HOW DOES FUTURE INCOME AFFECT CURRENT CONSUMPTION?*

CHRISTOPHER D. CARROLL

This paper tests a straightforward implication of the basic Life Cycle model of consumption: that current consumption depends on expected lifetime income. The paper projects future income for a panel of households and finds that consumption is closely related to projected current income, but unrelated to predictable changes in income. However, future income *uncertainty* has an important effect: consumers facing greater income uncertainty consume less. The results are consistent with "buffer-stock" models of consumption like those of Deaton [1991] or Carroll [1992a, 1992b], where precautionary motives greatly reduce the willingness of prudent consumers to consume out of uncertain future income.

INTRODUCTION

Common versions of the Life Cycle and Permanent Income Hypothesis (LC/PIH) models[1] assert that current consumption depends in large part on the present discounted value of future income. This paper tests that proposition by constructing estimates of future income for a sample of households whose consumption is directly observed. The results provide no support for the proposition that current consumption is influenced by predictable changes in income. On the other hand, further investigation finds that the degree of *uncertainty* in future income does have an important effect: consumers with greater income uncertainty, *ceteris paribus,* have lower current consumption. These two results seem contradictory, because the first suggests that consumers ignore the future and the second indicates that they prudently prepare against future contingencies. I argue, however, that both results are consistent with models proposed by Deaton [1991] and Carroll [1992a, 1992b] in which saving serves primarily as a "buffer-stock" against uncertainty.

The results of the first part of the paper can be interpreted as microeconomic confirmation of the "excess smoothness" results of

*This is a substantially revised version of my job market paper, which was titled "Uncertain Future Income, Precautionary Saving, and Consumption." Thanks to Olivier Blanchard, Lawrence Katz, Miles Kimball, James Poterba, Lawrence Summers, Jonathan Gruber, David Weil, David Wilcox, the members of the M.I.T. Public Finance and Money Lunches and Money Workshop, and seminar participants at departments of economics around the country for constructive comments. The usual disclaimer applies.

1. For example, the model used by Flavin [1981], Campbell [1987], Campbell and Deaton [1989], and many others.

The Quarterly Journal of Economics, February 1994

West [1988], Campbell and Deaton [1989], Viard [1993], and others, who have shown that current aggregate consumption is unresponsive to changes in expected future aggregate income. This paper finds that current household consumption is apparently not influenced by expected future changes in household income. Instead, in keeping with findings by Campbell and Mankiw [1989], current consumption is strongly linked to predictable current income. The results also provide further and more careful evidence for the existence of the "consumption/income parallel" documented graphically by Carroll and Summers [1991].

The evidence for an effect of uncertainty on consumption complements the recent theoretical work by Zeldes [1989], Kimball [1990], Deaton [1991], and Carroll [1992a, 1992b] exploring the nature and consequences of precautionary saving behavior. In fact, a theoretical measure of the intensity of the precautionary saving motive proposed by Kimball, the Equivalent Precautionary Premium, turns out to be a better predictor of the amount of precautionary saving than traditional but atheoretical measures like the variance or standard deviation of income.

The rest of the paper is organized as follows. The first section of the paper describes the methods used to construct estimates of future income for consumers in the 1960–1961 *Consumer Expenditure Survey* (CEX)[2] on the basis of characteristics like age, occupation, and education. I then show that these methods produce good predictions, by comparing the predictions to actual outcomes using the *Panel Study of Income Dynamics* (PSID). The next section shows that these predictable changes in income appear to have no influence on current consumption—a clear rejection of the simple certainty-equivalence (CEQ) LC/PIH models that have been popular recently, especially in the macroeconomics literature (e.g., Campbell [1987], Flavin [1981], and Campbell and Deaton [1989]). I consider several possible simple extensions of the model, including explicit treatment of family demographic structure, allowance for heterogeneous tastes, and the existence of a simple form of liquidity constraints, but none of these modifications provides a satisfactory explanation for the failure of future income to affect current consumption.

2. The 1960–1961 CEX (rather than a later survey) was used for a variety of reasons, the most important of which was the need for forward-compatibility with the 1968–1985 data from the *Panel Study of Income Dynamics* (PSID). It was desirable that the PSID data lie in the future of the consumption data, and the 1960–1961 CEX is the nearest consumption survey before 1968.

The following section uses the PSID to construct several measures of income uncertainty, and then adds these measures to the regressions of consumption on estimated current and future income. The uncertainty measures are highly statistically significant and of economically important size. Calculations suggest that for the typical household the effect of income uncertainty on consumption is large: the most plausible estimates suggest that a one-standard-deviation increase in uncertainty decreases consumption from 3 to 5 percent.

The last section of the paper considers the implications of these results for various models of saving, and concludes that although the results are inconsistent with the most common modifications of the usual CEQ LC/PIH model, they are consistent with the buffer-stock models of Deaton [1991] and Carroll [1992a, 1992b].

I. Methods for Estimating Future Income

This section presents two methods for estimating the future income of consumers participating in the 1960–1961 *Consumer Expenditures Survey* (CEX). The first method uses CEX income data to estimate age/income cross-sectional profiles by occupation, education, and other household characteristics. A particular household's expected future income is then assumed to be given by the average observed income of older households with similar education and occupation. The second method regresses actual 1969–1985 income on 1968 personal characteristics using data from the PSID. The coefficients from this regression are then transported back to the CEX where they are used to construct the second measure of expected future income.

Method 1: Forward Projection of a Point-in-Time Cross-Section Profile

I assume that the disposable labor income[3] of individual i at age a in year t can be described by an equation of the form,

$$Y_{i,t} = Z_{i,t} B_{y,t} + V_{y,t,i}, \tag{1}$$

where $V_{y,i,t}$ is the idiosyncratic component of income for household i

3. Labor income should be construed here to mean income from sources exogenous to the household's decisions about saving behavior. In particular, it includes transfer income such as unemployment compensation, social security, and disability payments.

in year t, and Z is given by

$$Z_{i,t} = [D_{i,t}, D_{i,t}a_{i,t}, D_{i,t}a_{i,t}^2], \tag{2}$$

where D is a set of dummy variables indicating the education, occupation, and race of the household head, and $a_{i,t}$ indicates the age of household i in year t. This structure allows a different age/income profile for each combination of occupation, education, and race.[4] Under the usual assumptions, this equation can be estimated by OLS, yielding $b_{y,t} = (Z_t'Z_t)^{-1} Z_t'Y_t$ (which is a consistent estimate of $B_{y,t}$), a predicted value of $y_{i,t} = Z_{i,t}b_{y,t}$, and $v_{y,i,t} = Y_{i,t} - y_{i,t}$, an estimate of the idiosyncratic component of income.

Table I presents the results from estimating equation (1) using the 1960–1961 CEX.[5] The main statistic of interest in this table is the corrected R^2 of .24, which indicates that an important fraction of the variation in income across consumers can be attributed to differences in personal characteristics. Assuming that an equation similar to (1) is expected to hold in the future, household i's expected income j periods in the future can be represented as

$$E_t Y_{i,t+j} = E_t Z_{i,t+j}B_{y,t+j} + E_t V_{y,i,t+j}. \tag{3}$$

A variety of further assumptions must be made in order to use such an equation to project future income. First, the relationship between income and the dummy variables must be expected to remain stable; i.e., $E_tB_{y,t+j} = B_{y,t}$.[6] Second, it is necessary to know the expected future value of the independent variables $E_tZ_{i,t+j}$. If, for a given household, the dummy variables never change over time (i.e., the occupation, education, and race of the household head never change), then $E_tZ_{i,t+j}$ is given by the Z that would apply now to a consumer with dummy variables D identical to consumer i's but j years older:

$$E_t Z_{i,t+j} = [D_{i,t}, D_{i,t}(a_{i,t} + j), D_{i,t}(a_{i,t} + j)^2]. \tag{4}$$

It is also necessary to make an assumption about the future value of the idiosyncratic component of the income. If $v_{y,i,t}$ represents a purely transitory deviation of income from $y_{i,t}$, the level that

4. In practice, the regressions used did not include interactions of the race dummy with age and age^2 because the coefficients on those interactions were insignificant.

