NBER WORKING PAPER SERIES

LIFE-CYCLE MODELS OF CONSUMPTION: IS THE EVIDENCE CONSISTENT WITH THE THEORY?

Angus Deaton

Working Paper No. 1910

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 1986

Presented at the Fifth World Congress of the Econometric Society, Cambridge, MA, August 16th - 23rd, 1985. Without implicating them, I should like to thank Alok Bhargava, Alan Blinder, Martin Browning, John Campbell, David Card, Rob Engle, Mark Gersovitz, Steve Goldfeld, Jerry Hausman, Fumio Hayashi, Tom McCurdy, John Muellbauer, Whitney Newey, Adrian Pagan, Larry Summers, Mark Watson, and Ken West, for helpful discussions during the preparation of this paper or for providing comments in response to the first draft. The research reported here is part of the NBER's research program in Economic Fluctuations. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper #1910
April 1986

Life-Cycle Models of Consumption: Is the Evidence Consistent with the Theory?

ABSTRACT

The paper considers a variety of evidence that casts light on the validity of the life-cycle model of consumer behavior. In the first part of the paper, simple non-parametric tests are used to examine representative agent models of consumption and labor supply. It seems extremely unlikely that post-war United States evidence can usefully be explained by such a model, at least if the assumption of intertemporal separability is maintained. Changes in aggregate consumption bear little relationship to after tax real interest rates, and consumption has tended to grow even during periods of negative real interest rates. Joint consideration of consumption and labor supply does nothing to resolve the problems that arise when consumption is taken by itself. It is argued that these results cast doubt, not on life-cycle theory itself, but on the representative agent assumption; there is little reason to suppose that changes in aggregate consumption *should* be related to the real interest rate. The second part of the paper is concerned with the time-series representation of disposable income and with its implications for the behavior of consumption under the assumptions of the life-cycle model. If real disposable income is truly a first-order autoregressive process in first differences, a process that fits the data well and is becoming increasing popular in the macro time-series literature, then the life-cycle model implies that changes in consumption should be *more* variable than innovations in income, a prediction that is manifestly false. Various possible resolutions of this problem are reviewed, including habit formation and alternative representations of disposable income. The paper concludes with some evidence on the excess sensitivity question, why it is that consumption responds to *anticipated* changes in income. Monte Carlo evidence supports the suggestion made by Mankiw and Shapiro that the presence of time trends can cause severe problems of inference in models containing variables with unit roots, but the results make it seem unlikely that this is the cause of the widespread excess sensitivity findings.

Angus Deaton,
311 Woodrow Wilson School,
Princeton University,
Princeton, N.J. 08540.

Introduction

In this paper I discuss some empirical evidence that reflects on the validity of life-cycle models of consumer behavior. I make no attempt to provide a survey but rather focus on a number of specific issues that seem to me to be important, or that seem to have been unreasonably neglected in the current literature. The paper has three sections. The first looks at the stylized facts. In particular I look at the non-parametric evidence with emphasis on both consumption and labor supply and the interaction between them. I present some aggregate time-series data from the U.S.; these suggest that simple representative agent models of the life cycle are unlikely to be very helpful, at least without substantial modification. It is particularly hard to come up with one explanation that is consistent both with these data and the wealth of evidence on consumption and labor supply from microeconomic information. However, I argue that the main problem here is not so much the theory as the aggregation; except under extremely implausible assumptions, including the supposition that consumers are immortal, life cycle theory does not predict the sort of aggregate relationships that are implied by representative agent models. In particular, it makes little sense to look for a simple relationship between the real rate of interest and the rate of growth of aggregate consumption. Section 2 is concerned with the estimation of parametric models on aggregate time series data. I review briefly the "excess sensitivity" issue, as well as some of the econometric problems associated with the nonstationarity of the income and consumption time series. My main point, however, is to argue that there are interactions between the time-series representation of income and the life-cycle model that have not been adequately recognized in the literature. In particular, if real disposable income can be adequately represented as a first-order autoregressive process in first differences, a formulation that is becoming increasingly popular in the macro literature, then consumption, far from being excessively sensitive, is not sensitive enough to innovations in current income. Indeed, the representative agent version of the permanent income hypothesis can be rejected because it fails to predict the fact that consumption is smooth, the very fact that it was invented to explain in the first place. I consider a number of possible explanations, and offer a menu of directions for escape: we can abandon the life-cycle theory, we can abandon intertemporal additivity, or we can abandon our time series description of income. There are attractions to both of the last two routes. The section concludes with some summary empirical evidence on the excess sensitivity issue as well as on the relevance of the distinction between anticipated and unanticipated variables. Some time series estimation problems are reviewed, and though they complicate inference, I nevertheless find convincing evidence for excess sensitivity, and more surprisingly for the view that there is little cost to ignoring the distinction between anticipated and unanticipated income. Section 4 is a brief summary of the main conclusions.

Section 1: Some non-parametric evidence.

In this section I first follow the lead of Martin Browning (1984) and discuss non-parametric tests of life-cycle behavior. Browning's work

$$W_0 = \Sigma\ \tilde{p}_t c_t - \tilde{w}_t h_t \tag{2}$$

In these equations, c_t and h_t are consumption and hours, p_t is the price of goods in t, and w_t the wage rate. Superimposed tildes imply that the price or wage is discounted back to period 0, so that W_0 is the initial wealth endowment at the beginning of life. Note that the subperiod utility functions are not indexed on t, and that for the moment, I have not allowed for a rate of time preference; this is more conveniently introduced later.

The empirical evidence consists of a finite set of pairs of observations on consumption and hours together with their associated discounted prices and wages. Browning shows that for these data to be consistent with the theory, it is necessary and sufficient that they satisfy the condition of "cyclical monotonicity", see also Green and Srivastava (1984). This is that for any "cycle" of observation indices, s, t, ..., v, say, it must be the case that

$$\begin{aligned} &\tilde{p}_v c_s + \tilde{p}_s c_t + \ldots\ldots\ldots\ldots + \tilde{p}_u c_v - \tilde{w}_v h_s - \tilde{w}_s h_t - \ldots - \tilde{w}_u h_v \\ &\geq \tilde{p}_s c_s + \tilde{p}_t c_t + \ldots + \tilde{p}_v c_v - \tilde{w}_s h_s - \tilde{w}_t h_t - \ldots - \tilde{w}_v h_v \end{aligned} \tag{3}$$

follows that of Sydney Afriat, for example, (1967), (1981), and Hal Varian, (1982),(1983), in seeking to confront a finite number of data points with the underlying theory without the intermediation of an arbitrarily chosen functional form. As we shall see, the tests are particularly simple for the life-cycle model, so that they provide an attractive way of organizing the evidence before moving on to more conventional analysis. The starting point is the (absurdly) extreme hypothesis that the evolution of aggregate behavior over time is simply the unfolding of the predetermined life-cycle plan of a single representative agent. This prescient individual, having returned from World War II, or let us say Korea, established a clear and subsequently correct view of the rest of the post-war period. This hypothesis is of interest only because, if it cannot be rejected, there is little point in testing weaker forms, like intertemporal choice under uncertainty. And Browning presents evidence from the U.S., the U.K., and Canada that suggests that there are in fact few obvious discrepancies. Let us see.

The theory begins with the assumption that preferences can be represented by the intertemporally additive utility function

$$u = \Sigma\, v(c_t, h_t) \qquad (1)$$

subject to a budget constraint, which under certainty, can be written decrease as the growth of real wages is greater than or less than the real rate of interest. If the latter exceeds the former, hours should be higher now than later in spite of the fact that wages are increasing; in terms of tomorrow's goods, the return on today's work is higher than that on tomorrow's work when the positive real rate of interest is taken into account. The reader can check that the addition of a rate of time preference to the utility function, so that each period's subutility is multiplied by the discount factor $1/(1+\delta)$ to the power of t, leaves these predictions unchanged except that the real rate of interest is simply replaced by the real rate of interest minus the rate of time preference. Note that if the data fail the conditions (5) and (6), i.e. if the schedules appear not to be monotone, then recourse can be had to the weaker condition (4). Provided labor supply behaves as it should, then there is some scope for a violation by consumption, and <u>vice versa</u>. However, note that some sign patterns are clearly inconsistent with <u>both</u> (4) on the one hand, and (5) and (6) on the other. In particular, if the real rate of interest is <u>negative</u> and consumption is <u>increasing</u>, then hours and discounted wages must move in the same <u>direction</u>.

