Chapter 7
CONSUMPTION

This chapter and the next investigate households' consumption choices and firms' investment decisions in more detail. Consumption and investment are important to both growth and fluctuations. With regard to growth, the division of society's resources between current consumption and various types of investment—in physical capital, human capital, and research and development—is central to standards of living in the long run. That division is determined by the interaction of households' allocation of their incomes between consumption and saving given the rates of return and other constraints they face, and firms' investment demand given the interest rates and other constraints they face. With regard to fluctuations, consumption and investment make up the vast majority of the demand for goods. Thus if we wish to understand how such forces as government purchases, technology, and monetary policy affect aggregate output, we must understand how consumption and investment are determined.

There are two other reasons for studying consumption and investment. First, they introduce some important issues involving financial markets. Financial markets affect the macroeconomy mainly through their impact on consumption and investment. In addition, consumption and investment have important feedback effects on financial markets. We will investigate the interaction between financial markets and consumption and investment both in cases where financial markets function perfectly and in cases where they do not.

Second, much of the most insightful empirical work in macroeconomics over the past twenty years has been concerned with consumption and investment. These two chapters therefore have an unusually intensive empirical focus.

7.1 Consumption under Certainty: The Life-Cycle/Permanent-Income Hypothesis

Assumptions

Although we have already examined aspects of individuals' consumption decisions in our investigations of the Ramsey and Diamond models in Chapter 2 and of real-business-cycle theory in Chapter 4, here we start with a simple case. Consider an individual who lives for T periods whose lifetime utility is

$$U = \sum_{t=1}^{T} u(C_t), \qquad u'(\bullet) > 0, \quad u''(\bullet) < 0, \tag{7.1}$$

where $u(\bullet)$ is the instantaneous utility function and C_t is consumption in period t. The individual has initial wealth of A_0 and labor incomes of $Y_1, Y_2, \ldots, Y_T$ in the T periods of his or her life; the individual takes these as given. The individual can save or borrow at an exogenous interest rate, subject only to the constraint that any outstanding debt must be repaid at the end of his or her life. For simplicity, this interest rate is set to zero.[1] Thus the individual's budget constraint is

$$\sum_{t=1}^{T} C_t \leq A_0 + \sum_{t=1}^{T} Y_t. \tag{7.2}$$

Behavior

Since the marginal utility of consumption is always positive, the individual satisfies the budget constraint with equality. The Lagrangian for his or her maximization problem is therefore

$$\mathcal{L} = \sum_{t=1}^{T} u(C_t) + \lambda\left(A_0 + \sum_{t=1}^{T} Y_t - \sum_{t=1}^{T} C_t\right). \tag{7.3}$$

The first-order condition for C_t is

$$u'(C_t) = \lambda. \tag{7.4}$$

[1] Note that we have also assumed that the individual's discount rate is zero (see [7.1]). Assuming that the interest rate and the discount rate are equal but not necessarily zero would have almost no effect on the analysis in this section and the next. And assuming that they need not be equal would have only modest effects.

Since (7.4) holds in every period, the marginal utility of consumption is constant. And since the level of consumption uniquely determines its marginal utility, this means that consumption must be constant. Thus $C_1 = C_2 = \cdots = C_T$. Substituting this fact into the budget constraint yields

$$C_t = \frac{1}{T}\left(A_0 + \sum_{\tau=1}^{T} Y_\tau\right) \qquad \text{for all } t. \tag{7.5}$$

The term in parentheses is the individual's total lifetime resources. Thus (7.5) states that the individual divides his or her lifetime resources equally among each period of life.

Implications

This analysis implies that the individual's consumption in a given period is determined not by income that period, but by income over his or her entire lifetime. In the terminology of Friedman (1957), the right-hand side of (7.5) is *permanent income,* and the difference between current and permanent income is *transitory income.* Equation (7.5) implies that consumption is determined by permanent income.

To see the importance of the distinction between permanent and transitory income, consider the effect of a windfall gain of amount Z in the first period of life. Although this windfall raises current income by Z, it raises permanent income by only Z/T. Thus if the individual's horizon is fairly long, the windfall's impact on current consumption is small. One implication is that a temporary tax cut may have little impact on consumption; as described in Chapter 6, this appears to be the case in practice.