5. See the Data Appendix for a description of the construction of disposable labor income in the 1960–1961 CEX.

6. Aggregate productivity growth can be accommodated by adding a common growth factor to this equation.

TABLE I

DISPOSABLE LABOR INCOME AS A FUNCTION OF PERSONAL CHARACTERISTICS IN THE 1960–1961 CEX

Independent variable	Estimated coefficient	Standard error	t-statistic
Intercept	−1609	3085	−0.5
White	927	136	6.8
Age	333	143	2.3
Age^2	−4	2	−2.4
o2	−8469	3569	−2.4
o3	678	3714	0.2
o4	−11	3048	0.0
o5	−1341	3109	−0.4
o6	160	3176	0.1
o7	1176	3163	0.4
o8	−3743	3623	−1.0
e2	1472	2193	0.7
e3	97	2038	0.0
e4	−4679	2735	−1.7
e5	−13366	3258	−4.1
e6	−15181	4170	−3.6
Age × o2	394	165	2.4
Age × o3	2	170	0.0
Age × o4	−55	144	−0.4
Age × o5	35	146	0.2
Age × o6	−64	149	−0.4
Age × o7	−143	148	−1.0
Age × o8	103	164	0.6
Age × e2	1	99	0.0
Age × e3	60	92	0.6
Age × e4	323	126	2.6
Age × e5	764	152	5.0
Age × e6	800	197	4.1
Age^2 × o2	−3.4	1.8	−1.8
Age^2 × o3	0.1	1.9	0.1
Age^2 × o4	0.8	1.6	0.5
Age^2 × o5	0.0	1.6	0.0
Age^2 × o6	1.1	1.7	0.7
Age^2 × o7	1.7	1.7	1.0
Age^2 × o8	−1.2	1.8	−0.6
Age^2 × e2	−0.4	1.1	−0.3
Age^2 × e3	−0.7	1.0	−0.7
Age^2 × e4	−3.5	1.4	−2.5
Age^2 × e5	−8.1	1.7	−4.8
Age^2 × e6	−7.3	2.2	−3.3
Number of observations			8362
R^2			0.25
Corrected R^2			0.24

Note. o2 through o8 and e2 through e6 are dummy variables for membership in each of the eight occupational or six education-level groups defined in Table II.

would be predicted by the wage equation, $E_t v_{y,i,t+j}$ would equal zero. If, on the other hand, $v_{y,i,t}$ represents a permanent difference between the income-earning capacity of household i and other households who are observationally identical, then $v_{y,i,t+j}$ might be expected to equal $v_{y,i,t}$. In practice, the v term is likely to be a mixture of both transitory and permanent effects, in unknown proportions. Following Hubbard, Skinner, and Zeldes [1993], I shall assume that the idiosyncratic error follows an AR(1) process:

$$E_t v_{y,i,t+j} = \rho^j v_{y,i,t}, \tag{5}$$

where the serial correlation parameter $\rho = 0.95$. The 0.95 parameter value from Hubbard, Skinner, and Zeldes was consistent with regressions I performed in the PSID; none of the results presented below is very sensitive to the assumption about the magnitude of ρ.

Human wealth from ages t through T (where T is usually taken to be the age of retirement, 65) is then given by

$$\begin{aligned} H_{i,t} &= \left[\sum_{j=1}^{T-a_{i,t}} R^{-j} E_t Y_{i,t+j} \right] \\ H_{i,t} &= \sum_{j=1}^{T-a_{i,t}} R^{-j} E_t Z_{i,t+j} b_{y,t} + \sum_{j=1}^{T-a_{i,t}} R^{-j} E_t v_{y,i,t+j}, \end{aligned} \tag{6}$$

where R is the gross interest rate and $E_t Z_{i,t+j}$ is constructed as in equation (4). Defining $h_{i,t}$ as the demographically predictable component of human wealth, and $v_{h,i,t}$ as the expected future idiosyncratic component of human wealth, we have $H_{i,t} = h_{i,t} + v_{h,i,t}$, where $h_{i,t}$ and $v_{h,i,t}$ are constructed from

$$h_{i,t} = \sum_{j=1}^{T-a_{i,t}} R^{-j} E_t Z_{i,t+j} b_{y,t} = Z_{i,t} b_h \tag{7}$$

$$v_{h,i,t} = \sum_{j=1}^{T-a_{i,t}} R^{-j} \rho^i v_{y,i,t} = \frac{1-(\rho/R)^{T-a_{i,t}+1}}{1-(\rho/R)} v_{y,i,t}. \tag{8}$$

It is not difficult to show that the linear quadratic structure of Z implies that, across households of the same age $a_{i,t}$, the summation in equation (7) reduces to a linear function of the Z's, i.e., $h_{i,t} = Z_{i,t} b_h$, where b_h is a function of remaining lifetime. Equation (8) indicates that $v_{h,i,t}$ is proportional to $v_{y,i,t}$ with a proportionality factor that is also solely a function of the remaining lifetime; this is an implication of the AR(1) process assumed for $v_{y,i,t}$.

It would be desirable to test all of the assumptions required to make this method of constructing human wealth correct. Probably the most important such assumption is that consumers expect the currently observed age/disposable-labor-income profiles to remain stable over time. Although there is no way to know what consumers expected, it is possible to examine whether the actual age/income profiles remained stable over time. Temporal stability would imply that when equation (1) is estimated using the 1960–1961 CEX, producing a coefficient $b_{y,\mathrm{CEX60}}$, and then estimated again using, say, the 1968 wave of the PSID (the first year of the PSID survey), producing a coefficient $b_{y,\mathrm{PSID68}}$, then $b_{y,\mathrm{CEX60}}$ and $b_{y,\mathrm{PSID68}}$ should be "close" in some relevant sense.

Unfortunately, however, the 1968 wave of the PSID does not contain information on taxes paid by the household, and contains less information than the 1960–1961 CEX on the split of income between capital income and labor income, so it was not possible to construct a direct measure of disposable labor income for 1968 in the PSID consistent with the measure in the 1960–1961 CEX. Of course, both surveys have information on total household income, so it is possible to test whether the age/total-household-income profiles were similar. If there were no major changes in the tax code or in the distribution of income between labor and capital income between 1961 and 1968, similar age/total-household-income profiles should imply similar age/disposable-labor-income profiles.

My method for judging whether the two age/total-household-income profiles are "close" was to estimate equation (1) for both data sets using total household income as the measure of income, producing coefficients $b_{y,\mathrm{CEX60}}$ and $b_{y,\mathrm{PSID68}}$, and then to construct the demographically predictable component of future total-household-income from equations (7) and (4), using first $b_{y,t} = b_{y,\mathrm{CEX60}}$ and then $b_{y,t} = b_{y,\mathrm{PSID68}}$. I then divided projected total future income h by the number of years in which that income was to be earned, producing average annual total future income, h'. Finally, I took the ratio of projected average annual future income to projected current income, h'/y. Table II presents the results from those calculations for consumers in three age brackets, 25–34, 35–44, and 45–54; six education brackets; and eight occupational categories. The results in this table match intuition: for instance, both $b_{y,\mathrm{CEX60}}$ and $b_{y,\mathrm{PSID68}}$ predict that young highly educated consumers and young professionals should expect a high rate of income growth, and that those with less education, and laborers, can expect lower income growth. Figure I plots the data in Table

TABLE II
RATIO OF PREDICTED FUTURE INCOME TO CURRENT INCOME, BY OCCUPATION AND EDUCATION GROUPS

		Age 25–34		Age 35–44		Age 45–54	
		1960–1961 CEX	1968 PSID	1960–1961 CEX	1968 PSID	1960–1961 CEX	1968 PSID
Occupation	Professional, technical, and kindred	1.34	1.35	1.02	1.05	0.93	0.94
	Managers, officials, and proprietors	1.38	1.32	1.09	1.30	0.99	1.18
	Self-employed businessmen	1.20	0.88	1.02	0.93	0.94	0.88
	Clerical and sales	1.20	1.14	1.00	0.99	0.94	0.96
	Craftsmen, foremen, and kindred	1.17	1.17	1.02	0.98	0.96	0.91
	Operatives and kindred	1.14	1.10	1.02	1.03	0.98	1.00
	Laborers and service	1.09	1.15	0.96	0.91	0.93	0.85
	Farmers and farm managers	1.10	1.23	0.91	1.03	0.84	0.89
Education	Less than or equal to 8 years	1.14	1.20	1.00	1.00	0.94	0.91
	9–11 years	1.05	1.11	0.95	0.93	0.90	0.88
	12 years	1.16	1.14	1.02	1.01	0.96	0.95
	13–15 years	1.23	1.14	1.00	1.01	0.91	0.98
	16 years	1.44	1.36	1.03	1.25	0.90	1.14
	Greater than 16 years	1.53	1.53	1.15	1.15	1.00	1.05

Notes. The figures represent predicted average annual income in every future year until retirement divided by predicted current income. For instance, in both surveys predicted annual future income for households with postgraduate education who are currently aged 25–34 is 1.53 times current income (the bottom figure in the first two columns).

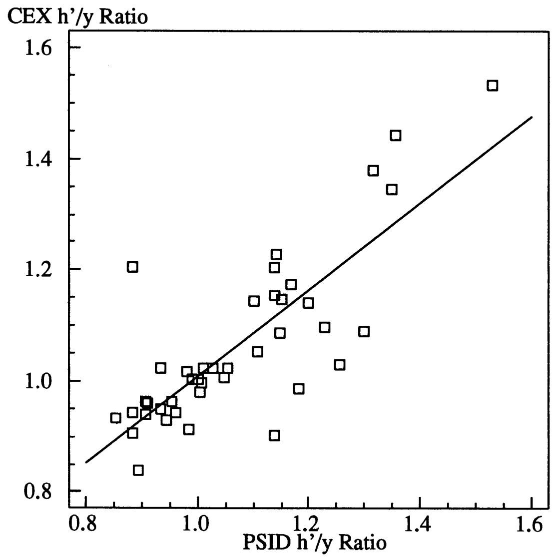

FIGURE I
Ratio of Future to Current Income by Group, Predicted Using PSID versus Using CEX

II.[7] For example, the dot in the upper right-hand corner is for households aged 25–34 in the highest education bracket: using both b's, the average projected annual income of professionals of age 35–65 was about 1.5 times the average income of professionals of age 25–34. Overall, the points lie fairly close to the 45° line, suggesting that the age/total-household-income profiles were indeed relatively stable over time (and between the two data sets).