Table 1 presents some data on annual changes in the aggregates for the U.S. from 1954 through 1984. Definitions of the series are shown at the foot of the Table. I have followed standard practice in excluding durable goods from the consumption concept. Note the importance of calculating real interest rates on an after tax basis; not only is there a prolonged period of negative rates through the 1970's, but there is also similar period from 1955 to 1959. The tax adjustment is made by comparing the yields on (tax-free) AAA municipal bonds with those on

Since I shall be attempting only to disprove this condition, I work with one of its implications. For any pair of indices, s and t, cyclical monotonicity implies that

$$(\tilde{p}_t - \tilde{p}_s).(c_t - c_s) \geqq (\tilde{w}_t - \tilde{w}_s).(h_t - h_s) \tag{4}$$

This condition can usefully be contrasted with those that are required if the utility function (1) is simultaneously <u>intra</u>- as well as <u>inter</u>-temporally additive, so that life-time utility is the sum of two sums, one involving hours alone, the other involving only consumption. In this case, cyclical consistency is equivalent to, for all t and s,

$$(\tilde{p}_t - \tilde{p}_s).(c_t - c_s) \leqq 0 \tag{5}$$

$$(\tilde{w}_t - \tilde{w}_s).(h_t - h_s) \geqq 0 \tag{6}$$

i.e. to the (obvious) requirement that the consumption demand function slope down and that the labor supply function slope up. Note that if t and s are successive observations in time, then the discounted price is falling if the real after-tax rate of interest is positive, while the discounted wage is rising if the proportional increase in real wages is greater than the real rate of interest. Hence, under simultaneous additivity, consumption should be increasing over periods when the real rate of interest is positive, while hours worked should increase or simultaneity, Kennan (1985), and so forth, but at this simple level there is clear evidence against the hypothesis.

Consumption and hours can also be considered jointly. Figure 1 is a scatter plot for the thirty-one years of real interest rates against the growth of real wages less the interest rate. Using the series DH1, points are marked as "1" for consumption growth and hours growth, "2" for consumption growth and hours decline, "3" for consumption decline and hours growth (one point, 1974), or "0" for 1963 when there was essentially no change in average hours. (There are no years in which consumption and hours both declined.) If we look at the second quadrant, where there is (discounted) wage growth but negative real interest rates, we know that increases in consumption must be associated with increases in hours for cyclical consistency to be satisfied. But of the 16 years represented in the quadrant in which consumption increased, in 5 of them hours declined. There are further problems in comparing the first and fourth quadrants. In 1983 and 1984, by a miracle of supply-side economics, hours increased (substantially) by both measures while discounted real wages were falling. Since consumption increased, no doubt in response to the unprecedented high real interest rates, such events could be explained by a high degree of complementarity between consumption and hours. However, in quadrant 1, where wage growth and interest rates are both positive, there are several years in which hours declined when consumption rose, presumably because hours and goods are substitutes! Once again, the addition of a positive rate of time preference is of little help. Subtracting $\delta>0$ from the real rate r similar (taxed) corporate bonds, and then applying the correction factor to the treasury-bill rate. This may not be correct, but it is better than making no correction at all.

Since consumption grows in every year except 1974, and since real interest rates are negative from 1955-59, 68-9, and 71-80, it is immediately clear that the simple monotonicity condition (5) is violated, indeed it is only satisfied for 15 out of the 31 years. Allowing for a positive rate of time preference is of little help in explaining why consumption <u>grows</u>; a <u>negative</u> rate of time preference would do much better.

The labor supply side is much more difficult to document, largely because, with different wage rates for different individuals, the aggregation makes even less sense. The series DH1 and DH2 are two of the many possible series for changes in hours; their sign patterns are similar but not identical. The real wage series is average hourly earnings of production workers in manufacturing deflated by the implicit price deflator of consumption. There is little relationship between the sign pattern in this series and those in either of the two series for changes in hours. This is what we would expect. The growth of real wages is close to being a random walk with drift, Ashenfelter and Card (1982),(1985), while hours move with the cycle, so that we can expect as many contradictions as confirmations of the monotonicity condition (6). Clearly, it is possible to take a more sophisticated view of these data, to allow for aggregation bias, see e.g. Barro and King (1984), for this, goods and hours have to be substitutes, which is inconsistent with the latter period of the typical life-cycle where hours, consumption and the wage rate decline together. One of the strengths of working on consumption behavior is the availability of many different types of data on which the same hypotheses can be tested, but one of the weaknesses of the theory is its inability to yield a simple explanation that appears to work for them all.

Many of these problems go away as soon as the fiction of a representative consumer is abandoned. In particular, imagine an economy with a stationary population and no secular real income growth, where each individual is a perfect life-cycle consumer. Suppose also that each consumer has the standard utility function, with identical parameters, and that the real rate of interest is greater than the rate of time preference. In consequence, the life-cycle consumption paths will be identical (except possibly for scale) for all individuals, and each will be characterized by a consumption stream that is rising over time. Furthermore, if all these consumers were to be transported to another economy with a higher real rate of interest, each of their consumption streams would be growing faster than in the original economy. Nevertheless aggregate consumption in this economy is constant, or at least it is so unless its inhabitants live for ever. Old people, whose consumption has been growing steadily over their lives and is thus higher than their permanent income, are continually dying off and being replaced by young people whose consumption is much less than their permanent income. In consequence, consumption in aggregate remains constant over time even though it is growing for each individual. Comparing otherwise simply moves the axes of Figure 1 down along the 45° line to the right; this does nothing for the "bad" points in quadrant 2.

These results seem to me to be entirely consistent with, and indeed almost an explanation for, the results that have been obtained in parametric studies that have modeled hours and consumption simultaneously in a life-cycle context. Mankiw, Rotemberg, and Summers (1985), using quarterly data from the U.S. find violations of concavity in their estimated utility function. With quite different data, a time series of cross sectional household surveys from the U.K., Browning, Deaton, and Irish (1985) also found inconsistencies between theory and evidence, with violations not only of concavity, but also of symmetry, so that goods and hours want to be both substitutes and complements for one another, as appears to be the case in Figure 1. In this evidence, it is clear that the life-cycle model cannot provide a single unified explanation that will cover both life-cycle and business cycle evidence. Data from other sources suggest that the difficulties are quite widespread. Information on individual households typically shows hump shaped life-cycle profiles of consumption, hours, and real wages, see Ghez and Becker (1975), Smith (1977), Thurow (1969), Browning, Deaton, and Irish (1985). This can be trivially explained by time varying preferences, and perhaps more interestingly by variations in household size over the life cycle. Alternatively, and since there is no obvious life-cycle shape to the real interest rate, the hump is consistent with the supposition that hours and goods are complements, as Heckman (1971), (1974) pointed out. However, in long-run time-series data, consumption increases, hours decline, real wages grow, and the real rate is (presumably) positive. To account for rather a typically humped shape, with declines in consumption in old age. Thirdly, casual evidence suggests that while children share much of their parents wealth and standards of living, it is not true that they immediately begin their life-cycle consumption paths at the same level as their parents.

Aggregation does not explain all of the contradictions given above, but it should teach us not to expect aggregate consumption to be growing or declining as the real rate of interest is greater than or less than the rate of time preference. Note that "Euler equation" models of consumption under uncertainty are also guilty of this sin. The Euler equation is the stochastic version of the first-order condition that equates the marginal utility of money across periods taking into account the real interest rate, and it is the equation that determines the rate of growth of consumption in relation to the real interest rate. Even if each consumer conforms to the Euler condition, there is no corresponding aggregate version without some version of the immortality assumption.

