Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence

Robert E. Hall

Center for Advanced Study in the Behavioral Sciences and National Bureau of Economic Research

Optimization of the part of consumers is shown to imply that the marginal utility of consumption evolves according to a random walk with trend. To a reasonable approximation, consumption itself should evolve in the same way. In particular, no variable apart from current consumption should be of any value in predicting future consumption. This implication is tested with time-series data for the postwar United States. It is confirmed for real disposable income, which has no predictive power for consumption, but rejected for an index of stock prices. The paper concludes that the evidence supports a modified version of the life cycle–permanent income hypothesis.

As a matter of theory, the life cycle–permanent income hypothesis is widely accepted as the proper application of the theory of the consumer to the problem of dividing consumption between the present and the future. According to the hypothesis, consumers form estimates of their ability to consume in the long run and then set current consumption to the appropriate fraction of that estimate. The estimate may be stated in the form of wealth, following Modigliani, in which case the fraction is the annuity value of wealth, or as permanent income, following Friedman, in which case the fraction should be very close to one. The major problem in empirical research based on the hypothesis has arisen in fitting the part of the model that relates current and past observed income to expected future income. The relationship almost always takes the form of a fixed distributed lag, though this practice has been very effectively criticized by Robert Lucas (1976). Further, the estimated distributed lag is usually puzzlingly short. Equations purporting to embody the life cycle–permanent income principle

are actually little different from the simple Keynesian consumption func-
tion where consumption is determined by contemporaneous income alone.

   Much empirical research is seriously weakened by failing to take proper
account of the endogeneity of income when it is the major independent
variable in the consumption function. Classic papers by Haavelmo (1943)
and Friedman and Becker (1957) showed clearly how the practice of
treating income as exogenous in a consumption function severely distorts
the estimated function. Even so, regressions with consumption as the depen-
dent variable continue to be estimated and interpreted within the life cycle–
permanent income framework.[1]

   Though in principle simultaneous-equations econometric techniques can
be used to estimate the structural consumption function when its major
right-hand variable is endogenous, these techniques rest on the hypothesis
that certain observed variables, used as instruments, are truly exogenous
yet have an important influence on income. The two requirements are often
contradictory, and estimation is based on an uneasy compromise where the
exogeneity of the instruments is uncertain. Furthermore, the hypothesis of
exogeneity is untestable.

   This paper takes an alternative econometric approach to the study of
the life cycle–permanent income hypothesis by asking exactly what can be
learned from a consumption regression where it is conceded from the outset
that none of the right-hand variables is exogenous. This proceeds from a
theoretical examination of the stochastic implications of the theory. When
consumers maximize expected future utility, it is shown that the conditional
expectation of future marginal utility is a function of today's level of con-
sumption alone—all other information is irrelevant. In other words, apart
from a trend, marginal utility obeys a random walk. If marginal utility is a
linear function of consumption, then the implied stochastic properties of
consumption are also those of a random walk, again apart from a trend.
Regression techniques can always reveal the conditional expectation of
consumption or marginal utility given past consumption and any other past
variables. The strong stochastic implication of the life cycle–permanent
income hypothesis is that only consumption lagged one period should have a
nonzero coefficient in such a regression. This implication can be tested
rigorously without any assumptions about exogeneity.

   Testing of the theoretical implication proceeds as follows: The simplest
implication of the hypothesis is that consumption lagged more than one
period has no predictive power for current consumption. A more stringent
testable implication of the random-walk hypothesis holds that consumption
is unrelated to *any* economic variable that is observed in earlier periods. In
particular, lagged income should have no explanatory power with respect
to consumption. Previous research on consumption has suggested that

---

[1] Examples are Darby 1972 and Blinder 1977.

lagged income might be a good predictor of current consumption, but this hypothesis is inconsistent with the intelligent, forward-looking behavior of consumers that forms the basis of the permanent-income theory. If the previous value of consumption incorporated all information about the well-being of consumers at that time, then lagged values of actual income should have no additional explanatory value once lagged consumption is included. The data support this view—lagged income has a slightly negative co-efficient in an equation with consumption as the dependent variable and lagged consumption as an independent variable. Of course, contempo-raneous income has high explanatory value, but this does not contradict the principal stochastic implication of the life cycle–permanent income hy-pothesis.

As a final test of the random-walk hypothesis, the predictive power of lagged values of corporate stock prices is tested. Changes in stock prices lagged by a single quarter are found to have a measurable value in predict-ing changes in consumption, which in a formal sense refutes the simple random-walk hypothesis. However, the finding is consistent with a modifi-cation of the hypothesis that recognizes a brief lag between changes in permanent income and the corresponding changes in consumption. The discovery that consumption moves in a way similar to stock prices actually supports this modification of the random-walk hypothesis since stock prices are well known to obey a random walk themselves.

The paper concludes with a discussion of the implications of the pure life cycle–permanent income hypothesis for macroeconomic forecasting and policy analysis. If every deviation of consumption from its trend is unexpected and permanent, then the best forecast of future consumption is just today's level adjusted for trend. Forecasts of future changes in income are irrelevant, since the information used in preparing them is already incorporated in today's consumption. In a forecasting model, consumption should be treated as an exogenous variable. For policy analysis, the pure life cycle–permanent income hypothesis supports the modern view that only unexpected changes in policy affect consumption—everything known about future changes in policy is already incorporated in present consump-tion. Further, unexpected changes in policy affect consumption only to the extent that they affect permanent income, and then their effects are expected to be permanent. Policies that have a transitory effect on income are in-capable of having a transitory effect on consumption. However, none of the findings of the paper implies that policies affecting income have no effect on consumption. For example, a permanent tax reduction generates an imme-diate increase in permanent income and thus an immediate increase in consumption. But the evidence that policies act only through permanent income certainly complicates the problem of formulating countercyclical policies that act through consumption.