There is a more direct way to judge whether the projection technique actually produces reasonably good forecasts of income: using the PSID, it is possible to regress the ex post sum of actual future income on predicted future income. Define

$$h_{i,t}^{17} = E_t\left[\sum_{j=1}^{17} R^{-i} Z_{i,t+j} b_{y,\mathrm{PSID68}}\right],$$

and

$$H_{i,t}^{17} = \left[\sum_{j=1}^{17} R^{-i} Y_{i,t+j}\right],$$

7. Cells for which either survey had five or fewer households were not plotted.

where $Y_{i,t+j}$ is actual observed total household income in year $t + j$ for household i in the PSID. Divide $h^{17}_{i,t}$ and $H^{17}_{i,t}$ by seventeen to put the data on a per year basis, drop the t subscript, and call the results h'_i and H'_i. Then we can estimate the equation[8]

$$H'_i = \underset{(2.49)}{3316} + \underset{(4.35)}{0.643\,y_i} + \underset{(3.40)}{0.41\,h'_i}, \tag{9}$$

where y_i is projected income in the current period from equation (1) and t-statistics (which are corrected for the fact that the independent variables are constructed) are in parentheses. If this estimation had yielded a coefficient on y of one and a h'_i coefficient near zero, the conclusion would have been that income is highly persistent and changes in income are entirely unpredictable. The highly significant coefficient on projected human wealth indicates instead that the projection technique captures significant and predictable differences between current and future income. The same point is made graphically by Figure II,[9] which plots H' against h' for the same groups of consumers represented in Table II and Figure I. There is a clear positive correlation between predicted and actual future income across groups.

Although Figure II and equation (9) indicate that the forward-projection method of predicting future income is successful, the method does have conceptual flaws. Perhaps the greatest is that it assumes that personal characteristics such as occupation and education do not change over time. If occupational transitions are common, however, a 25-year-old service worker might be a professional by age 30. Another potential problem is that if consumers possessed any knowledge of likely future shocks to the income of their occupation/education/racial groups, a procedure which merely projects current profiles forward could not capture such knowledge. These problems and others are addressed by the second technique for estimating future income.

Method 2: Backward Projection of Actual Experience onto Initial Characteristics

The second strategy for predicting future income involves estimating true ex post income in the PSID as a function of initial

8. Current income must be a regressor because the theory that will be tested below is driven by the proposition that there are predictable *differences* between current and future income. If income were constant for all consumers, or followed a random walk, h' might predict H' well but we would be unable to test the life-cycle model because there would be no predictable *changes* in income.

9. Cells for which there were five or fewer observations were not plotted.

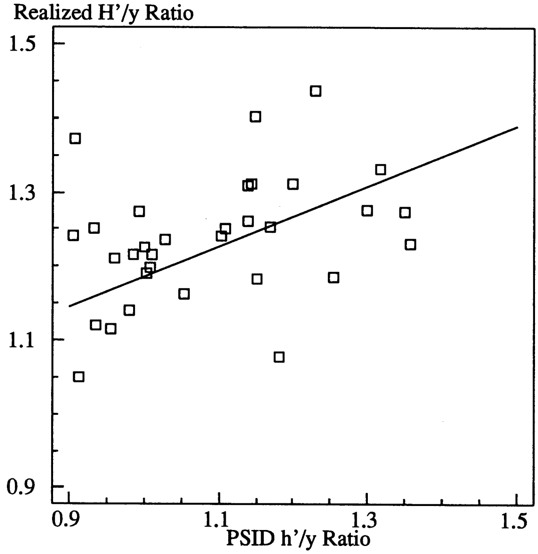

FIGURE II
Ratio of Future to Current Income by Group, Predicted Using PSID versus Actual Outcome

(1968) personal characteristics, i.e., estimating an equation of the form

$$H^{17}_{i,t} = Z_{i,t}b_h + v_{h,i,t}, \tag{10}$$

where $H^{17}_{i,t}$ is defined as above as the sum of actual 1969–1985 income, Z contains 1968 personal characteristics, and $v_{h,i,t}$ is the idiosyncratic error.[10] Such regressions generally have R^2's of 0.55 to 0.60, considerably higher than regressions like (1), which used a single year of income. This is unsurprising because transitory fluctuations that add variance to the point-in-time cross-section data should average out when income is cumulated over a long enough period.

Consider now our 25-year-old service worker who might be a professional by age 30. In estimating equation (10), the dependent variable would be the actual income that consumers who were

10. Estimated household tax payments become available in 1970, so it is possible to perform this regression using discounted after-tax income from 1970 to 1985 rather than discounted before-tax income from 1969 to 1985, as was done here. The results are much the same whichever measure of income is used, so the reported results will use before-tax income in order to maintain consistency with the results reported in Table II and Figure II.

service workers in 1968 earned over the next seventeen years; the measure of future income therefore automatically incorporates any changes in income due to changes after 1968 in occupational category or other personal characteristics. (Note that this method of estimating future income also automatically incorporates productivity growth.)

Once (10) has been estimated, the b_h coefficients on the Z variables can be transported over to the CEX to estimate future income for CEX consumers.[11] The results from estimating human wealth using this methodology were similar to the results using method 1, in the sense that the same groups had high or low levels and expected growth rates of income. For the sake of brevity, these results are not reported.

One possible objection to this technique is that post-1968 events that may have been unforeseeable in 1968, such as productivity shocks to occupational groups, will influence the estimated coefficients on personal characteristics. In effect, consumers are assumed to have a sneak preview into the true future for their cohort. If the consumers' method of estimating their future income is instead simply to project currently observed age/income profiles into the future as in method 1, then method 2 will produce a worse estimate of consumers' expectations of future income, even though it produces a better estimate of their actual future income. On the other hand, if consumers know in advance about upcoming changes in occupation, education, and cohort earnings, then this backward-projection method should produce better estimates of expected future earnings. Since it is difficult to know which method corresponds better to consumers' actual expectations, it seemed wisest to construct estimates of future income using both methods.

II. Does Future Income Affect Current Consumption?

The previous section demonstrated that future income differs predictably from current income. This section explores whether current consumption is influenced by these predictable changes in income. A preliminary look at the question is represented by Figures III and IV.

Each dot in Figure III represents the average current consumption and average current income of a group of consumers who are

11. The predicted income in the CEX using the PSID coefficients is rescaled so that the mean of the predicted PSID-basis income equals the mean of CEX income.

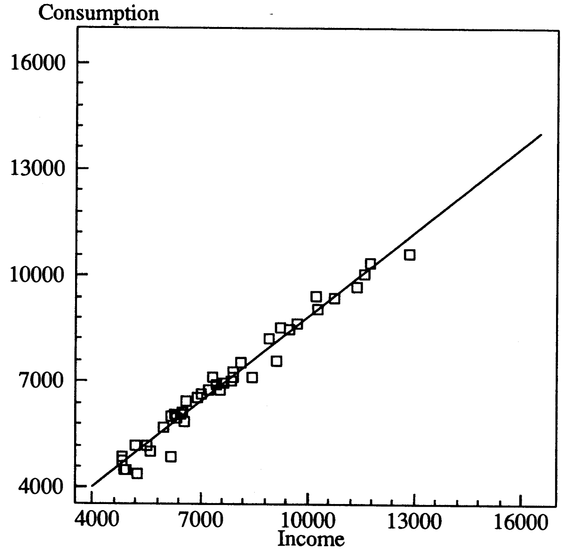

FIGURE III
Consumption versus Income by Group

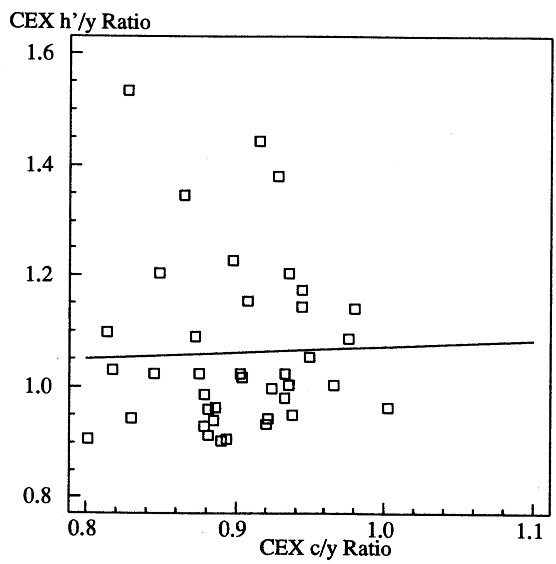

FIGURE IV
Ratio of Future to Current Income versus Ratio of Consumption to Current Income, by Group

of the same age bracket and either the same occupation or the same educational group. (The dots represent the same cells as in Table II and Figure I.) The figure reveals a strong relationship between average current income and average current consumption across groups. However, such a correlation may be uninformative about whether consumption depends on current income or lifetime income, because current income may be highly correlated with lifetime income. Hence Figure IV graphs the ratio of current consumption to predicted current income, c/y, against the ratio of predicted average future income to current income, h'/y.[12] If consumption were determined by expected lifetime income, this figure should show a strong positive relationship between the consumption ratio and the future-to-current income ratio. In fact, there is little apparent relationship at all between the consumption ratio and future income growth. These figures capture the essence of the more rigorous econometric results below: current consumption appears strongly related to predictable current income, and unrelated to predictable lifetime changes in income.

Econometric Tests

Simple certainty-equivalent (CEQ) versions of the Life Cycle/ Permanent Income Hypothesis[13] imply that consumption is given by some fraction of total (human and nonhuman) resources:

$$C_i = k_i[Y_i + H_i + W_i], \tag{11}$$

where k_i is a function of interest rates, the consumer's tastes and discount rate, and the consumer's age. W_i represents physical assets, Y_i is current period labor income, and H_i is human wealth as defined above (the t subscripts were dropped because these regressions will all be estimated across a set of consumers in the same year). If interest rates and discount rates are the same across households, k_i will be the same for all households of the same age. Given a sample restricted to households of the same age, therefore, the theory suggests an econometric specification like

$$C_i = \delta_0 + Y_i\delta_1 + H_i\delta_2 + W_i\delta_3 + e_i, \tag{12}$$

12. Future income was predicted using the first of the methods described in the previous section. Results were similar when income was predicted using the second method.