The aggregation assumptions are even more severe for labor supply than for consumption. Comparing (5) and (6), it is clear that if all consumers face the same prices and interest rates, and if the population of consumers is indeed immortal and thus identical at all points in time, then if each individual satisfies (5), so will the aggregate. However, different consumers face different wage rates, so that it is very easy to construct examples in which individual wages and hours move together, but the averages move contrarily.

stationary economies with different real interest rates will reveal no relationship between the real rates and the rate of growth of consumption, since the latter is always zero. It is only if there is population growth or growth in real income per capita that aggregate consumption will grow, and the rate of growth will be related to demographic factors on the one hand, and to a very long moving average of incomes (i.e. a time trend) on the other. Only if the economy is continuously on a "golden-age" growth path, with the real interest rate equal to the real rate of growth, and if it is capable of instantaneously jumping from one equilibrium growth path to another, only then will there be a stable relationship between the real rate of interest and the rate of growth of consumption. All of this was, of course, clearly worked out and carefully explained thirty years ago in the original papers on the life-cycle hypothesis by Modigliani and Brumberg (1955), (1979). It is a pity that the recent literature has lost sight of it.

One counter argument is that individuals may not live for ever, but that their descendants do, so that in aggregate we have "as if" immortality. This would require that as one consumer dies, his or her replacement, grandchild or great-grandchild perhaps, picks up the dead progenitor's consumption stream where it was momentarily interrupted, so that the growth path can continue. I find this hard to believe. Firstly, the discounted present value of future real incomes from here to eternity would almost certainly yield an eternal stream of real income far in excess of current income levels. If so, consumption ought to reflect it now, and it does not. Secondly, evidence on individual life-cycle consumption patterns does not show steady growth over the life-cycle, but

Section 2. Aggregate time series, sensitivity, and insensitivity.

2.0 Introduction

Compared with consumption function analysis of a decade or so ago, modern work is characterized by much greater attention to the time-series aspects of econometric analysis. This is partly because economists have become much more aware of the special econometric problems that arise with inference and estimation in time series models, but also because the applications of rational expectations theory to the consumption function has drawn attention to the essential part played by the time series properties of income in determining the form of the consumption function. This point was made in the context of the consumption function in Lucas's original critique (1976), but the basis for much of the subsequent work has been Hall's (1978) demonstration that, conditional on lagged consumption, expected future consumption should be independent of other lagged information. In his original paper, Hall found that such a formulation was surprisingly difficult to reject. Subsequent and more detailed analysis has led to something close to a consensus view that, contrary to the theory, aggregate consumption is responsive to anticipated changes in income, see in particular Flavin (1981), Hayashi (1982), and Hansen and Singleton (1982), though see also Campbell (1985), Bean (1985) and Blinder and Deaton (1985). Both of the last two studies find that as the consumption function is expanded to include a wider range of variables, e.g government expenditure, leisure, and various relative prices, it becomes more difficult to reject the theory. The

Before moving on to the more conventional parametric models, it is worth noting some of the advantages and disadvantages of the non-parametric technique. It will not have escaped the reader that the method shares much with the journalistic approach to economic events; if variables once move in the wrong direction relative to one another, a counter-example has been established, and the theory must be false. We spend a good deal of time, and for good reason, teaching undergraduates that this is not a good way to make inferences. Nevertheless, in the present case, the violations are more endemic than occasional, and an equally simple parametric approach would undoubtedly yield the same results. Additionally, it can be argued that since the parametric approach is <u>more</u> restrictive than the non-parametric, imposing a possibly restrictive functional form, then rejections of the theory can only get worse if parametric tests are carried out. However, there is an important sense in which non-parametric tests are too strong. By their nature, they focus on the relationship between two, possibly vector valued quantities. In these models, quantities respond to prices, and to nothing else. It is therefore straightforward to construct simple, exact models of consumption and labor supply that contain variables in addition to prices, and these simple models, though entirely consistent with the theory, could easily be made to fail the non-parametric tests.

aggregation over agents is unlikely to be the explanation, though I have so far been unable to rule out an explanation based on consumers' jointly modelling consumption, income, and other variables, for example, interest rates and unemployment. A simple model of habit formation is shown to be consistent with the evidence.

(iii) It is also possible to make the life-cycle model true by assumption, and to use its truth to infer backwards to the time series properties of income. I consider a number of possible time series representations for income, and argue that it is unlikely to be possible to discriminate between them on statistical grounds, even though each has very different implications, not only for the life-cycle model, but also for the way in which we think about the economy as a whole. As a consequence, and in contrast to the general promise of rational expectations models, time series analysis of income is unlikely to tell us very much about how consumption should behave, nor to yield testable implications for the consumption function. More concretely, the old problem of how consumption should respond to transitory income is essentially not resolvable since we cannot tell what are the implications of income innovations for the future of income.

(iii) On the excess sensitivity issue, I discuss briefly the issue raised by Mankiw and Shapiro (1984) of whether or not the non-stationarity of the joint consumption income process implies that tests of excess sensitivity are biased against the rational expectations model. I show that the "excess sensitivity" results on U.S. quarterly data cannot be explained away by this phenomenon. Further, there is little evidence "excess sensitivity" finding, which is usually attributed to the presence of a significant fraction of liquidity-constrained consumers, was also confirmed in an excellent study of food expenditures using the PSID by Hall and Mishkin (1982). Even so, there is much good work currently being done on panel data, and the final results are by no means in. For example, and again using the PSID, Altonji and Siow (1985) have found that allowing for measurement error can make a large difference to the results of tests of simple versus rational expectations type models.

Here I wish to discuss the following issues:

(i) I look at two standard time-series models for real disposable income. In the first, income is taken to be stationary around a deterministic trend, and the residuals after trend removal are modeled as a low order ARMA, typically AR(2). In the second, income is made stationary by first differencing, and the differences modeled as a low order ARMA, typically an AR(1) with positive persistence. Both procedures fit the data well, both have their determined adherents, and it is extremely difficult to tell them apart. However, if the second procedure is correct, it turns out that innovations to permanent income are <u>more</u> variable than innovations to current income, so that life-cycle theory predicts that changes in consumption should be larger than innovations in income. It seems fairly clear that such a conclusion is false.

(ii) Since the difference model of income is a plausible and attractive one, and since it is widely used, I consider some explanations for the smoothness of consumption that are consistent with it. I argue that value, it can ultimately be expected to return to the baseline trend. The differencing procedure may or may not imply that the effects of shocks eventually die out, depending on what sort of time series process is fitted to the differences. While the change in income is stationary and has a(n unconditional) mean that remains constant through time, there is in general nothing that acts to bring the level of income back to any particular path. In the simplest example of a random walk with drift, which turns out to be a good description of the growth of the real wage in the U.S., previous changes are immediately consolidated into the baseline, so that a good year with a large wage increase means that we can expect wage levels to be permanently higher with no expectation that at some future date, the good fortune will have to be paid for.

To compare the two approaches, I estimated simple representations using seasonally adjusted quarterly U.S. data from 1954,1 to 1984,4. I use two income concepts, real disposable labor income, y, and total income, including capital income, z. The latter concept is close to the published NIPA figure, and the former was constructed from other published data, see Blinder and Deaton (1985) for further details and for the series itself. It is important to make some attempt to exclude capital income, since it is the labor income concept that is the appropriate one for life-cycle models, and there is no presumption that the two series have the same time-series properties. The detrending procedure is best carried out by estimating the autoregression with time variables included. By the Frisch-Waugh theorem, the results are numerically identical, but the standard errors allow us to assess the relative that the distinction between anticipated and unanticipated income that is required by rational expectations theory has any basis in the empirical evidence.

2.1 The time-series processes of consumption and income

There is no clear consensus on how best to represent the process generating real disposable income. Everyone agrees that real income is not stationary, but that is far from enough to produce a consensus model. One approach is to remove a time trend from income, and then to fit a time-series representation to the residuals. On quarterly U.S. data, an AR(2) seems to do the job. The other approach, more favored by "professional" time-series analysts, see in particular Beveridge and Nelson (1981), is to difference either income or its logarithm, and to fit the time series model to the difference, rather than to the deviation from trend. More recently, Doan, Litterman and Sims (1985) have recommended first differencing the logarithms of virtually all trending macroeconomic time-series as a prelude to fitting vector autoregressive models. That it is rates of growth rather than levels that should be the appropriate object of our attention is an idea that has much to commend it.