I. Theory

Consider the conventional model of life-cycle consumption under uncertainty: maximize $E_t \sum_{\tau=0}^{T-t} (1 + \delta)^{-\tau} u(c_{t+\tau})$ subject to $\sum_{\tau=0}^{T-t} (1 + r)^{-\tau}(c_{t+\tau} - w_{t+\tau}) = A_t$. The notation used throughout the paper is:

E_t = mathematical expectation conditional on all information available in t;

δ = rate of subjective time preference;

r = real rate of interest $(r \geqq \delta)$, assumed constant over time;

T = length of economic life;

$u()$ = one-period utility function, strictly concave;

c_t = consumption;

w_t = earnings;

A_t = assets apart from human capital.

Earnings, w_t, are stochastic and are the only source of uncertainty. In each period, t, the consumer chooses consumption, c_t, to maximize expected lifetime utility in the light of all information available then. The consumer knows the value of w_t when choosing c_t. No specific assumptions are made about the stochastic properties of w_t except that the conditional expectation of future earnings given today's information, $E_t w_{t+\tau}$, exists. In particular, successive w_t's are not assumed to be independent, nor is w_t required to be stationary in any sense.[2]

The principal theoretical result, proved in the Appendix, is the following:

Theorem.—Suppose the consumer maximizes expected utility as stated above. Then $E_t u'(c_{t+1}) = [(1 + \delta)/(1 + r)]u'(c_t)$.

The implications of this result are presented in a series of corollaries.

Corollary 1.—No information available in period t apart from the level of consumption, c_t, helps predict future consumption, c_{t+1}, in the sense of affecting the expected value of marginal utility. In particular, income or wealth in periods t or earlier are irrelevant, once c_t is known.

Corollary 2.—Marginal utility obeys the regression relation, $u'(c_{t+1}) = \gamma u'(c_t) + \varepsilon_{t+1}$, where $\gamma = (1 + \delta)/(1 + r)$ and ε_{t+1} is a true regression disturbance; that is, $E_t \varepsilon_{t+1} = 0$.

Corollary 3.—If the utility function is quadratic, $u(c_t) = -\frac{1}{2}(\bar{c} - c_t)^2$ (where $\bar{c}$ is the bliss level of consumption), then consumption obeys the exact regression, $c_{t+1} = \beta_0 + \gamma c_t - \varepsilon_{t+1}$, with $\beta_0 = \bar{c}(r - \delta)/(1 + r)$. Again, no variable observed in period t or earlier will have a nonzero coefficient if added to this regression.

Corollary 4.—If the utility function has the constant elasticity of substitution form, $u(c_t) = c_t^{(\sigma-1)/\sigma}$, then the following statistical model describes the evolution of consumption: $c_{t+1}^{-1/\sigma} = \gamma c_t^{-1/\sigma} + \varepsilon_{t+1}$.

[2] An illuminating analysis of the behavior of consumption when income is stationary appears in Yaari (1976). Further aspects are discussed by Bewley (1976).

Corollary 5.—Suppose that the change in marginal utility from one period to the next is small, both because the interest rate is close to the rate of time preference and because the stochastic change is small. Then consumption itself obeys a random walk, apart from trend.[3] Specifically, $c_{t+1} = \lambda_t c_t + \varepsilon_{t+1}/u''(c_t) +$ higher-order terms where λ_t is $[(1 + \delta)/(1 + r)]$ raised to the power of the reciprocal of the elasticity of marginal utility

$$\lambda_t = \left(\frac{1 + \delta}{1 + r}\right)^{u'(c_t)/c_t u''(c_t)}$$

The rate of growth, λ_t, exceeds one because u'' is negative. It may change over time if the elasticity of marginal utility depends on the level of consumption. However, it seems likely that constancy of λ_t will be a good approximation, at least over a decade or two. Further, the factor $1/u''(c_t)$ in the disturbance is of little concern in regression work—it might introduce a mild heteroscedasticity, but it would not bias the results of ordinary least squares. From this point on, ε_t will be redefined to incorporate $1/u''(c_t)$ where appropriate.

This line of reasoning reaches the conclusion that the simple relationship $c_t = \lambda c_{t-1} + \varepsilon_t$ where ε_t is unpredictable at time $t - 1$, is a close approximation to the stochastic behavior of consumption under the life cycle–permanent income hypothesis. The disturbance, ε_t, summarizes the impact of all new information that becomes available in period t about the consumer's lifetime well-being. Its relation to other economic variables can be seen in the following way. First, assets, A_t, evolve according to $A_t = (1 + r)(A_{t-1} - c_{t-1} + w_{t-1})$. Second, let H_t be human capital, defined as current earnings plus the expected present value of future earnings: $H_t = \Sigma_{\tau=0}^{T-t} (1 + r)^{-\tau} E_t w_{t+\tau}$ where $E_t w_t = w_t$. Then H_t evolves according to $H_t = (1 + r)(H_{t-1} - w_{t-1}) + \Sigma_{\tau=0}^{T-t} (1 + r)^{-\tau}(E_t w_{t+\tau} - E_{t-1} w_{t+\tau})$. Let η_t be the second term, that is, the present value of the set of changes in expectations of future earnings that occur between $t - 1$ and t. Then by construction, $E_{t-1} \eta_t = 0$. Still, the first term in the expression for H_t may introduce a complicated intertemporal dependence into its stochastic behavior; only under very special circumstances will it be a random walk. The implied stochastic equation for total wealth is $A_t + H_t = (1 + r)(A_{t-1} + H_{t-1} - c_{t-1}) + \eta_t$. The evolution of total wealth then depends on the relationship between the new information about wealth, η_t, and the induced change in consumption as measured by ε_t. Under certainty equivalence, justified either by quadratic utility or by the small size of ε_t, the relationship is simple: $\varepsilon_t = [1 + \lambda/(1 + r) + \cdots + \lambda^{T-t}/(1 + r)^{T-t}]\eta_t = \alpha_t \eta_t$. This is the modified annuity value of the increment in wealth. The

