implicit in the above discussion was the assumption that individuals knew into which of the four transition states they belonged. More generally, we could allow for uncertainty about the state in which an individual might expect to find himself next period. Expectations about future levels of unemployment would enter the model yielding an optimal level of "precautionary savings" to act as a cushion against future unemployment. The path for consumption will be given by equation (3.3) modified to allow for the possibility that in period t+1 the individual may experience either a virtual budget constraint (w^*_{t+1}, y^*_{t+1}) or an actual budget constraint (w_t, y_t) depending on whether or not he will be unemployed next period.

Another set of constraints which render the basic life-cycle model inapplicable is capital market constraints. Because of the difficulty of borrowing against future labour earnings individuals who wish to consume more than their current income, and have run down their (liquid) assets to zero, may face a borrowing constraint.[15] The significance of such constraints is that the optimal intertemporal allocation of resources is no longer described by a stochastic Euler equation. Moreover, it is plausible to suppose that the incidence of borrowing constraints is highly correlated with constraints in the labour market, that is with unemployment. Formally, we could define the virtual nonstochastic interest rate and virtual lifetime wealth to be those values at which the individual would exactly choose the levels of full expenditure to which he or she were constrained. With a binding constraint on borrowing then full expenditure y_t is constrained to the value of the current endowment w_t. In a stochastic setting this approach does not lead to tractable specifications because only in very special cases is it possible to obtain an explicit solution for the optimal time path of full expenditure which is required in order to solve for the virtual interest rate and virtual wealth.

But we may illustrate the likely impact of capital market constraints on the behaviour of consumption. In period t there are two groups, those who face binding constraints on borrowing and those who do not. The consumption path of the unconstrained group is given by the appropriate stochastic Euler equation, equation (2.48) if they are free to choose their labour supply or equation (3.3) if they are rationed in the labour market. For the constrained group there are two possibilities. First, they may expect to be constrained again in period t+1 in which case the expectation at time t of their consumption in period t+1 is the expected value of labour earnings in period t+1 conditional on earnings in period t. Secondly, they may anticipate being unconstrained next period in which case consumption would be expected to return to a "normal" level reflecting lifetime opportunities. The value of this expectation can be determined only by solving the stochastic control problem facing the individual. But if he is very similar to a group of currently unconstrained consumers then we can use observations on their current consumption to form an expectation of his consumption next period.

A formal analysis of borrowing constraints is complicated by two factors. First, the timing of such constraints is endogenous to individuals' own savings plans and hence the identification of the "constrained group" is by no means obvious. Secondly, the optimal loan contract may contain a mixture of individual-specific interest rates and repayment periods, as well as upper limits on borrowing. Asymmetry in information about the risk of future unemployment will affect the outcome. In order to illustrate the nature of the consumption function when we allow for constraints in both capital and labour markets, we consider a very simple example. Individuals who do not experience unemployment are assumed never to be constrained in the capital market. For individuals who do experience unemployment we make the following assumptions. In period t it becomes known (both to the individual concerned and to lenders) that an individual will be unemployed in period t+1. Borrowing constraints are imposed only upon the current and prospective unemployed. If the person is already unemployed we assume that he has run down his assets by the end of period t and hence his consumption in period t+1 will be constrained to his receipt of unemployment benefit. If the person is employed we assume that he does not have such large debts that upon receipt of the "bad news" he is unable to reduce current consumption and save an optimal amount to provide a cushion against the future unemployment (a precautionary motive for savings). In other words, after the receipt of the bad news the optimal level of asset holdings is positive and the individual is not constrained in the capital market in period t.

We assume also that unemployment strikes randomly among the population. For the example analysed above in the context of labour market constraints, we have that expected per capita income in period t+1, conditional on information available in period t, in the four transition states is

(i) Group 1: employed in both periods:

$$E_t(C_{t+1}) = C_t \tag{3.11}$$

(ii) Group 2: unemployed in t but with a job in t+1: as someone who is currently unemployed we assume that an individual in this group is unable to borrow (even though he is also to return to work).[16] The Euler equation, in the form of an equality, does not apply. As an approximation we regard this individual in period t+1, once he is back at work as a random drawing from the population. For this individual also, therefore,

$$E_t(C_{t+1}) = C_t \tag{3.12}$$

where the C_t is an observation of the consumption of the currently employed. If the average earnings of those liable to unemployment were smaller than for the population as a whole then C_t would be replaced by, say, λC_t. Although (3.11) and (3.12) appear identical this is only because we have assumed for simplicity that λ equals unity.

(iii) Group 3; employed in t but facing unemployment in t+1: given our assumptions the Euler equation applies to this group and is given by equation (3.6)

$$E_t(C_{t+1}) = \frac{1}{1 - cw} C_t + h(O, w) \tag{3.13}$$

(iv) Group 4; unemployed in both periods: we assume that the average replacement ratio (ratio of unemployment benefit to current real earnings) is r.

$$E_t(C_{t+1}) = r\ E_t(w_{t+1}\ \ L_{t+1}) = rwL \tag{3.14}$$

Let the unemployment rate be u_t, and assume that a fraction Θ_t of the employed lose their jobs and a fraction δ_t of the unemployed find jobs in period t. This implies that

$$\Theta_t(1-u_t) + (1-\delta_t)u_t = u_{t+1} \tag{3.15}$$

Aggregate per capita consumption is given by

$$E_t(C_{t+1}) = \left[1 - (1-\delta_t)u_t + \frac{\theta_t cw(1-u_t)}{1 - cw}\right] C_t + \left[\theta_t(1-u_t)h(O,w) + rwu_t(1-\delta_t)L\right] \tag{3.16}$$

A regression model for consumption based on (3.16) would include not only lagged consumption, but also lagged earnings, the unemployment rate, and the change in the unemployment rate. There are some interesting special cases ($\theta = 0$, moving out of recession; $\delta = 1$, boom; $r = 1$, replacement ratio of unity) but these are left to the reader.

One of the implications of borrowing constraints is that for individuals whose consumption is constrained to current receipts, the marginal propensity to consume out of current receipts is unity. For unconstrained individuals the marginal propensity to consume out of current income is less than unity although its actual size depends on the stochastic process generating current incomes. This feature has been the focus of previous work on capital market constraints (such as Flemming 1973, Appelbaum and Harris 1978, Dolde 1978 and Koskela 1978). Of more interest in equation (3.16) is the fact that, except in special cases, the aggregate consumption function depends upon the distribution of individuals among the four transition states. The distribution of expectations about future unemployment will influence the level of aggregate consumption conditional upon the value of past consumption. The existence of constraints means that it is not possible to ignore aggregation problems when modelling the level of aggregate consumption. In his study of the effects of income distribution on the aggregate consumption function, Blinder (1975) found surprisingly little evidence of the effect of conventional measures of distribution on the level of aggregate consumption. This may well be because aggregation becomes important only when constraints are relevant. A simple example of this is the contrast between the linear Hall consumption function and Equation (3.16). The Hall function holds even if there is a distribution of wages among individuals. Anticipated changes in income distribution (resulting from, perhaps, changes in the degree of progression in the tax system) would be reflected in individuals current consumption plans. Conditional on the value of lagged consumption, no additional information is contained in a statistic measuring the income distribution. But of course it does not follow from this that changes in the degree of progression in the tax system have no effect on the level of consumption. In general, they will have significant effects but to model these in a stochastic setting is by no means straightforward given, as remarked above, the lack of explicit solutions to the decision rule relating consumption decisions to the current observed state of the economy.

