1999

Consumption Adjustment Under Time-Varying Income Uncertainty

Douglas G Steigerwald, *University of California, Santa Barbara*
Joon-Ho Hahm

CONSUMPTION ADJUSTMENT UNDER TIME-VARYING INCOME UNCERTAINTY

Joon-Ho Hahm and Douglas G. Steigerwald

Abstract—We study the effect of income uncertainty on consumption in a model that includes precautionary saving. In contrast to previous studies, we focus on time-series variation in income uncertainty. Our time-series measure of income uncertainty is constructed from a panel of forecasts. We find evidence of precautionary saving in that increases in income uncertainty are related to increases in aggregate rates of saving. We also find evidence that anticipated income growth rates have less explanatory power for consumption growth rates after conditioning on income uncertainty. The evidence indicates the presence of forward-looking consumers who gradually adjust precautionary savings in response to changing income uncertainty.

1. Introduction

The permanent income hypothesis (PIH) states that individuals base their consumption on the annuity value of current financial and human wealth.[1] Hall (1978) proposed a simple statistical test of the PIH that has spawned a large literature. Although Hall reported some evidence in favor of the PIH, researchers that followed Hall's methodology often failed to find evidence supporting the PIH.[2] Because statistical tests of the PIH always include ancillary assumptions, rejections could be due to misspecification of the ancillary assumptions.[3] We focus on the possible misspecification arising from the assumption that only the mean of future income affects individual consumption paths. If individual consumption decisions are influenced by uncertainty about future income, then the variance of future income should affect consumption. We posit that the degree of uncertainty about future income is time-varying and that incorrectly ignoring time-varying income uncertainty leads to rejection of the PIH. We find that time-varying income uncertainty does play a role in determining an individual's consumption path.

If the marginal utility of consumption is nonlinear, then individuals' consumption decisions do not depend only on the mean of future income.[4] With convex marginal utility, individuals accumulate precautionary savings, which are savings against uninsurable income risks. A test of the PIH that allows for income uncertainty is thus also a test of the precautionary saving theory. Recent theoretical work indicates that precautionary saving can provide answers to the consumption puzzles pointed out in the traditional certainty equivalent PIH literature.[5] In an effort to provide empirical support for precautionary saving, a number of authors have undertaken cross-section studies, which link household income uncertainty with household savings. While some support for precautionary saving has been found, the results are not conclusive.[6]

We focus on time-series, rather than cross-section, variation in income uncertainty. Our time series of income uncertainty is constructed from aggregate data because no reliable time-series data exist at the household level. Cochrane (1991) and Pischke (1995) argue that aggregate income uncertainty measures typically underestimate household level earnings uncertainty, which suggests that our measure of aggregate income uncertainty provides a lower bound for the total uninsurable income risk faced by households. In addition, if aggregate income fluctuations affect consumers unequally, then Blanchard and Mankiw (1988) show that aggregate income uncertainty may have a large effect on aggregate consumption.

In section II, we derive the optimal consumption path for individuals with convex marginal utility who face time-varying income uncertainty. The optimal consumption path leads to testable regression hypotheses for both consumption and savings. We discuss our measure of time-varying income uncertainty in section III. Our measure is novel in that it is constructed directly from a survey of professional forecasters rather than from a parametric model for time-varying conditional variances. Because survey data may be contaminated with measurement error, we do not rely only

Received for publication July 15, 1996. Revision accepted for publication May 8, 1998.

* Korea Development Institute and the University of California at Santa Barbara, respectively.

We thank Jim Stock, participants at the North American Summer Meeting of the Econometric Society, and two referees of this journal.

[1] The modern interpretation of permanent income, due to Hall (1978) and Flavin (1981), is consistent with intertemporal choice models but is less general than the original interpretation of Friedman (1957).

[2] Rejections of the simplest version of the PIH, which predicts a martingale property for consumption, are summarized by the terms *excess sensitivity* (Flavin, 1981) and *excess smoothness* (Deaton, 1987).

[3] Possible misspecifications include the failure to account for liquidity constraints (Flavin, 1985; Campbell & Mankiw, 1989); unrestricted information sets (West, 1988; Campbell & Deaton, 1989; Gali, 1991); finite horizons (Gali, 1990; Clarida, 1991); time aggregation bias (Christiano et al., 1991); durable goods (Mankiw, 1982); habit formation (Deaton, 1987); and stochastic real interest rates (Mankiw, 1981; Hall, 1988; Campbell & Mankiw, 1989; Hahm, 1998).

[4] If the marginal utility of consumption is a linear function of consumption, then an individual's plan for future consumption depends only on the mean of future income. This is often captured with the phrase *certainty equivalence.* Carroll and Kimball (1996) provide conditions under which the marginal utility of consumption is nonlinear.

[5] Zeldes (1989) and Caballero (1991) show that precautionary saving may explain the excess sensitivity and excess smoothness features of consumption. The importance of precautionary saving for government policy is studied in Barsky et al. (1986), Hubbard and Judd (1987), and Feldstein (1988).

[6] Several studies report evidence that supports precautionary saving: Guiso et al. (1991, with Italian survey data) find that consumption is lower for individuals with higher income uncertainty; and Carroll and Samwick (1992, with U.S. grouped data) find that the stock of wealth is higher for individual groups that have greater income uncertainty. Other studies report evidence that does not support precautionary saving: Skinner (1988) finds that saving rates are lower for occupations with higher income uncertainty; and Dynan (1993, with the U.S. Consumer Expenditure Survey) finds little evidence of precautionary saving.

The Review of Economics and Statistics, February 1999, 81(1): 32–40

on ordinary least-squares (OLS) estimates. We also construct estimators that are consistent in the presence of measurement error. We find that our results, which we report in section IV, are substantively similar across estimators, which indicates that measurement error is not driving the results. We find that, while time-varying income uncertainty has little role to play in explaining the instantaneous adjustment of consumption, income uncertainty is important in explaining the level of savings. Our results on instantaneous adjustment are echoed in Carroll (1992), who uses an unemployment expectations measure to capture income uncertainty. Together, these findings suggest that savings and (nondurable) consumption do not adjust completely in one period. We investigate the possibility that consumption adjustment is not completed within one period, and find that time-varying income uncertainty has a substantial role to play in explaining the adjustment of consumption over a longer horizon.

II. Model Specification

We model an infinitely lived representative consumer who maximizes the expected present value of lifetime utility. We assume that utility is additively separable through time and a function of consumption alone. Let C_t be the value of consumption in period t. We assume that the representative consumer has constant absolute risk aversion (CARA) utility of the following form

$$U(C_t) = -\frac{1}{\theta} e^{-\theta C_t}, \tag{2.1}$$

where $U(\cdot)$ is the representative consumer's utility function, and $\theta > 0$ is the coefficient of absolute risk aversion.