[3] Granger and Newbold (1976) present much stronger results for a similar problem but assume a normal distribution for the disturbance.

modification takes account of the consumer's plans to make consumption grow at proportional rate λ over the rest of his life. Then the stochastic equation for total wealth is $A_t + H_t = (1 + r)(1 - \alpha_{t-1})(A_{t-1} + H_{t-1}) + \eta_t$, which is a random walk with trend.

Consumers, then, process all available information each period about current and future earnings. They convert data on earnings, which may have large, predictable movements over time, into human capital, which evolves according to a combination of a highly predictable element associated with the realization of current earnings and an unpredictable element associated with changing expectations about future earnings. Taking account as well of financial assets accumulated from past earnings, consumers determine an appropriate current level of consumption. As shown at the beginning of this section, this implies that marginal utility evolves as a random walk with trend. As a result of consumers' optimization, wealth also evolves as a random walk with trend. Although it is tempting to summarize the theory by saying that consumption is proportional to wealth, wealth is a random walk, and so consumption is a random walk, this is not accurate. Rather, the underlying behavior of consumers makes both consumption and wealth evolve as random walks.

All of the theoretical results presented in this section rest on the assumption that consumers face a known, constant, real interest rate. If the real interest rate varies over time in a way that is known for certain in advance, the results would remain true with minor amendments—mainly, λ_t would vary over time on this account. The importance of known variations in interest rates depends on the elasticity of substitution between the present and future. If that elasticity is low, the influence would be unimportant. On the other hand, if the real interest rate applicable between periods t and $t + 1$ is uncertain at the time the consumption decision in period t is made, then the theoretical results no longer apply. However, there seems no strong reason for this to bias the results of the statistical tests in one direction or another.

II. Tests to Distinguish the Life Cycle–Permanent Income Theory from Alternative Theories

The tests of the stochastic implications of the life cycle–permanent income hypothesis carried out in this paper all have the form of estimating a conditional expectation, $E(c_t \mid c_{t-1}, x_{t-1})$, where x_{t-1} is a vector of data known in period $t - 1$, and then testing the hypothesis that the conditional expectation is actually not a function of x_{t-1}.[4] In all cases, the conditional expecta-

---

[4] The nature of the hypothesis being tested and the statistical tests themselves are essentially the same as in the large body of research on efficient capital markets (see Fama 1970). Sims (1978) treats the statistical problem of the asymptotic distribution of the regression coefficients of x_{t-1} in this kind of regression, with the conclusion that the standard formulas are correct.

tion is made linear in x_{t-1}, so the tests are the usual F-tests for the exclusion of a group of variables from a regression. Again, regression is the appropriate statistical technique for estimating the conditional expectation, and no claim is made that the true structural relation between consumption and its determinants is revealed by this approach.

What departures from the life cycle–permanent income hypothesis will this kind of test detect? There are two principal lines of thought about consumption that contradict the hypothesis. One holds that consumers are unable to smooth consumption over transitory fluctuations in income because of liquidity constraints and other practical considerations. Consumption is therefore too sensitive to current income to conform to the life cycle–permanent income principle. The second holds that a reasonable measure of permanent income is a distributed lag of past actual income, so the consumption function should relate actual consumption to such a distributed lag. A general consumption function embodying both ideas might let consumption respond with a fairly large coefficient to contemporaneous income and then have a distributed lag over past income. Such consumption functions are in widespread use and fit the data extremely well. But their estimation involves the very substantial issue that income and consumption are jointly determined. Estimation by least squares provides no evidence whether the observed behavior is consistent with the life cycle–permanent income hypothesis or not. Simultaneous estimation could provide evidence, but it would rest on crucial assumptions of exogeneity. Regressions of consumption on lagged consumption and lagged income *can* provide evidence without assumptions of exogeneity, as this section will show.

