

NBER WORKING PAPER SERIES

CONSUMPTION, INCOME, AND INTEREST RATES:
REINTERPRETING THE TIME SERIES EVIDENCE

John Y. Campbell

N. Gregory Mankiw

Working Paper No. 2924

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 1989

An earlier version of this paper was presented at the NBER Macroeconomics Conference, Cambridge, MA, March 10-11, 1989. We are grateful to Karen Dynan for research assistance, Alan Blinder, Olivier Blanchard, and Robert Hall for comments, and to the NBER and the National Science Foundation for financial support. This paper is part of NBER's research program in Economic Fluctuations. Any opinions expressed are those of the authors not of the National Bureau of Economic Research.

NBER Working Paper #2924
April 1989

CONSUMPTION, INCOME, AND INTEREST RATES:
REINTERPRETING THE TIME SERIES EVIDENCE

ABSTRACT

This paper proposes that the time-series data on consumption, income, and interest rates are best viewed as generated not by a single representative consumer but by two groups of consumers. Half the consumers are forward-looking and consume their permanent income, but are extremely reluctant to substitute consumption intertemporally. Half the consumers follow the "rule of thumb" of consuming their current income. The paper documents three empirical regularities that, it argues, are best explained by this model. First, expected changes in income are associated with expected changes in consumption. Second, expected real interest rates are not associated with expected changes in consumption. Third, periods in which consumption is high relative to income are typically followed by high growth in income. The paper concludes by briefly discussing the implications of these findings for economic policy and economic research.

John Y. Campbell
Woodrow Wilson School
Princeton University
Princeton, NJ 08544

N. Gregory Mankiw
National Bureau of Economic Research
1050 Massachusetts Avenue
Cambridge, MA 02138

Introduction

The study of aggregate consumption behavior has been profoundly altered by the rational expectations revolution in macroeconomics. The first example in Robert Lucas's (1976) influential critique of econometric policy evaluation involved consumption. Lucas argued that traditional consumption functions, no matter how well they fit the data, were not useful for evaluating the effects of alternative policies. Soon thereafter, Robert Hall (1978) proposed a new approach to studying consumption that was firmly founded on the postulate of rational expectations and that was immune to the problems Lucas pointed out. Hall suggested that aggregate consumption should be modelled as obeying the first-order conditions for optimal choice of a single, fully rational, and forward-looking representative consumer. The new style of research based on this assumption--sometimes called the "Euler equation approach"--has dominated work on consumption during the past decade.

In this paper we appraise what has been learned about aggregate consumption from this approach. We propose a simple, alternative characterization of the time-series data on consumption, income, and interest rates. We suggest that the data are best viewed as generated not by a single forward-looking consumer but by two types of consumers. Half the consumers are forward-looking and consume their permanent income, but are extremely reluctant to substitute consumption intertemporally in response to interest rate movements. Half the consumers follow the "rule of thumb" of consuming their current income. We document three empirical regularities that, we argue, are best explained by this model.

The first regularity is that expected changes in income are associated with expected changes in consumption. In contrast to the simplest version of the permanent income hypothesis, consumption is not a random walk: when income is expected to rise by 1 percent, consumption should be expected to rise by 0.5 percent. The strong connection between current income and consumption provides at least circumstantial evidence for "rule-of-thumb" behavior on the part of some consumers.

The second empirical regularity is that expected real interest rates are not associated with expected changes in consumption. This means that the predictable movements that we observe in consumption cannot be explained as a rational response to movements in real interest rates. It also means that forward-looking consumers do not adjust their consumption growth in response to interest rates, so their intertemporal elasticity of substitution in consumption must be close to zero. Hall (1988) also argues that the elasticity of substitution of permanent income consumers is small; but since he does not allow for current income consumers, he cannot explain the existence of any predictable movements in aggregate consumption.

The third empirical regularity is that periods in which consumption is high relative to income are typically followed by rapid growth in income. This finding suggests that at least some consumers are forward-looking: their knowledge of future income growth is reflected in current consumption. Yet we show that the magnitude of the association between consumption and future income growth is best explained by a model with both permanent income consumers and current income consumers.

Most of this paper is devoted to analyzing the data and documenting its consistency with the simple model we propose. In the final section, we

briefly discuss the broader implications for economic policy and economic research.

1. Is Consumption a Random Walk?

In this section we reexamine the evidence on the simplest version of the permanent income hypothesis, according to which consumption should follow a random walk. We begin by reviewing the basic model and discuss how it can be tested. Our approach differs from the standard one in two ways. First, we emphasize a specific alternative hypothesis under which some consumers follow the "rule of thumb" of consuming their current income rather than their permanent income. Second, we argue that more structural estimation using instrumental variables should be preferred over the standard tests for a random walk using the reduced form of the model. When we look at the data, we find that a substantial fraction of income accrues to rule-of-thumb consumers, indicating an economically important deviation from the permanent income hypothesis.

1.1. The Permanent Income Hypothesis and a "Rule-of-Thumb" Alternative

The permanent income hypothesis as usually formulated assumes that aggregate consumption can be modelled as the decisions of a representative consumer. The representative consumer maximizes

$$(1.1) \quad E_t \sum_{s=0}^{\infty} (1+\delta)^{-s} U(C_{t+s}) \quad U' > 0, U'' < 0$$

where C is consumption, δ is the subjective rate of discount, and E_t is the expectation conditional on information available at time t . If the

representative consumer can borrow and lend at the real interest rate r , then the first-order condition necessary for an optimum is

$$(1.2) \quad E_t U'(C_{t+1}) = \left(\frac{1+\delta}{1+r}\right) U'(C_t).$$

This says that marginal utility today is, up to a constant multiple, the best forecast of marginal utility tomorrow.

If we assume that $r=\delta$ and that marginal utility is linear, then we obtain the random walk result¹, $E_t C_{t+1} = C_t$. Consumption today is the optimal forecast of consumption tomorrow. This in turn implies

$$(1.3) \quad \Delta C_t = \epsilon_t$$

where ϵ_t is a rational forecast error, the innovation in permanent income. Thus, according to this formulation of the permanent income hypothesis, the change in consumption is unforecastable.

In evaluating how well this model fits the data, it is useful to keep in mind an explicit alternative hypothesis. We nest the permanent income hypothesis in a more general model in which some fraction of income λ accrues to individuals who consume their current income, while the remainder $(1-\lambda)$ accrues to individuals who consume their permanent income. If the incomes of the two groups are Y_{1t} and Y_{2t} , then total income is $Y_t = Y_{1t} + Y_{2t}$. Since the first group receives λ of total income,

¹ Obviously, these assumptions can be justified only as an approximation. One can obtain the random walk result with other sorts of approximations as well, e.g., the Taylor approximation in Mankiw (1981) or the log-normality assumption in Hansen and Singleton (1983). These other approximations may imply the log of consumption, rather than the level, is a random walk--a more appealing specification. They also often introduce other terms, such as the difference between δ and r and the variance of consumption growth; these other terms are usually included as part of the constant drift in consumption.

$Y_{1t} = \lambda Y_t$ and $Y_{2t} = (1-\lambda)Y_t$. Agents in the first group consume their current income, so $C_{1t} = Y_{1t}$, implying $\Delta C_{1t} = \Delta Y_{1t} = \lambda \Delta Y_t$. By contrast, agents in the second group obey the permanent income hypothesis, implying $\Delta C_{2t} = (1-\lambda)\epsilon_t$.

The change in aggregate consumption can now be written as

$$(1.4) \quad \Delta C_t = \Delta C_{1t} + \Delta C_{2t} = \lambda \Delta Y_t + (1-\lambda)\epsilon_t.$$

Under this alternative hypothesis, the change in consumption is a weighted average of the change in current income and the unforecastable innovation in permanent income. Equation (1.4) reduces to the permanent income hypothesis, equation (1.3), when $\lambda = 0$.²

Having set up the permanent income hypothesis as the null hypothesis and the existence of these "rule-of-thumb" consumers as the alternative hypothesis, there are two approaches to estimation and testing. The approach we advocate is to estimate λ directly and test the hypothesis that $\lambda=0$. It is important to note, however, that (1.4) cannot be estimated by Ordinary Least Squares, since the error term ϵ_t may be correlated with ΔY_t . The solution is to estimate (1.4) by instrumental variables. Any lagged stationary variables are potentially valid instruments since they are orthogonal to ϵ_t . Of course, good instruments must also be correlated with ΔY_t --therefore, one should choose lagged variables that can predict future income growth. Once such instruments are found, one can easily estimate

² This alternative model with some rule-of-thumb consumers is discussed briefly in Hall (1978). It is also a simpler version of the model proposed in Flavin (1981), in which the change in consumption responds not only to the contemporaneous change in current income, but also to lagged changes in current income. Flavin designs her model so that it is just-identified; by contrast, we view the over-identification of our model as one of its virtues. See also Bean (1986).

the fraction of income accruing to the rule-of-thumb consumers.