In each period, the representative consumer maximizes the expected present value of lifetime utility. To represent the consumer's problem, we follow much of the extant literature and assume that the real interest rate, $r > 0$, is constant and equal to the rate of time preference. Let E_t be the expectation operator conditional on all information available to the consumer in period t. To maximize the expected present discounted value of lifetime utility, the representative consumer solves

$$\max_{\{C_{t+i}\}} E_t \sum_{i=0}^{\infty} (1+r)^{-i} U(C_{t+i}), \tag{2.2}$$

subject to the budget constraint $C_{t+i} = Y_{t+i} + (1+r) A_{t+i-1} - A_{t+i}$, where Y_t is the period-t value of labor income, and A_t is the end of period-t value of nonhuman wealth that satisfies $\lim_{i\to\infty}(1+r)^{-i} A_{t+i} = 0$. Because future labor income is the only source of uncertainty for the consumer, labor income is the random variable that drives consumption.

The time path of consumption $\{C_t\}_{t=0}^{\infty}$ that solves equation (2.2) is given by the Euler equation

$$e^{-\theta C_t} = E_t e^{-\theta C_{t+1}}. \tag{2.3}$$

To understand the effect of uncertainty about future labor income on current consumption, we express the path of consumption that satisfies equation (2.3) in terms of the innovations to labor income. To capture time-varying uncertainty about future labor income, we allow labor income innovations to have time-varying conditional second moments.

To begin, we assume that labor income follows the unit-root process $Y_{t+1} = Y_t + W_{t+1}$, where W_{t+1} has a Gaussian conditional distribution that is centered at 0 with variance $E_t W_{t+1}^2$. We let $V_{t+1} = C_{t+1} - E_t C_{t+1}$ be the one-step-ahead forecast error for consumption. If the conditional distribution of V_{t+1} is Gaussian with mean zero, then equation (2.3) implies that

$$C_{t+1} = C_t + \frac{\theta}{2} E_t V_{t+1}^2 + V_{t+1}, \tag{2.4}$$

where $E_t V_{t+1}^2$ is the conditional variance of the consumption forecast error in period $t+1$.

To relate the optimal consumption path to the innovations to labor income, we must relate $\{V_{t+i}\}_{i=1}^{\infty}$ to $\{W_{t+i}\}_{i=1}^{\infty}$. To do so, we follow Caballero (1990) and rewrite the intertemporal budget constraint as

$$\begin{aligned}\sum_{i=1}^{\infty} \alpha^i \Bigg[C_t &+ \sum_{j=1}^{i} \frac{\theta}{2} E_t V_{t+j}^2 \\ &+ \sum_{j=1}^{i} \frac{\theta}{2} (E_{t+j-1} V_{t+j}^2 - E_t V_{t+j}^2) + \sum_{j=1}^{i} V_{t+j} \\ &- \sum_{j=1}^{i} W_{t+j} - E_t Y_{t+i} \Bigg] = A_t,\end{aligned} \tag{2.5}$$

where $\alpha = (1+r)^{-1}$. The algebraic steps that lead to equation (2.5) are not particularly enlightening and so are contained in the appendix. If we divide equation (2.5) into two components, then the basic prediction of the theory of precautionary saving falls out. The component that holds in period t is given by the period-t conditional expectation of equation (2.5)

$$\begin{aligned} C_t = \frac{1-\alpha}{\alpha} &\left(A_t + \sum_{i=1}^{\infty} \alpha^i E_t Y_{t+i} \right) \\ &- \frac{1-\alpha}{\alpha} \sum_{i=1}^{\infty} \alpha^i \left(\sum_{j=1}^{i} \frac{\theta}{2} E_t V_{t+j}^2 \right). \end{aligned} \tag{2.6}$$

Because $[(1-\alpha)/\alpha]\,(A_t + \Sigma_{i=1}^{\infty} \alpha^i E_t Y_{t+i})$ equals permanent income, the second term on the right-hand side of equation

(2.6) captures the amount that consumption is reduced in the face of uninsurable labor income uncertainty.[7]

The component of equation (2.5) that holds in future periods, that is, beyond period t, is obtained by substituting the right-hand side of equation (2.6) into equation (2.5):

$$\sum_{i=1}^{\infty} \alpha^i \left(\sum_{j=1}^{i} V_{t+j} + \sum_{j=1}^{i} \frac{\theta}{2} (E_{t+j-1} V_{t+j}^2 - E_t V_{t+j}^2) - \sum_{j=1}^{i} W_{t+j} \right) = 0. \tag{2.7}$$

If we let $X_{t+j} = (\theta/2)\,(E_{t+j-1}V_{t+j}^2 - E_t V_{t+j}^2)$, then equation (2.7) is rewritten as

$$\sum_{i=1}^{\infty} \alpha^i \left(\sum_{j=1}^{i} V_{t+j} + \sum_{j=1}^{i} X_{t+j} - \sum_{j=1}^{i} W_{t+j} \right) = 0,$$

which is satisfied for all t, if and only if

$$\sum_{i=1}^{\infty} \alpha^i (V_h + X_h - W_h) = 0 \text{ for all } h. \tag{2.8}$$

The solution to equation (2.8) is $V_h + X_h - W_h = 0$ for all h. In particular, for $h = t + 1$:

$$V_{t+1} = W_{t+1},$$

because $X_{t+1} = 0$. For periods further into the future, V_{t+i} equals W_{t+i} adjusted for future revisions in conditional expectations. For period $t + 2$, $V_{t+2} = W_{t+2} - (\theta/2)\,(E_{t+1} V_{t+2}^2 - E_t V_{t+2}^2)$.

Because V_{t+1} equals W_{t+1}, equation (2.4) becomes

$$C_{t+1} - C_t = \frac{\theta}{2} E_t W_{t+1}^2 + W_{t+1}. \tag{2.9}$$

For a representative consumer with precautionary savings, changing labor income uncertainty affects the value of consumption that solves the maximization problem in equation (2.2). The positive effect of labor income uncertainty on the change in consumption is due to the reduction in current consumption reflected in equation (2.6).

For mathematical simplicity, we derive equation (2.9) under the assumption that labor income follows a random walk. Because the first difference of labor income, W_{t+1}, may be serially correlated, it is natural to ask how serial correlation affects equation (2.9). If W_{t+1} is a serially correlated random variable, then the coefficient for labor income uncertainty is a function of θ and the squared sum of the serial correlation parameters for W_{t+1}. Although the coefficient is no longer interpretable as one-half the coefficient of absolute risk aversion, our essential result—that labor income uncertainty positively affects the change in consumption—is unaltered.