Our analysis also implies that although the time pattern of income is not important to consumption, it is critical to saving. The individual's saving in period t is the difference between income and consumption:

$$\begin{aligned} S_t &= Y_t - C_t \\ &= \left(Y_t - \frac{1}{T}\sum_{\tau=1}^{T} Y_\tau\right) - \frac{1}{T}A_0, \end{aligned} \tag{7.6}$$

where the second line uses (7.5) to substitute for C_t. Thus saving is high when income is high relative to its average—that is, when transitory income is high. Similarly, when current income is less than permanent income, saving is negative. Thus the individual uses saving and borrowing to smooth the path of consumption. This is the key idea of the life-cycle/permanent-income hypothesis of Modigliani and Brumberg (1954) and Friedman (1957).

What Is Saving?

At a more general level, the basic idea of the life-cycle/permanent-income hypothesis is a simple insight about saving: saving is future consumption. As long as an individual does not save just for the sake of saving, he or she saves to consume in the future. The saving may be used for conventional consumption later in life, or bequeathed to the individual's children for their consumption, or even used to erect monuments to the individual upon his or her death. But as long as the individual does not value saving in itself, the decision about the division of income between consumption and saving is driven by preferences between present and future consumption and information about future consumption prospects.

This observation suggests that many common statements about saving may be incorrect. For example, it is often asserted that poor individuals save a smaller fraction of their incomes than the wealthy do because their incomes are little above the level needed to provide a minimal standard of living. But this claim overlooks the fact that individuals who have trouble obtaining even a low standard of living today may also have trouble obtaining that standard in the future. Thus their saving is likely to be determined by the time pattern of their income, just as it is for the wealthy.

To take another example, consider the common assertion that individuals' concern about their consumption relative to others' tends to raise their consumption as they try to "keep up with the Joneses." Again, this claim fails to recognize what saving is: since saving represents future consumption, saving less implies consuming less in the future, and thus falling further behind the Joneses. Thus one can just as well argue that concern about relative consumption causes individuals to try to catch up with the Joneses in the future, and thus lowers rather than raises current consumption.[2]

Empirical Application: Understanding Estimated Consumption Functions

The traditional Keynesian consumption function posits that consumption is determined by current disposable income. Keynes (1936) argued that "the amount of aggregate consumption mainly depends on the amount of aggregate income," and that this relationship "is a fairly stable function." He claimed further that "it is also obvious that a higher absolute level of income...will lead, as a rule, to a greater *proportion* of income being saved" (Keynes, 1936, pp. 96–97; emphasis in original).

The importance of the consumption function to Keynes's analysis of fluctuations led many researchers to estimate the relationship between

[2] See Abel (1990) and Campbell and Cochrane (1995) for more on how individuals' concern about their consumption relative to others' affects saving once one recognizes that saving represents future consumption.

consumption and current income. Contrary to Keynes's claims, these studies did not demonstrate a consistent, stable relationship. Across households at a point in time, the relationship is indeed of the type that Keynes postulated; an example of such a relationship is shown in Panel (a) of Figure 7.1. But within a country over time, aggregate consumption is essentially proportional to aggregate income; that is, one sees a relationship like that in Panel (b) of the figure. Further, the cross-section consumption function differs across groups. For example, the slope of the estimated consumption function is similar for whites and blacks, but the intercept is higher for whites. This is shown in Panel (c) of the figure.

As Friedman (1957) demonstrates, the permanent-income hypothesis provides a straightforward explanation of all of these findings. Suppose that consumption is in fact determined by permanent income: $C = Y^P$. Current income equals the sum of permanent and transitory income: $Y = Y^P + Y^T$. And since transitory income reflects departures of current income from permanent income, in most samples it has a mean near zero and is roughly uncorrelated with permanent income.