The second approach to testing the permanent income hypothesis--used by Hall (1978) and in most of the subsequent literature--is to regress the change on consumption on lagged variables to see whether the change in consumption is forecastable. To see the relation between the two approaches, note that equation (1.4), estimated by instrumental variables, can be viewed as a restricted version of a more general two-equation system in which ΔC_t and ΔY_t are regressed directly on the instruments. If we have K instruments, X_{1t} through X_{Kt} , then the general system is

$$(1.5) \quad \begin{aligned} \Delta C_t &= \beta_0 + \beta_1 X_{1t} + \dots + \beta_K X_{Kt} + \eta_{Ct} = X_t \beta + \eta_{Ct} \\ \Delta Y_t &= \gamma_0 + \gamma_1 X_{1t} + \dots + \gamma_K X_{Kt} + \eta_{Yt} = X_t \gamma + \eta_{Yt} \end{aligned}$$

The permanent income hypothesis implies that the vector $\beta = 0$ (that is, $\beta_1 = \dots = \beta_K = 0$). This implication can be tested directly, without any need for considering the ΔY_t equation, by OLS estimation of the ΔC_t equation. When there is more than a single instrument, however, equation (1.4) places over-identifying restrictions on the two equation system (1.5): predictable changes in consumption and income, and therefore the vectors β and γ , are proportional to one another ($\beta = \lambda \gamma$, or $\beta_1/\gamma_1 = \dots = \beta_K/\gamma_K = \lambda$). The instrumental variables test that $\lambda=0$ is in essence a test that $\beta=0$ under the maintained hypothesis that these overidentifying restrictions are true.

Although estimating the reduced form equation for ΔC_t is more standard, there are compelling reasons to prefer the instrumental variables approach. One reason is power. Since there are many possible instruments, the instrumental variables procedure estimates far fewer parameters than are in the reduced form, thereby conserving on the degrees of freedom and

providing a more powerful test of the null hypothesis.

Perhaps more important, estimation of λ provides a useful metric for judging whether an observed deviation from the null hypothesis is economically important. As Franklin Fisher (1961) emphasized long ago, an economic model can be approximately true even if the strict tests of overidentification fail. It is therefore hard to interpret a rejection of the permanent income hypothesis in the reduced form framework. Indeed, Hall (1978) concluded that the evidence favors the permanent income hypothesis even though he reported formal rejections using stock prices. An estimate of λ is more informative about the economic importance of deviations from the theory.³ For example, if the estimate of λ is close to zero, then one can say the permanent income is approximately true--most income goes to consumers who obey the theory--even if the estimate of λ is statistically significant. Conversely, if the estimate of λ is large, then one must conclude that the evidence points away from the permanent income hypothesis.

One question that arises in interpreting a failure of the permanent income hypothesis is whether our rule-of-thumb alternative adequately captures the reason for the failure. The best way to answer the question is to consider explicitly other alternative hypotheses.⁴ Another way--more statistical and less economic--is to test the overidentifying restrictions that equation (1.4) imposes. This test is performed simply by regressing the residual from the instrumental variables regression on the instruments, and then to compare T times the R^2 from this regression, where T is the

³ Flavin (1981) also stresses this point.

⁴ For some examples, see Campbell and Mankiw (1987).

sample size, with the χ^2 distribution with $(K-1)$ degrees of freedom. We use this test below.

1.2. Two Specification Issues

Before we can estimate the model, we need to address two issues of specification that arise from the nature of the aggregate time series on consumption and income.

Our discussion so far has been couched in terms of levels and differences of the raw series C_t and Y_t . This is appropriate if these series follow homoskedastic linear processes in levels, with or without unit roots. Yet aggregate time series on consumption and income appear to be closer to log-linear than linear: the mean change and the innovation variance both grow with the level of the series. A correction of some sort appears necessary. The approach we take is simply to take logs of all variables. Although the parameter λ can no longer be precisely interpreted as the fraction of agents who consume their current income, one can view the model we estimate as the log-linear approximation to the true model. Thus, the interpretation of the results is not substantially affected. We use lower-case letters to denote log variables.⁵

A second data problem is that consumption and income are measured as quarterly averages rather than at points in time. If the permanent income hypothesis holds in continuous time, then measured consumption is the time average of a random walk. Therefore, the change in consumption will have a first-order serial correlation of 0.25, which could lead us to reject the

⁵ An alternative scaling method is to divide ΔC_t and ΔY_t by the lagged level of income, Y_{t-1} . In practice both scaling methods give very similar results.

model even if it is true.⁶ We deal with this problem by lagging the instruments more than one period, so that there is at least a two-period time gap between the instruments and the variables in equation (1.4). The time average of a continuous-time random walk is uncorrelated with all variables lagged more than one period, so by using twice-lagged instruments we obtain a test of the model that is valid for time-averaged data.

1.3. Another Look at U.S. Data

To estimate our model, we use standard U.S. quarterly time series data, obtained from the Data Resources, Inc. data bank. Y_t is measured as disposable personal income per capita, in 1982 dollars. C_t is consumption of non-durables and services per capita, in 1982 dollars. The sample period is 1953:1 to 1986:4.⁷

Table 1, which reports the results, has six columns. The first gives the row number and the second the instruments used⁸. The third and fourth columns give the adjusted R^2 statistics for OLS regressions of Δc_t and Δy_t , respectively, on the instruments. In parentheses we report the p-value for a Wald test of the hypothesis that all coefficients except the intercept are zero. The fifth column gives the instrumental variables estimate of λ , with an asymptotic standard error. The final column gives the adjusted R^2

⁶ See Working (1960).

⁷ In Campbell and Mankiw (1987) we discuss the importance of sample period and, in particular, the peculiar behavior of the first quarter of 1950, when there was a one-time National Service Life Insurance dividend payment to World War II veterans. The sample period of Table 1 extends the data used in Campbell and Mankiw (1987) by one year.

⁸ A constant term is always included as both an instrument and a regressor, but is not reported in the tables.

statistic for an OLS regression of the residual from the instrumental variables regression on the instruments. In parentheses we report the p-value for the corresponding test of the overidentifying restrictions placed by equation (1.4) on the general system (1.5). For reference, the first row of Table 1 shows the coefficient obtained when we estimate equation (1.4) by OLS.

Rows 2 and 3 of the table use lagged income growth rates as instruments. These are not strongly jointly significant in predicting consumption or income growth; in row 3, for example, lags two through six of income growth are jointly significant at the 21% level for consumption growth and at the 14% level for income growth. It appears that the univariate time series process for disposable income is close enough to a random walk that income growth rates are not well forecast by lagged income growth rates. Our instrumental variables procedure estimates λ at 0.506 with an asymptotic standard error of 0.176 in row 3; this rejects the permanent income hypothesis that $\lambda = 0$ at the 0.4% level. Yet instrumental variables procedures can be statistically unreliable when the instruments have only weak forecasting power for the right hand side variable.⁹ The rejection of the permanent income hypothesis in rows 2 and 3 should be interpreted cautiously.¹⁰

We obtain stronger results in rows 4 and 5 of the table, where we use lagged consumption growth rates as instruments. It is striking that lagged

⁹ See Nelson and Startz (1988) for an analysis of this issue.

¹⁰ These findings confirm the conclusions of Mankiw and Shapiro (1985): since disposable income is so close to a random walk, modelling income as a univariate process (e.g., Flavin (1981) or Bernanke (1985)) leads to tests with little power.

consumption forecasts income growth more strongly than lagged income itself does, and this enables us to estimate the parameter λ more precisely. This finding suggests that at least some consumers have better information on future income growth than is summarized in its past history and that they respond to this information by increasing their consumption. At the same time, however, the fraction of rule-of-thumb consumers is estimated at 0.523 in row 5 (and the estimate is significant at better than the 0.01% level). The OLS test also rejects the permanent income model in row 5.