Consider first the issue of excessive sensitivity of consumption to transitory fluctuations in income, which has been emphasized by Tobin and Dolde (1971) and Mishkin (1976). The simplest alternative hypothesis supposes that a fraction of the population simply consumes all of its disposable income, instead of obeying the life cycle–permanent income consumption function. Suppose this fraction earns a proportion μ of total income, and let $c'_t = \mu y_t$ be their consumption. The other part of consumption, say c''_t, follows the rule set out earlier: $c''_t = \lambda c''_{t-1} + \varepsilon_t$. The conditional expectation of total consumption, c_t, given its own lagged value, and, say, two lagged values of income, is $E(c_t \mid c_{t-1}, y_{t-1}, y_{t-2}) = E(c'_t \mid c_{t-1}, y_{t-1}, y_{t-2}) + E(c''_t \mid c_{t-1}, y_{t-1}, y_{t-2}) = \mu E(y_t \mid c_{t-1}, y_{t-1}, y_{t-2}) + \lambda(c_{t-1} - \mu y_{t-1})$. Suppose that disposable income obeys a univariate autoregressive process of second order, so $E(y_t \mid c_{t-1}, y_{t-1}, y_{t-2}) = \rho_1 y_{t-1} + \rho_2 y_{t-2}$. Then $E(c_t \mid c_{t-1}, y_{t-1}, y_{t-2}) = \lambda c_{t-1} + \mu(\rho_1 - \lambda) y_{t-1} + \mu \rho_2 y_{t-2}$. The life cycle–permanent income hypothesis will be rejected unless $\rho_1 = \lambda$ and $\rho_2 = 0$, that is, unless disposable income and consumption obey exactly the same stochastic process. If they do, permanent income and observed income are the same thing, and the liquidity-constrained fraction of the population is obeying the hypothesis anyway, so the hypothesis is confirmed. The proposed test

involving regressing c_t on c_{t-1}, y_{t-1}, and y_{t-2} will reject the life cycle–permanent income hypothesis in favor of the simple liquidity-constrained model whenever the latter is materially different from the former.

The distributed lag approximation to permanent income was first suggested by Friedman (1957, 1963) and has figured prominently in consumption functions ever since. Distributed lags are not necessarily incompatible with the life cycle–permanent income hypothesis—if income obeys a stable stochastic process, there should be a structural relation between the innovation in income and consumption (Flavin 1977).[5] Still, the theory of the consumer presented earlier rules out any extra predictive value of a distributed lag of income (excluding contemporaneous income) in a regression that contains lagged consumption. If consumers use a nonoptimal distributed lag in forming their estimates of permanent income, then this central implication of the life cycle–permanent income hypothesis is false. This proposition is easiest to establish for the simple Koyck or geometric distributed lag, $c_t = \alpha \Sigma_{i=0}^{\infty} \beta^i y_{t-i}$ or $c_t = \beta c_{t-1} + \alpha y_t$. Suppose, as before, that y_t obeys a second-order autoregressive process, $E(y_t \mid c_{t-1}, y_{t-1}, y_{t-2})$ $= \rho_1 y_{t-1} + \rho_2 y_{t-2}$. Then the conditional expectation is $E(c_t \mid c_{t-1}, y_{t-1}, y_{t-2}) = \beta c_{t-1} + \alpha \rho_1 y_{t-1} + \alpha \rho_2 y_{t-2}$. As long as income is serially correlated $(\rho_1 \neq 0$ or $\rho_2 \neq 0)$, this conditional expectation will not depend solely on c_{t-1} and the pure life cycle–permanent income hypothesis will be refuted. Discussion of the peculiarities of the case of uncorrelated income seems unnecessary since income is in fact highly serially correlated. With this slight qualification, the proposed test procedure will always detect a Koyck lag if it is present and thus refute the life cycle–permanent income hypothesis.

It is possible to show that the test also applies to the general distributed lag model used by Modigliani (1971) and others. If the lag in the underlying structural consumption function is nonoptimal, lagged income will have additional predictive power for current consumption beyond that of lagged consumption, so the life cycle–permanent income hypothesis will be rejected. Data generated by consumers who use an optimal distributed lag of current and past income in making consumption decisions will not cause rejection. This shows the crucial distinction between structural models which include contemporaneous income and the test regressions of this paper where the principle of the tests involves the inclusion of lagged variables alone.

This section has shown that simple tests of the predictive power of variables other than lagged consumption can detect departures from the pure life cycle–permanent income hypothesis in the two directions that have been widely suggested in previous research on consumption. Both excessive sensitivity to current income because of liquidity constraints and non-

---

[5] Lucas (1976) argues convincingly that the stochastic process for income will shift if policy rules change.

TABLE 1

REGRESSION RESULTS FOR THE BASIC MODEL, 1948–77

$$c_t^{-1/\sigma} = \gamma c_{t-1}^{-1/\sigma} + \varepsilon_t$$

Equation	σ	Constant	γ	SE	R^2	D-W Statistic
1.0	.2	. . .	.983 (.003)	.000735	.9964	2.06
1.2	1.0	. . .	.996 (.001)	.00271	.9985	1.83
1.3	−1.0	−.014	1.011 (.003)	.0146	.9988	1.70

NOTE.—The numbers in parentheses in these and subsequent regressions are standard errors.

optimal distributed lag behavior will give additional predictive power to lagged income beyond that of lagged consumption in a regression for current consumption. The discussion of this section focused on the possible role of lagged income because that role is so closely related to alternative theories of consumption. Valid tests can be performed with any variable that is known in period $t − 1$ or earlier. The additional tests presented in the next section use extra lagged values of consumption and lagged values of common stock prices. Both variables have plausible justifications, but are less closely related to competing theories of consumption.

III. The Data and Results for the Basic Model

The most careful research on consumption has distinguished between the investment and consumption activities of consumers by removing investment in consumer durables and adding the imputed service flow of the stock of durables to consumption. For the purposes of this paper, however, it is more satisfactory simply to examine consumption of nondurables and services. All of the theoretical foundations of the aggregate consumption function apply to individual categories of consumption as well. Dropping durables altogether avoids the suspicion that the findings are an artifact of the procedure for imputing a service flow to the stock of durables. The data on consumption used throughout the study, then, can be defined exactly as consumption of nondurables and services in 1972 dollars from the U.S. National Income and Product Accounts divided by the population. All data are quarterly.