Now consider a regression of consumption on current income:

$$C_i = a + bY_i + e_i. \tag{7.7}$$

In a univariate regression, the estimated coefficient on the independent variable is the ratio of the covariance of the independent and dependent variables to the variance of the independent variable. In this case, this implies

$$\begin{aligned}\hat{b} &= \frac{\text{Cov}(Y, C)}{\text{Var}(Y)} \\ &= \frac{\text{Cov}(Y^P + Y^T, Y^P)}{\text{Var}(Y^P + Y^T)} \\ &= \frac{\text{Var}(Y^P)}{\text{Var}(Y^P) + \text{Var}(Y^T)};\end{aligned} \tag{7.8}$$

here the second line uses the facts that current income equals the sum of permanent and transitory income and that consumption equals permanent income, and the last line uses the assumption that permanent and temporary income are uncorrelated. In addition, the estimated constant equals the mean of the dependent variable minus the estimated slope coefficient times the mean of the independent variable. Thus,

$$\begin{aligned}\hat{a} &= \overline{C} - \hat{b}\overline{Y} \\ &= \overline{Y}^P - \hat{b}(\overline{Y}^P + \overline{Y}^T) \\ &= (1 - \hat{b})\overline{Y}^P,\end{aligned} \tag{7.9}$$

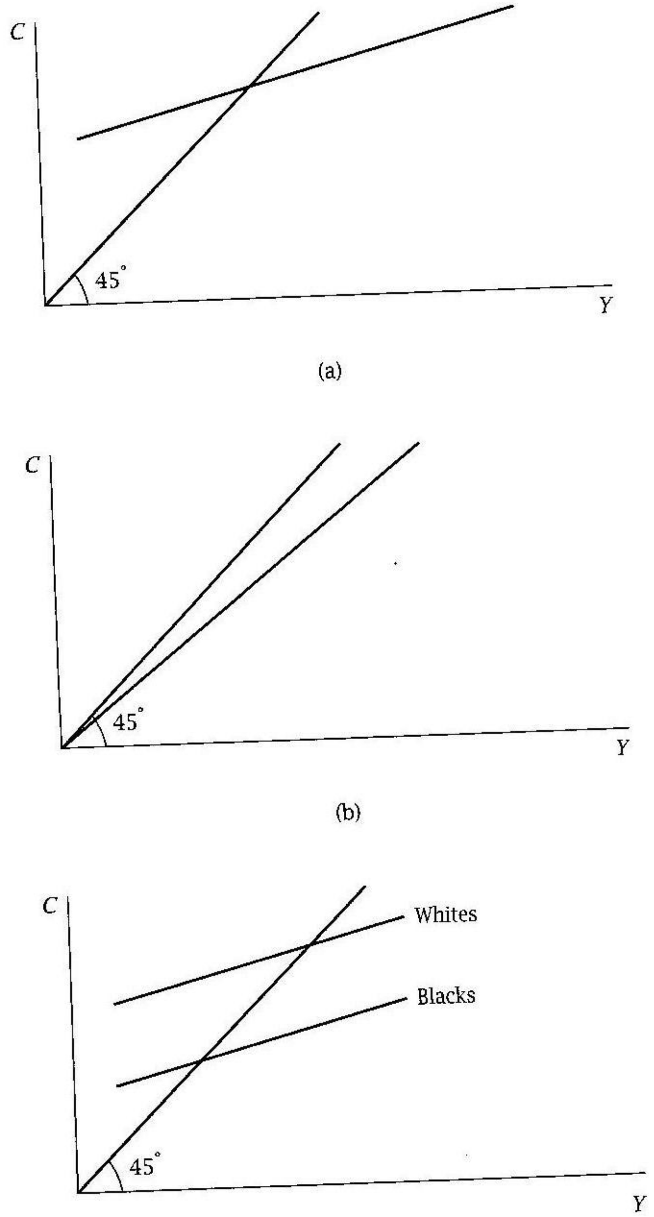

FIGURE 7.1 Some different forms of the relationship between current income and consumption

where the last line uses the assumption that the mean of transitory income is zero.

Thus the permanent-income hypothesis predicts that the key determinant of the slope of an estimated consumption function, $\hat{b}$, is the relative variation in permanent and transitory income. Intuitively, an increase in current income is associated with an increase in consumption only to the extent that it reflects an increase in permanent income. When the variation in permanent income is much greater than the variation in transitory income, almost all differences in current income reflect differences in permanent income; thus consumption rises nearly one-for-one with current income. But when the variation in permanent income is small relative to the variation in transitory income, little of the variation in current income comes from variation in permanent income, and so consumption rises little with current income.