In the analysis of constraints it was clear that virtual prices entered the consumption function, and since these are a function of quantities (constraints) then there is clearly a basis for a Keynesian consumption function in which quantities are relevant explanatory variables. But there are other non-Walrasian features which might usefully be incorporated into the model. We did not analyse the role of futures markets, and more particularly, the consequences of explicit or implicit contracts which may last for many periods. Such contracts have been the subject of much discussion in the literature on the labour market (see Hart 1983), and lead also to virtual prices which differ from observed prices. There is scope for future work on the impact of contracts on the behaviour of consumption. It is clear that one of the major empirical trends in the postwar period has been the rapid growth of 'contractual saving' and a large proportion of average savings bear a pre-determined relationship to the level of labour income. Where there are other opportunities for adjusting savings at the margin, then this fixed relationship may be of little consequence. But these alternative opportunities may receive less favourable tax treatment (and in general do), and if the allocation of savings among alternative assets is relevant then contracts will play a role.

There are several deeper objections which can be made to the life-cycle model. First, as mentioned above, Kurz (1981) has pointed out that if bequests are modelled then a concept of inter generational equilibrium rather than individual optimisation subject to exogenous constraints may be a more appropriate way of viewing the problem. Secondly, some individuals may either be rather poor at, or have no interest in, stochastic control theory. We return to this below. Thirdly, individuals may not act as 'rational planners' This objection is not helpful unless it leads to an alternative view of consumer behaviour. One possibility is to regard consumers as'myopic'. But if this means little more than the possibility that some consumers have high rates of time preference, then such behaviour can be analysed within the framework set out above. It is clear from the discussion in earlier sections that a combination of a high rate of time preference and borrowing constraints will lead to an outcome in which myopic individuals are generally constrained by current resources. Another view is that the problem of decision making under uncertainty is so great that the normal 'as if' justification for analysing optimal behaviour does not apply. Psychologists such as Tversky and Kahneman (1982) have argued that when constructing subjective probabilities people use a limited number of heuristic principles which can sometimes lead their subjective probabilities to violate Bayes' theorem. They provide empirical evidence to support this claim and in part the evidence derives from an inability of individuals to comprehend more than a certain part of the problem at any one time. Behaviour of this kind could lead to irrational expectations such that the current subjective probabilities used by individuals were unduly dominated by the most recent observed outcome. The omission of relevant prior information was, in fact, one of the empirical phenomena observed by Tversky and Kahneman (1982). Alternatively, the costs of decision-making under uncertainty may be so great that individuals devise rules of thumb for their consumption and saving behaviour.

Other approaches would see individuals behaving stochastically. In other words, instead of choosing a single action from a choice set based on an assessment of the likely outcomes, individual's choices in identical circumstances can be inconsistent. Barbera and Pattanaik (1981) have analysed the conditions under which stochastic choices can be represented as a lottery among a large number of different preference orderings. In a somewhat different vein, Shefrin and Thaler (1981,1982) analyse the savings decision as the outcome of a conflict between a split personality in which a 'higher' self, concerned with the long run consequences of current decisions, attempts to gain control over the 'lower' self which is concerned only with immediate pleasure. They formulate the problem as a principal-agent model. This does not seem appropriate. In their model it is possible for the higher self to ensure that the lower self always follows an optimal life-cycle savings plan. There is no asymmetry of information, and the essence of a principal-agent problem is lacking. More relevant would be a formulation in terms of a bargaining model in which the two opposing sets of values tussled for dominance in the individual's mind.

For any of these alternative models to be helpful, it must be possible to derive empirical predictions which differ from the predictions of a life-cycle model. As we have seen, very many types of behaviour can be rationalised by a life-cycle model. Nevertheless, for certain kinds of behaviour (such as low or negative wealth holding) to be explicable in a life cycle framework, then there are cross-equation restrictions for the behaviour of consumption, labour supply and asset holdings over the life cycle. It may be possible to test these restrictions in order to shed light on the extent to which the observed data are the result of forward-looking rational planning or represent the outcome of a family failing to cope with its own financial affairs. These different explanations for observed savings behaviour are not simply a matter of intellectual curiosity, but have a bearing on a wide range of policy issues such as the pensions system, redistribution over the life cycle, and income maintenance schemes. So-called 'paternalistic' intervention by government may be an appropriate second-best response to the high costs of obtaining information.

There is one final point to be made about the assumption of rational expectations in a life-cycle model. This assumption is most useful when it is reasonable to suppose that individuals have a sufficiently large number of observations of events for them to learn about the stochastic process generating the observations. Subjective probabilities can be related to observable frequencies. For this to be possible the stochastic processes must be stationary. But when planning over a life-cycle many of the relevant events will only be observed once. Most people expect to retire infrequently. When it comes to planning for retirement, learning by doing is somewhat difficult.

4. Empirical Studies of Savings Behaviour

In this section we shall survey some of the recent empirical work on consumption and savings behaviour. The focus will be on the light which these studies shed on the relevance of the models discussed above rather than on their respective merits as forecasting equations of consumption. At the outset three problems inherent in the empirical studies of consumption should be mentioned. All of the studies mentioned below are deficient with respect to at least one of these, and this should be borne in mind when interpreting the results. First, although we have ignored the influence of household composition on the optimal consumption profile, it is clear that changes in composition of families will in general lead to changes in the optimal pattern of consumption. The difficulties in identifying the effects of household composition on consumption patterns in cross-section data are well known, and existing results are far from robust. In studies using aggregate data, changes in demographic structure should appear in a fully specified model. Secondly, there are difficulties in defining the consumption variable used for empirical work. The models used above implicitly assume that the consumption good yields a flow of consumption services. The most natural variable to use in empirical work, therefore, is the value of consumption of non durables and services plus the imputed service flow from durable goods. But most studies exclude durables. This is partly because there are assumed to be costs of adjustment to durables although these are not explicitly modelled. An assumption about separability of preferences between consumption of non durables and the consumption of durables is made in order to justify a specification of consumption excluding durables. This strengthens the general conclusion that the models estimated are tests of the joint hypothesis of forward-looking behaviour and a particular parameterisation of preferences. Thirdly, in models using time-series data (whether aggregate or panel data) a choice must be made about the length of the period. Often this is dictated by the availability of data. But if the length of period used by consumers for planning purposes differs from that used in the estimation, time aggregation biases will be present.

4.1 Studies Using Aggregate Data

Any study based on aggregate data is obviously open to the charge that the existence of aggregation bias means that almost any set of aggregate data could be consistent with underlying rational individual behaviour. Clearly, one of the attractions of the simple linear Hall consumption function is that (leaving aside demographic changes) the model aggregates perfectly. A more interesting question is the use of such models to test the life-cycle hypothesis. This refers back to a point made in the introduction. The use of such tests (such as the significance of additional lagged variables in the consumption function) has been to accept or reject the life cycle model. An alternative inference with rather different implications for policy, is that part of the population behaves according to the life cycle model while another part does not. We shall argue below that this is not an unreasonable conclusion to draw from the data, and its investigation requires the use of microdata for individual households.

Before summarising the results of tests using aggregate data, it may be helpful to point to two stylised facts about aggregate postwar data which are in large part responsible for the results described below. The first is that shocks to the economy tend to result in positively correlated movements in consumption and labour supply. In other words, consumption and hours worked are procyclical. If we rule out the possibility that leisure is an inferior good, then this observation would be consistent with rational individual behaviour only if the real wage moved procyclically by enough to justify the changes in hours worked. But the aggregate real wage has not moved in this fashion. It might be argued that this reflects the fact that the aggregate real wage used in these studies is an average of the wage for those in employment. But whether we regard some households as unemployed in the sense that they face a constraint on hours worked or unemployed in the sense used above that their market wage is close to zero, the analysis above predicts that aggregate models will fail unless they deal explicitly with unemployment and constraints in the labour market. The second stylised fact concerns the lessons we might hope to draw about the intertemporal elasticity of substitution of consumption. This is relevant to those models which allow the interest rate to be time-dependent. Any model which uses factual observations on rates of return to compute predicted expected real rates of return (whether explicitly or as instrumental variables) will observe that there has been a fall in the postwar period in the expected rate of return whereas the growth rate of consumption has changed rather little. Hence it will be difficult to conclude from aggregate data that the response of consumption to changes in interest rates is anything other than rather small.