Because disposable income is equal to the sum of consumption and savings, a model of consumption is implicitly a model of savings. In fact, the precautionary saving theory argues that income uncertainty directly affects savings, and it is the effect on savings that feeds through to consumption. Let $Y_t^d = Y_t + rA_{t-1}$ be disposable income in period t, and let Y_t^p be permanent income in period t. If we substitute the equalities $Y_t^p = [(1-\alpha)/\alpha]\,(A_t + \Sigma_{i=1}^{\infty} \alpha^i E_t Y_{t+i})$ and $S_t + C_t = Y_t^d$ into equation (2.6), then

$$S_t = Y_t^d - Y_t^p + \frac{1-\alpha}{\alpha} \sum_{i=1}^{\infty} \alpha^i \left(\sum_{j=1}^{i} \frac{\theta}{2} E_t V_{t+j}^2 \right). \tag{2.10}$$

From equation (2.10) we see that S_t has two components. The first component, $Y_t^d - Y_t^p$, captures the standard role of saving in smoothing consumption, where saving anticipates future declines in income. The second component, $[(1-\alpha)/\alpha]\Sigma_{i=1}^{\infty}\alpha^i(\Sigma_{j=1}^{i}\,(\theta/2)E_t V_{t+j}^2)$, captures the amount of savings that is due to the riskiness of expected future labor income. That is, if uncertainty about expected future income increases, then savings increase.

A. Regression Specification

Our test of precautionary saving is based on the significance of the conditional variance of labor income shocks in a regression with a function of consumption or savings as the dependent variable. To ensure that our estimators are constructed from stationary random variables and that our estimates are comparable with those contained in previous studies, we use the consumption growth rate as a dependent variable rather than the first difference of consumption. Similarly, we use savings rates, rather than the level of savings, as a dependent variable.

We begin with specification of the consumption regression. To transform equation (2.9), we divide both sides by C_t and multiply and divide the right-hand side by Y_t^2 and estimate the consumption adjustment regression

$$\Delta \ln C_{t+1} = \beta_0 + \beta_1 \frac{Y_t^2}{C_t} E_t \frac{W_{t+1}^2}{Y_t^2} + \frac{W_{t+1}}{C_t}, \tag{2.11}$$

where (β_0, β_1) is a vector of parameters. Under rational expectations, $E_t(W_{t+1}/C_t) = 0$; so (W_{t+1}/C_t) is interpretable as the forecast error for the consumption growth rate, and $E_t(W_{t+1}^2/Y_t^2)$ is interpretable as the forecast error variance for the logarithm of labor income.[8] The precautionary saving

[7] Our result is distinct from that of Caballero (1990) in that equation (2.6) is not a closed-form solution as the right-hand side contains $E_t V_{t+j}^2$.

[8] Our results accord with previous results derived (under the assumption that the representative consumer has constant relative risk aversion utility) by Skinner (1988) and Zeldes (1989) who show that precautionary saving affects the optimal consumption path. Although both authors assume

theory implies that $\beta_1 = (\theta/2)$ is positive, although the magnitude of β_1 depends on the degree of risk aversion. The presence of the regressor $(Y_t^2/C_t)\, E_t\, (W_{t+1}^2/Y_t^2)$ in equation (2.11) indicates that, under precautionary saving, the conditional expectation of the consumption growth rate is not constant but is instead a function of income uncertainty.

Accounting for precautionary saving may also help explain the empirical finding that the expected income growth rate predicts the consumption growth rate. This finding, often referred to as the excess sensitivity of consumption, may simply arise from the fact that income uncertainty is incorrectly omitted from the regression and that income uncertainty is correlated with expected future income. To test for excess sensitivity of consumption under changing income uncertainty, we test the null hypothesis that β_2 equals zero in the consumption adjustment regression

$$\Delta \ln C_{t+1} = \beta_0 + \beta_1 \frac{Y_t^2}{C_t} E_t \frac{W_{t+1}^2}{Y_t^2} + \beta_2 E_t \Delta \ln Y_{t+1}^d + \frac{W_{t+1}}{C_t}, \tag{2.12}$$

where $E_t \Delta \ln Y_{t+1}^d$ is the expected growth rate of disposable income.[9]

To obtain a specification for the savings regressions, we begin with a savings adjustment regression that is analogous to equation (2.11). To transform equation (2.9), we replace C_t with $Y_t^d - S_t$ and divide both sides by Y_t^d:

$$\frac{\Delta S_{t+1}}{Y_t^d} = \beta_0 + \beta_1 Y_t^d E_t \frac{W_{t+1}^2}{Y_t^2} + \beta_2 \Delta \ln Y_{t+1}^d + \frac{W_{t+1}}{Y_t^d}, \tag{2.13}$$

where $\beta_1 = (\theta/2)$ and $\beta_2 = 1$. Once again, an increase in the conditional variance of the income growth rate reduces ΔS_{t+1} because it increases S_t, leaving S_{t+1} unchanged.

All of the specifications given above assume that adjustment in consumption and savings occurs completely within a given period as income uncertainty changes over time. As Carroll (1992) conjectures, instantaneous adjustment in consumption may be difficult, so that the level of consumption and savings may adjust incompletely within one period in the face of an increase in the level of income uncertainty.[10] To determine the level of empirical support for incomplete adjustment in consumption and savings, we first estimate a savings *level* regression. To derive a specification for a savings (rather than a savings adjustment) regression, we approximate equation (2.10). To approximate, we replace Y_t^p with Y_{t+1}^d, and we replace the third term on the right-hand side of equation (2.10) with $E_t W_{t+1}^2$. We then divide both sides by Y_t^d, and replace $E_t[W_{t+1}^2/(Y_t^d)^2]$ with $E_t(W_{t+1}^2/Y_t^2)$:

$$\frac{S_t}{Y_t^d} = \beta_0 + \beta_1 Y_t^d E_t \frac{W_{t+1}^2}{Y_t^2} + \beta_2 E_t \Delta \ln Y_{t+1}^d - \frac{W_{t+1}}{Y_t^d}. \tag{2.14}$$

The precautionary saving theory predicts that β_1 is positive and β_2 is negative; the latter implication follows because higher expected future income lowers saving, as in standard certainty-equivalence permanent income models (Campbell, 1987).

If savings do not adjust instantaneously to changes in income uncertainty, then an increase in income uncertainty in period t leads to an increase in both S_t and S_{t+1}. If both S_t and S_{t+1} increase, then β_1 in equation (2.14) is larger than β_1 in equation (2.13). As a result, a statistically significant estimate of β_1 in equation (2.14) and a statistically insignificant estimate of β_1 in equation (2.13) is evidence of incomplete adjustment. Because adjustments to savings are mirrored by adjustments to consumption, evidence of incomplete adjustment is also provided by a statistically insignificant estimate of β_1 in equation (2.12). To capture incomplete adjustment in response to changing income uncertainty we estimate the following consumption and savings adjustment regressions:

$$\ln C_{t+p} - \ln C_t = \beta_0 + \beta_1 \frac{Y_t^2}{C_t} E_t \frac{W_{t+1}^2}{Y_t^2} + U_{t+p}, \text{ and} \tag{2.15}$$

$$\ln C_{t+p} - \ln C_t = \beta_0 + \beta_1 \frac{Y_t^2}{C_t} E_t \frac{W_{t+1}^2}{Y_t^2} + \beta_2 E_t (\ln Y_{t+p}^d - \ln Y_t^d) + U_{t+p}, \tag{2.16}$$

where p is the number of quarters over which the adjustment process takes place. If p is greater than 1, the adjustment process is not completed within one period implying that an increase in income uncertainty in period t lowers consump-

constant income uncertainty, they also show, through comparative static analysis, that higher income uncertainty raises the expected consumption growth rate.