This analysis can be used to understand the estimated consumption functions in Figure 7.1. Across households, much of the variation in income reflects such factors as unemployment and the fact that households are at different points in their life cycles. As a result, the estimated slope coefficient is substantially less than 1, and the estimated intercept is positive. Over time, in contrast, almost all of the variation in aggregate income reflects long-run growth—that is, permanent increases in the economy's resources. Thus the estimated slope coefficient is close to 1, and the estimated intercept is close to zero.[3]

Now consider the differences between blacks and whites. The relative variances of permanent and transitory income are similar in the two groups, and so the estimates of b are similar. But blacks' average incomes are lower than whites'; as a result, the estimate of a for blacks is lower than the estimate for whites (see [7.9]).

To see the intuition for this result, consider a member of each group whose income equals the average income among whites. Since there are many more blacks with permanent incomes below this level than there are with permanent incomes above it, the black's permanent income is much more likely to be less than his or her current income than more. As a result, blacks with this current income have on average lower permanent income; thus on average they consume less than their income. For the white, in contrast, his or her permanent income is about as likely to be more than current income as it is to be less; as a result, whites with this current income on average have the same permanent income, and thus on average they consume

[3]In this case, although consumption is approximately proportional to income, the constant of proportionality is less than 1; that is, consumption is on average less than permanent income. As Friedman describes, there are various ways of extending the basic theory to make it consistent with this result. One is to account for turnover among generations and long-run growth: if the young generally save and the old generally dissave, the fact that each generation is wealthier than the previous one implies that the young's saving is greater than the old's dissaving.

their income. In sum, the permanent-income hypothesis attributes the different consumption patterns of blacks and whites to the different average incomes of the two groups, and not to any differences in tastes or culture.

7.2 Consumption under Uncertainty: The Random-Walk Hypothesis

Individual Behavior

We now extend our analysis to account for uncertainty. Continue to assume that both the interest rate and the discount rate are zero. In addition, suppose that the instantaneous utility function, $u(\bullet)$, is quadratic. Thus the individual maximizes

$$E[U] = E\left[\sum_{t=1}^{T} C_t - \frac{a}{2}C_t^2\right], \qquad a > 0. \tag{7.10}$$

We will assume that the individual's wealth is such that consumption is always in the range where marginal utility is positive. As before, the individual must pay off any outstanding debts at the end of his or her life. Thus the budget constraint is again given by equation (7.2), $\sum_{t=1}^{T} C_t \leq A_0 + \sum_{t=1}^{T} Y_t$.

To describe the individual's behavior, we use the Euler-equation approach that we employed in Chapters 2 and 4. Specifically, suppose that the individual has chosen first-period consumption optimally given the information available, and suppose that he or she will choose consumption in each future period optimally given the information then available. Now consider a reduction in C_1 of dC from the value the individual has chosen and an equal increase in consumption at some future date from the value he or she would have chosen. If the individual is optimizing, a marginal change of this type does not affect expected utility. Since the marginal utility of consumption in period 1 is $1 - aC_1$, the change has a utility cost of $(1 - aC_1)dC$. And since the marginal utility of period-t consumption is $1 - aC_t$, the change has an expected utility benefit of $E_1[1 - aC_t]dC$, where $E_1[\bullet]$ denotes expectations conditional on the information available in period 1. Thus if the individual is optimizing,

$$1 - aC_1 = E_1[1 - aC_t], \qquad \text{for } t = 2, 3, \ldots, T. \tag{7.11}$$

Since $E_1[1 - aC_t]$ equals $1 - aE_1[C_t]$, this implies

$$C_1 = E_1[C_t], \qquad \text{for } t = 2, 3, \ldots, T. \tag{7.12}$$

The individual knows that his or her lifetime consumption will satisfy the budget constraint, (7.2), with equality. Thus the expectations of the two sides of the constraint must be equal:

$$\sum_{t=1}^{T} E_1[C_t] = A_0 + \sum_{t=1}^{T} E_1[Y_t]. \tag{7.13}$$

Equation (7.12) implies that the left-hand side of (7.13) is TC_1. Substituting this into (7.13) and dividing by T yields

$$C_1 = \frac{1}{T}\left(A_0 + \sum_{t=1}^{T} E_1[Y_t]\right). \tag{7.14}$$

That is, the individual consumes $1/T$ of his or her expected lifetime resources.