We next consider using some financial variables as instruments. We tried using lagged changes in real stock prices (the quarterly percentage change in the real value of the Dow Jones Industrial Average), but found that this variable had no predictive power for consumption growth or income growth.¹¹ Results using lagged changes in quarterly average 3-month nominal Treasury bill rates (i_t) were more successful, and we report these in rows 6 and 7 of Table 1. The instruments are jointly significant for consumption growth at the 1.0% and 0.2% levels. The parameter λ is estimated at 0.698 in row 6 (significant at the 0.3% level), and at 0.584 in row 7 (significant at better than the 0.01% level).¹²

The final two rows of the table report restricted error-correction

¹¹ This finding contrasts with the positive results for stock prices reported by Hall (1978) and others. Yet close inspection of Hall's stock price regression (his equation (8), on p.984) suggests that almost all the explanatory power comes from the first lagged stock price change. When we include the first lag, we also find strong predictive power from stock price changes; but for the reasons discussed above, we regard this as an illegitimate test of the permanent income model.

¹² The spread between the yield on a long-term government bond and that on a three-month Treasury bill also provided a useful instrument. Using only the second lag of the yield spread, we obtained adjusted R^2 's of 0.094 for Δc and 0.048 for Δy , and an estimate of λ of 0.741 with a standard error of 0.235.

models for consumption and income. Row 8 has lags of consumption growth, income growth and the log consumption-income ratio as instruments; row 9 adds lagged interest rate changes. The results are broadly consistent with those in earlier rows.

Table 1 also tests the overidentifying restrictions of our model (1.4) on the unrestricted system (1.5). The test results are reported in the last column of the table. There is no evidence against our restrictions anywhere in this column.

Figures 1 and 2 illustrate what's going on in these instrumental variables estimates. Figure 1 is a scatterplot of ex post consumption growth against ex post income growth. The figure shows a positive relation, but not a tight one. Figure 2 is a scatterplot of expected consumption growth against expected income growth, where expectations were taken to be the fitted values from the reduced form equations estimated in row 9 of Table 1. Note that these points lie along a distinct line. In contrast to the permanent income hypothesis, expected increases in income are associated with expected increases in consumption.

The two lines shown in the figure are estimated by IV regression of Δc_t on Δy_t , as reported in Table 1, and by the reverse IV regression of Δy_t on Δc_t . It is apparent that the normalization of the IV regression makes little difference to the estimate of the slope λ ; this is what we would expect to find if our model is correctly specified and the true slope is not zero or infinite.¹³

¹³ Nelson and Startz (1988) point out that there are severe problems with the IV regression approach if the instruments do not forecast the right hand side variable. In our framework, this would occur in the IV regression of consumption growth on income growth if λ is infinite, and in the IV regression of income growth on consumption growth if λ is zero.

While the results in Table 1 follow most of the literature by examining consumer spending on non-durables and services, we have also examined two measures of consumption that include consumer durable goods. The results are potentially sensitive to the treatment of durable goods, because spending on them is so volatile. We therefore estimated equation (1.4) both using total consumer spending and using the sum of spending on non-durables and services and the imputed rent on the stock of consumer durables.¹⁴ The results obtained with these two measures turned out to be similar to those reported in Table 1.

In summary, we have found striking evidence against the permanent income hypothesis. The results from our instrumental variables test are particularly unfavorable to the permanent income model. When we use instruments that are jointly significant for predicting income growth at the 5% level or better, we get estimates of λ , the fraction of the population that consumes its current income, of about 0.5. The estimates are always strongly significant even though we have potentially lost some power by lagging the instruments two periods instead of one. The overidentifying restrictions of our model are not rejected at any reasonable significance level.

1.4. Evidence from Abroad

To examine the robustness of our findings for the United States, we now turn to examining data for several other countries. From various DRI

¹⁴ To calculate the stock of durables, we began with the Commerce Department's net stock of consumer durables for 1947 and then accumulated the spending flow assuming a depreciation rate of 5 percent per quarter. To calculate the imputed rent, we assumed a user cost of 6 percent per quarter

data banks, we obtained data on consumption and income to estimate equation (1.4) for the G-7 countries: Canada, France, Germany, Italy, Japan, the United Kingdom, and the United States.¹⁵

Two data issues arise. First, we found that long time series of quarterly consumption data are often available only for total spending, which includes spending on durables. Assuming exponential depreciation, however, durability should merely lead to the change in consumer spending being a first-order moving average process rather than white noise.¹⁶ Since we are using twice-lagged instruments, the inclusion of spending on durables does not change the implication of the permanent income hypothesis that forecastable changes in income should not lead to forecastable changes in consumption. We can therefore proceed as before.

The second data issue is that, for Canada, France, Italy, and Japan, we were unable to find a quarterly disposable personal income series and therefore used GDP as a proxy. The use of GDP to measure Y should still provide a valid test of the null hypothesis that the permanent income theory is correct. Yet real GDP is an imperfect proxy: in U.S. data, the correlation of real GDP growth and real disposable personal income growth is only 0.55. The use of this proxy can potentially reduce our test's power. It turns out, however, that loss of power appears not to be a problem.

Table 2 presents the estimates obtained for these seven countries.

¹⁵ Other studies that have used international data to test the permanent income hypothesis include Kormendi and LaHaye (1987) and Jappelli and Pagano (1988).

¹⁶ See Mankiw (1982). Matters become more complicated, however, if one allows more complicated forms of depreciation or the possibility of adjustment costs; see Heaton (1988).

The results from six of these seven countries tell a simple and consistent story. For Canada, France, Germany, Italy, Japan, and the United States, the estimate of the fraction of income going to rule-of-thumb consumers is significantly different from zero and not significantly different from 0.5. Moreover, the overidentifying restrictions imposed by our model are not rejected. The only exception is the United Kingdom, where neither the permanent income hypothesis nor our more general model appear to describe the data adequately. Taken as a whole, these results confirm the failure of the simple random-walk model for consumption and the apparent rule-of-thumb behavior of many consumers.

2. Consumption and the Real Interest Rate

The "random walk" theorem for consumption rests crucially on the assumption that the real interest rate is constant. Here we examine the Euler equation that allows for a varying and uncertain real interest rate.

There are two reasons we look at this extension of the basic model. First, a rejection of the theory might be attributable to the failure of this assumption, rather than to an important deviation from the permanent income hypothesis. In particular, variation through time in the real interest rate can make consumption appear excessively sensitive to income, even though individuals intertemporally optimize in the absence of borrowing constraints.¹⁷ We show, however, that the departure from the theory documented above--the apparent existence of rule-of-thumb consumers--is not an artifact of the assumed constancy of the real interest rate.

Second, we want to check whether Hall's (1988) conclusion that the

¹⁷ Michener (1984) makes this argument. See also Christiano (1987).

intertemporal elasticity of substitution is close to zero is robust to the presence of current-income consumers. Hall assumes that the underlying permanent income theory is correct and uses the absence of a relation between consumption growth and real interest rates as evidence for a small elasticity. In contrast, we argue that the underlying theory is not empirically valid. Unless one is willing to admit that a substantial fraction of income goes to rule-of-thumb consumers, the data cannot yield an answer on the intertemporal elasticity of substitution.

2.1. The Model with Only Permanent Income Consumers

We begin our examination of consumption and real interest rates by maintaining the hypothesis that the permanent income theory is correct. We will then go on to consider a more general model with some rule-of-thumb consumers.

The generalization of the consumer's Euler equation to allow for changes in the real interest rate is now well-known. The log-linear version of the Euler equation is¹⁸

$$(2.1) \quad \Delta c_t = \mu + \sigma r_t + \epsilon_t,$$

where r_t is the real interest rate contemporaneous with Δc_t , and as before the error term ϵ_t may be correlated with r_t but is uncorrelated with lagged

¹⁸ See, for example, Grossman and Shiller (1981), Mankiw (1981), Hansen and Singleton (1983), Hall (1988). Note that in the process of log-linearizing the first-order condition, the variance of consumption growth has been included in the constant term. Hence, heteroskedasticity is one possible reason for rejection of the model; see Barsky (1985) for a preliminary exploration of this issue.

variables. According to (2.1), high ex ante real interest rates should be associated with rapid growth of consumption. The coefficient on the real interest rate, σ , is the intertemporal elasticity of substitution.¹⁹

Equation (2.1) can be estimated using instrumental variables, just in the way we estimated equation (1.4). The nominal interest rate we use is the average three-month treasury bill rate over the quarter. The price index is the deflator for consumer non-durables and services. We assume a marginal tax rate on interest of 30%.