[9] Campbell and Mankiw (1989, 1990) construct a model in which there are two types of consumers: those who consume current income (due to liquidity constraints), and those who consume permanent income. The extension of their model into our framework implies that β_2 in equation (2.12) represents the share of consumers who consume current income. Campbell and Mankiw, who omit the uncertainty term, report point estimates of β_2 that cluster around 0.5, indicating that only half of income falls to consumers who follow the PIH. Of course, improperly omitting the uncertainty term biases the estimator of β_2. In fact, if $\beta_1 > 0$ in equation (2.12) and the uncertainty term and $\Delta \ln Y_{t+1}^d$ are positively correlated, then omitting the uncertainty term leads to an upward bias in the instrumental variables estimator of β_2, which understates the adequacy of the PIH to describe consumption.

[10] A number of authors have shown that consumption adjustment is incomplete within one period in their models: Constantinides (1990) and Heaton (1990) include habit formation; Goodfriend (1992) and Pischke (1995) include information lags; and Bertola and Caballero (1990) include other adjustment costs.

tion in both period t and period $t + 1$, so the sign on the coefficient of income uncertainty in the standard Euler equation is not clear. Clearly, if the adjustment process takes p periods or less, an increase in income uncertainty in period t leaves consumption in period $t + p$ unchanged, so the coefficient on income uncertainty in equation (2.15) is positive. To estimate the model, we set p equal to 4 to capture adjustment processes that are not instantaneous but are completed within one year.

III. Data

Our data set is novel in that we use a survey measure of income uncertainty. Recall that our regression model requires the conditional variance of the income growth rate as a regressor. Because we observe only one time series for income, we cannot construct a conditional variance of the income growth rate from the observed time series on income without parametric assumptions. One set of parametric assumptions, which is popular in the empirical finance literature, is to parameterize the conditional variance with a generalized autoregressive (GARCH) model. Yet any parametric model suffers from the weakness that there is little economic motivation for the specific parametric form of the model.[11] The problem is potentially serious as results often differ substantially over different parametric forms. We are able to avoid the problem by using what is, in effect, a nonparametric measure of conditional variance. That is, rather than trying to infer the conditional variance of income growth rates from past observations of income, we have a direct measure, namely the survey responses of forecasters of income.

Why do the surveyed income forecasts provide information on the conditional variance of income? To understand why, consider the model $\Delta \ln Y_{t+1} = \mu_t + V_{t+1}$, where μ_t is the conditional mean given all information available in period t and V_{t+1} is the forecast error. The conditional variance of the income growth rate—where use of the word *conditional* means conditional on period-t information—is the conditional variance of the forecast error. To measure the conditional variance of the forecast error we first construct the one-quarter forecast of the income growth rate for each of the forecasters in the panel, and, second, construct the variance, across forecasters, of the one-quarter forecast. Since our measure of the conditional variance of the forecast error is the variance of the individual point forecasts of the quarterly income growth rate, we refer to our measure as the point forecast measure of the conditional variance.

In addition to the point forecasts, in each survey the respondents are also asked to provide additional information about the distribution of their annual forecast. Specifically, a range of intervals is provided (each interval is of the form "Income will increase between 2.0 and 2.9 percent"), and the respondents assign probabilities to the intervals. The range of probabilities assigned to the intervals leads to a measure of the conditional variance of the forecast error for each respondent. The average across respondents, of the conditional variance for each respondent, is an interval forecast measure of the conditional variance of the forecast error. If respondents exercise as much care in assigning probabilities to the listed intervals as they do in constructing their point forecasts, then the interval forecast measure of the conditional variance should be at least as accurate as the point forecast measure of the conditional variance.

Unfortunately, the interval forecasts are only surveyed at the annual frequency, so there is no consistent interval forecast at the quarterly frequency. To determine the adequacy of our point forecast measure of the conditional variance of the forecast error in quarterly income growth rates, we turn to the correspondence between the interval forecast measure and the point forecast measure for the conditional variance of annual income growth rates. Zarnowitz and Lambros (1987) study the relation between these two measures of the conditional variance for the annual income growth rates contained in the survey. They find that the correlation coefficient between the two measures is 0.71, which suggests that our point forecast measure is an informative measure of the conditional variance of the forecast errors.

Of course, survey measures are not without drawbacks. Our survey was initially gathered by the American Statistical Association, in conjunction with the National Bureau of Economic Research, and begun in 1968. Over time, responsibility for gathering the survey data has shifted to the Federal Reserve Bank of Philadelphia. Each of the institutions could potentially survey different groups. (In fact, the Federal Reserve Bank of Philadelphia has attempted to continue the original survey design used by the ASA/NBER, and the survey group is restricted to professional forecasters.) Note that we do not have a panel of forecasters in which we track a specific group of forecasters over a given time period. Rather, our individual forecasters exit and enter the survey, and they occasionally fail to respond. Thus, we have different numbers of responses in different quarters, and we do not track a consistent group of forecasters. Further, the number of responses differs from period to period. Because we do not sample the exact same group of forecasters in each period—and, even if we did, because our group of professional forecasters does not cover the entire population—our regressor is measured with error. In the econometric work we describe below, we take care to treat measurement error and calculate estimators of the parameters that are consistent in the presence of measurement error. Details on the construction of a measurement-error consistent estimator are contained in the appendix.

[11] As evidence of the problems of statistical modeling without theory, we estimate a GARCH(1,1) model for disposable income in which the conditional mean is a first-order autoregression. A GARCH(1,1) model parameterizes the conditional variance for disposable income in period t, J_t, as $J_t = \gamma_0 + \gamma_1(Y^d_{t-1} - E_{t-2}Y^d_{t-1})^2 + \gamma_2 J_{t-1}$. The estimate of γ_2 is negative, which implies that forecasts of the conditional variance of disposable income can be negative.

To construct our measure of income uncertainty, $E_t(W^2_{t+1}/Y^2_t)$, we use survey responses for real GDP. Although the model in section II is constructed from labor income, the closest measure to labor income in our survey is real GDP.[12] Each survey response contains forecasts of the level of real GDP for four quarters into the future. From the forecast level of GDP, we construct the implicit forecast of the growth rate of real GDP.

Our measurements of the other variables (namely consumption, disposable income, and savings) are drawn from the U.S. National Income and Product Accounts. For C_t, we use quarterly consumption of nondurables and services; for Y^d_t, we use quarterly disposable personal income; and, for S_t, we use quarterly personal savings, where all series are per capita, seasonally adjusted, and measured in 1987 dollars.[13] Because survey data for real GDP is gathered beginning in the third quarter of 1981, our sample period begins in the third quarter of 1981 and continues through the fourth quarter of 1994.