13. In this section of the paper, analysis will be restricted to the CEQ version of the LC/PIH models. The last section of the paper will discuss a more sophisticated model in which income uncertainty is important.

and implies that $\delta_1 = \delta_2 = \delta_3$ (and, interpreted rigidly, that $\delta_0 = 0$). My focus will be on whether the propensity to consume out of current and future income is the same, i.e., whether $\delta_1 = \delta_2$, and whether either δ_1 or δ_2 is close to the values predicted by the theory.

Unfortunately, the CEX does not contain a direct measure of physical wealth W, so a measure had to be constructed using information on the value of owner-occupied housing, mortgage interest payments, and capital income (see the Data Appendix for details). Because this method undoubtedly measures wealth with error, equation (12) had to be estimated using instrumental variables in order to obtain consistent estimates of δ_3. The instruments were the same Z variables used in estimating equation (1) and in constructing the demographically predictable component of human wealth, $h_{i,t}$.

The equation was estimated using the Two Sample Two-Stage Least Squares (TS2SLS) technique described in Carroll and Weil [1994]. This procedure allows estimation of the first stage of a 2SLS regression on one sample and estimation of the second stage on a different sample. In the case at hand, the first-stage regressions are estimated on the entire sample of people aged 25–65, while the second-stage regressions are estimated separately for households aged 25–34, 35–44, and 45–54. Carroll and Weil show that this technique produces consistent estimates of the coefficients if the usual assumptions required for 2SLS estimation hold. The calculation of standard errors is slightly more complex; details of the estimation procedure are contained in an econometric appendix available upon request from the author.

The second-stage regression is given by

$$C_i = \delta_0 + y_i\delta_1 + h_i\delta_2 + w_i\delta_3 + u_i, \tag{13}$$

where y_i and w_i are the predicted values of Y and W from the first-stage regressions of Y and W on the instruments Z in the full sample, and h_i is given either as in equation (7) when method 1 is used to project future income or as in equation (10) when method 2 is used. The age of retirement T is assumed to be 65, and the annual real interest rate is 3 percent.[14] These second-stage regressions are run for consumers in the three age brackets defined above.

Before presenting the results, it is necessary to consider what the CEQ model predicts. Under the simplest version of the model,

14. Results were similar for interest rates 0 percent and 8 percent.

consumption is constant over the remainder of the lifetime. Assuming a life span of 80 years, this means that an average member of the 25–34 year old group who is 30 years old should have $\delta_1 = \delta_2 \approx (1/50) = 0.02$, an average member of the 35–44 year old group should have $\delta_1 = \delta_2 \approx (1/40) = 0.025$, and an average member of the 45–54 year old group should have $\delta_1 = \delta_2 \approx (1/30) \approx 0.0333$.

A final issue is the appropriate definition of consumption. Equation (11) is derived from a model in which all consumption is perfectly nondurable. I therefore subtracted purchases of durable goods (vehicles, household durables, and radios and televisions) from total consumption. The appropriate treatment of insurance payments is also unclear, but I removed them as well. I then scaled the resulting measure of nondurables and services consumption by the average ratio of total consumption to nondurables and services consumption, on the assumption that the flow of durables services is proportional to the amount of nondurables and services consumption.[15]

Table III presents the empirical estimates of equation (13) using each of the two methods of predicting income. The results are overwhelmingly unfavorable to the model. The coefficient on current income is always considerably more than an order of magnitude too large, and is overwhelmingly significant, while the coefficient on human wealth is always far too small, is sometimes negative, and is only significantly positive in two regressions.[16] In

15. To be concrete, the average ratio of total consumption to nondurables and services consumption was 1.220 for the consumers in my sample. My measure of estimated total consumption was therefore defined as 1.220 times the observed consumption of nondurables and services. This adjustment would be exactly right if the elasticity of substitution between durables and nondurables were fixed, and if there were no life-cycle pattern in the proportion of consumption accounted for by each category of consumption. The latter assumption can be tested, and is supported: the ratio of total consumption to nondurables and services consumption was 1.222 for the 25–34 age group, 1.223 for the 35–44 age group, and 1.215 for the 45–54 age group. I also performed all the regressions in the paper using total actual expenditures including durables; the results were similar to those reported, and are available upon request.

16. The failure is, if anything, worse for the regressions when income is predicted using the second method, because in that case the coefficient on the h term should be biased upward. The reason is that the second method produced a projection of income, $h_{i,t}^{17}$, only for the next seventeen years of life. The correct variable is $h_{i,t}$ through the end of life. Thus, there is an omitted variable in the regression, the income to be earned from years $t + 18$ through retirement, $(h_{i,t} - h_{i,t}^{17})$. If income earned in years $t + 18$ through retirement is positively correlated with income earned in the next seventeen years (as one would expect), then the usual logic of omitted variable bias shows that the coefficient on the included variable should be biased upward.

TABLE III
CEX REGRESSIONS OF CONSUMPTION ON CURRENT INCOME AND FUTURE INCOME

Regression	Income projection method	Age group	Constant	Current income y	Future Income h	Wealth w	Number of obs.
1	1	25–34	1174	0.705*,†	0.0019†	−0.007	1788
			[275]	[0.062]	[0.0020]	[0.009]	
2	1	35–44	669	1.004*,†	−0.0084**,†	−0.025	2518
			[325]	[0.095]	[0.0034]	[0.015]	
3	1	45–54	385	0.745*,†	0.0080**,†	0.004	2237
			[252]	[0.084]	[0.0044]	[0.011]	
4	2	25–34	1254	0.611*,†	0.0087**,††	−0.009	1788
			[273]	[0.088]	[0.0049]	[0.009]	
5	2	35–44	922	0.878*,†	−0.0046†	−0.017	2518
			[311]	[0.112]	[0.0078]	[0.014]	
6	2	45–54	361	0.873*,†	−0.0022†	0.002	2237
			[272]	[0.048]	[0.0020]	[0.012]	

Notes. Estimated using Two Sample Two-Stage Least Squares, as described in Carroll [1993a, econometric appendix]. Standard errors are in brackets. Variable construction is described in the text and in the data appendix. Theoretically predicted values, calculated in the text, are 0.02 for 25–34 age group 0.025 for the 35–44 age group, and 0.0333 for the 45–54 age group.

*Significantly different from zero at the 1 percent level of significance.

**Significantly different from zero at the 10 percent level of significance.

†Significantly different from theoretically predicted value at the 1 percent level of significance.

††Significantly different from theoretically predicted value at the 10 percent level of significance.

every regression the coefficient on current income is significantly different from the theoretically predicted value at much better than the 1 percent level of significance. In all regressions but one, the coefficient on human wealth is significantly different from the theoretical value at the 1 percent level; in that one regression the estimated coefficient is different from the theoretical value at the 10 percent level. The coefficient on nonhuman wealth is estimated and is never statistically different from zero, although it is usually significantly less than the theoretically predicted value, and is negative more often than positive. Given that the wealth measure had to be constructed from data on asset income and home ownership status, it may not be surprising that the coefficients are unstable and statistically insignificant.

Put simply, these regressions suggest that, for example, a 25–34 year old professional does not consume significantly more than a 25–34 year old salesman with a similar current income, even though the professional can expect much higher lifetime income.[17]

It would have been preferable (and easier) to estimate the regressions using method 2 in Table III using the PSID alone, in order to eliminate the difficulties caused by attempting to meld information from two different data sets.[18] Unfortunately, however, this cannot be done because the PSID does not contain comprehensive data on consumption. It does contain data on food consumption, and many researchers have attempted to use the PSID food consumption variable as a proxy for total consumption or nondurables consumption. Frankly, I am skeptical about this strategy. The PSID is designed as an income survey and collects food consumption data almost as an afterthought. Shapiro [1982] has estimated that 95 percent of the variance in the PSID food consumption variable is noise; Runkle [1991] estimates that about 70 percent is noise. Discussions with economists at the Bureau of Labor Statistics, which conducts the Consumer Expenditures Surveys, have led me to believe that a survey that casually asks questions about consumption without training survey respondents to track their consumption is unlikely to elicit even remotely accurate answers. Furthermore, it is not even clear whether the measure of food consumption available in the first wave of the PSID, food consumed at home, is a normal or inferior good, and hence even its theoretical relation to lifetime income is unclear.[19] Finally, even if strong relationships between current food consumption and future income could be discovered, the implications for the deeper validity of the Life Cycle/Permanent Income Hypothesis would be murky because it would be unclear how to translate

17. There is an important subtlety in interpreting the results of Table III: the high coefficient on the "projected current income" term does not necessarily correspond to a high marginal propensity to consume out of transitory income. Imagine a model in which consumers always set consumption in a given year equal to *expected* income in that year. The marginal propensity to consume out of unanticipated transitory shocks to income would be zero, yet regressions like those in Table III would find a coefficient of one on the projected current income term and a coefficient near zero on the predicted future income term. (This description, in fact, may not be far from what Friedman [1957] had in mind in his original presentation of the Permanent Income Hypothesis.)