Implications

Equation (7.12) implies that $E_1[C_2]$ equals C_1. More generally, reasoning analogous to what we have just done implies that in each period, expected next-period consumption equals current consumption. This implies that changes in consumption are unpredictable. By the definition of expectations, we can write

$$C_t = E_{t-1}[C_t] + e_t, \tag{7.15}$$

where e_t is a variable whose expectation as of period $t-1$ is zero. Thus, since $E_{t-1}[C_t] = C_{t-1}$, we have

$$C_t = C_{t-1} + e_t. \tag{7.16}$$

This is Hall's famous result that the life-cycle/permanent-income hypothesis implies that consumption follows a random walk (Hall, 1978). The intuition for this result is straightforward: if consumption is expected to change, the individual can do a better job of smoothing consumption. Suppose, for example, that consumption is expected to rise. This means that the current marginal utility of consumption is greater than the expected future marginal utility of consumption, and thus that the individual is better off raising current consumption. Thus the individual adjusts his or her current consumption to the point where consumption is not expected to change.

In addition, our analysis can be used to find what determines the change in consumption, e. Consider for concreteness the change from period 1 to period 2. Reasoning parallel to that used to derive (7.14) implies that C_2 equals $1/(T-1)$ of the individual's expected remaining lifetime resources:

$$
\begin{aligned}
C_2 &= \frac{1}{T-1}\left(A_1 + \sum_{t=2}^{T} E_2[Y_t]\right) \\
&= \frac{1}{T-1}\left(A_0 + Y_1 - C_1 + \sum_{t=2}^{T} E_2[Y_t]\right),
\end{aligned}
\tag{7.17}
$$

where the second line uses the fact that $A_1 = A_0 + Y_1 - C_1$. We can rewrite the expectation as of period 2 of income over the remainder of life, $\sum_{t=2}^{T} E_2[Y_t]$, as the expectation of this quantity as of period 1, $\sum_{t=2}^{T} E_1[Y_t]$, plus the information learned between period 1 and period 2, $\sum_{t=2}^{T} E_2[Y_t] - \sum_{t=2}^{T} E_1[Y_t]$. Thus we can rewrite (7.17) as

$$
C_2 = \frac{1}{T-1}\left\{A_0 + Y_1 - C_1 + \sum_{t=2}^{T} E_1[Y_t] + \left(\sum_{t=2}^{T} E_2[Y_t] - \sum_{t=2}^{T} E_1[Y_t]\right)\right\}. \tag{7.18}
$$

From (7.14), $A_0 + Y_1 + \sum_{t=2}^{T} E_1[Y_t]$ equals TC_1. Thus (7.18) becomes

$$
\begin{aligned}
C_2 &= \frac{1}{T-1}\left\{TC_1 - C_1 + \left(\sum_{t=2}^{T} E_2[Y_t] - \sum_{t=2}^{T} E_1[Y_t]\right)\right\} \\
&= C_1 + \frac{1}{T-1}\left(\sum_{t=2}^{T} E_2[Y_t] - \sum_{t=2}^{T} E_1[Y_t]\right).
\end{aligned}
\tag{7.19}
$$

Equation (7.19) states that the change in consumption between period 1 and period 2 equals the change in the individual's estimate of his or her lifetime resources divided by the number of periods of life remaining.

Finally, note that the individual's behavior exhibits certainty equivalence: as (7.14) shows, the individual consumes the amount he or she would if his or her future incomes were certain to equal their means; that is, uncertainty about future income has no impact on consumption.