Table 1 presents descriptive statistics of the data used in the empirical analysis. In the summary statistics panel, the first row contains the mean, and the second row contains the standard deviation of each variable. Below the summary statistics panel, we report the matrix of correlations between the variables. The estimated correlations support the theory that our income uncertainty measure, $E_t(W^2_{t+1}/Y^2_t)$, is positively correlated with both real per capita consumption growth rates and savings rates, and is negatively correlated with the ratio of the change in savings to disposable income.

We present the mean forecast growth rate and the actual growth rate of real GDP in figure 1. As the figure indicates, the growth rate forecasts are unbiased. (The average deviation over the entire sample is −0.0436% per year with a standard deviation of 2.5773.) The correlation between the mean forecast growth rate and the actual growth rate is 0.51. In figure 2, we plot the conditional standard deviation of the income growth rate from the survey. As can be seen from figure 2, there is a marked change in the conditional variance of the forecasts over time. This may reflect the high variability of inflation in the early 1980s.

TABLE 1.—DATA DESCRIPTION, 1981:III–1994:IV

Summary Statistics

$\Delta \ln C_t$	$\Delta \ln Y_t$	(S_t/Y_t)	$(\Delta S_t/Y_t)$	$E_t(W^2_{t+1}/Y^2_t)$
1.502	1.543	5.644	−0.187	3.478
(1.546)	(3.595)	(1.746)	(3.726)	(4.790)

Correlation Matrix

	$\Delta \ln C_t$	(S_t/Y_t)	$(\Delta S_t/Y_t)$	$\Delta \ln Y_t$
(S_t/Y_t)	0.245			
$(\Delta S_t/Y_t)$	−0.187	−0.313		
$\Delta \ln Y_t$	0.287	−0.162	0.809	
$E_t(W^2_{t+1}/Y^2_t)$	0.198	0.658	−0.175	−0.020

Note: Growth rates for C_t and Y_t are measured in percent per year.

IV. Empirical Results

For each regression specification, we test the significance of income uncertainty. More precisely, we test the null hypothesis that the PIH without income uncertainty holds against the alternative hypothesis that the PIH with income uncertainty holds. For each specification, the null and alternative hypotheses have precise implications for the coefficient on the uncertainty terms: namely, $\beta_1 = 0$ under the null hypothesis, and $\beta_1 > 0$ under the alternative hypothesis for the consumption adjustment and savings rate regressions; and $\beta_1 = 0$ under the null hypothesis, and $\beta_1 < 0$ under the alternative hypothesis for the savings adjustment regressions. Because the alternative hypothesis for β_1 is one sided, we construct one-sided significance tests for β_1. Rejection of the null hypothesis is thus support for our model.

Many of the specifications contain an additional regressor that is a function of expected disposable income. The coefficient on this additional regressor, β_2, is assumed to have the same value under both the null and alternative hypotheses for β_1. As a result, it is not a simple matter to construct joint significance tests. We proceed by constructing separate significance tests. Tests of β_2 are general tests of the adequacy of the PIH with income uncertainty. Because violations of the PIH imply a two-sided rejection region for β_2, we use two-sided significance tests for β_2.

A. Instantaneous Adjustment

To determine the adequacy of the prediction that consumption adjusts instantaneously to changes in income uncertainty, in table 2 we report estimates from equation (2.11) and equation (2.12). In table 2 (as in each of the remaining

FIGURE 1.—REAL GDP GROWTH RATES: REALIZED AND MEAN FORECAST VALUE

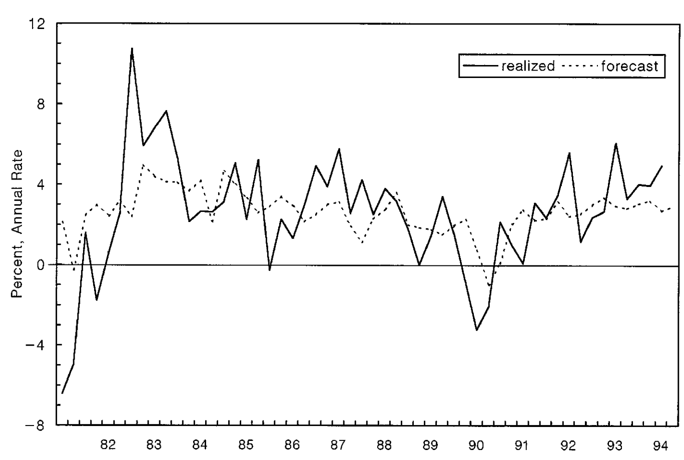

[12] For the sample period, the standard deviation of the growth rate of disposable income is 3.6%, while the standard deviation of the growth rate of GDP is 3.0%.

[13] The survey response is a forecast of real GDP rather than per capita real GDP. We constructed the conditional variance of surveys for the growth rate of real GDP and per capita real GDP. Because our conditional variance measures were virtually identical (which reflects the stability of the population growth rate over our sample), we use the conditional variance of the growth rate of real GDP because that is the quantity reported in the survey.

FIGURE 2.—STANDARD DEVIATION OF REAL GDP GROWTH RATE FORECASTS

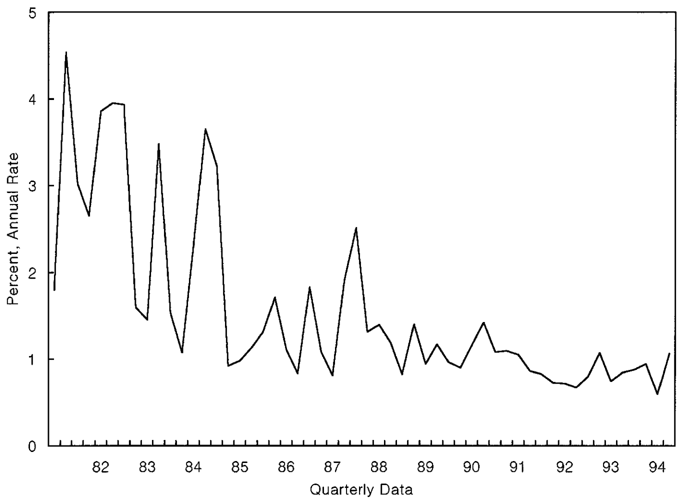

TABLE 2.—CONSUMPTION AND INCOME UNCERTAINTY, 1981:III–1994:IV

A. Consumption Adjustment Regression

$$\Delta \ln C_{t+1} = \beta_0 + \beta_1(Y_t^2/C_t)\, E_t(W_{t+1}^2/Y_t^2) + U_{t+1}$$

$\hat{\beta}_1$	$\hat{\beta}_1^{me}$
0.0042	0.0029
(0.0032)	$F_{.05} = -0.0013$

B. The Permanent Income Hypothesis under Changing Income Uncertainty

$$\Delta \ln C_{t+1} = \beta_0 + \beta_1(Y_t^2/C_t)\, E_t(W_{t+1}^2/Y_t^2) + \beta_2 E_t \Delta \ln Y_{t+1}^d + U_{t+1}$$