18. Carroll [1993, econometric appendix] describes the statistical difficulties. Furthermore, the definitions of occupational and educational categories were not identical in the two data sets, so I had to develop new definitions of occupational and educational categories that could be constructed using the data incorporated in each data set.

19. Later PSID survey waves have data on food at home and away from home (restaurant meals). Results using those data were qualitatively similar to the reported results, although the coefficient on the current income term was larger.

TABLE IV
PSID REGRESSIONS OF FOOD CONSUMPTION ON CURRENT INCOME AND FUTURE INCOME

Regression	Income projection method	Age group	Constant	Current income y	Future income h	Wealth w	Number of obs.
1	1	25–34	560	0.098*	−0.0015*	0.004	283
			[225]	[0.018]	[0.0006]	[0.022]	
2	1	35–44	1329	0.046*	−0.0003	−0.030	308
			[226]	[0.017]	[0.0006]	[0.025]	
3	2	25–34	498	0.074**	−0.0006	0.015	283
			[254]	[0.030]	[0.0017]	[0.025]	
4	2	35–44	1328	0.004	0.0021	−0.034	308
			[224]	[0.027]	[0.0016]	[0.025]	

Notes. Estimated using Two Sample Two-Stage Least Squares, as described in Carroll [1993a, econometric appendix]. Standard errors are in brackets. Variable construction is described in the text and in the data appendix.
*Significantly different from zero at the 1 percent level of significance.
**Significantly different from zero at the 10 percent level of significance.

information about changes in food consumption into implications about overall lifetime consumption and spending patterns.

Despite these reservations, the results of regressions using the PSID food consumption variable may interest some readers, so they are presented in Table IV. The sample size was too small to run the regressions for the oldest age group, so results are presented only for the two younger groups. The PSID results are qualitatively similar to those in Table III: current income is usually highly significant, and future income is usually insignificant (and, perversely, usually negative). In fact, the coefficient on current income is usually not far from 8 percent, the simple ratio of food consumption to total income in the PSID.

III. Is There a Simple Modification to the CEQ LC/PIH Model Which Can Explain These Results?

A. *Education/Occupation Choice Is Correlated with the Degree of Impatience*

Suppose that individuals are identical in every respect except their willingness to defer gratification. Those with low discount rates might invest heavily in education while young, enduring scholarly poverty for a few years, and would reap the rewards of a more positively sloped (and higher) income profile over the rest of their lives. These same people would desire a more positively sloped consumption path, because that is optimal for people with low discount rates. Hence a low discount rate would "cause" both a high growth rate of income and a high growth rate of consumption.

The discount-rate-selection argument can be tested empirically by excluding education from the variables predicting income. The hope is that this removes, or at least attenuates, the mechanism through which discount rates affect income, but leaves intact other factors which can predict income. Regressions 1–3 in Table V show the results when education is excluded from the set of variables used to explain income growth. The results are essentially the same as in Table III: consumption is very closely positively related to predictable current income, but future income does not have a reliable positive effect on current consumption. The proposition that the coefficients equal their theoretically predicted

TABLE V
REGRESSIONS USING LIMITED INSTRUMENT SETS

Regression	Instrument restrictions	Age group	Constant	Current income y	Future income h	Wealth w	Number of obs.
1	Excluding education	25–34	1125 [306]	0.673*,† [0.074]	0.0034† [0.0023]	0.001 [0.011]	1788
2	Excluding education	35–44	717 [341]	0.881*,† [0.102]	−0.0030† [0.0035]	−0.017 [0.019]	2518
3	Excluding education	45–54	−303 [374]	0.900*,† [0.119]	0.0063† [0.0060]	−0.026 [0.019]	2237
4	Excluding occupation	25–34	1154 [401]	0.723*,† [0.079]	0.0019† [0.0020]	−0.025 [0.017]	1788
5	Excluding occupation	35–44	1060 [500]	0.907*,† [0.129]	−0.0081**,† [0.0032]	−0.006 [0.025]	2518
6	Excluding occupation	45–54	1339 [423]	0.416*,† [0.134]	0.0156*,† [0.0055]	0.060 [0.021]	2237

Notes. Estimated using Two Sample Two-Stage Least Squares, as described in Carroll [1993a, econometric appendix]. Standard errors are in brackets. Variable construction is described in the text and in the data appendix. Theoretically predicted values, calculated in the text, are 0.02 for 25–34 age group 0.025 for the 35–44 age group, and 0.0333 for the 45–54 age group.

*Significantly different from zero at the 1 percent level of significance.

**Significantly different from zero at the 5 percent levelof significance.

†Significantly different from theoretically predicted value at the 1 percent level of significance.

††Significantly different from theoretically predicted value at the 5 percent level of significance.

values can be rejected at the 1 percent level in every regression. The foregoing argument about educational choice can be repeated with occupational choice by arguing that people with a low discount rate will chose occupations with steeply sloped income paths. Again, when the appropriate regressions are performed (Table V, regressions 4–6) the results in Table III are essentially duplicated. Therefore, to the extent that the discount-rate-selection hypothesis can be tested, it is not supported. (Of course, if both occupation and education are associated with tastes and the degree of patience, we would like to exclude *both* from the variables predicting future income. Unfortunately, this cannot be done because occupation and education are the main variables we are using to predict future income, so we cannot exclude both simultaneously.)

B. Children or Family Structure

A more sophisticated version of the life-cycle model that maintained the essential approach of CEQ intertemporal optimization but used a utility function that subsumed the consumption needs of children could predict that the presence of children would affect family consumption. Marital status also changes the group of people over which optimization is performed and so would presumably affect consumption. Race, too, might be imagined to have an effect on current consumption, either through associated cultural or taste differences or through relative income effects as in Duesenberry [1949]. These points can be addressed by adding demographic variables to the consumption equations being estimated.

Consider children. The presence of children should increase the fraction of lifetime resources devoted to current family consumption. Under the null that the life-cycle model is true, therefore, a dummy variable interacting the number of children with the amount of lifetime income should be included on the right-hand side of any equation of the form (1). The coefficient on this variable would then indicate how much an additional child would raise current consumption as a fraction of lifetime income.

Note that it would be inappropriate here to make estimates of future income using the first methodology described above, in which we assume that demographic characteristics other than age stay the same throughout the future, because the number of children at home will change, and marital status may change, over time. The backwards-projection method (method 2), however, is appropriate: we can estimate actual future income in the PSID as a function of 1968 personal characteristics, this time including 1968 marital status and number of children. The appropriate equation to estimate, therefore, is equation (1) using these PSID-based estimates of human wealth and with the addition of the human-wealth-interacted number of children, marital status, and race.

When such regressions were estimated, the qualitative results are essentially the same as those in Tables III and V. These results are omitted for brevity.

C. Liquidity Constraints

Maybe the reason consumption typically remains close to income is that most people would like to be spending more than their current income, but are unable to do so because of liquidity constraints. Since predicted current income generally receives a coefficient of at least 0.6 in most of the regressions, direct liquidity constraints could explain the results only if a very large proportion of consumers are liquidity constrained, and if liquidity-constrained consumers have a marginal propensity to consume near one.[20] If most of the population were liquidity constrained in this primitive sense, most of the population should have zero liquid assets. In fact, data from the 1983 *Survey of Consumer Finances* show that only about 12 percent of households held zero liquid assets [Kennickell and Shack-Marquez 1992].[21]

Some tests of whether this kind of simple, binding liquidity constraint can explain the results are reported in Table VI. The first test (regressions 1–3) restricts the sample to households whose estimated liquid assets were greater than one-quarter of their income. Although both estimated assets and projected income are subject to measurement error, this group should at least be substantially less liquidity constrained than the remainder of the sample. Although the coefficient on projected current income δ_1 remains at least ten times larger than predicted by the CEQ theory in all regressions, the smaller sample size results in a larger standard error, and the hypothesis that δ_1 equals the predicted

20. A consumer who is liquidity constrained today and will also be liquidity constrained tomorrow may have an MPC far below one because she may wish to spread the extra dollar over the whole period during which the constraint applies. If anything, the presumption should probably be that the MPC of liquidity-constrained consumers is far less than one, since most liquidity-constrained consumers probably remain constrained for long periods.

21. A more sophisticated model of liquidity constraints, like that of Deaton [1991], does not necessarily imply that consumers should hold zero assets. Indeed, the last section of the paper will argue that the empirical results of the paper are consistent with the kind of "buffer-stock saving" behavior that emerges from a model like Deaton [1991] or Carroll [1992a, 1992b]. The current discussion is targeted only at the most naive models of liquidity constraints.

TABLE VI
TESTS OF LIQUIDITY CONSTRAINTS

Regression	Sample restriction	Age group	Constant	Current income y	Future income h	Wealth w	Number of obs.
1	$A > Y/4$	25–34	1191	0.577**,††	0.0079	−0.001	174
			[1388]	[0.266]	[0.0084]	[0.037]	
2	$A > Y/4$	35–44	891	0.929*,†	−0.0083†	0.011	445
			[930]	[0.264]	[0.0093]	[0.042]	
3	$A > Y/4$	45–54	545	0.574**,††	0.0160	0.055	554
			[978]	[0.307]	[0.0161]	[0.038]	
4	Homeowners	25–34	978	0.610**,††	−0.0016††	0.046	88
			[1046]	[0.299]	[0.0105]	[0.043]	
5	Homeowners	35–44	−100	0.794*,†	0.0021††	−0.007	333
			[796]	[0.261]	[0.0102]	[0.038]	
6	Homeowners	45–54	−159	0.459**,††	0.0294*	0.039	630
			[565]	[0.191]	[0.0104]	[0.025]	

Notes. Estimated using Two Sample Two-Stage Least Squares, as described in Carroll [1993a, econometric appendix]. Standard errors are in brackets. Variable construction is described in the text and in the data appendix. Theoretically predicted values, calculated in the text, are 0.02 for 25–34 age group, 0.025 for the 35–44 age group, and 0.0333 for the 45–54 age group. $A > Y/4$ indicates that estimated liquid assets were greater than one-quarter of income. Homeowner indicates that the household owns the house outright, with no mortgage debt.