To see the intuition for this certainty-equivalence behavior, consider the Euler equation relating consumption in periods 1 and 2. With a general instantaneous utility function, this condition is

$$
u'(C_1) = E_1[u'(C_2)]. \tag{7.20}
$$

When utility is quadratic, marginal utility is linear. Thus the expected marginal utility of consumption is the same as the marginal utility of expected consumption. That is, since $E_1[1 - aC_2] = 1 - aE_1[C_2]$, for quadratic utility (7.20) is equivalent to

$$
u'(C_1) = u'(E_1[C_2]). \tag{7.21}
$$

This implies $C_1 = E_1[C_2]$.

This analysis shows that quadratic utility is the source of certainty-equivalence behavior: if utility is not quadratic, marginal utility is not linear, and so (7.21) does not follow from (7.20). We return to this point in Section 7.6.[4]

7.3 Empirical Application: Two Tests of the Random-Walk Hypothesis

Hall's random-walk result ran strongly counter to existing views about consumption.[5] The traditional view of consumption over the business cycle implies that when output declines, consumption declines but is expected to recover; thus it implies that there are predictable movements in consumption. Hall's extension of the permanent-income hypothesis, in contrast, predicts that when output declines unexpectedly, consumption declines only by the amount of the fall in permanent income; as a result, it is not expected to recover.

Because of this divergence in the predictions of the two views, a great deal of effort has been devoted to testing whether predictable changes in income produce predictable changes in consumption. The hypothesis that consumption responds to predictable income movements is referred to as *excess sensitivity* of consumption (Flavin, 1981).[6]

Campbell and Mankiw's Test Using Aggregate Data

The random-walk hypothesis implies that the change in consumption is unpredictable; thus it implies that no information available at time $t-1$ can

[4]Although the specific result that the change in consumption has a mean of zero and is unpredictable (equation [7.16]) depends on the assumption of quadratic utility (and on the assumption that the discount rate and the interest rate are equal), the result that departures of consumption growth from its average value are not predictable arises under more general assumptions. See, for example, Problem 7.3.

[5]Indeed, it is said that when Hall first presented the paper deriving and testing the random-walk result, one prominent macroeconomist told him that he must have been on drugs when he wrote the paper.

[6]The permanent-income hypothesis also makes predictions about how consumption responds to unexpected changes in income. In the model of Section 7.2, for example, the response to news is given by equation [7.19]. The hypothesis that consumption responds less than the permanent-income hypothesis predicts to unexpected changes in income is referred to as *excess smoothness* of consumption. Since excess sensitivity concerns expected changes in income and excess smoothness concerns unexpected changes, it is possible for consumption to be excessively sensitive and excessively smooth at the same time. For more on excess smoothness, see Campbell and Deaton (1989); West (1988); Flavin (1993); and Problem 7.4.

be used to forecast the change in consumption from $t-1$ to t. Thus one approach to testing the random-walk hypothesis is to regress the change in consumption on variables that are known at $t-1$. If the random-walk hypothesis is correct, the coefficients on the variables should not differ systematically from zero.

This is the approach that Hall took in his original work. He was unable to reject the hypothesis that lagged values of either income or consumption cannot predict the change in consumption. He did find, however, that lagged stock-price movements have statistically significant predictive power for the change in consumption.

The disadvantage of this approach is that the results are hard to interpret. Hall's result that lagged income does not have strong predictive power for consumption, for example, could arise not because predictable changes in income do not produce predictable changes in consumption, but because lagged values of income are of little use in predicting income movements. Similarly, it is hard to gauge the importance of the rejection of the random-walk prediction using stock-price data.

Campbell and Mankiw (1989a) therefore use an instrumental-variables approach to test Hall's hypothesis against a specific alternative. The alternative they consider is that some fraction of consumers simply spend their current income, and the remainder behave according to Hall's theory. This alternative implies that the change in consumption from period $t-1$ to period t equals the change in income between $t-1$ and t for the first group of consumers, and equals the change in estimated permanent income between $t-1$ and t for the second group. Thus if we let λ denote the fraction of consumption that is done by consumers in the first group, the change in aggregate consumption is

$$\begin{aligned} C_t - C_{t-1} &= \lambda(Y_t - Y_{t-1}) + (1-\lambda)e_t, \\ &\equiv \lambda Z_t + v_t, \end{aligned} \tag{7.22}$$

where e_t is the change in consumers' estimate of their permanent income from $t-1$ to t.