We obtained the results in Table 3. We find fairly small values for the coefficient on the real interest rate. Hall interprets evidence of this sort as indicating that the intertemporal elasticity of substitution is close to zero--that is, consumers are extremely reluctant to substitute intertemporally.

In our view, however, the equation estimated in Table 3 is misspecified because it does not allow for the presence of rule-of-thumb consumers. This misspecification shows up in several ways in Table 3. First, the hypothesis that consumption growth is unpredictable is rejected at the 1% level or better in 5 out of 8 rows of Table 3, and at the 5% level or better in 7 rows. This is inconsistent with Hall's interpretation of the data: if the permanent income theory were true and σ were zero, consumption should be a random walk. Second, the overidentifying restrictions of equation (2.1) are rejected at the 5% level or better whenever lagged real interest rates are included in the set of instruments.

¹⁹ If the representative agent has power utility, then σ is the reciprocal of the coefficient of relative risk aversion. Epstein and Zin (1987a, 1987b) and Giovannini and Weil (1989) have shown that the same Euler equation can be obtained in a more general model in which risk aversion and the intertemporal elasticity of substitution are decoupled.

Third, the estimates of σ are highly unstable; while they are generally small, they do exceed one when nominal interest rate changes are used as instruments.

Perhaps the most telling check on the specification comes from reversing the Hall IV regression. Table 4 shows the IV regression of the real interest rate on the change in consumption. We do not find that the estimates of $1/\sigma$ are extremely large, as would be predicted by the Hall hypothesis; instead, they cluster around one.²⁰

Figure 3 shows graphically why the results are so sensitive to normalization. We regressed Δc and r on the instruments in row 9 of Table 3 and then plotted the fitted values as estimates of the expected change in consumption and the real interest rate. The figure shows that there is substantial variation in these two variables over time. Yet contrary to the predictions of the theory, the fitted values do not lie along a line. The two lines in this figure correspond to the two regressions estimated with the two normalizations. Because the fitted values are not highly correlated, the estimated regression is crucially dependent on which variable is on the left-hand side. Hence, this scatterplot does not imply that the elasticity of substitution is small. Instead, it suggests that the model underlying the Euler equation (2.1) should be rejected.

2.2. Including Rule-of-Thumb Consumers

We now reintroduce our rule-of-thumb consumers into the model. That is, we consider a more general model in which a fraction λ of income goes

²⁰ This cannot be explained by small-sample problems of the Nelson and Startz (1988) variety, since consumption growth is fairly well predicted by the instruments in Table 3.

to individuals who consume their current income and the remainder goes to individuals who satisfy the general Euler equation (2.1). We estimate by instrumental variables

$$(2.2) \quad \Delta c_t = \mu + \lambda \Delta y_t + \theta r_t + \epsilon_t,$$

where $\theta = (1-\lambda)\sigma$. We thus include actual income growth and the ex post real interest rate in the equation, but instrument using twice lagged variables. The results are in Table 5.

The first implication of the results is that the rule-of-thumb consumers cannot be explained away by allowing for fluctuations in the real interest rate. The coefficient on current income remains substantively and statistically significant.

The second implication of the results in Table 5 is that there is no evidence that the ex ante real interest rate is associated with the growth rate of consumption after allowing for the rule-of-thumb consumers. The coefficient on the real interest rate is consistently less than its standard error. The small estimated coefficients on the real interest rate indicate that the intertemporal elasticity of substitution for the permanent income consumers is very small. In addition, there is no evidence of any misspecification of the sort found when the rule-of-thumb consumers were excluded. The overidentifying restrictions are never close to being rejected.

Figure 4 illustrates the finding of a small elasticity of substitution by plotting the expected real interest rate and the expected change in consumption for the permanent income consumers assuming $\lambda=0.5$. This figure

is exactly analogous to Figure 3, except that Δc has been replaced by $\Delta c - 0.5\Delta y$. These fitted values lie almost along a horizontal line, as is required for an elasticity near zero. The figure also includes the regression line of the expected consumption change on the expected real interest rate, and it is near horizontal. Note that we cannot estimate the reverse normalization: we have been unable to find any instruments that forecast $\Delta c - 0.5\Delta y$ (as must be the case if $\lambda = 0.5$ and $\sigma = 0$).

In summary, the data show little or no correlation between expected changes in consumption and ex ante real interest rates. Yet this finding should not be interpreted as implying that the permanent income hypothesis holds with a small intertemporal elasticity of consumption: that hypothesis would require that expected changes in consumption are small and linearly dependent on the ex ante real interest rate. Instead, it seems that expected changes in consumption are dependent on expected changes in income, which can be explained by the existence of some rule-of-thumb consumers. Once these rule-of-thumb consumers are admitted into the model, the data become consistent with an elasticity of substitution near zero for the permanent income consumers.

3. From Euler Equation to Consumption Function

Modern empirical work on consumption behavior has focused almost exclusively on the Euler equations implied by optimizing models of intertemporal choice. Our own work is no exception. Yet it seems that something has been lost in this change of emphasis. The Euler equation determines only the level of consumption today, relative to the level of consumption tomorrow. We would like to be able to determine the absolute

level of consumption, given either wealth and expected future interest rates, or expected future income flows and interest rates. For this we need a traditional consumption function, that is, a closed-form solution for consumption given exogenous variables.

Of course, there are considerable technical difficulties in deriving a consumption function from an optimizing model. In fact, closed-form solutions are available only in a very few special cases, the best-known being log utility or power utility with independently and identically distributed asset returns.²¹ The problem is that a closed-form solution is obtained by combining an Euler equation with the intertemporal budget constraint. But even when the Euler equation is linear or log-linear, the budget constraint is always non-linear when asset returns are random. Consumption is subtracted from wealth to give the amount invested, and this amount is then multiplied by a random rate of return to give tomorrow's level of wealth.

In this section we explore a class of approximate consumption functions obtained by log-linearizing the intertemporal budget constraint. These approximate consumption functions give considerable insight into the implications of alternative models, and they offer an alternative way to confront the models with the data.²²

²¹ See Samuelson (1969) or Ingersoll (1987).

²² Our log-linearization is similar to the one used by Campbell and Shiller (1988) to study stock prices, dividends and discount rates. It differs slightly because we define wealth inclusive of today's consumption, which is analogous to a cum-dividend asset price. There is also an interesting parallel between our approach and the continuous-time model of Merton (1971). Merton was able to ignore the product of random returns and consumption flows, since this becomes negligible in continuous time. See also Hayashi (1982), who examines a similar model under the maintained

3.1. The intertemporal budget constraint

To see the way our approach works, consider the budget constraint of a consumer who invests his wealth in a single asset with a time-varying risky return R_t . We do not explicitly model income at this stage; this is legitimate provided that all the consumer's income flows (including his labor income) are capitalized into marketable wealth. The period-by-period budget constraint is

$$(3.1) \quad W_{t+1} = R_{t+1}(W_t - C_t).$$

Solving forward with an infinite horizon and imposing the transversality condition that the limit of discounted future wealth is zero, we obtain

$$(3.2) \quad W_t = C_t + \sum_{i=1}^{\infty} C_{t+i} / \left(\prod_{j=1}^i R_{t+j} \right).$$

This equation says that today's wealth equals the discounted value of all future consumption.

We would like to approximate the nonlinear equations (3.1) and (3.2) in such a way that we obtain linear relationships between log wealth, log consumption, and log returns, measured at different points of time. To do this, we first divide equation (3.1) by W_t , take logs and rearrange. The resulting equation expresses the growth rate of wealth as a nonlinear function of the log return on wealth and the log consumption-wealth ratio.

assumption of a constant real interest rate.

In the Appendix we show how to linearize this equation using a Taylor expansion. We obtain

$$(3.3) \quad \Delta w_{t+1} \approx k + r_{t+1} + (1-1/\rho)(c_t - w_t).$$

In this equation lower-case letters are used to denote the logs of the corresponding upper-case letters. The parameter ρ is a number a little less than one, and k is a constant.²³ This equation says that the growth rate of wealth is a constant, plus the log return on wealth, less a small fraction $(1-1/\rho)$ of the log consumption-wealth ratio. In the Appendix we solve equation (3.3) forward to obtain

$$(3.4) \quad c_t - w_t = \sum_{j=1}^{\infty} \rho^j (r_{t+j} - \Delta c_{t+j}) + \rho k / (1-\rho).$$

Equation (3.4) is a log-linear version of the infinite-horizon budget constraint (3.2). It states that a high log consumption-wealth ratio today must be associated either with high future rates of return on invested wealth, or with low future consumption growth.