At a conceptual level, these two procedures are quite different, and they have different implications for the estimation and interpretation of the relationship between consumption and income. The detrending procedure assumes that trend around which real income moves is fixed and deterministic; even if income is temporarily shocked above or below its trend insufficient to induce stationarity. After deleting insignificant terms, including the time trends, the differenced models are

$$\hat{\Delta y}_t = \underset{(3.3)}{8.4} + \underset{(5.4)}{0.435}\, y_{t-1} \qquad R^2=0.1918,\ Q(33)=37.7,\ ESE= 25.3 \tag{9}$$

$$\hat{\Delta z}_t = \underset{(4.2)}{12.5} + \underset{(4.5)}{0.377}\, z_{t-1} \qquad R^2=0.1428,\ Q(33)=35.7,\ ESE= 27.7 \tag{10}$$

These are simple and parsimonious descriptions of two nonstationary time series; the change in labor income is somewhat more autoregressive and has a less variable innovation than does the change in total income and this is what one would expect. The closeness of (9) and (10) to the corresponding unrestricted forms (7) and (8) suggests that formal tests are unlikely to be able to separate them, and this is in fact the case. Following Dickey and Fuller (1981), the adequacy of the unit-root models (9) and (10) can be tested by calculating an "F-test" for these models against the more general alternatives in which y (or z) is regressed on its first lag, its lagged first-difference, and a time trend. The calculated statistics are 1.97 for labor income and 3.95 for total income compared with the critical values given by Dickey and Fuller of 6.49 at 5% and 5.47 at 10%.

contributions of both elements. For labor income, y, the regression equation is

$$\hat{y}_t = \underset{(2.3)}{84.9} + \underset{(15.2)}{1.39}\, y_{t-1} \underset{(-2.9)}{-0.46}\, y_{t-2} + \underset{(1.1)}{0.16}\, y_{t-3} \underset{(-1.5)}{-0.14}\, y_{t-4} + \underset{(2.1)}{0.63}\, t$$

R^2=0.9983, Q(33) =31.7, ESE =25.1 (7)

while, for total real disposable income, z, the regression is

$$\hat{z}_t = \underset{(2.8)}{139.7} + \underset{(14.7)}{1.30}\, z_{t-1} \underset{(-2.3)}{-0.35}\, z_{t-2} + \underset{(0.5)}{0.08}\, z_{t-3} \underset{(-1.0)}{-0.09}\, z_{t-4} + \underset{(2.8)}{1.35}\, t$$

R^2=0.9987, Q(33) =31.6, ESE =27.1. (8)

Note that these time-series representations are broadly similar, though the absolute values of the coefficients on the first and second lags are larger for labor than for total income, while the standard deviation of the innovation is about 20% larger for the broader concept. The time trends are both significant, though much less important than the autoregressive components, and if the regressions are rerun in logarithmic form, neither time trend is significant at conventional levels. Note that at the parameter values shown, both autoregressions are estimated to be stationary, but in each case by a very small margin. Indeed, the estimated coefficients strongly suggest a differenced model in which y is modelled as an AR(1). This would be the immediate reaction of most time-series analysts; for example, Beveridge and Nelson (1981) do not quote results like (7) and (8) but simply state that detrending is is, the estimated autoregressions suggest that, not only are shocks permanent, but they are positively autoregressive, so that a better than average change in any given quarter can be expected to lead to further good fortune in subsequent quarters. This implies that the permanent income value of an unanticipated change in income is greater than the change itself. Clearly, the choice between the two time series processes considered here, though almost impossible to make on statistical grounds, is of considerable importance in calculating permanent income.

How then can we choose between, on the one hand, the (marginally) stationary representation of the deviations of income from trend, and on the other, the AR(1) representation of the differences? Since the data cannot discriminate between them, the choice must be made on theoretical grounds. For myself, I find it hard to believe that real per capita income is centered round a deterministic trend. As emphasized in a recent paper by Campbell and Mankiw (1986), such a representation tends to assume that income shocks originate from the demand side rather than from supply. If some particularly unpleasant negative "epsilon" was the result, for example, of the destruction of part of the capital stock, then I see no reason for believing that income will eventually get back on the old trend. It is much more credible that it would begin to grow again from the new lower base, which is what the first difference model says it will do. Of course, the AR(1) first-difference model is far from being the only time-series representation of income that allows this sort of shock persistence, but it is certainly one of the most parsimonious and widely used.

Consider then a representative consumer who is calculating his or her permanent income and whose representative income is generated by one of these stochastic processes. For the general case, where income (or deviations of income around a deterministic trend) are generated by an ARMA(p,q) process, the change in permanent income from t-1 to t is given, see Flavin (1981), by

$$\Delta y_t^p = \frac{r\{1 + \Sigma(1 + r)^{-s}\pi_s\}}{\{1 - \Sigma(1 + r)^{-s}\rho_s\}} . u_t \tag{11}$$

where the π's are the MA coefficients, the ρ's the AR coefficients, u_t is the current innovation, and r is the interest rate, here assumed to be fixed to avoid complications that are irrelevant to my current concerns. Equation (11) is valid whether or not the process is non-stationary, see Hansen and Sargent (1981). If we re-estimate the stationary labor income equation above excluding the longest two lags, the autoregressive parameters are 1.41 and -0.45, so that, with a real rate of 1%, the factor multiplying the innovation in the above equation is 0.22, a sizeable but reasonable figure. However, if we move to equation (9), which as we have seen cannot be rejected against the model with a time trend, the corresponding parameters are 1.435 and -0.435, so that the response of consumption to a unit innovation is now 1.8, i.e. eight times as large.

For the non-stationary model (9), such a result is intuitively obvious. If income were a random walk, permanent and measured income would be the same; an income shock is expected to be sustained indefinitely. As it

If the arguments in favor of the differenced model are accepted, then the permanent income model of the representative consumer is in some trouble. Innovations in permanent income should be magnified versions of the innovations in measured income, so that if, to take the simplest case, consumption is equal to permanent income, then the variance of the change in consumption should be larger than the variance of the innovation in the income process. Equation (9) above gives the standard deviation of the income innovation process as $25.27 per capita in 1972 prices. The standard deviation of the change in aggregate consumption of goods and services over the same period, 1954,1 to 1984,4 is $12.08. Even if it is only required that consumption be proportional to permanent income, identical problems arise. The standard deviation of the rate of growth of consumption is still only a half of the standard deviation of the innovation in the AR(1) process describing the first difference of the logarithm of income. Even this understates the difficulties, since no allowance has been made for any other shocks to consumption. The addition of shocks that are independent of innovations to income, "transitory consumption," will further increase the expected variation in consumption, and further deepen the puzzle.

Of course, the representative consumer may be able to predict income changes better than can a simple autoregression. Indeed, if equation (9) is supplemented by other variables, including several lags of consumption, wealth, interest rates, and inflation rates, the multiple correlation coefficient can be doubled from 0.20 to 0.40. But even if the R^2 were to be 0.50, the standard deviation of the innovation would still be $20.2, which is still very much larger than the corresponding figure for changes in consumption. However, this is not the relevant comparison. If the consumer uses a number of other variables to predict income, then those variables must themselves be predicted if they are to be a guide to future values of incomes, so that the innovation in permanent income will ultimately be a function, not only of the income innovation, but also of the innovations in all of these other variables. West (1985) has shown that, in spite of these complexities, expansion of the information set cannot <u>increase</u> the variance of the innovations to permanent income, at least provided that the discount rate is positive and that permanent income is discounted over an infinite horizon. Discounting is the key to this result. If, for example, today's unemployment rate helps predict tomorrow's income, knowledge of the fact can help shift some of today's uncertainty to tomorrow at which date unemployment itself will have to be forecast in order to get the next day's income. But, with a positive discount rate, the postponement of shocks helps reduce the innovation variance in the estimate of permanent income. I have not made any progress in attempts to quantify the size of this reduction, and remain doubtful the that more complex forecasting equation can resolve the paradox. Even so, it remains an important avenue for further research.