Z_t and v_t are almost surely correlated. Times when income increases greatly are usually also times when households receive favorable news about their total lifetime incomes. But this means that the right-hand-side variable in (7.22) is positively correlated with the error term. Thus estimating (7.22) by ordinary least squares (OLS) leads to estimates of λ that are biased upward.

The solution to correlation between the right-hand-side variable and the error term is to use instrumental variables (IV) rather than OLS. The intuition behind IV estimation is easiest to see using the two-stage least squares interpretation of instrumental variables. What one needs are variables correlated with the right-hand-side variables but uncorrelated with the residual. Once one has such *instruments*, the first-stage regression is a regression of the right-hand-side variable, Z_t, on the instruments. The second-stage regression is then a regression of the left-hand-side variable, $C_t - C_{t-1}$, on the fitted value of Z_t from the first-stage regression, $\hat{Z}_t$. That is, we estimate:

$$\begin{aligned} C_t - C_{t-1} &= \lambda\hat{Z}_t + \lambda(Z_t - \hat{Z}_t) + v_t \\ &\equiv \lambda\hat{Z}_t + \tilde{v}_t. \end{aligned} \tag{7.23}$$

The residual in (7.23), $\tilde{v}_t$, consists of two terms, v_t and $\lambda(Z_t - \hat{Z}_t)$. By assumption, the instruments used to construct $\hat{Z}$ are not systematically correlated with v_t. And since $\hat{Z}$ is the fitted value from a regression, by construction it is uncorrelated with the residual from that regression, $Z - \hat{Z}$. Thus regressing $C_t - C_{t-1}$ on $\hat{Z}$ yields a valid estimate of λ.[7]

The usual problem in using instrumental variables is finding valid instruments: it is often hard to find variables that one can be confident are uncorrelated with the residual. But in cases where the residual reflects new information between $t-1$ and t, theory tells us that there are many candidate instruments: any variable that is known as of time $t-1$ is uncorrelated with the residual.

We can now turn to the specifics of Campbell and Mankiw's test. They measure consumption as real purchases of consumer nondurables and services per person, and income as real disposable income per person. The data are quarterly, and the sample period is 1953–1986. They consider various sets of instruments. They find that lagged changes in income have almost no predictive power for future changes. This suggests that Hall's failure to find predictive power of lagged income movements for consumption is not strong evidence against the traditional view of consumption. As a base case, they therefore use lagged values of the change in consumption as instruments. When three lags are used, the estimate of λ is 0.42, with a standard error of 0.16; when five lags are used, the estimate is 0.52, with a standard error of 0.13. Other specifications yield similar results.

Thus Campbell and Mankiw's estimates suggest quantitatively large and statistically significant departures from the predictions of the random-walk model: consumption appears to increase by about fifty cents in response to an anticipated 1-dollar increase in income, and the null hypothesis of no effect is strongly rejected. At the same time, the estimates of λ are far below

[7]The fact that $\hat{Z}$ is based on estimated coefficients causes two complications. First, the uncertainty about the estimated coefficients must be accounted for in finding the standard error of the estimate of λ; this is done in the usual formulas for the standard errors of instrumental-variables estimates. Second, the fact that the first-stage coefficients are estimated introduces some correlation between $\hat{Z}$ and v in the same direction as the correlation between Z and v. This correlation disappears as the sample size becomes large; thus IV is consistent but not unbiased. If the instruments are only moderately correlated with the right-hand-side variable, however, the bias in finite samples can be substantial. See, for example, Nelson and Startz (1990).

1. Thus the results also suggest that the permanent-income hypothesis is important to understanding consumption.[8]

Shea's Test Using Household Data

Testing the random-walk hypothesis with aggregate data has several disadvantages. Most obviously, the number of observations is small. In addition, it is difficult to find variables with much predictive power for changes in income; it is therefore hard to test the key prediction of the random-walk hypothesis that predictable changes in income are not associated with predictable changes in consumption. Finally, the theory concerns individuals' consumption, and additional assumptions are needed for the predictions of the model to apply to aggregate data. Entry and exit of households from the population, for example, can cause the predictions of the theory to fail in the aggregate even if they hold for each household individually.