3.2. Wealth-based and income-based consumption functions

So far we have merely manipulated a budget constraint, without stating any behavioral restrictions on consumer behavior. We now assume that the consumer satisfies the log-linear Euler equation discussed earlier in Section 2:

²³ The parameter ρ can also be interpreted as the average ratio of invested wealth, $W-C$, to total wealth, W .

$$(3.5) \quad E_t \Delta c_{t+1} = \mu + \sigma E_t r_{t+1}.$$

Equation (3.5) can be combined with equation (3.4) to give a consumption function relating consumption, wealth, and expected future returns on wealth. Take conditional expectations of equation (3.4), noting that the left hand side is unchanged because it is in the consumer's information set at time t . Then substitute in for expected consumption growth from (3.5). The resulting expression is

$$(3.6) \quad c_t - w_t = (1-\sigma) E_t \sum_{j=1}^{\infty} \rho^j r_{t+j} + \rho(k-\mu)/(1-\rho).$$

This equation generalizes Paul Samuelson's (1969) results for independent and identically distributed asset returns. It says that the log consumption-wealth ratio is a constant, plus $(1-\sigma)$ times the expected present value of future interest rates, discounted at the rate ρ . When $\sigma = 1$, the consumer has log utility and we get the well-known result that consumption is a constant fraction of wealth. When $\sigma > 1$, an increase in interest rates lowers the log consumption-wealth ratio because substitution effects outweigh income effects; when $\sigma < 1$, income effects are stronger and high interest rates increase consumption. Whatever the sign of the effect, persistent movements in interest rates have a stronger impact on the level of consumption than transitory movements do.

Traditional macroeconomic consumption functions usually determine consumption in relation to income flows rather than wealth. We can move from the wealth-based consumption function (3.6) to an income-based

consumption function by expressing the market value of wealth in terms of future expected returns and the future expected income flows from wealth. A full derivation is given in the Appendix. The resulting consumption function is

$$(3.7) \quad c_t - y_t = E_t \sum_{j=1}^{\infty} \rho^j (\Delta y_{t+j} - \sigma r_{t+j}) - \rho\mu/(1-\rho),$$

where y_{t+j} is the income at time $t+j$ generated by the wealth held at time t . The log consumption-income ratio depends on the expected present value of future income growth, less σ times the expected present value of future interest rates. As σ falls towards zero, interest rates have less and less effect on the consumption-income ratio and the model becomes a log-linear version of the standard permanent income model which ignores interest rate variation.

Two aspects of (3.7) are worthy of special mention. First, the interest rate terms in (3.7) capture the effects of changes in interest rates holding future income constant (while the market value of wealth is allowed to vary). By contrast, the interest rate terms in (3.4) capture the effects of changes in interest rates holding wealth constant (while future income is allowed to vary). When one holds future income constant, higher interest rates lower the market value of wealth; when one holds the market value of wealth constant, higher interest rates increase future income flows. As Lawrence Summers (1981) has emphasized, higher interest rates reduce consumption more when income flows are held fixed, since there is no positive income effect to offset the negative substitution effect of interest rates on consumption. With fixed income flows, the impact of

interest rates on consumption approaches zero as σ approaches zero.

Second, the income growth terms in (3.7) represent the influence of expected growth in income on current wealth, that is, net of the effects of further wealth accumulation. This complicates the use of (3.7) in empirical work, although the component of measured income growth that is due to wealth accumulation may be small in practice.²⁴

The analysis of this section has so far ignored the possibility that some fraction λ of income accrues to individuals who consume their current income rather than obeying the consumption function (3.7). But it is straightforward to generalize (3.7) to allow for these consumers. We obtain

$$(3.8) \quad c_t - y_t = (1-\lambda) E_t \sum_{j=1}^{\infty} \rho^j (\Delta y_{t+j} - \sigma r_{t+j}) - (1-\lambda)\rho\mu/(1-\rho).$$

The presence of current-income consumers reduces the variability of the log consumption-income ratio. The model of Hall (1988) sets $\sigma = \lambda = 0$ and thus has the consumption-income ratio responding fully to expected income growth but not at all to expected interest rates. By contrast, our model with $\lambda = 0.5$ has a reduced response of the consumption-income ratio to expected future income growth.

3.3. Empirical implementation

Since equation (3.8) shows that both the permanent income model and our more general model with rule-of-thumb consumers can be written as a

²⁴ For a discussion of this issue, see Flavin (1981).

present value relation, all the econometric techniques available for examining present value relations can be used to test and estimate these models. Applying these techniques is beyond the scope of this paper. To see what such exercises are likely to find, however, we take an initial look at the data from the perspective of this present value relation.

If we assume the intertemporal elasticity of substitution is small and set $\sigma = 0$, equation (3.8) says that the log of the average propensity to consume ($c-y$) is the optimal forecast of the present value of future income growth. To see if in fact there is any relation between these variables, Figure 5 plots the log of the average propensity to consume (computed using spending on non-durables and services) and the present value of realized income growth (computed using personal disposable income per capita). We assume a quarterly discount factor of 0.99, and set the out-of-sample income growth rates at the sample mean. As the theory predicts, the figure shows a clear positive relationship between these variables. When consumption is high relative to current income, income will tend to grow faster than average. When consumption is low relative to current income, income will tend to grow slower than average.²⁵

We can obtain an estimate of λ , the fraction of income going to rule-of-thumb consumers, by regressing the present value of realized income growth on the log of the average propensity to consume. Since the error in this relationship is an expectations error, it should be uncorrelated with currently known variables--in particular, $c-y$. The coefficient on $c-y$ is therefore a consistent estimate of $1/(1-\lambda)$. We can see from Figure 5 that

²⁵ This figure thus confirms the findings using vector autoregressions in Campbell (1987).

the estimate is likely to be greater than one: the present value of future income growth seems to respond more than one-for-one to fluctuations in c_y , which suggests that λ is greater than zero.

Table 6 shows the regression results for three measures of consumption: spending on non-durables and services, total consumer spending, and the sum of spending on non-durables and services and the imputed rent on the stock of consumer durables. We present the results with and without a time trend.²⁶ The implied estimates of λ in Table 6 vary from 0.233 to 0.496, which are similar to those obtained in Table 1.²⁷ These findings lead us to believe that more sophisticated examinations of the present value relation will likely yield a conclusion similar to the one we reached examining the Euler equation: a model with some permanent income consumers and some rule-of-thumb consumers best fits the data.

4. Conclusions

We have argued that aggregate consumption is best viewed as generated not by a single representative consumer but rather by two groups of consumers--one consuming their permanent income and the other consuming

²⁶ We include a time trend to proxy for mismeasurement in the average propensity to consume due to the treatment of consumer durables. The ratio of spending on consumer durables to spending on consumer non-durables and services has grown over time. Therefore, a failure to include consumer durables or an incorrect imputation is likely to cause mismeasurement in c_y that is correlated with time. We confess that inclusion of a time trend is a crude correction at best.

²⁷ We have somewhat more confidence in the estimates of λ obtained from Euler equation estimation. In Table 6, measurement error in consumption biases downward the estimate of λ (as does the inability to observe the out-of-sample values of future income growth.) Yet such measurement error does not affect the Euler equation estimates if this measurement error is uncorrelated with the instruments.

their current income. We have estimated that each group of consumers receives about fifty percent of income and that the intertemporal elasticity of substitution for the permanent income consumers is close to zero. This alternative model can explain why expected growth in consumption accompanies expected growth in income, why expected growth in consumption is unrelated to the expected real interest rate, and why periods in which consumption is high relative to income are typically followed by high growth in income.

Our model also has the potential to explain the "excess smoothness" of aggregate consumption pointed out by Angus Deaton (1987).²⁸ Deaton shows that if income follows a persistent time series process, then the variance of the innovation in permanent income exceeds the variance of the change in current income. According to the permanent income model, the change in consumption should then be more variable than the change in income; but in fact consumption is considerably smoother than income. Our model can resolve this puzzle because it makes the change in consumption a weighted average of the change in current income and the change in permanent income. If these two income changes are not perfectly correlated, then a weighted average of them can be less variable than either one considered in isolation. Aggregate consumption is smooth in our model because it is a "diversified portfolio" of the consumption of two groups of agents.²⁹

²⁸ See also Campbell and Deaton (1989), Christiano (1987), and West (1988).