*Significantly different from zero at the 1 percent level of significance.
**Significantly different from zero at the 10 percent level of significance.
†Significantly different from theoretically predicted value at the 1 percent level of significance.
††Significantly different from theoretically predicted value at the 10 percent level of significance.

value is now only rejected at the 10 percent significance level for the oldest and youngest age groups, although it is still rejected at the 1 percent level for the middle age group. The increased standard errors have a more potent effect for the coefficient on human wealth: the estimated coefficient is now significantly different from the predicted value only for the middle age group, although the point estimates are not much different from those in previous tables.

The second test for the importance of liquidity constraints, in regressions 4–6, selects only those households who own their homes, on the assumption that homeowners are less likely to be liquidity constrained than nonhomeowners. The results for δ_1 essentially duplicate the results from the first test of liquidity constraints in regressions 1–3. However, the hypothesis that δ_2 equals its theoretical value is now rejected at the 10 percent level in two of the three age groups.

On balance, these results suggest that a simple formulation of liquidity constraints is not the explanation for the close association between consumption and predictable current income. The foregoing analysis, however, is aimed primarily at disposing of the very simplest liquidity constraints models. The final section of the paper will argue that these same empirical results can be explained by a more sophisticated model in which people would like to borrow if future income were known with certainty, but prudence in the face of income uncertainty causes consumers to maintain positive assets.

IV. Does Uncertainty About Future Income Affect Current Consumption?

The previous section argued that there is little evidence in the CEX or PSID that predictable long-horizon changes in the level of income affect current consumption. This section examines whether and how income uncertainty affects current consumption. The results indicate not only that uncertainty matters, but that it matters in just the way that modern precautionary saving theory suggests it should.

A. Literature Review

Using the 1972–1973 *Consumer Expenditure Survey,* Skinner [1988] found that saving rates were less for the self-employed and sales workers. He had no explicit measure of income uncertainty, but asserted that these two groups have greater income uncertainty than other groups,[22] and so should have had greater saving. This is an extremely crude test, and was only meant to be a very preliminary first step.

Dynan's [1993] recent test is more sophisticated. She notes that the Euler equation for consumption implies that those who face greater uncertainty should have greater consumption growth (because they are depressing the level of consumption now in order to do precautionary saving). Dynan calculates quarterly consumption growth and the variance of quarterly consumption growth using data from the 1985 *Consumer Expenditure Survey,* and regresses actual consumption growth on the variance of consumption growth, instrumented by occupation, education, age, etc. She finds no evidence that consumption growth is higher for those whose predicted consumption variance is higher, and interprets this as evidence against the existence of precautionary saving.

For a careful discussion of the theory behind Dynan's test and some subtle theoretical objections to her test, see Carroll [1992a]. An empirical objection is that her measure of uncertainty, the variance of *quarterly* consumption growth, probably bears very little relationship to actual uncertainty. She finds that quarterly consumption growth has a standard deviation of about 20 percent in her sample; this is surely far too high to represent quarterly reevaluations of the level of lifetime income. Instead, it probably reflects vacation expenses, schooling costs, and even the number of shopping trips one has managed to make in a given quarter. These things are probably completely unrelated to uncertainty, and therefore should be unrelated to consumption growth. Although Dynan's test is a clever theoretical idea, in practice she probably asks too much of quarterly household consumption data.

Guiso, Jappelli, and Terlizzese [1991] use data from an Italian household survey to see whether consumption is related to a self-reported expected variance of next year's income. It is clear that many households did not understand the survey question: a very large proportion of households reported point expectations for the next year's income. Nevertheless, Guiso, Jappelli, and Terlizzese do find a statistically significant, though small, relationship between self-reported income uncertainty and consumption. Finally, two recent papers by Carroll and Samwick [1993a, 1993b]

22. A proposition for which there is some support, at least regarding the self-employed, in my Table VII.

regress wealth holdings in the PSID in 1984 on measures of income uncertainty and do find statistically significant evidence for economically important precautionary saving.

B. My Tests

The PSID extract used in this paper contains eighteen years of income for a relatively large panel of households. Traditional measures of income uncertainty are the standard deviation or the variance of income. After removing the predictable life-cycle component of income changes,[23] variance and standard deviation measures were constructed for each household. Unfortunately, except under very special circumstances, neither of these is "the right" measure of uncertainty, in the sense of a measure which theory is a sufficient statistic for the amount of precautionary saving that will be induced by a given income distribution.

Kimball [1990], however, has derived a measure called the "equivalent precautionary premium" which, if the standard theory of precautionary saving is correct, might be a better measure of uncertainty than variance or the standard deviation. Kimball's equivalent precautionary premium is given by the amount $\psi(c)$ such that for given consumption c, $u'(c - \psi) = E\, u'(c + \zeta)$, where ζ is a random error term in consumption. Kimball shows that the equivalent precautionary premium is, in essence, a direct measure of the intensity of the precautionary saving motive at the point of zero precautionary saving.

Under certain assumptions it is possible to construct a measure corresponding to ψ for each household in the PSID. Suppose that each household had to consume exactly its income each year. Consider a household i who in each period t consumed amount $Y_{i,t}$ which was distributed i.i.d. with mean μ_i, $Y_{i,t} = \mu_i + \zeta_{i,t}$. Define the mean of $Y_{i,t}$ over the observed period for this consumer as $\hat{\mu}_i = 1/18 \sum_{t=1968}^{1985} Y_{i,t}$. If households really had to consume their income, actual marginal utility in each period would be $u'(Y_{i,t})$. Define the estimated income shock in year t as $\zeta_{i,t} = Y_{i,t} - \hat{\mu}_i$. Assuming that, on average, expected marginal utility equals actual marginal utility (rational expectations), and further assuming that the distribution of ζ's does not change over time, we have

$$E[u'(Y_1)] = 1/18 \sum_{t=1968}^{1985} u'(Y_{i,t}) + \gamma_i,$$

23. This was done by expressing income as a fraction of predicted income generated using wage equations like those described in Section I, including a term for aggregate productivity growth (assumed to be 2 percent a year).

where γ_i represents the deviation of actual experience over the 1968–1985 period from true underlying expectations. Rational expectations would imply that on average γ_i should be zero, so it should be acceptable to proxy for $E\,[u'(Y_i)]$ using $1/18\ \Sigma_{t=1968}^{1985}\ u'(Y_{i,t})$.

Suppose that this consumer had a utility function of the form $u(c) = c^{1-\alpha}/1 - \alpha$, where α is the coefficient of relative risk aversion, implying that $u'(c) = c^{-\alpha}$. Again assuming that consumption c must equal income in each period, it is possible to solve for an estimate of ψ:

$$u'(\hat{\mu}_i - \psi_i) = E[u'(Y_i)]$$

(14)
$$\hat{\psi} = \hat{\mu}_i - \left[\frac{1}{18}\sum_{t=1968}^{1985}(Y_{i,t})^{-\alpha}\right]^{-1/\alpha}.$$

In this case $\hat{\psi}_i$ will be an unbiased estimator for the true ψ which corresponds to the true distribution of Y.[24] The only further assumption needed in order to compute $\hat{\psi}$ for each household in the PSID is an assumption about α.

Table VII presents the normalized standard deviation, variance, and equivalent precautionary premium for income assuming that $\alpha = 3$, by occupational and educational group for the consumers in the PSID. In general, the patterns in the table seem sensible: farmers and self-employed businessmen had the highest income uncertainty, while professionals and highly educated workers had low income uncertainty.[25]

The remaining task is to determine how consumption is related to uncertainty. Dropping the household subscript i, write the uncertainty-augmented consumption model as

(15)
$$C = \delta_0 + Y\delta_1 + H\delta_2 + W\delta_3 + S\delta_4 + e,$$

where S represents uncertainty, as measured by any of the three measures constructed above. This equation cannot be estimated within the PSID, because, as argued above, the PSID does not have an adequate measure of consumption. It also cannot be estimated directly in the CEX, because the CEX does not contain the panel

24. Recall that the actual income series used here has been detrended to remove both aggregate productivity growth and predictable life-cycle changes in income, because those changes in income are predictable and we are attempting to construct a measure of uncertainty.

25. Note that the standard deviation is not exactly equal to the square root of the variance. For any individual household the standard deviation will equal the square root of the variance, but this does not hold for the group average square roots and variances because the square of the average is not equal to the average of the square.