$\hat{\beta}_1$	$\hat{\beta}_1^{me}$	$\hat{\beta}_2$	$\hat{\beta}_2^{me}$
0.0049*	0.0029	0.2133	0.1832
(0.0029)	$F_{.05} = -0.0014$	(0.1164)	(0.1325)

Notes: An * indicates rejection of the null hypothesis at the 5% significance level. Because the alternative hypothesis is $\beta_1 > 0$, we reject the null hypothesis that $\beta_1 = 0$ at the 5% significance level if the fractile corresponding to the lower 5% of the empirical distribution of $\hat{\beta}_1^{me}$ (which is $F_{0.05}$) exceeds zero.

tables), $\hat{\beta}$ denotes the least-squares estimator (two-stage least squares for specifications that contain $E_t \Delta \ln Y_{t+1}^d$, and OLS for all other specifications), and $\hat{\beta}^{me}$ denotes the measurement-error consistent estimator.[14] In parentheses below $\hat{\beta}$, we report estimated (serial-correlation and hetero-skedasticity consistent) standard errors, and below $\hat{\beta}^{me}$ we report either the appropriate fractiles from the empirical distribution of $\hat{\beta}^{me}$ for one-sided tests or the standard error from the empirical distribution for two-sided tests. We use the empirical distribution for $\hat{\beta}^{me}$ constructed from bootstrap resampling, because the asymptotic covariance matrix is not easily obtained. (Details of the bootstrap algorithm are in the appendix.)[15] Again, because the alternative hypothesis is one sided, we reject the null hypothesis if the fractile corresponding to the lower 5% of the empirical distribution of $\hat{\beta}^{me}$ exceeds 0.

[14] To ensure that our results are driven by uncertainty about the income growth rate and not by variables such as Y_t^2/C_t or Y_t^d, we estimate all regressions with the uncertainty regressor set equal to $E_t(W_{t+1}^2/Y_t^2)$. The results from these regressions, which are available on request, are essentially the same as the results we report.

[15] The fractiles from the empirical distribution for $\hat{\beta}^{me}$ are denoted $F_{.05}$ and $F_{.95}$, where five percent of the bootstrap values are less than $F_{0.05}$ and 95% of the bootstrap values are less than $F_{0.95}$.

TABLE 3.—SAVING AND INCOME UNCERTAINTY, 1981:III–1994:IV

A. Savings Adjustment Regression

$$(\Delta S_{t+1}/Y_t^d) = \beta_0 + \beta_1 Y_t^d E_t(W_{t+1}^2/Y_t^2) + \beta_2 \Delta \ln Y_{t+1}^d + U_{t+1}$$

$\hat{\beta}_1$	$\hat{\beta}_1^{me}$	$\hat{\beta}_2$	$\hat{\beta}_2^{me}$
−0.0086	−0.0136*	0.6962	0.7369
(0.0066)	$F_{0.05} = -0.0076$	(0.3017)	(0.2384)

B. Savings Rate Regression

$$(S_t/Y_t^d) = \beta_0 + \beta_1 Y_t^d E_t(W_{t+1}^2/Y_t^2) + \beta_2 \Delta \ln Y_{t+1}^d + U_{t+1}$$

$\hat{\beta}_1$	$\hat{\beta}_1^{me}$	$\hat{\beta}_2$	$\hat{\beta}_2^{me}$
0.0177*	0.0170*	−0.4141	−0.4980
(0.0021)	$F_{0.05} = 0.0144$	(0.1188)	(0.1061)

Notes: An * indicates rejection of the null hypothesis at the 5% significance level. Because the alternative hypothesis is $\beta_1 < 0$ for the specification in panel A, we reject the null hypothesis that $\beta_1 = 0$ at the 5% significance level if the fractile corresponding to the upper 95% of the empirical distribution of $\hat{\beta}_1^{me}$ (which is $F_{0.95}$) is less than zero.

From panel A, which contains estimates for equation (2.11), we see that both estimates of β_1 are positive as theory predicts. However, both estimates are also insignificant at the 5% significance level. (The estimated value of $\hat{\beta}^{me}$ is not significant because the upper 95% range includes zero as indicated by the negative value for the 5% fractile.) From panel B, which contains estimates for equation (2.12), both estimates of β_1 are again positive. Further, $\hat{\beta}_1$ is significant (recall that the appropriate critical value is 1.645), although $\hat{\beta}_1^{me}$ is insignificant at the 5% significance level.[16] Because the PIH implies $\beta_2 = 0$ for the specification in panel B, the statistically insignificant estimates of β_2 support the PIH. Thus, our consumption adjustment regressions provide evidence that income uncertainty affects consumption (although measurement error could be driving the effect) and that income uncertainty may remove the effect of expected income growth on consumption as indicated by the insignificant estimates of β_2 in panel B.

Next, we turn to the savings regressions in table 3. In panel A, we report results on the savings adjustment regression equation (2.13). Both $\hat{\beta}_1$ and $\hat{\beta}_1^{me}$ are negative, as predicted by the precautionary saving theory. While $\hat{\beta}_1$ is insignificant at the 5% significance level, $\hat{\beta}_1^{me}$ is significantly less than zero, as indicated by the negative 95% fractile. Correcting for measurement error results in evidence in support of the precautionary saving theory, as an increase in income uncertainty raises current savings S_t and lowers $\Delta S_{t+1}/Y_t^d$. Because the PIH implies $\beta_2 = 1$ for the specification in panel A, the result that both estimates of β_2 are not significantly different from one supports the PIH.

Further evidence in support of the precautionary saving theory is contained in panel B, where we report results for the savings rate regression (equation (2.14)). Both $\hat{\beta}_1$ and $\hat{\beta}_1^{me}$ are significantly greater than zero, which implies that increasing income uncertainty immediately increases savings. To assess the magnitude of the effect, we use $\hat{\beta}_1^{me}$ to

[16] We use two-stage least squares to estimate regressions with income growth rates. The regressors for the first-stage regression are a constant, the first four lags of both income growth rates and consumption growth rates, and the first lag of the logarithm of C_t/Y_t.

TABLE 4.—UNCERTAINTY EFFECT UNDER INCOMPLETE ADJUSTMENT, 1981:III–1994:I

A. Consumption Adjustment Regression

$$\ln C_{t+4} - \ln C_t = \beta_0 + \beta_1 (Y_t^2/C_t)\, E_t(W_{t+1}^2/Y_t^2) + U_{t+4}$$

$\hat{\beta}_1$	$\hat{\beta}_1^{me}$
0.0056*	0.0041*
(0.0022)	$F_{0.05} = 0.0003$

B. Excess Sensitivity of Consumption to Current Income

$$\ln C_{t+4} - \ln C_t = \beta_0 + \beta_1 (Y_t^2/C_t)\, E_t(W_{t+1}^2/Y_t^2) + \beta_2 E_t(\ln Y_{t+4}^d - \ln Y_t^d) + U_{t+4}$$

$\hat{\beta}_1$	$\hat{\beta}_1^{me}$	$\hat{\beta}_2$	$\hat{\beta}_2^{me}$
0.0058*	0.0041*	0.2268	0.1983
(0.0019)	$F_{0.05} = 0.0002$	(0.1365)	(0.1773)

Note: An * indicates rejection of the null hypothesis at the 5% significance level.

infer that an increase of one standard deviation in the conditional variance of the income growth rate, given that Y_t is set to the sample mean value of 13.58 (in thousands of 1987 dollars), leads to an increase in the savings rate of approximately 1.1 percentage points. Note that the estimates of β_2 (which is the coefficient on the instrumentized income growth rate) are significantly less than zero. The negative estimates are consistent with the savings behavior of forward-looking consumers under the PIH, as noted by Campbell (1987).