²⁹ As an example, consider the case in which income is a random walk but is known one period in advance. In this case, since the change in permanent income and the change in current income are contemporaneously uncorrelated, our model implies that the variance of the change in consumption will be one half the variance of the change in income. For more discussion of excess smoothness in our model, see the 1989 version of

Although our emphasis in this paper has been on characterizing the aggregate data rather than on analyzing economic policies, our findings are suggestive regarding the effects of policies. In particular, if current income plays as central a role in consumption as our alternative model suggests, economists should not turn so readily to the permanent income hypothesis for policy analysis. An important application of this conclusion is in the debate over the national debt. Since the Ricardian equivalence proposition relies on the permanent income hypothesis, the failure of the permanent income hypothesis casts doubt on this proposition's empirical validity. Rule-of-thumb consumers are unlikely to increase private saving and bequests in response to government deficits. The old-fashioned Keynesian consumption function may therefore provide a better benchmark for analyzing fiscal policy than does the model with infinitely-lived consumers.

Our alternative model with rule-of-thumb consumers is very different from the alternative models considered in much recent work on Ricardian equivalence.³⁰ Those alternatives are forward-looking, but involve finite horizons or wedges between the interest rates that appear in private sector and government budget constraints. We believe that such effects may be present, but are hard to detect because they have much more subtle effects than the rule-of-thumb behavior we document here. Thus, the tests in the

Campbell and Mankiw (1987).

³⁰ For example, see Evans (1988), which tests Ricardian equivalence within the framework of Blanchard (1985).

literature may have low power.³¹

The failures of the representative consumer model documented here are in some ways unfortunate. This model held out the promise of an integrated framework for analyzing household behavior in financial markets and in goods markets. Yet the failures we have discussed are not unique. The model is also difficult to reconcile with the large size of the equity premium, the cross-sectional variation in asset returns, and time-series fluctuations in the stock market.³² The great promise of the representative consumer model has not been realized.

One possible response to these findings is that the representative consumer model examined here is too simple. Some researchers have been attempting to model the aggregate time series using a representative consumer model with more complicated preferences. Non-time-separabilities and departures from the von Neumann-Morgenstern axioms are currently receiving much attention.³³ It is also possible that there are non-separabilities between non-durables and services consumption and other

³¹ An exception is the study by David Wilcox (1989) which reports that consumer spending rises when Social Security benefits are increased. This finding provides evidence against the infinite-horizon model of the consumer. Moreover, since these benefit increases were announced in advance, this finding also provides evidence against models with forward-looking, finite-horizon consumers.

³² See Mehra and Prescott (1985), Mankiw and Shapiro (1986), and Campbell and Shiller (1988).

³³ For models with non-time-separability, see Constantinides (1988) and Heaton (1988). For departures from the von Neumann-Morgenstern axioms, see Epstein and Zin (1987a, 1987b) and Giovannini and Weil (1989).

contemporaneous variables.³⁴

Alternatively, some have argued that random shocks to the representative consumer's utility function may be important.³⁵ This contrasts with the standard assumption in the consumption literature that fluctuations arise from shocks to other equations, such as productivity shocks or changes in monetary and fiscal policy. If there are shocks to the utility function and if they are serially correlated, then they enter the residual of the Euler equation and may be correlated with lagged instruments, invalidating standard test procedures.³⁶

Unlike our model with rule-of-thumb consumers, these approaches remain in the spirit of the permanent income hypothesis by positing forward-looking consumers who do not face borrowing constraints. We believe that such modifications of the standard model are worth exploring, but we doubt that they will ultimately prove successful. We expect that the simple model presented here--half of income going to permanent income consumers

³⁴ In Campbell and Mankiw (1987), we looked at cross-effects with labor supply, government spending, and durable goods; we found no evidence for these types of non-separabilities. There is perhaps more evidence for non-separability with the stock of real money balances; see Koenig (1989). Nason (1988) proposes a model in which the marginal utility of consumption depends on current income. His model is observationally equivalent to ours, and has the same implications for policy; it is a way to describe the same facts in different terms.

³⁵ See Garber and King (1983) and Hall (1986).

³⁶ One response to this point is to try to find instruments that are uncorrelated with taste shocks. We have experimented with several instrument sets, including lagged growth of defense spending and political party dummies, but these did not have much predictive power for income. On the other hand, the change in the relative price of oil had significant predictive power two quarters ahead. When we used lags 2 through 6 as instruments, we estimated the fraction of current income consumers to be 0.28 with a standard error of 0.09. These instruments, however, did not have significant predictive power for real interest rates, so we were unable to estimate the more general Euler equation.

and half going to current income consumers--will be hard to beat as a description of the aggregate data on consumption, income, and interest rates.

Appendix: Derivation of Approximate Consumption Functions

We first divide equation (3.1) by W_t and take logs. The resulting equation is

$$(A.1) \quad w_{t+1} - w_t = r_{t+1} + \log(1 - C_t/W_t) - r_{t+1} + \log(1 - \exp(c_t - w_t)).$$

The last term in equation (A.1) is a nonlinear function of the log consumption-wealth ratio, $c_t - w_t = x_t$. The next step is to take a first-order Taylor expansion of this function, $\log(1 - \exp(x_t))$, around the point $x_t = x$. The resulting approximation is

$$(A.2) \quad \log(1 - \exp(c_t - w_t)) \approx k + (1 - 1/\rho)(c_t - w_t),$$

where the parameter $\rho = 1 - \exp(x)$, a number a little less than one, and the constant $k = \log(\rho) - (1 - 1/\rho)\log(1 - \rho)$. The parameter ρ can also be interpreted as the average ratio of invested wealth, $W - C$, to total wealth, W . Substituting (A.2) into (A.1), we obtain (3.3).

The growth rate of wealth, which appears on the left hand side of equation (3.3), can be written in terms of the growth rate of consumption and the change in the consumption-wealth ratio:

$$(A.3) \quad \Delta w_{t+1} = \Delta c_{t+1} + (c_t - w_t) - (c_{t+1} - w_{t+1}).$$

Substituting (A.3) into (3.3) and rearranging, we get a difference equation relating the log consumption-wealth ratio today to the interest rate, the consumption growth rate, and the log consumption-wealth ratio tomorrow:

$$(A.4) \quad c_t - w_t = \rho(r_{t+1} - \Delta c_{t+1}) + \rho(c_{t+1} - w_{t+1}) + \rho k.$$

Solving forward, we obtain (3.4).

To obtain an income-based consumption function, we suppose that total wealth W_t consists of N_t shares, each with ex-dividend price P_t and dividend payment Y_t in period t :

$$(A.5) \quad W_t = N_t(P_t + Y_t).$$

The return on wealth can be written as

$$(A.6) \quad R_{t+1} = (P_{t+1} + Y_{t+1})/P_t.$$

Combining (A.5) and (A.6) and rearranging, we get

$$(A.7) \quad W_{t+1}/N_{t+1} = R_{t+1}(W_t/N_t - Y_t),$$

where $W_t/N_t = P_t + Y_t$ is the cum-dividend share price at time t . This equation is in the same form as (3.1) and can be linearized in the same way. The log-linear model is

$$(A.8) \quad y_t - w_t = -n_t + E_t \sum_{j=1}^{\infty} \rho^j (r_{t+j} - \Delta y_{t+j}) + \rho k/(1-\rho).$$

(Implicitly we are assuming that the mean dividend-price ratio equals the mean consumption-wealth ratio since the same parameter ρ appears in (A.8) and in (3.4)). Normalizing $N_t=1$ ($n_t=0$) and substituting (A.8) into (3.6), we obtain (3.7).