It is not surprising, therefore, to find that tests of the stochastic life cycle model have been rather disappointing. The study nearest to the theoretical model set out above is that of Mankiw, Rotemberg and Summers (1982). They estimated the stochastic Euler equations for consumption and labour supply using aggregate quarterly data for the United States in the period 1947-80. To do so they used the instrumental variables technique discussed in Section 2.2. The two major findings were that the orthogonality restrictions implied by the first order conditions were rejected, and that the implied utility function was often not concave. Some of their estimates allow for a non separable preference ordering and in this respect their study is unique. The preference ordering which they assume is of the following rather special form.

$$U = \frac{1}{1-\gamma}\left[\frac{C_t^{1-\alpha} - 1}{1-\alpha} + d.\,\frac{(1-L_t)^{1-\beta} - 1}{1-\beta}\right]^{1-\gamma} \tag{4.1}$$

When $\gamma = 0$ then preferences are separable. Given the difficulties with aggregate data mentioned above, it may well be that no other preference ordering would improve matters, but there is something to be said for experimenting with a functional form consistent with the existing results on labour supply as suggested in Section 2.1 above. Perhaps the main conclusion to be drawn from this study, however, is that it is important to incorporate constraints in the labour and capital markets into an empirical specification. This is closely related to the proposition that over the business cycle changes in hours worked are not spread evenly among the working population but tend to fall disproportionately on a minority. The non linear labour supply function of Section 2.1 has the property that for any given change in the aggregate real wage of the population as a whole, a change in hours worked is much smaller if the change in wage is equal for all than if the change in wages takes the form of a small proportion receiving a zero wage offer.

Hansen and Singleton (1982) also estimate the Euler equations and report a rejection of the restrictions implied by the model using postwar US data. Although Hall (1981) does not estimate the Euler equations as such, he makes assumptions about the stochastic processes underlying the model such that he is able to derive an explicit form for the time path of consumption. He assumes that both the interest rate factor and the marginal utility of consumption are lognormally distributed, and so the reduced form is a bivariate regression for consumption and the real interest rate. This is because the product of two log-normally distributed random variables is itself lognormally distributed, and so (2.39) reduces to a regression model. Preferences between consumption and leisure are assumed to be separable. The nature of the stochastic processes for wage and interest rates required for the distribution of marginal utility to be log normal is unclear. The model was estimated using postwar annual data for the US and the coefficient on the expected interest rate was small.

The remaining studies do not estimate the stochastic Euler equations derived from an explicit model of optimising behaviour, but specify a regression model. A formal justification of the specification employed requires the assumption of a nonstochastic interest rate. In fact all but two studies assume a constant interest rate. Wickens and Molana (1982) and Muellbauer (1983) both deal with time-dependent interest rates, but their specifications hold only under the assumption of point expectations. All studies except the two last mentioned use seasonally adjusted data (as do Mankiw et al) despite the autocorrelation of residuals which this practice may induce. All of the papers assume separable preferences between consumption and leisure. Hayashi (1982) is the only study to use data for consumption both on non durables and also the imputed flows from durable goods. Hayashi also uses annual rather than quarterly data for the US from 1948-78 and concludes that disposable income has some explanatory power in addition to its role in the determination of permanent income. Hayashi interprets his point estimates as implying that 17.1% of the population is liquidity constrained. In contrast, Hall (1978) used quarterly data for the US to estimate the linear model discussed above and concluded that lagged values of disposable income were insignificant. He did find, however, that the lagged value of the stock market index had some explanatory power. The papers by Sargent (1978), Flavin (1981) and

Muellbauer (1983) all use the additional information that if labour income follows an autoregressive process then there must be cross-equation restrictions between the innovations in consumption and the stochastic process of income. Flavin corrects Sargent's error in including capital income in his definition of the relevant stochastic process, and concludes that consumption is more sensitive to current income innovations than is implied by the observed auto regressive process for income. This conclusion is conditional on the belief that real interest rates are constant. As Bernanke (1982) and Wickens and Molana (1982) have pointed out, much of this excess sensitivity of consumption to current income may well be due to the omission of changes in real interest rates. Using data for the UK, Wickens and Molana (1982) found that the Hall model adjusted for time-varying interest rates fitted well and provided a good explanation of aggregate consumption. Muellbauer (1983) however, found that although the basic Hall model fitted quarterly UK data for the period 1955-79 surprisingly well, the model was clearly rejected when it was tested over two sub periods within the postwar period. In a study of the joint behaviour of expenditure on non durables and durables, Bernanke (1982) found that the life cycle model could be rejected using quarterly data for the US. Given the assumptions of quadratic utility and a non stochastic interest rate, Bernanke found explicit solutions for the behaviour of consumption in terms of lagged consumption and past stocks of durable goods.

Finally, there is the voluminous literature on econometric studies of the Keynesian consumption function which relates consumption to current and distributed lags of disposable income. Much of this work in recent years has been directed at exploring the role of inflation in the consumption function; see, for example Hendry and von Ungern-Sternberg (1980), Davidson et al (1978), Davidson and Hendry (1981), and Pesaran and Evans (1982). These studies confirm that the best forecasting equation for consumption is one which includes variables additional to lagged consumption. The conclusions we might draw from the studies are the following. First, if one's main aim is to obtain the best forecast of future consumption then a rather eclectic view of the appropriate consumption function seems appropriate. Any lagged variable which improves the predictive power could be included although it is likely that the equation will exhibit parameter instability if the stochastic processes determining the exogenous variables relevant to consumption decisions are changing. But most forecasters are in the habit of re-estimating their equations on a regular basis. Such an approach does not tell us very much about the underlying model of behaviour. For this we need to contrast empirical predictions of competing models. Our analysis of the stochastic life cycle model with constraints in capital and labour markets suggests that it is unlikely that aggregate data will enable us to distinguish between competing explanations of observed behaviour. The results to be obtained from aggregate models are largely dictated by the two stylised facts mentioned at the beginning of this section. Moreover, very few of the above studies estimated a model based on explicit optimising behaviour under uncertainty. Nevertheless it is clear that tests using aggregate data generally support the view that consumption is too sensitive to current income for us to accept the life-cycle model without some modification, such as the introduction of borrowing constraints for at least part of the population.

4.2 Studies Using Household Data

In principle the use of individual microdata gives us many more observations with which to test the implications of the life cycle model. It allows us to estimate the proportion of the population for which this model appears relevant. But it draws attention also to two problems which are often overlooked at the aggregate level. First, there is the question of the quality of the data. With the development of econometric techniques it has become increasingly common (if not even fashionable) to use panel data for which we have observations on the same household for several periods. Such data do allow us to test the stochastic life cycle model using the formulations derived above. This is an advantage over cross-section survey data. But panel data incur problems of attrition as households drop out of the survey, and often the data for later years in the panel are obtained in a less reliable fashion, such as by interviews on the telephone. Some of the cross-section surveys, in contrast, are based on detailed interviews where the respondents are asked to provide documentary evidence of their replies. In assessing the quality of data it is helpful to understand how the data were collected and to adopt an appropriate estimation procedure. The second problem, sometimes used to justify the use of aggregate data, is that once individual data are used it is difficult to deny the existence of heterogeneous preferences. Where these preference parameters are correlated with observable variables, then it may be very hard to identify the model. This problem, however, provides no justification for the use of aggregate data. If preferences are indeed correlated with observable variables then any aggregate model will exhibit chronic parameter instability. It suggests that models be formulated in such a way that it is not unreasonable that the preference parameters be assumed to be uncorrelated with the explanatory variables.