B. Incomplete Adjustment

The results from the savings rate regression indicate that income uncertainty has an effect on current savings, but the effect is not easily detected in savings and consumption adjustment regressions. The reason may be that the adjustment process takes more than one period. If this is the case, then a change in income uncertainty today affects consumption and savings both today and tomorrow, so the measured effect on the difference is reduced. To measure the impact of incomplete adjustment within one period, in table 4 we report empirical results for equation (2.15) and (2.16). In panel A, both $\hat{\beta}_1$ and $\hat{\beta}_1^{me}$ are significantly greater than zero, which provides evidence that income uncertainty affects consumption but that the adjustment requires more than one period.

In panel B, we report estimates from equation (2.16), which includes the expected income growth rate as a regressor. For both the 2SLS and the consistent estimator, the same result emerges. Income uncertainty has a significantly positive effect on the change in consumption. Further, the estimated coefficient on the expected income growth rate is insignificant.[17] Because a significantly positive coefficient on the income growth rate is often referred to as the excess sensitivity of consumption to current income, our results suggest that previous findings of excess sensitivity may be due to the incorrect assumption that the conditional variance of income is constant.

[17] Because the shock to consumption is the shock to income, the income growth rate regressor ($\ln Y_{t+p} - \ln Y_t$) is correlated with the regression error. To remove the correlation, we use two-stage least squares, in which the instruments used in the first-stage regression are ($\ln Y_t - \ln Y_{t-p}$), ($\ln Y_{t-p} - \ln Y_{t-2p}$), ($\ln C_t - \ln C_{t-p}$), ($\ln C_{t-p} - \ln C_{t-2p}$), and $\ln (C_t/Y_t)$.

V. Conclusion

We derive and estimate a simple framework in which consumers optimally revise their intertemporal consumption plan, not only in response to changes in the level of permanent income but also to changes in their uncertainty about future income. We find that our measure of income uncertainty changes significantly over time, indicating that the popular assumption of constant income uncertainty over time is misleading. Further, the precautionary savings response to changing income uncertainty is a significant source of observed changes to both consumption and savings, and, the higher the uncertainty level, the more precautionary savings consumers accumulate. However, the adjustment does not seem to occur instantaneously, possibly due to information lags or adjustment costs. The estimates from the incomplete adjustment model indicate that the excess sensitivity of consumption to current income may be partially explained by the role of time-varying income uncertainty operating through precautionary savings.

The overall evidence indicates that there exist forward-looking consumers who adjust precautionary savings in response to changing income uncertainty. Although our research focuses on consumption of nondurable goods and services, our results also have implications for consumption of durable goods. Because durable consumption is believed to be quite volatile over the business cycle and sensitive to consumer sentiment, future models of the optimal consumption of durable goods should include time-varying income uncertainty. The consumption response to changes in uncertainty about future income, which is the optimal precautionary savings response, is a potentially important and previously overlooked component of adjustment in both consumption and savings.

REFERENCES

Barsky, Robert B., N. Gregory Mankiw, and Stephen P. Zeldes, ''Ricardian Consumers with Keynesian Propensities,'' *American Economic Review* 76 (September 1986), 676–691.

Bertola, G., and R. Caballero, ''Kinked Adjustment Costs and Aggregate Dynamics,'' in O. Blanchard and S. Fischer (eds.), *NBER Macroeconomics Annual* (Cambridge: MIT Press, 1990), 237–295.

Blanchard, Olivier J., and N. Gregory Mankiw, ''Consumption: Beyond Certainty Equivalence,'' *American Economic Review* 78 (May 1988), 173–177.

Caballero, Ricardo J., ''Consumption Puzzles and Precautionary Savings,'' *Journal of Monetary Economics* 25 (January 1990), 113–136.

——— ''Earnings Uncertainty and Aggregate Wealth Accumulation,'' *American Economic Review* 81 (September 1991), 859–871.

Campbell, John Y., ''Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis,'' *Econometrica* 55 (November 1987), 1249–1273.

Campbell, John Y., and Angus Deaton, ''Why is Consumption so Smooth?,'' *Review of Economic Studies* 56 (July 1989), 357–374.

Campbell, John Y., and N. Gregory Mankiw, ''Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence,'' in O. Blanchard and S. Fischer (eds.), *NBER Macroeconomics Annual* (Cambridge: MIT Press, 1989).

——— ''Permanent Income, Current Income, and Consumption,'' *Journal of Business and Economic Statistics* 8 (July 1990), 265–279.

Carroll, Christopher D., ''The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,'' *Brookings Papers on Economic Activity,* (Issue 2, 1992), 61–127.

Carroll, C., and M. Kimball, ''On the Concavity of the Consumption Function,'' *Econometrica* 64 (July 1996), 981–992.

Carroll, Christopher D., and Andrew A. Samwick, ''The Nature and Magnitude of Precautionary Wealth,'' Federal Reserve Board of Governors, Working Paper (March, 1992).

Christiano, Lawrence J., Martin Eichenbaum, and David Marshall, ''The Permanent Income Hypothesis Revisited,'' *Econometrica* 59 (March 1991), 397–423.

Clarida, Richard H., ''Aggregate Stochastic Implications of the Life-cycle Hypothesis,'' *Quarterly Journal of Economics* 106 (August 1991), 851–867.

Cochrane, John, H., ''A Simple Test of Consumption Insurance,'' *Journal of Political Economy* 99 (October 1991), 957–976.

Constantinides, G., ''Habit Formation: A Resolution of the Equity Premium Puzzle,'' *Journal of Political Economy* 98 (June 1990), 519–543.

Deaton, Angus, ''Life-Cycle Models of Consumption: Is the Evidence Consistent with the Theory?'' in T. Bewley (ed.), *Advances in Econometrics: Fifth World Congress* (Cambridge: Cambridge University Press, 1987).

Dynan, Karen, ''How Prudent Are Consumers?,'' *Journal of Political Economy* 101 (December 1993), 1104–1113.

Feldstein, Martin, ''The Effects of Fiscal Policies When Incomes Are Uncertain: A Contradiction to Ricardian Equivalence,'' *American Economic Review* 78 (March 1988), 14–23.

Flavin, Marjorie, ''The Adjustment of Consumption to Changing Expectations About Future Income,'' *Journal of Political Economy* 89 (October 1981), 974–1009.