2.2 Why is consumption so smooth? Aggregation and separability.

One plausible explanation for the smoothness of consumption is that individuals have a great deal of personal, idiosyncratic information about the likely future course of their labor income, so that even if their income path looked very noisy to an observer, it would contain few

$$u_t = u_t^{sh} + u_t^{eh} \tag{13}$$

$$\varepsilon_t = \varepsilon_t^{sh} + \varepsilon_t^{eh} \tag{14}$$

$$\Delta c_t^h = u_t^{sh} + \varepsilon_t^{sh} \tag{15}$$

If, as seems reasonable, it is assumed that each sub-component of averages to zero over the population, then aggregate average consumption Δc and income Δy are given by

$$\Delta c_t = \bar{u}_t^{s} \tag{16}$$

$$\Delta y_t = u_t \tag{17}$$

so that consumption can be smoother than income if the variance of the average surprise component of the <u>common</u> innovation is less than the variance of the aggregate innovation. If u_t is genuinely unpredictable at the aggregate level, it is hard to see why it should be predictable by individuals; it is by definition an <u>aggregate</u> surprise. The ability of individuals to anticipate a large component of their own incomes is largely irrelevant to the relationship between aggregate (common) shocks in income and consumption, at least in this simple model.

If not aggregation, then what? One simple possibility is to abandon the intertemporal additivity assumption with which I began. I think that surprises for the individual so that consumption would nevertheless be very smooth. The question is whether this explanation can carry through to the aggregate.

Unfortunately, relatively little is known for certain about the relationship between individual and aggregate income fluctuations over time. There is considerable and accumulating evidence that the income changes observed in panel data such as the PSID are heavily contaminated by errors of measurement. For example, by comparing two different wage indicators in the PSID, Altonji (1985) estimates that 72.2 of the observed variance of wages is measurement error, see also Abowd and Card (1985) and Ashenfelter (1985) for other evidence that is consistent with this finding. In consequence, I have no direct evidence to offer, only a balance of argument.

Consider the simplest case where, for each consumer h, labor income as perceived by an outsider follows a random walk. Write this as

$$\Delta y_t^h = u_t + \varepsilon_t^h \tag{12}$$

where u_t is a common component across all agents, and ε_t^h is the idiosyncratic component which averages to zero over the population as a whole. Decompose each of these into their (by individual h) anticipated and unanticipated components, with consumption responding only to the latter taking $\tilde{c}_t = c_t - \alpha c_{t-1}$ as instruments, and then shows how to

rewrite the budget constraint so as to define corresponding prices that reflect not only market prices of the goods, but also the costs or benefits of consumption now in terms of pleasure foregone later. Doing this gives a budget constraint under certainty of

$$\Sigma \ \tilde{p}_t \tilde{c}_t = W_t - \alpha c_{t-1}\{1 - \alpha/(1+r)\}^{-1} \tag{19}$$

$$\tilde{p}_t = (1+r)^{-t} \ \Sigma \ \{\alpha/(1+r)\}^k \tag{20}$$

where W is the usual discounted present value of human and nonhuman wealth. If preferences are such that $\tilde{c}$ is constant over time, then the consumption function has the form

$$c_t = \alpha c_{t-1} + \beta \ [y_t^p - \alpha r c_{t-1}\{1 - \alpha/(1+r)\}^{-1} \ ] \tag{21}$$

where $\beta = (1+r-\alpha)/(1+r)$ and permanent income is as conventionally defined. If r is small enough, $\beta=(1-\alpha)$ and the second term in square brackets is small so that for the differenced version, we have

$$\Delta c_t = \alpha \Delta c_{t-1} + (1-\alpha)\Delta y_t^p \tag{22}$$

this ought to be done with some reluctance, since almost any lag pattern between income and consumption can be modeled by a insertion of suitable lags of consumption and hours into the current period subutility function. In consequence, the hypothesis loses much of its sharpness and predictive power. Nevertheless, if the evidence is against additivity, then alternatives must be explored. If lagged hours affects the enjoyment of current leisure, and if current consumption and leisure are not separable within the period, there is no problem in explaining the "excess sensitivity" finding, that the changes in consumption depends on last period's income, and there is a growing body of work that interprets the evidence in this way, see e.g. Kydland and Prescott (1982), Eichenbaum, Hansen, and Singleton (1984), and Singleton (1984). There is also a long tradition of modeling "habit-formation" or "stock-effects" in the demand analysis literature, see Phlips (1972), Houthakker and Taylor (1970), and especially the elegant work of Spinnewyn (1979a), (1979b). Such models can also explain the smoothness of consumption in relation to income.

Consider the simplest model of non-separability recently discussed and applied to British data by Muellbauer (1985). The utility function is

$$u = (1+\delta)^{-t} v(c_t - \alpha c_{t-1}), \quad \alpha > 0 \tag{18}$$

where α is a measure of habit formation. Spinnewyn suggests taking hence the relationship between innovations to income and changes in permanent income. Note also that the first difference of y is a stationary process and can easily be seen to be a (restricted) ARMA(2,2).

Watson contrasts this unobserved components model with the AR(1) first differences model as a description of the logarithm of real GDP, a quantity that has very similar time series properties to those of real disposable income. Not surprisingly, the AR(2) for the second component looks very similar to the AR(2) implied by the ARIMA(1,1,0) in the differenced model, and the growth term g is the mean of y. Moreover, both models perform about equally well over the sample (1949 through 1984), though the unobservable components model does somewhat better in tests of longer-run forecasting ability. However, because the unobservable components model attributes only a fraction of the current innovation to shifts in the underlying trend, the effects of innovations on present discounted values are much less than in the ARIMA, so that such a model is entirely consistent with the change in permanent income being substantially less than the innovation in income. Watson's analysis therefore provides us with a model that allows permanent supply-side shocks but is nevertheless consistent with consumption being smoother than income.

In a very recent paper, Campbell and Mankiw (1986) have directly addressed the question of the extent to which innovations in GDP are permanent. They use exact maximum likelihood methods to estimate a wide range of ARMA processes for the first differences of the logarithm of GDP, and use the results to calculate the fraction of an innovation that

If we assume that 80% of permanent income eventually gets consumed as non-durable goods and services, then the variance results given above require an estimate of α of 0.78, so that current consumption exerts a very high price in terms of future enjoyment.

2.3 Alternative specifications for the income process.

Each of the alternative income generation processes can be modified in a number of different ways. The deterministic trend model is an example of an unobserved components model, in which the time series is decomposed into two or more components. A model that is very close to this has recently been proposed by Watson (1986). In his model, the trend is not deterministic but is modelled as a random walk with drift, but as in our original model, there is a low order stationary process in addition to the trend. The model can therefore be written in the form

$$y_t = x_t + z_t \tag{23}$$

$$x_t = g + x_{t-1} + v_{1t} \tag{24}$$

$$z_t = \alpha_1 z_{t-1} + \alpha_2 z_{t-2} + v_{2t} \tag{25}$$

where v_1 and v_2 are two mutually orthogonal white noise processes. Note that, if the variance of v_1 is zero, Watson's model is exactly the same as an AR(2) around a deterministic trend, but the addition of v_1 permits the trend itself to shift. In consequence, innovations in y will contain innovations in both v_1, which are permanent, and innovations in v_2, which are not. The ratio of the variances of the two types of innovations determines the fraction of any given innovation that will persist and unanticipated events is surely one that makes a great deal of sense. A simple form of a "surprise" consumption function estimated on the same U.S. quarterly data is shown in the first column of Table 2. Once again, the income variable is the Blinder and Deaton (1975) series for labor income, not the conventional NIPA magnitude. The dependent variable is the change in the logarithm of consumption of non-durable goods and services. Lagged consumption is included in order to capture any possible habit-formation effects, while lagged income is included largely to test for "excess sensitivity". The change in income is decomposed with reference to a supplementary equation (not shown) in which the change is explained by two lags each of income and consumption together with a quadratic time trend. Time is included to model the very long-run moving average of income that is predicted by the aggregation theory of Section 1 above. The econometric procedure is to first run the supplementary regression followed by the main regression with the expectations and surprises replaced by their calculated values from the first stage. The standard errors are calculated by a straightforward application of Newey's (1984) method; for all the models shown here, the results of Pagan (1984) and Bean (1985) apply, so that the procedure is asymptotically fully efficient, and the standard errors are identical to the raw OLS standard errors for the "surprise" terms, and to the two-stage least squares standard errors for the others. All standard errors and results could have been equivalently obtained by applying three stage least squares to the system.