Given the wealth of information obtained in micro data sets, it is often helpful to use this in two ways. The first is to test explicit models such as the stochastic life cycle model, and the second is to display the correlations between recorded variables in the survey. We shall discuss examples of both. The stochastic life cycle model relates to the rate of change of consumption. None of the studies using individual household data had access to sufficiently long time series of changes in consumption for individual households to enable the model to be tested as it has been on aggregate data. But the change in consumption can be related to the innovations in the stochastic process generating incomes, and by estimating these processes one can test the cross-equation restrictions on the innovations in consumption and income implied by the life-cycle model. This was attempted by Hall and Mishkin (1982) who examined data for 2,309 families from the University of Michigan's panel study of income dynamics. The measure of consumption was annual expenditure on food used at home and expenditure in restaurants. Data on the rates of change of consumption for five years were available. Hall and Mishkin assumed that labour income consisted of three components, (1) a deterministic component which was age-dependent, (2) a first stochastic component which measured changes in lifetime prospects and was assumed to follow a random walk, and (3) a second stochastic component measuring transitory fluctuations and assumed to be described by a moving average time series process. They assumed also that the utility function in consumption (but not defined over labour supply) was quadratic and a nonstochastic interest rate equal to the rate of time preference. With these assumptions then the change in consumption is given by the following equation.

$$\Delta C = \varepsilon_t + \beta_t \eta_t \qquad (4.2)$$

where ε_t is the innovation in the first stochastic component and η_t is the innovation in the transitory stochastic component. Because the first stochastic component is assumed to follow a random walk then its innovations will be expected to be permanent and the marginal propensity to consume out of this component is equal to the marginal propensity to consume out of permanent income. Innnovations in the transitory stochastic component will have a much smaller impact on current consumption and the value of β_t is the annuity value of an innovation in the transitory component. Because this is assumed to follow a moving average time-series process, the value of β_t depends both upon interest rates, expected lifetime and also the parameters of the moving average process. But it will clearly be less than unity. The variance of changes in consumption will be a function of the variances of the stochastic processes generating income. The ratio of the marginal propensity to consume out of transitory income relative to the marginal propensity to consume out of lifetime income was estimated at 0.29. Given the estimated moving average process for income, this estimate would be consistent with life cycle behaviour only at real interest rates of around 30%. Hence the model appears to be rejected. By augmenting the model to allow for the possibility that a certain fraction of consumption is constrained to be equal to current income, Hall and Mishkin find that about 80% of consumption can be explained by the lifecycle model and 20% appears to be determined by current income. They do not however test the more interesting model in which a certain fraction of individuals follow the life cycle model and the remaining fraction are constrained. Instead, the same model is imposed for all individuals and the mixture model proposed above has yet to be tested. Using the same model for the stochastic processes generating income, Bernanke (1981) found that for data on expenditures on automobiles of 1,434 house holds the life cycle model could be accepted and the test that consumption was constrained by current income was rejected. When integrating the demand for durables with that for nondurables, however, Bernanke (1982) rejected the life cycle model (although this test used aggregate data).

The studies cited above examined the behaviour of consumption. An alternative approach is to examine the pattern of wealth holdings and their relationship with age in a cross-section of households. An ambitious attempt to estimate the age profile of asset holdings was attempted by Blinder et al (1980) who, assuming perfect certainty and a specific functional form for the additively separable utility function of consumption, estimated the deterministic path for wealth holdings given the individual household's earnings profile. This was obtained using data for 4,133 white males from the Longitudinal Retirement History Survey. The optimising model implies a very specific non linear function of asset holdings over the life cycle, and because of the tight parameterisation the model proved hard to estimate with imprecisely determined parameters. One obvious difficulty with such a model is that there is no reason to believe that optimal behaviour under uncertainty can be proxied by adding a stochastic error term to a model of optimal behaviour under certainty. An example of this is the low rate of decumulation of wealth one would expect to find after retirement in the presence of uncertainty about date of death and an absence of sufficient real annuity markets. Nevertheless, one would expect to find a 'hump-shaped' pattern of wealth holding for a majority of the population if the life cycle model were an adequate description of their behaviour. The existence of such a pattern has proved hard to find. Lydall (1955) found no tendency for wealth to decline with age using data from the 1953 Oxford Institute of Statistics Survey, and a similar result was found by Mirer (1979) in the United States for a sample of 2713 couples. But this finding was obtained using a simple regression of wealth on age. In these studies no attempt was made to control the wealth-age relationship for the effect of differences in permanent income, the amount of wealth held in the form of pension rights, and the fact that the life cycle model implies a highly non propensity to consume out of lifetime income was estimated at 0.29. Given the estimated moving average process for income, this estimate would be consistent with life cycle behaviour only at real interest rates of around 30%. Hence the model appears to be rejected. By augmenting the model to allow for the possibility that a certain fraction of consumption is constrained to be equal to current income, Hall and Mishkin find that about 80% of consumption can be explained by the lifecycle model and 20% appears to be determined by current income. They do not however test the more interesting model in which a certain fraction of individuals follow the life cycle model and the remaining fraction are constrained. Instead, the same model is imposed for all individuals and the mixture model proposed above has yet to be tested. Using the same model for the stochastic processes generating income, Bernanke (1981) found that for data on expenditures on automobiles of 1,434 house holds the life cycle model could be accepted and the test that consumption was constrained by current income was rejected. When integrating the demand for durables with that for nondurables, however, Bernanke (1982) rejected the life cycle model (although this test used aggregate data).

The studies cited above examined the behaviour of consumption. An alternative approach is to examine the pattern of wealth holdings and their relationship with age in a cross-section of households. An ambitious attempt to estimate the age profile of asset holdings was attempted by Blinder et al (1980) who, assuming perfect certainty and a specific functional form for the additively separable utility function of consumption, estimated the deterministic path for wealth holdings given the individual household's earnings profile. This was obtained using data for 4,133 white males from the Longitudinal Retirement History Survey. The optimising model implies a very specific non linear function of asset holdings over the life cycle, and because of the tight parameterisation the model proved hard to estimate with imprecisely determined parameters. One obvious difficulty with such a model is that there is no reason to believe that optimal behaviour under uncertainty can be proxied by adding a stochastic error term to a model of optimal behaviour under certainty. An example of this is the low rate of decumulation of wealth one would expect to find after retirement in the presence of uncertainty about date of death and an absence of sufficient real annuity markets. Nevertheless, one would expect to find a 'hump-shaped' pattern of wealth holding for a majority of the population if the life cycle model were an adequate description of their behaviour. The existence of such a pattern has proved hard to find. Lydall (1955) found no tendency for wealth to decline with age using data from the 1953 Oxford Institute of Statistics Survey, and a similar result was found by Mirer (1979) in the United States for a sample of 2713 couples. But this finding was obtained using a simple regression of wealth on age. In these studies no attempt was made to control the wealth-age relationship for the effect of differences in permanent income, the amount of wealth held in the form of pension rights, and the fact that the life cycle model implies a highly non linear relationship between the ratio of wealth to permanent income and age. Using data for 12,734 Canadian families in 1977, King and Dicks-Mireaux (1982) found that a hump-shaped profile for asset holdings over the life cycle could be detected even for assets excluding rights to pensions and social security payments. The estimated profile was plotted in Figure 2 above. There was evidence of decumulation after retirement, although the rate of decumulation is much less than would be predicted by a life cycle model. We have commented in Section 2 on the possibility that low rates of decumulation can be explained by uncertainty about date of death. We suggested there an empirical test to allow for this possibility by estimating separately the wealth and insurance effects of pensions. Slightly higher rates of decumulation, of up to 5% per annum, were found by Diamond and Hausman (1982) in a study based on data from the National Longitudanal Survey of Mature Men. An interesting observation, common to both the Diamond and Hausman, and King and Dicks-Mireaux studies, is the fact that a substantial proportion of the sample reported extremely low levels of asset holdings. Diamond and Hausman found that 19.7% of the sample reported net worth below $1,000 in 1966 and King and Dicks-Mireaux found that 23.4% of their sample had net worth in 1977 of less than $2,500 Canadian (a comparable figure allowing for inflation and the exchange rate). It is conceivable that such low levels of asset holdings are consistent with the life cycle model if individuals have somewhat higher than average rates of time preference and expect to receive pensions comparable with their earnings before retirement. But King and Dicks Mireaux found no correlation between the incidence of low wealth holdings and the ratio of pensions to permanent income. Although further study is required, there does seem to be evidence that the population consists of groups with different types of savings behaviour. One question which follows naturally is, therefore, what are the characteristics which determine whether households behave according to the life cycle model or according to some other model? King and Dicks Mireaux found that there was a distinct difference between households who reported negative net worth and those who reported positive but small amounts of net worth. Those with negative net worth were significantly younger than the rest of the population and this finding is consistent with the view that young households incur debts, which are recorded in the survey, to purchase relatively inexpensive durable goods which are often excluded from such surveys. But households with positive though very low asset holdings were not significantly younger than the rest of the population, but did share certain characteristics. The households with small amounts of assets were much more likely to have unemployed members of the household or members receiving extremely low levels of current earnings. This is consistent with the view of constraints in labour and capital markets discussed in Section 3. Another finding was that even controlling for the level of permanent income and earnings, education had additional explanatory power in determining whether a household held 'adequate' assets. Households where the head had received high levels of education were more likely to hold large amounts of assets. It would probably be unwise to speculate but there may be something in the view that the ability to manage one's own financial affairs and make provision for retirement is correlated with the level of education.