Table 1 presents the results of fitting the basic regression relation between current and lagged marginal utility predicted by the pure life cycle–permanent income theory. Equations 1.1 and 1.2 are for the constant-elasticity utility function, with $\sigma = 0.2$ and 1.0, respectively. Equation 1.3 is for the quadratic utility function exactly, or for any utility function approximately, and is simply a regression of consumption on its own lagged value and a constant. All three equations show that the predictive value of

lagged marginal utility for current marginal utility is extremely high; that is, the typical information that becomes available in each quarter, as measured by ε_t, has only a small impact on consumption or marginal utility. Of course, this is no more than a theoretical interpretation of the well-known fact that consumption is highly serially correlated. The close fit of the regressions in table 1 is not itself confirmation of the life cycle–permanent income hypothesis, since the hypothesis makes no prediction about the variability of permanent income and the resultant variance of ε_t. The theory is compatible with any amount of unexplained variation in the regression.

There is no usable statistical criterion for choice among the three equations in table 1. The transformation of the dependent variable rules out the simple principle of least squares. Under the assumption of a normal distribution for ε_t, there is a likelihood function with an extra term, the Jacobian determinant, to take account of the transformation. However, for this sample, it proved to be an increasing function of σ for all values, so no maximum-likelihood estimator is available. This seems to reflect the operation of corollary 5—the ε_t's are small enough that any specification of marginal utility is essentially proportional to consumption itself, and the effective content of the life cycle–permanent income theory is to make consumption itself evolve as a random walk with trend. From this point on, the paper will discuss only equation 1.3 and its extensions to other variables.

The principal stochastic implication of the life cycle–permanent income hypothesis is that no other variables observed in quarter $t - 1$ or earlier can help predict the residuals from the regressions in table 1. Before formal statistical tests are used, it is useful to study the residuals themselves. The pattern of the residuals is extremely similar in the three regressions, but the residuals themselves are easiest to interpret for equation 3, where they have the units of consumption per capita in 1972 dollars. These residuals appear in table 2.

The standard error of the residuals in 14.6, so roughly six of the observations should exceed 29.2 in magnitude. There are in fact six. Three are drops in consumption, and of these, one coincides with the standard dating of recessions: 1974:4. Five milder recessions contribute drops of less than two standard deviations: 1949:3, 1953:4, 1958:1, 1960:3, and 1970:4. The other major decline in consumption is associated with the Korean War, in 1950:4. Most of the drops in consumption occurred quickly, in one or two quarters. The only important exception was in the period from 1973:4 to 1975:1, when six straight quarters of consecutive decline took place. On the expansionary side, there is little consistent evidence of any systematic tendency for consumption to recover in a regular pattern after a setback. The largest single increase occurred in 1965:4. This, together with three successive increases in 1964, accounts for all of the increase in consumption relative to trend associated with the prolonged boom of the mid- and late sixties.

TABLE 2

RESIDUALS FROM REGRESSION OF CONSUMPTION ON LAGGED CONSUMPTION, 1948–77 ($)

1948:		1956:		1964:		1972:	
1....	.5	1....	2.8	1....	17.8	1....	20.0
2....	8.0	2....	−10.1	2....	20.0	2....	32.7
3....	−15.5	3....	−6.1	3....	14.6	3....	8.4
4....	3.5	4....	1.2	4....	−4.4	4....	21.1
1949:		1957:		1965:		1973:	
1....	−8.6	1....	−10.8	1....	5.8	1....	8.0
2....	−8.5	2....	−6.1	2....	5.2	2....	−15.6
3....	−27.2	3....	2.4	3....	10.2	3....	3.0
4....	−.1	4....	−13.6	4....	38.7	4....	−32.8
1950:		1958:		1966:		1974:	
1....	5.6	1....	−29.1	1....	−1.3	1....	−27.3
2....	23.8	2....	8.3	2....	3.7	2....	−23.4
3....	15.5	3....	14.3	3....	−1.9	3....	−16.6
4....	−31.0	4....	−1.9	4....	−12.4	4....	−42.8
1951:		1959:		1967:		1975:	
1....	16.0	1....	15.1	1....	10.2	1....	−5.8
2....	−24.8	2....	3.8	2....	2.0	2....	25.6
3....	9.0	3....	−2.7	3....	−2.8	3....	−21.3
4....	−6.1	4....	1.1	4....	−7.5	4....	.9
1952:		1960:		1968:		1976:	
1....	−15.1	1....	−5.6	1....	15.6	1....	24.4
2....	18.0	2....	8.8	2....	7.9	2....	8.3
3....	10.0	3....	−24.3	3....	22.2	3....	2.3
4....	8.6	4....	−10.1	4....	−8.1	4....	30.4
1953:		1961:		1969:		1977:	
1....	−1.6	1....	1.8	1....	.8	1....	−1.8
2....	−1.0	2....	10.1	2....	−5.3		
3....	−22.1	3....	−16.9	3....	−2.4		
4....	−27.0	4....	15.8	4....	.3		
1954:		1962:		1970:			
1....	−.1	1....	−1.7	1....	4.0		
2....	−2.4	2....	3.4	2....	−12.3		
3....	11.9	3....	−.8	3....	−.3		
4....	7.4	4....	.6	4....	−21.5		
1955:		1963:		1971:			
1....	6.1	1....	−8.4	1....	1.5		
2....	7.0	2....	−1.3	2....	−2.4		
3....	−.8	3....	11.7	3....	−14.6		
4....	22.3	4....	−6.3	4....	−6.6		

The data contain no obvious refutation of the unpredictability of the residuals from the basic model, but, just as a study of stock prices will never convince the "chartist" that it is futile to try to predict their future, the confirmed believer in regular fluctuations in consumption will not be swayed by the data alone. More powerful methods for summarizing the data are required.