Column (1) of Table 2 shows the basic regression with only income and consumption included. It shows little evidence that the rational is expected to persist indefinitely. Their results, though not inconsistent with the belief that such a question is hard to answer, suggest that innovations <u>are</u> persistent, frequently with a long-run effect that is larger than their immediate impact. If these results survive further testing, we are back with the original paradox. Life-cycle theory predicts that changes in consumption should be larger than innovations in income, and the prediction is false.

At the present, it is probably safest to reserve judgment. It is clear that the nature of real disposable income or of GDP is such that it is extremely difficult to discriminate between a wide range of different time series representations, all of which have similar short-run properties, but which differ radically in the way in which they respond in the long-run to short-run shocks. And not surprisingly, the data are much more informative about short-run than about long-run properties. Unfortunately however, the relationship between consumption and income depends on the long-run properties of the income generation process, the very properties about which we know very little. If further research cannot modify this conclusion, then in the end, the study of the income process will have told us very little about the structure of the consumption function, and the promise that is held out by the rational expectations approach will remain unfulfilled.

2.4 Surprise consumption functions and excess sensitivity

Whatever the final verdict on rational expectations consumption functions, the distinction that they recognize between anticipated and surprises matter is resoundingly rejected, whether or not lagged consumption is included in the regression, though now the coefficients on the two wealth terms are sufficiently different to permit a rejection of the traditional formulation without the decomposition. Note also that if the coefficients on lagged consumption and income are equal and opposite, consumption has a long run unit elasticity with respect to income and the model is of the partial adjustment variety that has recently been much recommended in Britain by Hendry and his collaborators, Davidson, Hendry, Srba and Yeo (1978), Davidson and Hendry (1981), and Hendry (1983). In both models this restriction is rejected at the 0.5% level or less.

More and more detailed results of this kind can be found in Blinder and Deaton (1985), and there are more elaborate models in which it is possible to reject the traditional in favor of the surprise version. However, here I am more concerned with evaluating an objection to this type of testing that has recently been raised by Mankiw and Shapiro (1986). Mankiw and Shapiro note, as in Section 2.1 above, that real income is non-stationary, and they model it by a random walk. They point out that if the random walk formulation is correct, permanent income and measured income are identical, so that if consumption is equal to permanent income, all three series are the same random walk. Hence, if the change in consumption is regressed on lagged income to test for excess sensitivity, the regression is essentially the regression of a random walk on its own first lag so that the excess sensitivity test is a t-test for the unit root in a first-order autoregression. But as can be checked from the tables provided by Dickey and Fuller (1981) for this situation, the t-distribution is a poor guide to the actual distribution expectations model can help explain the data. The coefficient on the unanticipated income term is almost identical with the coefficient on the anticipated change in income, and the lagged value of income also attracts a statistically significant coefficient. The chi-squared test for the joint exclusion of lagged and anticipated income has a p-value that is less than one in 10,000, and even if the test is repeated with lagged consumption excluded from the regression (so as to conform to other tests such as that by Flavin), the "surprises-only" hypothesis can still be rejected at around the one percent level, the "variant" figure in the table. However, the test against what I call the traditional model, one in which there are no surprises, is accepted. The decomposition of the change in income into its components has no significant effect in the regression, and offers no improvement over regressing consumption on income, income lagged, consumption lagged, and time. This "traditional" formulation has graced (or disgraced) standard Keynesian macroeconometric models for decades, and has been widely attacked on rational expectations grounds. And while it may be difficult to defend from a theoretical perspective, it is entirely consistent with the data.

Column two show the results of including wealth terms, again decomposed into anticipated and unanticipated terms. As before, prediction equations are fitted for these variables, and the VAR's are extended to include two lags of wealth. The results for the wealth variable are much more satisfactory, with anticipated changes in wealth having no significant effect, while wealth surprises have a highly significant positive effect on consumption. However, the hypothesis that <u>only</u>

formulation, consumption responds one for one to the innovation in income and has no other source of variation. The model that was estimated was

$$\Delta \log c_t = b_0 + b_1 \log y_{t-1} + b_2 \mathrm{time}_t + \mathrm{error}_t \qquad (31)$$

In a hundred replications, errors were drawn from independent normal distributions, and the income and consumption series calculated afresh each time, with their actual value in 1953 serving as a base for the calculations. The technical problem here is, as in Mankiw and Shapiro, the non-stationarity of the income series in the regression, so that standard results cannot be called on to derive the distribution of the calculated "t-statistics." We would expect the difficulties to be most severe when the consumption series is most closely linked to the income series, i.e. in the baseline case where $\tau=1$ and there is no independent innovation in consumption.

Table 3 presents the results. The baseline case, though for a different model than that used by Mankiw and Shapiro, and one that is more like the actual data, essentially reproduces their results. The mean value of the absolute value of the t-statistics is 2.25 with a standard error of 0.075, and the true null hypothesis that lagged income is absent from the regression is rejected 62% of the time. The rest of the experiments differ from the first in that by setting $\sigma_1 = .003842$, an independent source of variation is introduced into the change of consumption. In the case where $\tau=0.3$, this setting replicates the observed variance of the change in log consumption. This makes little difference to the results of the so-called t-statistic in this situation, even asymptotically. Mankiw and Shapiro generate artificial series for income and consumption on the assumption that consumption equals permanent income and that the random walk model is true. They then find in a Monte Carlo experiment in which the change in consumption is regressed on lagged income and a time trend, that the median value of the t-statistic on the latter is -2.2 with a consequent "excess sensitivity" finding 61% of the time, even though the rational expectations permanent income model is true.

Mankiw and Shapiro's procedures are somewhat different from those in Table 2, and the time series analysis presented above does not suggest that income is a pure random walk. However, the parallels are close enough to be worrying, and it seemed worth discovering how sensitive the Mankiw and Shapiro results would be to variations of their assumptions in the directions of the models estimated here. I work with the following artificial model:

$$\Delta \log c_t = \alpha + \tau u_{2t} + u_{1t} \qquad E(u_{1t}) = 0, \quad Var(u_{1t}) = \sigma_1^2 \tag{26}$$

$$\Delta \log y_t = \mu + \rho \Delta \log y_{t-1} + u_{2t} \qquad E(u_{2t}) = 0, \quad Var(u_{2t}) = \sigma_2^2 \tag{27}$$

Equation (27) replicates reality, and in all the experiments reported below I use the estimated parameter values and variances to generate the data. For equation (26), I always take α to be the sample mean of $\Delta \log c$ over 1954,1 to 1984,4. My base case, corresponding as closely as possible to that of Mankiw and Shapiro, is where $\tau=1$ and $\sigma_1=0$. In this coefficients, particularly those on income and wealth, are more or less unchanged, and the tests for the surprises only hypothesis and for the traditional formulation give the same results as before. The test for the long-run unit elasticity are, however, extremely sensitive to the inclusion of the time trend. Of course, these results do not tell us whether or not the time trend ought to be in the regression, but that given the non-stationarity of income and consumption, it is very difficult to tell. However, these time-series problems, although real and potentially misleading, do not appear to be the source of the finding of excess sensitivity.

Section 3. Summary and conclusions.