Table 1
United States 1953-1986

$$\Delta c_t = \mu + \lambda \Delta y_t$$

Row	Instruments	First-stage regressions Δc equation	Δy equation	λ estimate (s.e.)	Test of restrictions
1	None (OLS)	---	---	0.316 (0.040)	---
2	$\Delta y_{t-2}, \dots, \Delta y_{t-4}$	-0.005 (0.500)	0.009 (0.239)	0.417 (0.235)	-0.022 (0.944)
3	$\Delta y_{t-2}, \dots, \Delta y_{t-6}$	0.017 (0.209)	0.026 (0.137)	0.506 (0.176)	-0.034 (0.961)
4	$\Delta c_{t-2}, \dots, \Delta c_{t-4}$	0.024 (0.101)	0.045 (0.028)	0.419 (0.161)	-0.009 (0.409)
5	$\Delta c_{t-2}, \dots, \Delta c_{t-6}$	0.081 (0.007)	0.079 (0.007)	0.523 (0.131)	-0.016 (0.572)
6	$\Delta i_{t-2}, \dots, \Delta i_{t-4}$	0.061 (0.010)	0.028 (0.082)	0.698 (0.235)	-0.016 (0.660)
7	$\Delta i_{t-2}, \dots, \Delta i_{t-6}$	0.102 (0.002)	0.082 (0.006)	0.584 (0.137)	-0.025 (0.781)
8	$\Delta y_{t-2}, \dots, \Delta y_{t-4},$ $\Delta c_{t-2}, \dots, \Delta c_{t-4},$ $c_{t-2} - y_{t-2}$	0.007 (0.341)	0.068 (0.024)	0.351 (0.119)	-0.033 (0.840)
9	$\Delta y_{t-2}, \dots, \Delta y_{t-4},$ $\Delta c_{t-2}, \dots, \Delta c_{t-4},$ $\Delta i_{t-2}, \dots, \Delta i_{t-4},$ $c_{t-2} - y_{t-2}$	0.078 (0.026)	0.093 (0.013)	0.469 (0.106)	-0.029 (0.705)

Note: The columns labelled "First-stage regressions" report the adjusted R^2 for the OLS regressions of the two variables on the instruments; in parentheses is the p-value for the null that all the coefficients except the constant are zero. The column labelled " λ estimate" reports the IV estimate of λ and, in parentheses, its standard error. The column labelled "Test of restrictions" reports the adjusted R^2 of the OLS regression of the residual on the instruments; in parentheses is the p-value for the null that all the coefficients are zero.

Table 2
Evidence from Abroad

$$\Delta c_t = \mu + \lambda \Delta y_t$$

	Country (sample period)	First-stage regressions		λ estimate (s.e.)	Test of restrictions
		Δc equation	Δy equation		
1	Canada (1963-1986)	0.047 (0.127)	0.090 (0.030)	0.616 (0.215)	0.007 (0.263)
2	France (1970-1986)	0.083 (0.091)	0.155 (0.015)	1.095 (0.341)	-0.055 (0.714)
3	Germany (1962-1986)	0.028 (0.211)	0.086 (0.031)	0.646 (0.182)	-0.030 (0.639)
4	Italy (1973-1986)	0.195 (0.013)	0.356 (0.000)	0.400 (0.094)	-0.034 (0.488)
5	Japan (1959-1986)	0.087 (0.020)	0.205 (0.000)	0.553 (0.096)	0.018 (0.178)
6	United Kingdom (1957-1986)	0.092 (0.012)	0.127 (0.002)	0.221 (0.153)	0.086 (0.010)
7	United States (1953-1986)	0.040 (0.092)	0.079 (0.014)	0.478 (0.158)	0.004 (0.269)

Note: For all countries, the consumption data are total spending. The set of instruments is: $\Delta y_{t-2}, \dots, \Delta y_{t-4}, \Delta c_{t-2}, \dots, \Delta c_{t-4}, c_{t-2} \cdot y_{t-2}$. Also see note, Table 1.

Table 3
United States, 1953-1986

$$\Delta c_t = \mu + \sigma r_t$$

Row	Instruments	First-stage regressions Δc equation	r equation	σ estimate (s.e.)	Test of restrictions
1	None (OLS)	---	---	0.276 (0.079)	---
2	r_{t-2}, \dots, r_{t-4}	0.063 (0.009)	0.431 (0.000)	0.270 (0.118)	0.031 (0.029)
3	r_{t-2}, \dots, r_{t-6}	0.067 (0.014)	0.426 (0.000)	0.281 (0.118)	0.034 (0.050)
4	$\Delta c_{t-2}, \dots, \Delta c_{t-4}$	0.024 (0.101)	-0.021 (0.966)	-0.707 (2.586)	0.000 (0.215)
5	$\Delta c_{t-2}, \dots, \Delta c_{t-6}$	0.081 (0.007)	0.007 (0.316)	0.992 (0.478)	0.008 (0.189)
6	$\Delta i_{t-2}, \dots, \Delta i_{t-4}$	0.061 (0.010)	0.024 (0.105)	1.263 (0.545)	-0.021 (0.918)
7	$\Delta i_{t-2}, \dots, \Delta i_{t-6}$	0.102 (0.002)	0.028 (0.119)	1.213 (0.445)	-0.022 (0.700)
8	$r_{t-2}, \dots, r_{t-4},$ $\Delta c_{t-2}, \dots, \Delta c_{t-4},$	0.062 (0.026)	0.455 (0.000)	0.204 (0.114)	0.047 (0.033)
9	$r_{t-2}, \dots, r_{t-4},$ $\Delta c_{t-2}, \dots, \Delta c_{t-4},$ $\Delta i_{t-2}, \dots, \Delta i_{t-4}$	0.103 (0.006)	0.476 (0.000)	0.150 (0.111)	0.100 (0.005)

Note: See Table 1.

Table 4

United States, 1953-1986

$$r_t = \mu + 1/\sigma \Delta c_t$$

Instrument	First-stage regressions Δc equation	r equation	$1/\sigma$ estimate (s.e.)	Test of restrictions
None (OLS)	---	---	0.304 (0.087)	---
r_{t-2}, \dots, r_{t-4}	0.063 (0.009)	0.431 (0.000)	1.581 (0.486)	0.086 (0.001)
r_{t-2}, \dots, r_{t-6}	0.067 (0.014)	0.426 (0.000)	1.347 (0.390)	0.113 (0.001)
$\Delta c_{t-2}, \dots, \Delta c_{t-4}$	0.024 (0.101)	-0.021 (0.966)	-0.341 (0.428)	-0.021 (0.878)
$\Delta c_{t-2}, \dots, \Delta c_{t-6}$	0.081 (0.007)	0.007 (0.316)	0.419 (0.258)	-0.010 (0.440)
$\Delta i_{t-2}, \dots, \Delta i_{t-4}$	0.061 (0.010)	0.024 (0.105)	0.768 (0.334)	-0.021 (0.919)
$\Delta i_{t-2}, \dots, \Delta i_{t-6}$	0.102 (0.002)	0.028 (0.119)	0.638 (0.249)	-0.024 (0.747)
$r_{t-2}, \dots, r_{t-4},$ $\Delta c_{t-2}, \dots, \Delta c_{t-4}$	0.062 (0.026)	0.455 (0.000)	1.034 (0.333)	0.236 (0.000)
$r_{t-2}, \dots, r_{t-4},$ $\Delta c_{t-2}, \dots, \Delta c_{t-4},$ $\Delta i_{t-2}, \dots, \Delta i_{t-4}$	0.103 (0.006)	0.476 (0.000)	0.521 (0.220)	0.455 (0.000)

Note: See Table 1.

Table 5
United States, 1953-1986

$$\Delta c_t = \mu + \lambda \Delta y_t + \theta r_t$$

Row	Instruments	First-stage regressions			λ (s.e.)	θ (s.e.)	Test of restrictions
		Δc	Δy	r			
1	None (OLS)	---	---	---	0.294 (0.041)	0.150 (0.070)	---
2	$\Delta y_{t-2}, \dots, \Delta y_{t-4}$ r_{t-2}, \dots, r_{t-4}	0.045 (0.061)	0.030 (0.125)	0.471 (0.000)	0.438 (0.189)	0.080 (0.123)	-0.010 (0.441)
3	$\Delta c_{t-2}, \dots, \Delta c_{t-4}$ r_{t-2}, \dots, r_{t-4}	0.062 (0.026)	0.046 (0.060)	0.455 (0.000)	0.467 (0.152)	0.089 (0.110)	-0.006 (0.391)
4	$\Delta i_{t-2}, \dots, \Delta i_{t-4}$ r_{t-2}, \dots, r_{t-4}	0.092 (0.005)	0.034 (0.106)	0.431 (0.000)	0.657 (0.212)	0.016 (0.146)	-0.022 (0.665)

Note: See Table 1.