IV. Can Consumption Be Predicted from Its Own Past Values?

The simplest testable implication of the pure life cycle–permanent income hypothesis is that only the first lagged value of consumption helps predict current consumption. This implication would be refuted if consumption had a definite cyclical pattern described by a difference equation of second

or higher order.[6] Intelligent consumers ought to be able to offset any such cyclical pattern and restore the noncyclical optimal behavior of consumption predicted by the hypothesis. The following regression tests this implication by adding additional lagged values of consumption to equation 1.3:

$$c_t = 8.2 + 1.130c_{t-1} - 0.040c_{t-2} + 0.030c_{t-3} - 0.113c_{t-4};$$
$$(8.3)\quad(0.092)\qquad(0.142)\qquad(0.142)\qquad(0.093)$$
$$R^2 = .9988;\qquad s = 14.5;\qquad \text{D-W} = 1.96.$$

The contribution of the extra lagged values is to increase the accuracy of the forecast of current consumption by about 10 cents per person per year. The F-statistic for the hypothesis that the coefficients of c_{t-2}, c_{t-3}, and c_{t-4} are all zero is 1.7, well under the critical point of the F-distribution of 2.7 at the 5 percent level. Only very weak evidence against the pure life cycle–permanent income hypothesis appears in this regression. In particular, there are no definite signs that consumption obeys a second-order difference equation capable of generating stochastic cycles. In this respect, consumption differs sharply from other aggregate economic measures, which do typically obey second-order autoregressions.

V. Can Consumption Be Predicted from Disposable Income?

If lagged income has substantial predictive power beyond that of lagged consumption, then the life cycle–permanent income hypothesis is refuted. As discussed in Section II, this evidence would support the alternative views that consumers are excessively sensitive to current income, or, more generally, that they use an ad hoc, nonoptimal distributed lag of past income in making consumption decisions.

Table 3 presents a variety of regressions testing the predictive power of real disposable income per capita, measured as current dollar disposable income from the national accounts divided by the implicit deflator for consumption of nondurables and services and divided by population. Equation 3.1 shows that a single lagged level of disposable income has essentially no predictive value at all. The coefficient of y_{t-1} is slightly negative, but this is easily explained by sampling variation alone. The F-statistic for the exclusion of all but the constant and c_{t-1} is 0.1, far below the critical F of 3.9. Equation 3.2 tries a year-long distributed lag estimated without constraint. The first lagged value of disposable income has a slight positive coefficient, but this is more than outweighed by the three negative coefficients for the longer lags. The long-run "marginal propensity to consume," measured by the sum of the coefficients, is actually negative, though again this could easily result from sampling variation. The F-statistic for the joint predictive value of all four lagged income variables is 2.0, somewhat less than the critical value of 2.4 at the 5 percent level. Note

[6] Fama (1970) calls the similar test for asset prices a "weak form" test.

TABLE 3

EQUATIONS RELATING CONSUMPTION TO LAGGED CONSUMPTION AND PAST LEVELS OF REAL DISPOSABLE INCOME

Equation No. and Equation	R^2	s	D-W	F	F*
3.1 .... $c_t = -16 + 1.024\, c_{t-1} - .010\, y_{t-1}$ (11) (.044) (.032)	.9988	14.7	1.71	.1	3.9
3.2 .... $c_t = -23 + 1.076\, c_{t-1} + .049\, y_{t-1} - .051\, y_{t-2}$ (11) (.047) (.043) (.052) $-.023\, y_{t-3} - .024\, y_{t-4}$ (.051) (.037)	.9989	14.4	2.02	2.0	2.4
3.3 .... $c_t = -25 + 1.113\, c_{t-1} + \sum_{i=1}^{12} \beta_i\, y_{t-i}$ $\Sigma\beta_i = .077$ (11) (.054) (.040)	.9988	14.6	1.92	2.0	2.7

that the pure life cycle–permanent income hypothesis would be rejected if the size of the test were 10 percent or higher.

Equation 3.3 fits a 12-quarter Almon lag to see if a long distributed lag can compete with lagged consumption as a predictor for current consumption. Again, the sum of the lag coefficients is slightly negative, now almost significantly so. The F-statistic for the hypothesis of no contribution from the complete distributed lag on income is again close to the critical value.

The sample evidence of the relation between consumption and lagged income seems to say the following: There is a statistically marginal and numerically small relation between consumption and very recent levels of disposable income. The sum of the lag coefficients is slightly negative. Further, there is no evidence at all supporting the view that a long distributed lag covering several years helps to predict consumption. This evidence casts just a little doubt on the life cycle–permanent income hypothesis in its purest form but is not at all destructive to a somewhat more flexible interpretation of the hypothesis, to be discussed shortly.