Because of these considerations, many investigators have examined consumption behavior using data on individual households. Shea (1995) takes particular care to identify predictable changes in income. He focuses on households in the Panel Study of Income Dynamics (commonly referred to as the PSID) with wage-earners covered by long-term union contracts. For these households, the wage increases and cost-of-living provisions in the contracts cause income growth to have an important predictable component.

Shea constructs a sample of 647 observations where the union contract provides clear information about the household's future earnings. A regression of actual real wage growth on the estimate constructed from the union contract and some control variables produces a coefficient on the constructed measure of 0.86, with a standard error of 0.20. Thus the union contract has important predictive power for changes in earnings.

Shea then regresses consumption growth on this measure of expected wage growth; the permanent-income hypothesis predicts that the coeffi-

---

[8]In addition, the instrumental-variables approach has *overidentifying restrictions* that can be tested. If the lagged changes in consumption are valid instruments, they are uncorrelated with v. This implies that once we have extracted all of the information in the instruments about income growth, they should have no additional predictive power for the left-hand-side variable: if they do, that means that they are correlated with v, and thus that they are not valid instruments. This implication can be tested by regressing the estimated residuals from (7.22) on the instruments and testing whether the instruments have any explanatory power. Specifically, one can show that under the null hypothesis of valid instruments, the R^2 of this regression times the number of observations is asymptotically distributed χ^2 with degrees of freedom equal to the number of overidentifying restrictions—that is, the number of instruments minus the number of endogenous variables.

In Campbell and Mankiw's case, this TR^2 statistic is distributed χ^2_2 when three lags of the change in consumption are used, and χ^2_4 when five lags are used. The values of the test statistic in the two cases are only 1.83 and 2.94; these are only in the 59th and 43rd percentiles of the relevant χ^2 distributions. Thus the hypothesis that the instruments are valid cannot be rejected.

cient should be zero.[9] The estimated coefficient is in fact 0.89, with a standard error of 0.46. Thus Shea also finds a quantitatively large (though only marginally statistically significant) departure from the random-walk prediction.

Recall that in our analysis in Sections 7.1 and 7.2, we assumed that households can borrow without limit as long as they eventually repay their debts. One reason that consumption might not follow a random walk is that this assumption might fail—that is, that households might face *liquidity constraints.* If households are unable to borrow and their current income is less than their permanent income, their consumption is determined by their current income. In this case, predictable changes in income produce predictable changes in consumption.

Shea tests for liquidity constraints in two ways. First, following Zeldes (1989) and others, he divides the households according to whether they have liquid assets. Households with liquid assets can smooth their consumption by running down these assets rather than by borrowing. Thus if liquidity constraints are the reason that predictable wage changes affect consumption growth, the prediction of the permanent-income hypothesis will fail only among the households with no assets. Shea finds, however, that the estimated effect of expected wage growth on consumption is essentially the same in the two groups.

Second, following Altonji and Siow (1987), Shea splits the low-wealth sample according to whether the expected change in the real wage is positive or negative. Individuals facing expected declines in income need to save rather than borrow to smooth their consumption. Thus if liquidity constraints are important, predictable wage increases produce predictable consumption increases, but predictable wage decreases do not produce predictable consumption decreases.

Shea's findings are the opposite of this. For the households with positive expected income growth, the estimated impact of the expected change in the real wage on consumption growth is 0.06 (with a standard error of 0.79); for the households with negative expected growth, the estimated effect is 2.24 (with a standard error of 0.95). Thus there is no evidence that liquidity constraints are the source of Shea's results.

7.4 The Interest Rate and Saving

An important issue concerning consumption involves its response to rates of return. For example, many economists have argued that more favorable

[9] An alternative would be to follow Campbell and Mankiw's approach and regress consumption growth on actual income growth by instrumental variables, using the constructed wage growth measure as an instrument. Given the almost one-for-one relationship between actual and constructed earnings growth, this approach would be likely to produce similar results.