Low ratios of wealth to permanent income were found also in a study of Japanese data by Shimono and Tachibanaki (1982) who found that for workers with 'lifetime' contracts the ratio of wealth to income at age 55 (the normal age of retirement from the main job in Japan) was on average about 4 whereas for workers employed without such contracts the ratios were significantly lower. Again, there are educational differences between the two groups, and it would be interesting also to model the expected differences in wealth holdings resulting from the different types of labour contracts under which the two operate.

We have contrasted studies which examine either the change in consumption over time or the level of asset holdings over the life cycle. It is also possible to examine the cross-section variations in consumption and examine its relationship with age. Using data for over 1,900 married couples from the Retirement History Survey, Kotlikoff, Spivak and Summers (1982) computed the constant consumption stream between the current age and anticipated age at death which could be financed from current assets and expected future returns. They computed also a constant consumption stream for the entire lifetime of each couple which could be financed from lifetime earnings and pension receipts (using matching records on earnings from the Social Security Earnings File). A comparison of the two consumption streams indicates the ability of households to maintain in retirement a consumption level close to that which could have been financed throughout the entire working life. A substantial proportion of households has ratios of potential constant consumption over remaining lifetime to constant consumption over entire lifetime greater than unity. This is adduced as evidence that given the present social security system in the US, there is little evidence of inadequate savings for retirement by the aged. Hamermesh(1982) has argued, however, that it may be more appropriate to compare potential consumption during retirement with consumption immediately after retirement. Using this standard, rather than consumption over entire lifetime, he found that many individuals in the Retirement History Survey experienced declines in consumption during retirement. To the extent that it is optimal for consumption to decline in retirement, either because of greater leisure or ill health, then this pattern may not be disturbing. It is also the case that the elderly consume the services of durables to a greater extent than those of young households. But further study of the consumption patterns of the retired would throw some light on the issue.

The main conclusion of this section must be that support for the life cycle model is mixed. The phenomenon worthy of further investigation is the identification of the characteristics which determine whether households are more likely to follow the life cycle model or an alternative model of savings behaviour. This approach has, as yet, not been formally modelled but it appears to be fruitful and to test it one would like data on individual families for two or three years with emphasis on the quality of the data on incomes, consumption, and assets. It would not be inappropriate to sacrifice the length of the panel to improve the quality of the survey responses.

5. Bequests

The introduction of bequests raises a number of important questions which we shall do no more than mention, such as the way in which the bequest motive enters the donor's utility function, the division of legacies among recipients, the timing of gifts, and the allocation of expenditures on children's education. But whatever the optimal level of bequests which emerges from these considerations, a potential donor will still allocate consumption over his or her life cycle using the same criteria as analysed in preceding sections. The existence of bequests does not, therefore, alter the first-order conditions determining the rate of change of consumption over the life cycle. It does, however, introduce some new dimensions into the discussion of savings.

If the utility of an individual is defined over the utility level of his heirs, then he belongs to an infinite chain of families. This means that there can be no intergenerational externalities and the existence and size of government debt and social security, whether funded or pay-as-you-go, has no real effect (Barro, 1974, 1978). This is because individuals can alter their bequests so as to exactly offset the intergenerational redistribution implied by policy and attain the original optimum (although taxes may lead to additional efficiency losses). In this way private transfers will offset the public transfers implied by debt or social security policy. The result depends, however, on the ability of individuals to reach an interior optimum. If the private transfers required to offset the effects of policy lead to a corner solution (for example, negative bequests may not be feasible) then the neutrality result no longer obtains.

To test the validity of the view that individuals care about the welfare of their descendants we could examine the empirical significance of the neutrality proposition. There are, however, two problems with this test. First, social security wealth is unobservable and empirical proxies are subject to the errors in variables bias which will lead to estimated coefficients on social security wealth which may be below their true value. The insignificance of social security wealth in a regression may be the result of errors in measurement rather than the optimal behaviour of an infinitely-lived family. Secondly, the infinitely-lived family may behave in exactly the same way as a completely myopic family. The introduction of an unfunded pension scheme would lead pure life cycle savers to reduce their current saving and hence social security would depress their capital accumulation. But infinitely lived families would continue to accumulate assets so that future generations would be able to meet the cost of social security schemes. Equally, myopic families would ignore the future entitlements to social security and its introduction would leave savings patterns unaffected. This observational equivalence can be settled only by a priori reasoning or alternative empirical tests.

There is a deeper question about whether the decision to give or leave bequests should be regarded as the outcome of an individual optimisation or whether we should take into account also the response of the recipient. Kurz (1982) has argued that gifts and bequests should be seen as the outcome of an intergenerational game in which the interdependence in utilities is explicitly recognised and embodied in the appropriate equilibrium concept. One example of such an equilibrium would be a Nash equilibrium in which the bequests of each generation were optimal given the bequests made by all other generations. There are two problems with this view. First, as pointed out by Kurz, the notion of an equilibrium of a game in which most of the players are absent is unclear. Secondly, there are some difficulties in defining equilibrium for an economy in which bequests play a key role when some families have no heirs (Bewley 1981).

Viewing bequests as the outcome of an intergenerational game is only one of a number of ways of modelling the phenomenon. For example, we could assume that parents care about their children's opportunities rather than their utility as such, and define parental utility over own consumption and the budget sets of heirs. There is no game-theoretic aspect here and bequests enter the utility function in a way which depends upon children's talents and endowments. Alternatively, we could view bequests as constituting an inducement for children to spend time with their parents. This approach has been modelled by Bernheim, Shleifer and Summers (1983).