TABLE VII
MEASURES OF THE UNCERTAINTY OF FUTURE INCOME IN THE PSID BY OCCUPATION AND EDUCATIONAL GROUP

Category of consumer		Number of households	Variance of normalized income	Std. dev. of normalized income	Equivalent precautionary premium†
Occupation	Professional, technical, and kindred	141	0.08	0.33	0.14
	Managers, officials, and proprietors	77	0.07	0.28	0.14
	Self-employed businessmen	16	0.10	0.40	0.21
	Clerical and sales	90	0.09	0.35	0.17
	Craftsmen, foremen, and kindred	122	0.08	0.30	0.14
	Operatives and kindred	101	0.09	0.31	0.17
	Laborers and service	44	0.12	0.37	0.21
	Farmers and farm managers	14	0.30	0.58	0.59
Education	Less than or equal to 8 years	65	0.12	0.39	0.25
	9–11 years	101	0.10	0.36	0.16
	High school grad	202	0.08	0.31	0.16
	Some college	105	0.08	0.29	0.15
	Finished college	77	0.08	0.32	0.13
	Some postgraduate	55	0.08	0.34	0.14

Source. Calculations by the author using 1969–1985 incomes of consumers in 1985 wave of the PSID.
†Calculated assuming a coefficient of relative risk aversion of three.
All three uncertainty variables are normalized to render them scaleless. The standard deviation and the EPP are normalized by mean income, and the variance is normalized by the square of mean income.

data on income necessary to construct S. What can be done, however, is Two Sample Two-Stage Least Squares (TS2SLS) estimation, in which the relationship between the instruments and uncertainty is estimated in the PSID, and predicted values of uncertainty are constructed using the same instruments in the CEX.

To be concrete, in the PSID we can estimate the equation,

$$S_p = Z_p b_s + v_{p,s}, \tag{16}$$

where Z_p is a set of instruments identical to the instruments used in the CEX estimation for Y and W above, and $v_{p,s}$ is the error from estimating the S equation in the PSID. We can then move to the CEX data set and construct $s = Zb_s$. Proofs in Carroll and Weil

[1994] can be applied to show that OLS estimation of the equation,

$$C = \delta_0 + y\,\delta_1 + h\,\delta_2\,w\,\delta_3 + s\,\delta_4 + u, \tag{17}$$

will give consistent estimates of δ. Unfortunately, under TS2SLS it is not possible to calculate a consistent estimate for the true standard errors in the case where the instrumented variable (in this case, s) is measured with error. It was possible, however, to derive a measure that is an upper bound for the correct standard errors; the derivation is contained in the econometric appendix, available from the author.

Table VIII presents the results, where the reported standard

TABLE VIII
ADDING MEASURES OF INCOME UNCERTAINTY TO THE REGRESSIONS

Regression	Uncertainty measure	Age group	Constant	Current income y	Future income h	Wealth	Uncertainty	Number of obs.
1	STD	25–34	1387	0.696	0.0027	−0.003	−0.136	1788
			[452]	[0.098]	[0.0031]	[0.014]	[0.090]	
2	VAR	25–34	1286	0.673	0.0031	−0.002	−2.27E-05	1788
			[430]	[0.098]	[0.0032]	[0.014]	[1.50E-05]	
3	EPP	25–34	1749	0.698	0.0007	0.007	−0.348**	1788
			[482]	[0.096]	[0.0031]	[0.015]	[0.140]	
4	STD	35–44	709	1.012	−0.0087	−0.024	−0.034	2518
			[610]	[0.173]	[0.0061]	[0.028]	[0.176]	
5	VAR	35–44	669	1.005	−0.0085	−0.025	−1.59E-06	2518
			[573]	[0.171]	[0.0061]	[0.028]	[3.00E-05]	
6	EPP	35–44	2327	0.763	−0.0076	0.033	−0.447*	2518
			[676]	[0.158]	[0.0051]	[0.028]	[0.126]	
7	STD	45–54	250	0.852	0.0045	−0.008	−0.054	2237
			[549]	[0.199]	[0.0099]	[0.026]	[0.042]	
8	VAR	45–54	67	0.891	0.0019	−0.007	−4.54E-06	2237
			[585]	[0.207]	[0.0104]	[0.026]	[3.03E-06]	
9	EPP	45–54	462	0.825	0.0060	−0.018	−0.127***	2237
			[540]	[0.185]	[0.0095]	[0.027]	[0.069]	

Notes. Estimated using Two Sample Two-Stage Least Squares, as described in Carroll [1993a, econometric appendix]. Upper-bound estimates for standard errors are in brackets. Variable construction is described in the text and in the data appendix. STD indicates standard deviation, VAR indicates variance, and EPP indicates Equivalent Precautionary Premium.

*Significantly different from zero at the 1 percent level of significance.

**Significantly different from zero at the 5 percent level of significance.

***Significantly different from zero at the 10 percent level of significance.

errors are actually the upper bound estimate of the standard errors derived in the econometric appendix. When added to the regression equations, all three measures of income uncertainty enter negatively: more uncertainty implies less consumption. However, using the upper-bound standard errors, only the equivalent precautionary premium is statistically significant in any of these regressions: it is significant at the 5 percent level for the youngest households; at the 1 percent level for middle-aged households; and at the 10 percent level for the oldest households. The correct, smaller standard errors would likely substantially increase the statistical significance of these estimates, and might make the other uncertainty measures statistically significant, but would not modify the conclusion that the EPP performs better than the other measures of uncertainty.

It is interesting to note that the estimated coefficient on income uncertainty is substantially smaller for the oldest age group than for the younger groups. This is consistent with a buffer-stock model of saving in which young and middle-aged households are trying to build up a buffer stock, but by the time they have reached their peak earning years, 45–54, they have achieved a large enough buffer and so do not need to continue depressing consumption to continue building up the stock further.

Even if the coefficients are statistically significant, if they were so small as to imply negligible effects on saving, then the results of this table might be uninteresting. A crude way to judge the magnitude of these effects is by asking what would happen to consumption if each measure of income uncertainty increased by a certain amount. A natural experiment is to increase each measure by an amount equal to its own cross-sectional standard deviation. This experiment is performed in Table IX. To be perfectly clear, take the case of the EPP for young households. Recall that an estimated EPP was constructed for each household. The cross-sectional standard deviation of the EPP for young consumers was \$595. The coefficient of -0.348 on the EPP in Table VIII implies that if the EPP were increased by \$595 for the typical consumer, consumption would fall by $0.348 \times \$595 = \$207$. This \$207 represents 3.2 percent of the average level of consumption for young consumers, \$6447. Thus, the table indicates that a young consumer with an EPP one standard deviation above the mean should have a saving rate about 3.2 percent higher (a consumption ratio 3.2 percent lower). These effects are quite large relative to an aggregate personal saving rate in 1960–1961 of about 6 percent.

TABLE IX
INCREASE IN SAVING RATE CAUSED BY A ONE-STANDARD-DEVIATION INCREASE IN UNCERTAINTY

	Uncertainty measure		
	Coefficient of variation	Variance	Equivalent Precautionary Premium
Age group	Increase in saving as a fraction of income (percent)		
25–34	2.0	2.0	3.2
35–44	0.3	0.1	4.8
45–54	2.4	2.9	3.5

Notes. Calculated using the coefficients from Table VIII combined with calculated standard deviations for each of the measures of uncertainty. For example, the 3.2 percent figure in the upper right corner of the table is calculated as follows. The standard deviation of the EPP for the young age group is \$595. The coefficient on the EPP from Table VIII is −0.348. Thus, a one-standard-deviation increase in the EPP would reduce consumption and increase saving by \$595 × 0.348 = \$207. The average income for the young consumers was \$6447, so the effect on the saving rate out of income is 207/6447 = 3.2 percent.

One apparent implication of these results may seem problematic. Table VII shows that professionals and highly educated consumers have lower income uncertainty. These latest results suggest that *ceteris paribus* lower uncertainty should mean higher consumption and lower saving. Yet casual experience suggests that professionals and highly educated consumers save more, not less, than other groups. How can they save more if they have proportionally lower income uncertainty? The answer is that the equations in Table VIII also allow people with high incomes to save more, independent of the effect of uncertainty. If consumption $c = a + by$, then the consumption ratio $c/y = b + a/y$ will fall as income rises, as long as $a > 0$. Since the constant term a is always estimated to be positive, we can attribute the higher saving of professionals and highly educated people to their higher income, and argue that they would save even *more* if they faced relatively as much income uncertainty as manual laborers.

A host of other measures of the variability of the income stream were tried,[26] but none performed consistently better than the three measures reported here.

26. For example, the number of times income fell to less than 25 percent or 50 percent of its mean value, or the number of times it fell by more than 25 percent or 50 percent from one year to the next. Although some of these measures sometimes outperformed variance or standard deviation, none consistently outperformed the EPP.

Taken as a whole, these results cast substantial doubt on the relevance of CEQ models which imply that income uncertainty has no effect on consumption.

CONCLUSION

This paper performs a direct test of perhaps the central implication of the simplest Life Cycle model—that current consumption should depend on expected future income. In a wide variety of tests, little evidence is found to support this proposition. Instead, projected current income is found to be overwhelmingly important in determining current consumption.

This might seem to recommend a simple Keynesian consumption model in which consumers spend their current income. However, the latter part of the paper shows that consumption responds strongly to uncertainty in future income—implying that consumers are rational and forward-looking (and hence not Keynesian).

Although neither of these results is consistent with the CEQ LC/PIH model of consumption, both are broadly consistent with "buffer-stock" models of saving such as those of Deaton [1991] or Carroll [1992a, 1992b]. In these models consumers are "impatient" in the sense that they would like to borrow against future labor earnings in order to finance current consumption. However, they either face explicit liquidity constraints [Deaton] or are prevented from borrowing by a particular form of precautionary saving behavior [Carroll]. Consumers hold a buffer stock of assets that serves to shield their consumption against high frequency fluctuations in income, but on average they are neither accumulating nor decumulating assets. Thus, on average, consumption equals income. However, for a given level of wealth and income, consumers with greater income uncertainty will consume less and save more—as found in Table VIII.