The argument in this paper may be briefly summarized as follows:

(i) Non-parametric tests using aggregate time series data suggest that the representative consumer life-cycle model is likely to have difficulty in explaining the evidence on the joint movement of the real after tax interest rate and changes in consumption. Taking consumption and labor supply jointly helps very little. This is hardly surprising since the labor supply model does even worse at explaining the evidence on hours and wages than does the consumption model at explaining consumption and interest rates. These tests are very naive and hardly constitute a convincing rejection of the life-cycle model. But the question arises as to how much intellectual effort is worth expending in the attempt to rescue a model that seems to be so blatantly at odds with the raw, untreated data.

if remains as high as unity. However, as is reduced, so that the change in consumption varies less with the shock to income, the t-values become smaller, though even in the (realistic) case where $\tau=0.3$, the mean is still 1.46 and the null would be rejected in 26% of the cases. If the t-values were normally distributed, the expected values would be $2/\sqrt{(2\pi)}=0.80$, and the rejection frequency would be 5%. Experiments 1-3b reproduce the three previous ones but with the time trend excluded. There is a remarkable improvement with the t-values much closer to their theoretical values and rejection frequencies of 8%, 7%, and 5% respectively. Clearly, it is the presence of the time trend that causes the greatest difficulty and without it there are no great difficulties in making inferences in the usual way. The final experiment sets $\tau=2.12$ which is the theoretically correct value if the rational expectations model is true given the process determining income. These results are essentially the same as when $\tau=1$; things do not get any worse. (Note that these results do <u>not</u> suggest that a model of this kind could have generated the <u>actual</u> data. If it had, we would indeed be likely to make incorrect inferences, but the data would look very different, and in particular changes in consumption would have a very much larger variance than changes in income.)

The results of these experiments suggest that the results discussed above should be recalculated without the time trends. The new results are reported on the right hand side of Table 2, and the differences are not as large as suggested by the Monte Carlo results. Here, the main interaction is between the time trend and the lagged consumption term, which does not figure in the excess sensitivity tests. Other contain non-stationary variables. However, these difficulties do not yield an adequate explanation of why consumption appears to respond to information that is not previously unanticipated.

List of works cited:

Abowd, John M.,and David Card (1985),"The covariance structure of earnings and hours changes in three panel data sets," Princeton University, processed.

Afriat, Sydney N. (1981),"On the constructability of consistent price indices between several periods simultaneously", in Deaton, Angus S.(ed), Essays in the theory and measurement f consumer behaviour in honour of Sir Richard Stone, Cambridge. Cambridge University Press,pp 133-61

Afriat, Sydney N. (1967),"The construction of a utility function from expenditure data", International Economic Review, 8, 67-77.

Altonji, Joseph (1986),"Intertemporal substitution in labor supply:evidence from micro data", Journal of Political Economy (forthcoming)

Altonji, Joseph ,and Aloysius Siow (1985),"Testing the response of consumption to income changes with (noisy) panel data", Industrial Relations Section Working Paper No.186, Princeton University, processed

Ashenfelter, Orley (1984),"Macroeconomic analyses and microeconomic analyses of labor supply", presented to Carnegie-Rochester Conference, Bal Harbor, Florida, November 1983.

Ashenfelter, Orley,and David Card (1985),"Why have unemployment rates in Canada and the U.S. diverged?" Industrial Relations Section, Princeton University, processed.

Barro, Robert J., and Robert G.King (1984),"Time-separable preferences and intertemporal-substitution models of business cycles", Quarterly Journal of Economics, 99, 817-39

Bean, Charles R. (1986),"The estimation of surprise models and the surprise consumption function", Review of Economic Studies, forthcoming.

(ii) Unless consumers are immortal, or unless their children simply pick up the eternal dynastic consumption track where their parents left off, the life-cycle model does not predict that aggregate data should be explicable by a representative consumer model. In particular, the rate of growth of average consumption should not be directly related to the real rate of interest. This casts doubt on the rationale for "Euler equation" type consumption models which rest heavily on the relationship between these two variables.

(iii) The evolution of real disposable income over time, whether including or excluding capital income, is well described by an AR(1) in the first difference with a positive autoregressive term. Given this, changes in aggregate consumption ought to have a variance that is greater than the variance of the innovation in the income process. This prediction is falsified by the evidence, and it seems unlikely that inappropriate aggregation can explain this conflict with life cycle theory. The finding can be explained by habit formation, and/or by the adoption of different time-series processes for income, though the question of whether these last are appropriate is very far from being settled.

(iv) Surprise consumption functions estimated on quarterly U.S. aggregate time-series show clear evidence of excess sensitivity of consumption to predictable events. This excess sensitivity is consistent with a simple minded traditional view in which the distinction between anticipated and unanticipated income terms is ignored. There are very real problems associated with inference in models such as these which

Green, R.C.,and S. Srivastava (1984),"Expected utility maximization and observable demand behavior", Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, processed.

Hall, Robert E. (1978),"Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence, Journal of Political Economy, 86, 971-87

Hall, Robert E.,and Frederic S.Mishkin (1982),"The sensitivity of consumption to transitory income: estimates from panel data on households", Econometrica, 50, 461-81

Hansen, Lars Peter and Kenneth J.Singleton (1982),"Generalized instrumental variables estimation of non-linear rational expectations models", Econometrica, 50, 1269-86

Hayashi, Fumio (1982),"The permanent income hypothesis: estimation and testing by instrumental variables", Journal of Political Economy, 90, 895-916

Heckman, James J. (1971),"Three essays on the supply of labor and the demand for goods", unpublished Ph.D. thesis, Princeton University.

Heckman, James J. (1974),"Life cycle consumption and labor supply: an explanation of the relationship between consumption and income over the life cycle", American Economic Review, 64 188-99.

Hendry, David F. (1983),"Econometric modelling: the consumption function in retrospect", Scottish Journal of Political Economy, 30, 193-220.

Houthakker, Hendrik S.,and Lance D. Taylor (1970), Consumer demand in the United States: Analysis and projections, Cambridge. Harvard University Press.

Kennan, John (1985),"An econometric analysis of aggregate fluctuations in labor supply and demand", Department of Economics, University of Iowa, processed.

Kydland, Finn E.,and Edward C.Prescott (1982),"Time to build and aggregate fluctuations", Econometrica, 50, 1345-70.

Lucas, Robert E. (1976), "Econometric policy evaluation: a critique", in Brunner, Karl & Alan Meltzer, eds, The Phillips curve and labor markets, Carnegie-Rochester Conference Series on Public Policy, 1, Amsterdam, North-Holland. pp 19-46

Mankiw, N.Greg, Julio J.Rotemberg, & Lawrence H.Summers (1985),"Intertemporal substitution in macroeconomics", Quarterly Journal of Economics,

Mankiw, N.Greg,and Matthew Shapiro (1986),"Trends, random walks, and tests of the permanent income hypothesis", Economics Letters, forthcoming.

Beveridge, S.,and Charles R.Nelson (1981),"A new approach to decomposition of economic time series into permanent and transitory components", Journal of Monetary Economics, 7,151-74

Blinder, Alan S.,and Angus Deaton (1985),"The time-series consumption function revisited", Brookings Papers on Economic Activity,2, pp 465-521.

Browning, Martin J.(1984) "A non-parametric test of the life-cycle rational expectations hypothesis", Department of Economics, McMaster University, processed

Browning, Martin J., Deaton, Angus S.,and Margaret J.Irish (1985),"A profitable approach to labor supply and commodity demands over the life cycle", Econometrica, 53, 503-43.

Campbell, John Y.,(1985),"Does saving anticipate declining labor income? An alternative test of the permanent income hypothesis," Princeton University, Department of Economics, processed.

Campbell, John Y.,and N.Greg Mankiw,(1986),"Are output fluctuations transitory?" Princeton University, Department of Economics, processed.