VI. Wealth and Consumption

Of the many alternative variables that might be included on the right-hand side of a regression to test the pure life cycle–permanent income hypothesis, some measure of wealth is one of the leading candidates. Theory and prevailing practice agree that contemporaneous wealth has a strong influence on consumption, so lagged wealth is a logical variable to test. Again, the hypothesis implies that wealth measured in earlier quarters should have no predictive value with respect to this quarter's consumption. All information contained in lagged wealth should be summarized in lagged consumption.

Reliable quarterly data on property values are not available for most categories of property. For one major category, however, essentially perfect data are available at any frequency, namely, the market value of corporate stock. Tests of the random-walk hypothesis do not require a comprehensive wealth variable, so a test based on stock prices is appropriate, even though the resulting equation does not describe the structural relation between wealth and consumption. The tests reported here are based on Standard and Poor's comprehensive index of the prices of stocks deflated by the implicit deflator for nondurables and services and divided by population. This variable will be called s. It makes a statistically unambiguous contribution to prediction of current consumption:

$$c_t = -22 + 1.012c_{t-1} + 0.223s_{t-1} - 0.258s_{t-2} + 0.167s_{t-3} - 0.120s_{t-4}$$
$$(8) \quad (0.004) \quad (0.051) \quad (0.083) \quad (0.083) \quad (0.051)$$
$$R^2 = .9990; \quad SE = 14.4; \quad D\text{-}W = 2.05.$$

The F-statistic for the hypothesis that the coefficients of the lagged stock prices are all zero is 6.5, well above the critical value of 2.4 at the 5 percent

level. Further, each coefficient considered separately is clearly different from zero according to the usual t-test. However, the improvement in the predictive power of the regression, while statistically significant, is not numerically large. The standard error of the regression is about 20 cents per person per year smaller in this equation compared with the basic model of equation 1.3 ($\$14.40$ against $\$14.60$). Most of the predictive value of the stock price comes from the *change* in the price in the immediately preceding quarter. A smaller contribution is made by the change in the price 3 quarters earlier. Use of the Almon lag technique for both levels and differences in the stock price failed to turn up any evidence of a longer distributed lag.

VII. Implications of the Empirical Evidence

The pure life cycle–permanent income hypothesis—that c_t cannot be predicted by any variable dated $t - 1$ or earlier other than c_{t-1}—is rejected by the data. The stock market is valuable in predicting consumption 1 quarter in the future. Most of the predictive power comes from Δs_{t-1}. But the data seem entirely compatible with a modification of the hypothesis that leaves its central content unchanged. Suppose that consumption does depend on permanent income, and that marginal utility indeed does evolve as a random walk with trend, but that some part of consumption takes time to adjust to a change in permanent income. Then any variable that is correlated with permanent income in $t - 1$ will help in predicting the change in consumption in period t, since part of that change is the lagged response to the previous change in permanent income. Both the finding that consumption is only weakly associated with its own past values and that immediate past values of changes in stock prices have a modest predictive value are compatible with this modification of the life cycle–permanent income hypothesis.

Whatever problems remain in the consumption function, there seems little reason to doubt the life cycle–permanent income hypothesis. Within a framework in which permanent income is treated as an unobserved variable the data seem fully compatible with the hypothesis, provided a short lag between permanent income and consumption is recognized. Of course, acceptance of the hypothesis does not yield a complete consumption function, since no equation for permanent income has been developed. The evidence against the ad hoc distributed-lag model relating permanent income to actual income seems fairly strong. The task of further research is to create a more satisfactory model for permanent income, one that recognizes that consumers appraise their economic well-being in an intelligent way that involves looking into the future.

It is important not to treat any of the equations of this paper as structural relations between consumption and the variables that are used to predict it.

For example, table 3 should not be read as implying that income has a negative effect on consumption. The effect of a particular change in income depends on the change in permanent income it induces, and this can range anywhere from no effect to a dollar-for-dollar effect, depending on the way that consumers evaluate the change. In any case, the regressions understate the true structural relation between the change in income and the change in consumption because they omit the contemporaneous part of the relation.

VIII. Implications for Forecasting and Policy Analysis

Under the pure life cycle–permanent income hypothesis, a forecast of future consumption obtained by extrapolating today's level by the historical trend is impossible to improve. The results of this paper have the strong implication that beyond the next few quarters consumption should be treated as an exogenous variable. There is no point in forecasting future income and then relating it to income, since any information available today about future income is already incorporated in today's permanent income. Forecasts of consumption next quarter can be improved slightly with current stock prices, but no further improvement can be achieved in this way in later quarters.

With respect to the analysis of stabilization policy, the findings of this paper go no further than supporting the view that policy affects consumption only as much as it affects permanent income. In the analysis of policies that are known to leave permanent income unchanged, consumption may be treated as exogenous. Further, only *new* information about taxes and other policy instruments can affect permanent income. Beyond these general propositions, the policy analyst must answer the difficult question of the effect of a given policy on permanent income in order to predict its effect on consumption. Regression of consumption on current and past values of income are of no value whatsoever in answering this question.

Appendix

1. Theorem

If a consumer maximizes $E_t \sum_{t=0}^{T} (1 + \delta)^{-t} u(c_t)$, subject to $\sum_{t=0}^{T} (1 + r)^{-t} (c_t - w_t)$ $= A_0$, sequentially determining c_t at each t, then $E_t u'(c_{t+1}) = [(1 + \delta)/(1 + r)]$ $\times u'(c_t)$.