——— ''Excess Sensitivity of Consumption to Current Income: Liquidity Constraints or Myopia?,'' *Canadian Journal of Economics* 18 (February 1985), 117–136.

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

Gali, Jordi, ''Finite Horizons, Life Cycle Savings and Time Series Evidence on Consumption,'' *Journal of Monetary Economics* 26 (November 1990), 433–452.

——— ''Budget Constraints and Time Series Evidence on Consumption,'' *American Economic Review* 81 (December 1991), 1238–1253.

Goodfriend, M., ''Information Aggregation Bias,'' *American Economic Review* 82 (June 1992), 508–519.

Guiso, Luigi, Tullio Jappelli, and Danielle Terlizzese, ''Earnings Uncertainty and Precautionary Savings,'' Bank of Italy, Working Paper (April 1991).

Hall, Robert E., ''Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence,'' *Journal of Political Economy* 86 (December 1978), 971–987.

——— ''Intertemporal Substitution in Consumption,'' *Journal of Political Economy* 96 (April 1988), 339–357.

Hahm, Joon-Ho, ''Consumption Adjustment to Real Interest Rates: Intertemporal Substitution Revisited,'' *Journal of Economic Dynamics and Control* 22 (February 1998), 293–320.

Heaton, J., ''The Interaction Between Time-Nonseperable Preferences and Time Aggregation,'' MIT, Working Paper (1990).

Hubbard, R. Glenn, and Kenneth L. Judd, ''Social Security and Individual Welfare: Precautionary Saving, Liquidity Constraints, and the Payroll Tax,'' *American Economic Review* 77 (September 1987), 630–646.

Mankiw, N. Gregory, ''The Permanent Income Hypothesis and The Real Interest Rate,'' *Economics Letters* 7 (1981), 307–311.

——— ''Hall's Consumption Hypothesis and Durable Goods,'' *Journal of Monetary Economics* 10 (1982), 417–425.

Pal, M., ''Consistent Moment Estimators of Regression Coefficients in the Presence of Errors in Variables,'' *Journal of Econometrics* 14 (1980), 349–364.

Pischke, J., ''Individual Income, Incomplete Information, and Aggregate Consumption,'' *Econometrica* 63 (July 1995), 805–840.

Skinner, Jonathan, ''Risky Income, Life Cycle Consumption, and Precautionary Savings,'' *Journal of Monetary Economics* 22 (1988), 237–255.

West, Kenneth D., ''The Insensitivity of Consumption to News about Income,'' *Journal of Monetary Economics* 21 (January 1988), 17–34.

Zarnowitz, V., and L. Lambros, ''Consensus and Uncertainty in Economic Prediction,'' *Journal of Political Economy* 95 (June 1987), 591–621.

Zeldes, Stephen P., ''Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence,'' *Quarterly Journal of Economics* 104 (May 1989), 275–298.

APPENDIX

Derivation of the Budget Constraint

To derive equation (2.5), we add and subtract the conditional expectation of Y_{t+i} and substitute recursively for A_{t+i} in the period-t budget constraint, which yields

$$\sum_{i=1}^{\infty} \alpha^i C_{t+i} - \sum_{i=1}^{\infty} \alpha^i (Y_{t+i} - E_t Y_{t+i}) - \sum_{i=1}^{\infty} \alpha^i E_t Y_{t+i} = A_t. \tag{5.1}$$

Because $\{Y_t\}$ is a unit-root process, $Y_{t+i} - E_t Y_{t+i} = \Sigma_{j=1}^{i} W_{t+j}$. Because equation (2.4) implies that $C_{t+i} = C_{t+i-1} + (\theta/2) E_{t+i-1} V_{t+i}^2 + V_{t+i}$ for any value of i,

$$\begin{aligned} C_{t+i} = C_t &+ \sum_{j=1}^{i} \frac{\theta}{2} E_t V_{t+j}^2 + \sum_{j=1}^{i} \frac{\theta}{2} (E_{t+j-1} V_{t+j}^2 - E_t V_{t+j}^2) \\ &+ \sum_{j=1}^{i} V_{t+j}. \end{aligned} \tag{5.2}$$

The constraint (equation (2.5)) then follows by substituting into equation (5.1) the right-hand side of equation (5.2) for C_{t+i}, and $\Sigma_{j=1}^{i} W_{t+j}$ for $Y_{t+i} - E_t Y_{t+i}$.

Measurement-Error Consistent Estimator

As is well known, if a regressor is measured with error, OLS estimators are inconsistent. The measurement error problem is essentially an identification problem. To restore identification, and hence consistency, we turn to sufficient assumptions to separately identify the coefficients of the variables that are measured with error. To simplify notation, let $H_t = Y_t^2/C_t\, E_t W_{t+1}^2$ and let G_t generically denote the dependent variable. We do not observe H_t because it contains $E_t W_{t+1}^2$; rather we observe a contaminated version $H_t^c = H_t + M_t$, where M_t denotes measurement error. If we assume that mismeasurement of the conditional variance is not systematic—so that $EM_t = 0$, and that the regression error is symmetrically distributed, so that $EM_t^3 = 0$—then we are able to derive a consistent and asymptotically normal method of moments estimator. (See Pal (1980) for more details.) Our measurement-error consistent estimator of β_1 is $\hat{\beta}_1^{me} = m_{21}/m_{30}$, where $m_{21} = T^{-1}\Sigma_t (H_t - \overline{H})^2 (G_t - \overline{G})$ and $m_{30} = T^{-1}\Sigma_t (H_t - \overline{H})^3$.

To construct standard errors for $\hat{\beta}_1^{me}$, we use a bootstrap method. For the bootstrap simulations, we first estimate the regression model to obtain $\hat{\beta}^{me}$ and the residuals $\{\hat{U}_t\}_{t=1}^n$. We then estimate an AR(1) process for the residuals, $\hat{U}_t = \rho \hat{U}_{t-1} + V_t$, to obtain $\hat{\rho}$ and a second set of residuals, $\{\hat{V}_t\}_{t=1}^n$. (We use the AR(1) specification throughout, because there is little evidence of serial correlation in $\{\hat{V}_t\}_{t=1}^n$.)We sample with replacement from the white-noise sequence $\{\hat{V}_t\}_{t=1}^n$ to obtain a bootstrap sample $\{\hat{V}_t^*\}_{t=1}^n$. For each bootstrap sample, we use $\hat{\rho}$ to construct $\{\hat{U}_t^*\}_{t=1}^n$. and then use $(\beta^{me}, \{H_t\}_{t=1}^n, \{\hat{U}_t^*\}_{t=1}^n)$ to construct $\{G_t^*\}_{t=1}^n$. For each bootstrap sample $\{G_t^*, H_t\}_{t=1}^n$, we obtain a bootstrap value of the estimator $\hat{\beta}^{me*}$. We repeat the procedure 1,000 times to create a bootstrap distribution $F^*(\hat{\beta}^{me*})$.