Davidson, James E.H.,and David F.Hendry (1981),"Interpreting econometric evidence: consumers expenditure in the U.K. European Economic Review, 16, 177-92

Davidson, James E.H.,et al. (1978),"Econometric modelling of the aggregate time-series relationship between consumers expenditure and income in the United Kingdom", Economic Journal, 88, 661-92

Dickey, David A.,and Wayne A.Fuller (1981),"Likelihood ratio statistics for autoregressive time series with a unit root", Econometrica, 49, 1057-72

Doan, Thomas, Robert Litterman,and Christopher A. Sims, (1985), "Forecasting and conditional projection using realistic prior distributions," Econometric Reviews, 3, 1-100

Eichenbaum, Martin S., Lars Peter Hansen, & Kenneth J. Singleton (1984),"A time series analysis of representative agent models of consumption and leisure choice under uncertainty", Graduate School of Industrial Administration, Carnegie-Mellon University,Pittsburgh, processed.

Flavin, Marjorie (1981),"The adjustment of consumption to changing expectations about future income", Journal of Political Economy, 89, 974-1009

Ghez, Gilbert, and Gary S.Becker (1975), The allocation of time and goods over the life-cycle, New York. Columbia University Press.

TABLE 1
Changes in consumption, hours, wages, and real after tax interest rates

Date	DC	DLNC	RR	DH1	DH2	DLNW-RR
1954	7.57	0.41	0.22	-5.1	-2.1	1.0
1955	53.0	2.85	-0.19	1.8	2.6	3.4
1956	43.9	2.30	-0.38	0.6	-0.7	3.7
1957	18.6	0.96	-0.97	-2.5	-1.5	2.4
1958	9.82	0.50	-0.12	-4.4	-1.6	0.9
1959	53.5	2.69	-0.06	1.2	2.8	2.3
1960	18.2	0.90	0.80	-0.5	-1.4	0.5
1961	27.8	1.36	0.88	-2.2	0.3	0.5
1962	43.6	2.10	0.34	0.2	1.4	1.2
1963	38.5	1.82	0.45	0.0	0.3	0.8
1964	74.2	3.41	1.03	0.5	0.5	0.6
1965	78.4	3.48	0.80	1.9	1.2	0.7
1966	77.2	3.31	0.16	1.0	0.4	0.9
1967	58.5	2.44	0.85	0.6	-1.9	0.7
1968	83.3	3.37	-0.80	0.7	0.4	3.1
1969	72.6	2.85	-0.41	0.9	-0.2	1.9
1970	60.6	2.32	0.80	-1.2	-2.0	-0.2
1971	35.2	1.32	-0.29	-0.8	0.1	2.0
1972	88.0	3.23	-0.73	2.5	1.7	4.1
1973	56.0	2.00	-2.09	2.6	0.3	3.2
1974	-3.51	-0.12	-4.57	0.3	-1.6	2.9
1975	40.5	1.42	-2.46	-3.0	-1.5	3.8
1976	100.0	3.44	-1.14	2.6	1.7	3.8
1977	92.3	3.07	-2.46	2.8	0.5	5.2
1978	78.0	2.52	-2.99	3.9	0.2	4.6
1979	61.9	1.96	-3.50	1.6	-0.6	3.1
1980	23.3	0.73	-3.15	-1.8	-1.3	1.8
1981	15.8	0.49	0.04	-0.3	0.3	1.0
1982	20.9	0.65	3.82	-3.0	-2.2	-3.5
1983	78.5	2.38	3.27	1.2	3.0	-3.1
1984	79.4	2.36	3.31	4.4	1.4	-2.5

Notes:- DC and DLNC are the backward first differences in the levels and logarithms, the latter x 100, of consumption of nondurable goods and services, excluding clothing and footwear. RR is the real after tax rate of interest from t-1 to t. It is computed from the treasury bill rate by applying the implicit rate of tax revealed by comparing AAA municipal with AAA corporate bonds. The change in the logarithm of the price deflator of full NIPA consumption is subtracted on a quarterly basis, and converted to annual rates of return. DH1 is the first difference of average hours per employee from the Citibase tape. DH2 is the first difference of average weekly hours of production workers in manufacturing. DLNW is 100 times the annual change in the real value of average hourly earnings, not corrected for overtime, of production workers in manufacturing.

Modigliani, Franco, and Richard Brumberg (1954),"Utility analysis and the consumption function: an interpretation of cross-section data", in Kurihara, Kenneth K., ed., Post-Keynesian Economics, New Brunswick, Rutgers University Press.

Modigliani, Franco, and Richard Brumberg (1979),"Utility analysis and aggregate consumption functions: an attempt at integration", in Abel, Andrew, ed., The Collected Papers of Franco Modigliani, Vol 2, Cambridge, Mass. MIT Press. 128-98

Muellbauer, John (1985),"Habits, rationality and the life-cycle consumption function", Nuffield College, Oxford, processed.

Newey, Whitney K. (1984),"A method of moments interpretation of sequential estimators", Economics Letters, 14, 201-6

Pagan, Adrian R. (1984),"Econometric issues in the analysis of regressions with generated regressors", International Economic Review, 25, 221-47.

Phlips, Louis (1972),"A dynamic version of the linear expenditure model" Review of Economics and Statistics, 64, 450-8

Singleton, Kenneth J. (1984),"Testing specifications of economic agents' intertemporal optimum problems against non-nested alter natives", Graduate School of Industrial Administration, Carnegie Mellon University, processed.

Smith, James P. (1977),"Family labor supply over the life cycle", Explorations in Economic Research, 4, 205-76

Spinnewyn, Franz (1979a),"Rational habit formation", European Economic Review, 15, 91-109

Spinnewyn, Franz (1979b), "The cost of consumption and wealth in a model with habit formation", Economics Letters, 2, 145-8

Thurow, Lester (1969),"The optimum lifetime distribution of consumption expenditures", American Economic Review, 59, 324-30

Varian, Hal R. (1983),"Non-parametric tests of consumer behavior", Review of Economic Studies, 50, 99-110

Varian, Hal R. (1982),"The non-parametric approach to demand analysis", Econometrica, 50, 945-73

Watson, Mark W. (1986), "Univariate detrending method with stochastic trends," Journal of Monetary Economics, forthcoming.

West, Kenneth D. (1984), "Speculative bubbles and stock price volatility," Department of Economics, Princeton University, Financial Research Center Memorandum No. 54, processed.

TABLE 3

Monte Carlo experiments on surprise consumption functions

	τ	σ_1	abs(t)	rej	s.d.	trend
baseline	1.00	0	2.25	0.62	0.075	yes
experiment 1a	1.00	.003842	2.10	0.60	0.082	yes
experiment 2a	0.65	.003842	1.92	0.50	0.085	yes
experiment 3a	0.30	.003842	1.46	0.26	0.087	yes
experiment 1b	1.00	.003843	.949	0.08	0.067	no
experiment 2b	0.65	.003843	.962	0.07	0.075	no
experiment 3b	0.30	.003843	.874	0.05	0.666	no
PIH true	2.12	.003843	2.28	0.64	0.079	yes

Table 2

Regression results for surprise consumption functions

constant	0.743 (3.2)	0.660 (2.7)	0.006 (0.4)	- 0.047 (1.3)
log c(-1)	-0.178 (3.5)	-0.190 (3.6)	-0.018 (1.4)	-0.057 (2.1)
log y(-1)	-0.217 (2.7)	-0.190 (2.3)	-0.273 (3.3)	-0.223 (2.7)
Elogy	0.294 (3.6)	0.273 (3.5)	0.295 (3.3)	0.258 (2.7)
logy-Elogy	0.277 (6.0)	0.228 (4.7)	0.277 (5.7)	0.228 (4.6)
Elogw	-	0.013 (1.0)	-	0.022 (1.6)
logw-Elogw	-	0.119 (3.1)	-	0.119 (3.0)
time	.0006 (3.2)	.0005 (3.3)	-	-
all surprises	0.0000	0.0000	0.0018	0.0007
variant	0.0125	0.0159	0.0035	0.0058
traditional	0.8619	0.0316	0.8943	0.0635
unit elastic	0.0012	0.0034	0.6702	0.8182
e.s.e.	.00387	.00373	.00402	.00386
R^2	0.3380	0.3951	0.2799	0.3470

Figure 1

Changes in real wqges less after tax real rate versus real rate