Proof.—At time t, the consumer chooses c_t so as to maximize $(1 + \delta)^{-t} u(c_t) +$ $E_t \sum_{\tau=t+1}^{T} (1 + \delta)^{-t} u(c_\tau)$ subject to $\sum_{\tau=t}^{T} (1 + r)^{(t-\tau)}(c_\tau - w_\tau) = A_t$. The optimal sequential strategy has the form $c_t = g_t(w_t, w_{t-1}, \ldots, w_0, A_0)$. Consider a variation from this strategy: $c_t = g_t(w_t, \ldots) + x$: $c_{t+1} = g_{t+1}(w_{t+1}, w_t, \ldots) - (1 + r)x$.

Note that the new consumption strategy also satisfies the budget constraint. Now consider $\max_x \{(1 + \delta)^{-t} u(g_t + x) + E_t [(1 + \delta)^{-t-1} u(g_{t+1} - (1 + r)x) + \sum_{\tau=t+2}^{T} (1 + \delta)^{-\tau} u(g_\tau)]\}$.

The first-order condition is $(1 + \delta)^{-t} u'(g_t + x) - E_t (1 + \delta)^{-t-1}(1 + r)$ $u'(g_{t+1} - (1 + r)x) = 0$ as asserted.

2. Proof of Corollary 5

Recall that $u'(c_{t+1}) = [(1 + \delta)/(1 + r)]u'(c_t) + \varepsilon_{t+1}$ and $\lambda_t = [(1 + \delta)/(1 + r)]^{u'(c_t)/[c_t u''(c_t)]}$. Expand the implicit equation for c_{t+1} in a Taylor series at the point $\lambda_t = 1$ ($r = \delta$) and $\varepsilon_{t+1} = 0$: $c_{t+1} = c_t + (\lambda_t - 1)(\partial c_{t+1}/\partial \lambda_t) + \varepsilon_{t+1}(\partial c_{t+1}/\partial \varepsilon_{t+1})$ + higher-order terms. At the point $\lambda_t = 1$ and $\varepsilon_{t+1} = 0$, c_{t+1} equals c_t, and it is not hard to show that $\partial c_{t+1}/\partial \lambda_t = c_t$ and $\partial c_{t+1}/\partial \varepsilon_{t+1} = 1/u''(c_t)$. Thus $c_{t+1} = c_t + (\lambda_t - 1)c_t + \varepsilon_{t+1}/u''(c_t) = \lambda_t c_t + \varepsilon_{t+1}/u''(c_t)$, as asserted.

References

Bewley, Truman. "The Permanent Income Hypothesis: A Theoretical Formulation." Mimeographed. Cambridge, Mass.: Harvard Univ., Dept. Econ., September 1976.

Blinder, Alan. "Temporary Taxes and Consumer Spending." Mimeographed. Jerusalem: Inst. Advanced Studies, April 1977.

Darby, Michael R. "The Allocation of Transitory Income among Consumers' Assets." *A.E.R.* 62 (December 1972): 928–41.

Fama, Eugene F. "Efficient Capital Markets: A Review of Theory and Empirical Work." *J. Finance* 25 (May 1970): 383–417.

Flavin, Marjorie. "The Adjustment of Consumption to Changing Expectations about Future Income." Unpublished paper, Massachusetts Inst. Tech., February 1977.

Friedman, Milton. *A Theory of the Consumption Function*. Princeton, N.J.: Princeton Univ. Press, 1957.

———. "Windfalls, the 'Horizon,' and Related Concepts in the Permanent-Income Hypothesis." In *Measurement in Economics: Studies in Mathematical Economics and Econometrics in Memory of Yehuda Grunfeld*, edited by Carl Christ et al. Stanford, Calif.: Stanford Univ. Press, 1963.

Friedman, Milton, and Becker, Gary S. "A Statistical Illusion in Judging Keynesian Models." *J.P.E.* 65, no. 1 (February 1957): 64–75.

Granger, C. W. J., and Newbold, P. "Forecasting Transformed Series." *J. Royal Statis. Soc.* ser. B, 38, no. 2 (1976): 189–203.

Haavelmo, Trygve. "The Statistical Implications of a System of Simultaneous Equations." *Econometrica* 11 (January 1943): 1–12.

Lucas, Robert E., Jr. "Econometric Policy Evaluation: A Critique." In *The Phillips Curve and Labor Markets*, edited by Karl Brunner and Allan H. Meltzer. Carnegie-Rochester Conference Series on Public Policy. Amsterdam: North-Holland; New York: American Elsevier, 1976.

Mishkin, Frederic. "Illiquidity, Consumer Durable Expenditure, and Monetary Policy." *A.E.R.* 66 (September 1976): 642–54.

Modigliani, Franco. "Monetary Policy and Consumption." In *Consumer Spending and Monetary Policy: The Linkages*. Conference Series No. 5. Boston: Federal Reserve Bank of Boston, 1971.

Sims, Christopher A. "Least Squares Estimation of Autoregressions with Some Unit Roots." Discussion Paper no. 78–95, Univ. Minnesota, Center Econ. Res., March 1978.

Tobin, James, and Dolde, Walter. "Wealth, Liquidity, and Consumption." In *Consumer Spending and Monetary Policy: The Linkages*. Conference Series No. 5. Boston: Federal Reserve Bank of Boston, 1971.

Yaari, Menahem E. "A Law of Large Numbers in the Theory of Consumer's Choice under Uncertainty." *J. Econ. Theory* 12 (April 1976): 202–17.