One pressing question is whether the buffer-stock models of saving that appear consistent with these micro data can explain a variety of puzzles in the macroeconomic literature. Such models at least have the potential to explain the "excess smoothness" puzzle mentioned in the introduction, because they imply that consumers are reluctant to spend on the basis of uncertain future income. And Carroll [1992a, 1992b] argues that a buffer-stock model may also be consistent with Campbell and Mankiw's [1989] finding that consumption growth is closely associated with predictable income growth. It does not appear unreasonable to hope that buffer-stock models may be capable not only of explaining the empirical results of this paper but also of resolving a wide range of puzzles in both the micro and macro consumption literatures.

Data Appendix

This appendix provides details about the construction of the data used in the paper, and about sample selection and variable definitions. The data sets and programs used to produce the results in this paper are available from the author.

The Panel Study of Income Dynamics

The PSID data were taken from Wave XVIII of the study, which contained data on income from 1968 to 1985 for 8129 individuals who were heads of a household at some time between 1968 and 1985. Characteristics of the PSID poverty subsample differed very substantially from those of the rest of the sample, so all the PSID regressions excluded households which were part of the poverty subsample. The regressions are unweighted, because the PSID weights apply only to the full sample that contains the poverty subsample. All results are also restricted to individuals who were the head of their household in 1968, for whom there were valid income data for all years, for whom there were valid occupation and education data for 1968, and who were aged 25 to 65 in 1968. These restrictions reduced the sample to 1033. In calculating the uncertainty measures, the sample was further restricted to individuals who were always the head of their household from 1968 to 1985, and who were less than 63 years old in 1985. These restrictions further reduced the sample to 577. Other sample restrictions used in producing individual results are detailed in the text and in the tables.

The theoretically desirable definition of income is total disposable labor income of the household. However, the PSID does not contain sufficient data to construct total labor income before 1976, or to construct total disposable income before 1970. The regressions for constructing actual future income using method 2 therefore use total household income, as explained in the text.

Several additional issues arise when calculating measures of income uncertainty. One issue was how to treat households whose head changed over the period. Because changes of head are often associated with dramatic changes in family status, such as marriage or divorce or departure from the nuclear family, I thought it would be inappropriate to treat changes in income associated with changes in the household head as a reflection of income uncertainty. I therefore restricted the sample to households whose head remained the same over the entire period. Another problem is that the progressivity of the tax code could have potentially had a very important effect on the magnitude of income uncertainty. I therefore wanted to use disposable household income in calculating the income uncertainty figures, but as noted above, disposable income is available only starting in 1970. For 1968 and 1969 the tax rate for each household was assumed to be equal to that household's average tax rate from 1970 through 1972. Despite the potential insurance effect of the tax code, however, the empirical results were not much different when the uncertainty measures were calculated using total household income rather than disposable household income. Because disposable household income is closer to the theoretical concept, the empirical results presented use uncertainty figures calculated using disposable household income.

The Consumer Expenditure Survey, 1960–1961

The 1960–1961 CEX surveyed 13,728 households in 1960 and 1961. My sample was restricted to households for whom reported consumption was positive, valid education and occupation information were available for the household head, the household composition did not change during the course of the survey year, the head was aged 25 to 65, and data on capital income were present. These restrictions narrowed the sample to 8364. The CEX regressions, like the PSID regressions, were unweighted. Results from weighted regressions were very similar to those for unweighted regressions, but the statistical theory detailed in the econometric appendix becomes even more cumbersome when performed with weights, so the simpler unweighted regressions were described and used.

Income variables reported in the survey include total household income, total household taxes, and capital income from interest, dividends, and rent. Total labor income was constructed as total household income minus total capital income. In order to construct disposable labor income, I assumed that total taxes were divided between labor income and capital income taxes in proportion to their shares in total income (i.e., if labor income was three-quarters of total income, three-quarters of total taxes were assumed to be labor income taxes).[27]

My measure of household physical wealth in the CEX was constructed from data on interest income, interest expenses, and the reported value of the consumer's home. The estimated value of mortgage debt was imputed by dividing mortgage interest payments by an assumed mortgage interest rate of 5.25 percent. I was unable to locate data on the average rate on outstanding mortgages, but did find data on the average interest rate on new fixed-rate 30-year mortgages beginning in 1972. A regression of this mortgage interest rate series on the interest rate on ten-year Treasury bills produced an R^2 of about 0.97. An equation that performed nearly as well simply set mortgage interest rates equal to ten-year T-bill interest rates plus 1.85 percent. Data on interest rates on ten-year Treasury bills were available back to 1953, so I assumed that the average interest rate on outstanding mortgages in 1960 and 1961 was given by the average T-bill rate from 1953–1961 plus 1.85 percent, yielding my 5.25 percent assumption. To impute a value of interest-bearing assets, I just divided interest income received by the 1961 yield on ten-year T-bills, 3.88 percent. For the value of the home, the survey contained a coded variable indicating whether the value of the home was less than $5000, between $5000 and $10,000, etc. I used data from the 1962–1963 *Survey of Consumer Finances* to calculate an average value of homes in each of the given ranges (deflating back to 1961), and set the value of the home for each consumer to the average value of homes in the price range indicated by their code. The measure of net worth was given this value of the home plus the imputed value of liquid assets minus the imputed value of debt. As noted in the text, this estimate of wealth undoubtedly contains very substantial measurement error, so all regressions had to be estimated using instrumental variables.

My measure of consumption was total household expenditures excluding principal and interest payments on debt. This definition includes insurance payments, gifts to individuals outside the household, and out-of-pocket expenditures on durable goods such as cars. (If a car was purchased with a $1000 down payment and

27. The exception to this rule was for the few households who reported negative income, usually as a result of negative self-employment income. For these households I assumed that disposable labor income was equal to disposable total income; results were not sensitive to alternative assumptions.

interest payments in the first year were $100, motor vehicle consumption would be $1100.) Results using this measure of consumption were very similar to the reported results, which are for consumption excluding durable goods, as described in the text.

BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM

REFERENCES

Campbell, John Y., "Does Saving Anticipate Declining Labor Income? A New Test of the Life Cycle-Permanent Income Hypothesis," *Econometrica*, LV (1987), 1249–73.

Campbell, John Y., and Angus Deaton, "Why is Consumption So Smooth?" *Review of Economic Studies*, LVI (1989), 357–74.

Carroll, Christopher D., "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity* (1992a), 61–156.

——, "Buffer Stock Saving and the Permanent Income Hypothesis," Board of Governors of the Federal Reserve System, unpublished paper, 1992b.

——, "How Does Future Income Affect Current Consumption: Econometric Appendix" (available from the author upon request), mimeo, Board of Governors of the Federal Reserve System, 1993.

Carroll, Christopher D., and Andrew A. Samwick, "The Nature of Precautionary Saving," Board of Governors of the Federal Reserve System, unpublished manuscript, 1993a.

Carroll, Christopher D., and Andrew A. Samwick, "How Important Is Precautionary Saving?" Board of Governors of the Federal Reserve System, Economic Activity Working Paper No. 145, 1993b.

Carroll, Christopher D., and Lawrence H. Summers, "Consumption Growth Parallels Income Growth: Some New Evidence," in B. Douglas Bernheim and John B. Shoven, eds., *National Saving and Economic Performance* (Chicago: University of Chicago Press for NBER, 1991), pp. 305–43.

Carroll, Christopher D., and David N. Weil, "Saving and Growth: A Reinterpretation," in Allan H. Meltzer and Charles I. Plosser, eds., *The Carnegie-Rochester Conference Series on Public Policy*, Volume 40 (Amsterdam: North-Holland, 1994), forthcoming.

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, LIX (1991), 1221–48.

Duesenberry, James S., *Income, Saving, and the Theory of Consumer Behavior* (Cambridge MA: Harvard University Press, 1949).

Dynan, Karen E., "How Prudent Are Consumers?" Board of Governors of the Federal Reserve System, unpublished paper, 1993.

Flavin, Marjorie, "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy*, LXXXIX (1981), 974–1009.

Friedman, Milton, *A Theory of the Consumption Function* (Princeton: Princeton University Press, 1957).

Guiso, Luigi, Tullio Jappelli, and Danielle Terlizzese, "Earnings Uncertainty and Precautionary Saving," Bank of Italy, unpublished paper, 1991.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," mimeo, Bradley Policy Research Center, April 1993.

Kennickell, Arthur B., and Janice Shack-Marquez, "Changes in Family Finances from 1983 to 1989: Evidence from the Survey of Consumer Finances," *Federal Reserve Bulletin*, LXXVIII (January 1993), 1–18.

Kimball, Miles S., "Precautionary Saving in the Small and in the Large," *Econometrica*, LVIII (1990), 53–73.

Runkle, David E., "Liquidity constraints and the Permanent Income Hypothesis," *Journal of Monetary Economics*, XXVII (1991), 73–98.

Shapiro, Matthew, "A Note on Tests of the Permanent Income Hypothesis in Panel Data," mimeo, University of Michigan, 1982.
Skinner, Jonathan, "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics,* XXII (1988), 237–55.
Viard, Alain D., "The Productivity Slowdown and the Savings Shortfall: A Challenge to the Permanent Income Hypothesis," *Economic Inquiry,* XXXI (October 1993), 549–63.
West, Kenneth D., "The Insensitivity of Consumption to News about Income," *Journal of Monetary Economics,* XXI (1988), 17–33.
Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics,* CIV (1989), 275–98.