Within a life-cycle framework the timing of gifts and bequests clearly influences the age profile of asset holdings (Atkinson, 1971). We have no good theory concerning the timing of gifts. Bequests also have implications for the composition of household portfolios. With uncertainty over date of death savings for post-retirement consumption will be in the form of annuities (if indexed annuities are available) and the riskiness of legacies can be reduced by the purchase of life insurance. Differential mortality rates and the valuation of life insurance are a major problem in interpreting cross-section evidence on wealth-age profiles (Shorrocks, 1975).

In two studies based on data collected for the US President's Commission on pension policy, Kurz (1981)(1982) derived specifications for asset holdings within a life cycle model and for transfers (whether gifts made or received in a model of inter-generational equilibrium. Optimal transfers were assumed to be a linear function of current assets, the wage rate, net social security wealth, public transfers, and a vector of socioeconomic characteristics. The parameters of the model are shown to be generation-dependent. This causes a problem for empirical work because in the formal model each generation lives for only two periods, and hence there is no confusion between the age of an individual and his generation. It would appear from the model that the relevant explanatory variables in Kurz's equation should be those for a generation, and hence should be lifetime equivalent values of assets and the wage rate rather than the values at any particular age. In principle, we need to distinguish between generational differences in the explanatory variables and differences which result simply because we are observing individuals at different stages of their life cycle. This distinction was not made in the empirical work. Leaving this to one side, Kurz proposes three tests of the model. First, the effect of social security wealth on transfers should be such that any increment to net social security wealth of a generation will result in an equal increment to transfers to the next generation. Changes in public policy cannot affect the intergenerational equilibrium, as in the Barro model. Secondly, the structural equation determining transfers should hold at all stages of life and hence is relevant for both recipients and donors. Thirdly, the effect of social security wealth on private savings should be such that private savings will be unaffected by changes in social security wealth, again leaving the intergenerational equilibrium undisturbed. The results of the first test were inconclusive. The second test led to a rejection of the model of intergenerational equilibrium. The third test on an equation for asset holdings led to support for the model. One problem with these results is that they could all be affected by errors in variables if the social security wealth variable is measured with error. This would lead to an unnecessary rejection of the model in test A and acceptance in test C.

Despite the difficulties in modelling bequests empirically, it is clear that the subject merits further work both at the theoretical and empirical level. For some households, although probably a minority, bequests are likely to represent one of the main motives for accumulation of wealth.

6. Conclusions

We may briefly summarise the findings of the Survey as follows. The literature has drawn a number of false inferences about the relevance of the life cycle model from empirical studies. We have tried to sort out those empirical tests which would throw light on the issue from those which merely reflect assumptions about the parameterisation of preferences. The proposed tests rely on the cross-equation constraints imposed by the life cycle model (for example, those between consumption and labour supply profiles) and on the way in which constraints should enter the model.

There is an unfortunate tendency for studies based on aggregate data to regard the life cycle model as one which applies to everyone or no-one. In fact, there is a remarkable conformity in the findings using cross-section data, panel data, and aggregate time-series data that about twenty or twenty-five per cent of the population behaves in a way which may well be inconsistent with the life cycle model, whereas for the rest of the population observed savings behaviour is not inconsistent with the model. The findings of Diamond and Hausman and King and Dicks-Mireaux with cross- section data, the panel study of Hall and Mishkin and the numerous aggregate time-series studies all produce coefficient estimates which point towards this conclusion. The problem with aggregate data is that we know the stylised facts are inconsistent with the view that aggregate behaviour can be explained as if it were generated by a representative individual following a life cycle model (see section 4.1 above). Hence studies based on aggregate data are bound to find significant coefficients on some variables which appear to violate the predictions of a pure life cycle model. But the interesting question is whether these violations are common to all households or whether they reflect the fact that changes over the business cycle fall disproportionately on a minority of the population. My conjecture is that it will prove fruitful to explore the latter point of view and the issue is important for policy purposes.

The life cycle model allows individuals to separate their consumption in any year from their current income. In this framework the stage of the lifecycle at which taxes or transfers are levied is a matter of indifference. If individuals are forward-looking but are constrained in any period, then their lifetime welfare cannot be measured by the present value of lifetime endowment. There is a role for policy in either relaxing the constraints or using transfers to compensate for the reduction in the utility which results from the existence of binding constraints. Obviously, unemployment insurance is an example of this type of policy. For individuals who are backward-looking, 'myopic', or whose behaviour is described by an alternative model, then there may be scope for paternalistic intervention. It is important for future empirical work to try to identify the characteristics of individuals who do not follow a life cycle pattern of behaviour. In theoretical terms the difficulty is that we need a coherent alternative model to analyse such behaviour. Such a model would need to take into account the following considerations. First, in a life cycle context there is little possibility of learning by doing (or at least by personal experience). Decisions about retirement and savings for the future can not be compared with decisions about purchases of breakfast cereal or soap powder. Convergence to rational expectations is a process which depends upon observations of past errors and their incorporation into revised beliefs about the underlying process generating observed data. Over the life cycle it is unlikely that the underlying processes are stationary. Given that those currently retired have experienced at least one major war and can now purchase products which could not have been conceived in their youth, decisions about retirement raise again the old issue of risk versus uncertainty. A second consideration is that the cost of thinking about the future may well exceed the expected benefit from devoting more time to planning retirement portfolios. This is either because of the complications of decision making under uncertainty (which may lead to bounded rationality) or because thoughts of the future may imply contemplating unpleasant outcomes such as ill health. The implications of such behaviour for policy are unclear. Paternalistic intervention to help individuals who behave in this way runs the risk that other individuals will, by anticipating the policy response, be able to benefit substantially. Moral hazards of this sort might justify some element of compulsory saving for all individuals, and the analysis of optimal public policy in an economy with different types of savings behaviour appears to be a fruitful topic. Empirical work can make a contribution here by quantifying the relative importance of different models. Just as Friedman's(1957) study of the consumption function evolved out of empirical work on family budgets by Dorothy Brady and Margaret Reid, it is the interplay between empirical findings and theoretical models which is likely to take us a step forward. From this point of view, the aim of empirical work should be less to estimate parameters or obtain the best predictions for the future than to fulfil the hopes of Keynes (1973, p.296 in a letter to Roy Harrod) that 'the object of statistical study is not so much to fill in missing variables with a view to prediction as to test the relevance and validity of the model'.

Footnotes

[1] For further discussion see Becker (1981) and Deaton and Muellbauer (1980, Chapter 8). Tobin (1967) presents simulations based on an explicit life-cycle pattern of family composition.

[2] We are abstracting from changes in tax rates. Such changes would mean that the normalization would be appropriate only if the relevant tax rates were introduced explicitly into the model.

[3] The fixed effect is the multiplier λ and so must be a function of the parameters of the individual's optimization problem.

[4] This is true only if we can assume that all individuals face the same real interest rate. Conversely, if interest rates vary across individuals (because of different marginal tax rates, for example) then the equation could be estimated using micro data.

[5] This assumption is made in all of the recent work on consumption and also in the labour supply study of Heckman and MaCurdy (1981).

[6] For a further discussion of the derivation of the indirect utility function from estimated labour supply functions see Hausman (1981a)

[7] For the Slutsky condition to hold the expenditure function corresponding to (2.24) must be concave in the wage rate. This implies that $c(1 - L) - b.\gamma \exp(-\gamma w) \leq 0$. This condition is always satisfied in our simulations, and would only be violated at very high wage rates.

[8] Although U is defined only up to a positive monotonic transformation, the cardinalization of v does matter because it determines the intertemporal allocation of total expenditure. The cardinalization assumed in (2.24) and (2.25) was chosen to yield age profiles for labour supply consistent with the empirical evidence cited in the text (see Figure 1).