Table 6

United States, 1953-1986

$$\sum_{j=1}^{\infty} \rho^j \Delta y_{t+j} = \mu + [1/(1-\lambda)](c_t - y_t)$$

<u>Consumption Measure</u>	<u>1/(1-λ)</u>	<u>time</u>	<u>R²</u>	<u>Implied λ</u>
Non-durables and Services	1.306 (0.233)		0.690	0.234
Non-durables and Services	1.983 (0.221)	0.0005 (0.0001)	0.792	0.496
Total Consumer Spending	1.455 (0.408)		0.302	0.313
Total Consumer Spending	1.303 (0.256)	-0.0004 (0.0002)	0.463	0.233
Non-durables, Services, and Imputed Rent on Durables	1.576 (0.225)		0.740	0.366
Non-durables, Services, and Imputed Rent on Durables	1.937 (0.203)	0.0003 (0.0001)	0.776	0.484

Note: These regressions were estimated using Ordinary Least Squares. The present value of future growth was computed assuming $\rho=0.99$; out-of-sample growth rates were set at the sample mean. Standard errors in parentheses were computed using the Newey-West (1987) correction for serial correlation; these standard errors use a lag length of 20, although lag lengths of 10 and 30 yielded similar results.

Bibliography

Barsky, Robert B., Three Interest Rate Paradoxes, PhD Dissertation, M.I.T., 1985.

Bean, Charles R., "The Estimation of 'Surprise' Models and the 'Surprise' Consumption Function," Review of Economic Studies 53: 497-516, August 1986.

Bernanke, Ben S., "Adjustment Costs, Durables and Aggregate Consumption", Journal of Monetary Economics 15:41-68, January 1985.

Blanchard, Olivier J., "Debt, Deficits, and Finite Horizons", Journal of Political Economy 93:223-247, April 1985.

Campbell, John Y., "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis", Econometrica 55:1249-1273, November 1987.

Campbell, John Y., and Angus S. Deaton, "Why is Consumption So Smooth?" Review of Economic Studies, April 1989, forthcoming.

Campbell, John Y. and N. Gregory Mankiw, "Permanent Income, Current Income, and Consumption", NBER Working Paper No. 2436, November 1987; revised January 1989.

Campbell, John Y. and Robert J. Shiller, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors", Review of Financial Studies 1:195-228, Fall 1988.

Christiano, Lawrence J., "Why Is Consumption Less Volatile than Income?", Federal Reserve Bank of Minneapolis Quarterly Review 2-20, Fall 1987.

Constantinides, George M., "Habit Formation: A Resolution of the Equity Premium Puzzle", unpublished paper, Graduate School of Business, University of Chicago, 1988.

Deaton, Angus S., "Life-Cycle Models of Consumption: Is the Evidence Consistent with the Theory?" in T.F. Bewley ed. Advances in Econometrics, Fifth World Congress, Vol. 2, Cambridge: Cambridge University Press, 1987.

Epstein, Larry G. and Stanley E. Zin (1987a), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Prices I: A Theoretical Framework", Econometrica, forthcoming.

Epstein, Larry G. and Stanley E. Zin (1987b), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Prices II: An Empirical Analysis", unpublished paper, 1987.

Evans, Paul, "Are Consumers Ricardian? Evidence for the United States", Journal of Political Economy 96:983-1004, October 1988.

Fisher, Franklin, "On the Cost of Approximate Specification in Simultaneous Equation Estimation," Econometrica 29, April 1961, 139-170.

Flavin, Marjorie A., "The Adjustment of Consumption to Changing Expectations about Future Income", Journal of Political Economy 89:974-1009, October 1981.

Garber, Peter M., and Robert G. King, "Deep Structural Excavation? A Critique of Euler Equation Methods," NBER Technical Paper 31, 1983.

Giovannini, Alberto, and Philippe Weil, "Risk Aversion and Intertemporal Substitution in the Capital Asset Pricing Model," NBER Working Paper 2824, 1989.

Grossman, Sanford J. and Robert J. Shiller, "The Determinants of the Variability of Stock Market Prices", American Economic Review 71:222-227, May 1981.

Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence", Journal of Political Economy 86:971-987, October 1978.

Hall, Robert E., "The Role of Consumption in Economic Fluctuations," The American Business Cycle: Continuity and Change, ed. Robert J. Gordon, Chicago: University of Chicago Press, 1986, 237-266.

Hall, Robert E., "Intertemporal Substitution in Consumption", Journal of Political Economy 96:339-357, April 1988.

Hansen, Lars P. and Kenneth J. Singleton, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns", Journal of Political Economy 91:249-265, April 1983.

Hayashi, Fumio, "The Permanent Income Hypothesis: Estimation and Testing by Instrumental Variables," Journal of Political Economy 90:895-916, October 1982.

Heaton, John, "The Interaction Between Time-Nonseparable Preferences and Time Aggregation", unpublished paper, University of Chicago, November 1988.

Ingersoll, Jonathan E., Jr., Theory of Financial Decision Making, Totowa: Rowman and Littlefield, 1987.

Jappelli, Tullio and Marco Pagano, "Consumption and Capital Market Imperfections: An International Comparison", CEPR Discussion Paper No. 244, June 1988.

Koenig, Evan, "Real Money Balances and the Timing of Consumption: An Empirical Investigation", Discussion Paper, Federal Reserve Bank of Dallas, March 1989.

Kormendi, Roger C. and Laura LaHaye, "Cross-Country Evidence on the Permanent Income Hypothesis", unpublished paper, University of Michigan, revised December 1987.

Lucas, Robert E., Jr., "Econometric Policy Evaluation: A Critique", in Carnegie-Rochester Conference Series on Public Policy 1:19-46, 1976.

Mankiw, N. Gregory, "The Permanent Income Hypothesis and the Real Interest Rate", Economics Letters 7:307-311, 1981.

Mankiw, N. Gregory, "Hall's Consumption Hypothesis and Durable Goods", Journal of Monetary Economics 10:417-426, November 1982.

Mankiw, N. Gregory and Matthew D. Shapiro, "Trends, Random Walks and Tests of the Permanent Income Hypothesis", Journal of Monetary Economics 16:165-174, 1985.

Mankiw, N. Gregory and Matthew D. Shapiro, "Risk and Return: Consumption Beta versus Market Beta", Review of Economics and Statistics 68:452-459, August 1986.

Mehra, Rajnish and Edward C. Prescott, "The Equity Premium: A Puzzle", Journal of Monetary Economics 15:145-161, 1985.

Merton, Robert C., "Optimum Consumption and Portfolio Rules in a Continuous-Time Model", Journal of Economic Theory 3:373-413, 1971.

Michener, Ron, "Permanent Income in General Equilibrium", Journal of Monetary Economics 13: 297-306, 1984.

Nason, James M., "Permanent Income, Current Income, Consumption, and Changing Tastes", unpublished paper, Board of Governors, Federal Reserve System, December 1988.

Nelson, Charles R. and Richard Startz, "The Distribution of the Instrumental Variables Estimator and its t-Ratio When the Instrument is a Poor One", Discussion Paper 88-07, University of Washington, May 1988.

Newey, Whitney, and Kenneth West, "A Simple, Positive Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix", Econometrica 55:703-708, 1987.

Samuelson, Paul A., "Lifetime Portfolio Selection by Dynamic Stochastic Programming", Review of Economics and Statistics 51:239-246, 1969.

Sargan, J.D., "The Estimation of Economic Relationships Using Instrumental Variables", Econometrica 26:393-415, 1958.

Summers, Lawrence H., "Capital Taxation and Accumulation in a Life Cycle Growth Model", American Economic Review 71:533-544, September 1981.

West, Kenneth D., "The Insensitivity of Consumption to News about Income," Journal of Monetary Economics 21:17-33, January 1988.

Wilcox, David, "Social Security Benefits, Consumption Expenditure, and the Life Cycle Hypothesis", Journal of Political Economy, April 1989.

Working, Holbrook, "Note on the Correlation of First Differences of Averages in a Random Chain", Econometrica 28:916-918, October 1960.

FIGURE 1

Scatterplot of Changes in Consumption and Income




FIGURE 2

Scatterplot of Expected Changes in Consumption and Income




FIGURE 3




FIGURE 4

Scatterplot of Expected Change in Consumption for "Permanent Income" Consumers and the Expected Real Interest Rate




FIGURE 5

The Average Propensity to Consume as a Forecast of Future Income Growth