[9] Although the parameter values have been chosen to replicate observed life-cycle profiles of labour supply and earnings, it is interesting to note that the labour supply function is consistent also with results from cross-section studies. If we take a first-order expansion of the exponential term in equation (2.20), then we obtain the linear model of labour supply estimated by Hausman (1981 a, b). The coefficient values are remarkably similar. For example, the coefficient on the hourly wage estimated by Hausman was 495.1 compared with 585.0 implied by our parameters. Both the intercept term and the "income" coefficient are

similar also, although in our specification "income" is replaced by the current period's "full expenditure" on consumption and leisure. This is endogenous to the optimal life-cycle plan. If we include an extra second-order term in the exponential expansion then we obtain the restricted quadratic specification proposed by Stern (1983) in which labour supply is quadratic in the wage but linear in other variables.

10 Stochastic Euler equations in the case of additive separability over time have been derived by Le Roy (1973) Lucas (1978) and Hall (1978) among others.

11 This is because in the dual form there is only a single stochastic Euler equation to be estimated regardless of the number of commodities over which preferences are defined.

12 Constraints at positive quantities of labour supply can be handled in a straightforward manner.

13 The variance of first differences of consumption will, however, differ in the constrained case. Explicit examples can be obtained by using the preference ordering used for the simulation in section 2 for which the virtual wage may be obtained in implicit form. It is easy to confirm that when $r = \rho$ and $w_t = w_{t-1}$, both y_t and C_t follow a random walk in both Cases 1 and 4.

14 In the simulations in Section 2.1 we found that a value for c of 0.1 fitted the stylized facts and a similar coefficient was found by Hausman (1981a) for the linear approximation ot our nonlinear labour supply function.

15 Of couse many individuals do borrow against future earnings to purchase durable goods such as owner-occupied houses because the durables may be used as collateral. But it is more difficult to borrow for the purpose of consumption of nondurables and this is another reason for believing that borrowing constraints are likely to be more relevant to those who have suffered substantial falls in earnings, such as the unemployed.

16 Perhaps because it is easier for lenders to monitor announcements or rumours about redundancies (which tend to affect many easily identified individuals) than anticipated returns to work.

References

Aaron, H.J. (1982) Economic Effects of Social Security, Brookings Institution, Washington D.C.

Appelbaum, E. and Harris, R. (1978) "Imperfect Capital Markets and Life-Cycle Savings" Canadian Journal of Economics, 11, 319-24.

Atkinson, A.B. (1971), "The Distinction of Wealth and the Individual Life-Cycle", Oxford Economic Papers 23, 239-254.

Atkinson, A.B. and Sandmo, A. (1980), "Welfare Implications of the Taxation of Savings", Economic Journal 90, 529-49.

Auerbach, A.J. and Kotlikoff, L.J. (1981) "National Savings, Economic Welfare and the Structure of Taxation", National Bureau of Economic Research Working Paper No. 729.

Barbera, S. and Pattanaik, P.K. (1981) "Rationalizability of Stochastic Choices in Terms of Random Orderings", University of Birmingham mimeo.

Barro, R.J. (1974) "Are Government Bonds Net Wealth?", Journal of Political Economy 82, 1095-1117.

Barro, R.J. (1978) "The Impact of Social Security on Private Saving: Evidence from the US Time Series", American Enterprise Institute, Washington D.C.

Becker, G.S. (1981) A Treatise on the Family, Harvard University Press.

Bernanke, B. (1981) "Permanent Income, Liquidity, and Expenditure on Automobiles: Evidence from Panel Data", National Bureau of Economic Research Working Paper No. 756, mimeo.

Bernanke, B. (1982) "Adjustment Costs, Durables and Aggregate Consumption", National Bureau of Economic Research Working Paper No. 1038.

Bernheim, B.D., Shleifer, A. and Summers, L.H. (1983), "The Manipulative Bequest Motive", mimeo.

Bewley, T. (1981) "The Indeterminacy of Interest Rates", Northwestern University Discussion Paper No. 491, mimeo.

Blinder, A.S. (1975) "Distribution Effects and the Aggregate Consumption Function", Journal of Political Economy, 83, 447-75.

Blinder, A.S. (1982) "Private Pensions and Public Pensions: Theory and Fact", National Bureau of Economic Research Working Paper No. 902.

Blinder, A.S., Gordon, R.H. and Wise, D.E. (1980) "Life-Cycle Savings and Bequests: Cross-Sectional Estimates of the Life Cycle Model", National Bureau of Economic Research Working Paper 619.

Champernowne, D.G. (1969) Uncertainty and Estimation in Economics, Vol. 3, Holden-Day, San Francisco.

Courant, P., Gramlich, E. and Laitner, J. (1982) "A Dynamic Micro Estimate of the Life Cycle Model", presented to the Brookings Institution Conference on Retirement and Aging, Washington, D.C., mimeo.

Davidson, J.E.H. and Hendry, D.F. (1981) "Interpreting Econometric Evidence: The Behaviour of Consumers' Expenditure in the UK", European Economic Review 15, 177-92.

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

Davies, J.B. (1980) "Uncertain Lifetime, Consumption and Dissaving in Retirement", Journal of Political Economy, 89, 561-77.

Deaton, A.S. and Muellbauer, J. (1980) Economics and Consumer Behaviour, Cambridge University Press.

Diamond, P.A. (1977) "A Framework for Social Security Analysis", Journal of Public Economics, 8, 275-98.

Diamond, P.A. and Hausman, J.A. (1982) "Individual Retirement and Savings Behaviour", presented to the NBER/SSRC Conference on Microdata and Public Economics, Oxford.

Dicks-Mireaux, L. and King, M.A. (forthcoming) Pension Wealth and Household Savings: Tests of Robustness", Journal of Public Economics.

Dolde, W. (1978) "Capital Markets and the Short-Run Behavior of Life-Cycle Savers", Journal of Finance 33, 413-28.

Duesenberry, J.S. (1948) "Income-Consumption Relations and their Implications" in Income, Employment and Public Policy, Essays in Honor of Alvin H. Hansen, Norton, New York.

Duesenberry, J.S. (1949) Income, Saving and the Theory of Consumer Behavior, Harvard University Press, Cambridge.

Eden, B. and Pakes, A. (1981) "On Measuring the Variance-Age Profile of Lifetime Earnings", Review of Economic Studies, 48, 385-94.

Fisher, I. (1907) The Rate of Interest: Its Nature, Determination and Relation to Economic Phenomena, Macmillan, New York.

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

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

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

Grossman, S.J. and Shiller, R.J. (1981) "The Determinants of the Variability of Stock Market Prices", American Economic Review 71, Papers and Proceedings, 222-27.

Hall, R.E. (1978) "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence", Journal of Political Economy 86, 971-87.

Hall, R.E. (1981) "Intertemporal Substitution in Consumption", National Bureau of Ecnomic Research Working Paper No. 720.

Hall, R.E. and Mishkin, F.S. (1982) "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households" Econometrica 50, 461-81

Hamermesh, D.S. (1982) "Consumption during Retirement: The Missing Link in the Life Cycle", University of Michigan, mimeo.

Hansen, L.P. and Sargent, T.J. (1980) "Formulating and Estimating Dynamic Linear Rational Expectations Models", Journal of Economic Dynamics and Control, Vol. 2, 7-46.

Hansen, L.P. and Sargent, T.J. (1981) "Linear Rational Expectations Models for Dynamically Interrelated Variables", Chapter 8 in (eds.) Lucas, R.E. and Sargent, T.J. Rational Expectations and Econometric Practice, Allen and Unwin, London.

Hansen, L.P. and Singleton, K.J. (1982) "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models" Econometrica Vol. 50, 1269-86.

Hart, O.D. (1983) "Optimal Labour Contracts under Asymmetric Information: An Introduction", Review of Economic Studies, 50, 3-35.

Hausman, J.A. (1981a) "Exact Consumer's Surplus and Dead-weight Loss", American Economic Review 71, 662-76.

Hausman, J.A. (1981b) "Labour Supply" in (eds.) Aaron, H.J. and Pechman, J.A. How Taxes Affect Economic Behaviour Brookings Institution, Washington, D.C.

Hayashi, F. (1982) "The Permanent Income Hypothesis: Estimation and Testing by Instrumental Variables", Journal of Political Economy, 90,895-918.

Heckman, J.J. (1974) "Life Cycle Consumption and Labour

Supply: An Explanation of the Relationship Between Income and Consumption over the Life Cycle", American Economic Review 64, 188-194.

Heckman, J.J. and MaCurdy, T.E. (1980) "A Life Cycle Model of Female Labor Supply" Review of Economic Studies 47, 47-74.

Hendry, D.F. and von Ungern-Sternberg, T. (1980) "Liquidity and Inflation Effects on Consumers' Expenditure", in (ed.) Deaton, A.S. Essays in the Theory and Measurement of Consumers' Behaviour, Cambridge University Press.

Hicks, J.R. (1939) Value and Capital, Oxford University Press.

Hurd, M. (1971) "Changes in Wage Rates Between 1959 and 1967", Review of Economics and Statistics, 53, 189-99.

Keynes, J.M. (1936) The General Theory of Employment, Interest and Money (page reference to The Collected Writings of John Maynard Keynes, Vol. VII, Macmillan, London).

Keynes, J.M. (1973) The Collected Writings of John Maynard Keynes XIV The General Theory and After Part II Defence and Development, Macmillan, London.

King, M.A. (1980) "Savings and Taxation", in (eds.) Hughes, G.A. and Heal, G.M. Public Policy and the Tax System, Allen and Unwin London.

King, M.A. and Dicks-Mireaux, L. (1982) "Asset Holdings and the Life Cycle", Economic Journal 92, 247-67.

King, M.A. and Fullerton, D. (eds.)(forthcoming) The Taxation of Income from Capital: A Comparative Study of the US, UK, Sweden and West Germany, Chicago University Press.

Koskela, E. (1978) "On the Theory of Consumption Function with Borrowing Constraints", Schweizerische Zeitschrift für Volkswirtschaft und Statistik 114, 739-52.

Kotlikoff, L.J. (1973), "Taxation and Savings", mimeo.

Kotlikoff, L.J., Spivak, A. and Summers, L.H. (1982) "The Adequacy of Savings", American Economic Review 72, 1056-69.

Kurz, M. (1981) "The Life-Cycle Hypothesis and the Effects of Social Security and Private Pensions on Family Savings", Stanford University Institute for Mathematical Studies in the Social Sciences Technical Report No. 335 (revised), mimeo.

Kurz, M. (1982) "Capital Accumulation and the Characteristics of Private Intergenerational Transfers" Stanford University Institute for Mathematical Studies in the Social Sciences Technical Report No. 364, mimeo.

LeRoy, S.F. (1973) "Risk Aversion and the Martingale Property of Stock Prices", International Economic Review 14, 436-46.

Levhari, D. and Mirman, L.J. (1977) "Savings and Consumption with an Uncertain Horizon", Journal of Political Economy 85, 265-81.

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

Lydall, H. (1955) "The Life Cycle in Income, Saving and Asset Ownership", Econometrica 23, 131-50.

MaCurdy, T.E. (1981) "An Empirical Model of Labor Supply in a Life-Cycle Setting", Journal of Political Economy, 89, 1059-85.

Mankiw, N.G., Rotemberg, J.J. and Summers, L.H. (1982) "Intertemporal Substitution in Macroeconomics", National Bureau of Economic Research Working Paper No. 898.

Mann, N.E.W. (1981) "A Bibliography of Life Cycle and Related Studies", QSEP Research Report No. 9, McMaster University Hamilton, mimeo.

Mincer, J. (1976) "Progress in Human Capital Analysis of the Distribution of Earnings" in (ed.) Atkinson, A.B. The Personal Distribution of Incomes, Allen and Unwin, London.

Mirer, T.W. (1979) "The Wealth-Age Relationship Among the Aged", American Economic Review, 69, 435-43.

Modigliani, F. and Ando, A. (1957) "Tests of the Life Cycle Hypothesis of Savings", Bulletin of the Oxford University Institute of Economics and Statistics 19, 99-124.

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

Muellbauer, J. (1983) "Surprises in the Consumption Function", Economic Journal, Supplement of Conference Papers.

Muth, J.F. (1960) "Optimal Properties of Exponentially Weighted Forecasts", Journal of the American Statistical Association 55, 299-306.

Pesaran, M.H. and Evans, R.A. (1982) "An Inflation-Adjusted Life-Cycle Explanation of UK Personal Savings: 1953-80" University of Cambridge, mimeo.

Psacharopoulos, G. and Layard, R. (1979) "Human Capital and Earnings: British Evidence and a Critique", Review of Economic Studies, 46, 485-503.

Ramsey, F.P. (1928) "A Mathematical Theory of Saving", Economic Journal, 38, 543-59.

Sandmo, A. (1970) "The Effects of Uncertainty on Saving Decisions", Review of Economic Studies, 37, 353-60.

Sandmo, A. (Forthcoming) "The Effects of Taxation on Savings and Risk-Taking", in (eds.) Auerbach, A.J. and Feldstein, M.S. Handbook of Public Economics, North-Holland.

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

Schumpeter, J.A. (1954) History of Economic Analysis, Allen and Unwin, London.

Shefrin, H.M. and Thaler, R.H. (1982) "A Self-Control Based Theory of Personal Saving", Cornell University Working Paper No. 82-12, mimeo.

Shiller, R.J. (1981) "Do Stock Prices Move Too much to be Justified by Subsequent Changes in Dividends?", American Economic Review 71, 421-36.

Shimono, K. and Tachibanaki, T. (1982) "Savings and Life-Cycle: A Cohort Analysis", Kyoto University, mimeo.

Shorrocks, A.F. (1975), "The Age-Wealth Relationship: A Cross-Section and Cohort Analysis", Review of Economics and Statistics 57, 155-163.

Stern, N.H. (1981) "Hours of Work in the UK in the 1970's" SSRC Programme on Taxation, Incentives and the Distribution of Income Discussion Paper No. 15, mimeo, London.

Stern, N.H. (1983) "On the Specification of Labour Supply Functions", SSRC Programme on Taxation, Incentives and the Distribution of Income Discussion Paper, mimeo London.

Summers, L.H. (1982) "Tax Policy, The Rate of Return and Savings" National Bureau of Economic Research Working Paper No. 995.

Thaler, R.H. and Shefrin, H.M. (1981) "An Economic Theory of Self-Control", Journal of Political Economy, 89, 392-406.

Thurow, L. (1969) "The Optimum Lifetime Distribution of Consumption Expenditure", American Economic Review, 59, 324-30.

Tobin, J. (1967) "Life Cycle Saving and Balanced Growth" in Fellner, W. et al. Ten Economic Studies in the Tradition of Irving Fisher, John Wiley and Sons, London.

Tversky, A. and Kahneman, D. (1982) "Judgement under Uncertainty:Heuristics and Biases", Chapter 1 in (eds) Kahneman, D., Slovic, P. and Tversky, A. Judgement under Uncertainty: Heuristics and Biases, Cambridge University Press.

US (1976), Report of the Consultant Panel on Social Security, Government Printing Office, Washington DC.

Wickens, M.R. and Molana, H. (1982) "Stochastic Life Cycle Theory with Varying Interest Rates and Prices", University of Southampton Discussion Paper, No. 8224, mimeo.

Yaari, M.E. (1965) "Uncertain Lifetime, Life Insurance and the Theory of the Consumer", Review of Economic Studies, 32, 137-50.