

Chapter 7

Consumption

This chapter focuses on household behavior in response to economic fluctuations. In the traditional Keynesian model and the IS/MP/IA framework, the so-called Keynesian consumption function associates changes in consumption with *current* disposable income, by some constant fraction (the marginal propensity to consume). The Keynesian consumption function is given by:

$$C(Y - T) = a + b(Y - T)$$

where $0 < b < 1$ is the marginal propensity to consume and a is consumption that is independent of current income (such as wealth or expectations).¹ Keynes (1936) claimed that this function was relatively stable. Before examining the implications of the consumption models presented in this chapter, it is useful to consider those from the textbook Keynesian consumption function.

A key implication of this consumption function fueled later debate discussed further below. Assume no taxes and let the average propensity to consume, APC is defined as:

$$\begin{aligned} APC &\equiv C/Y = \frac{a + bY}{Y} \\ APC &= \frac{a}{Y} + b \end{aligned}$$

The above expression implies that at higher levels of income, the APC will be lower.²

We can examine the empirical evidence to see whether the data support this description of consumption behavior.

• Time series implications

¹We could incorporate expectations of future income into the IS/MP/IA model by thinking of a change as an IS shock. However, this model does not allow for potential feedback into future consumption decisions. Since the IS/MP/IA model is fundamentally static, it is not well-suited for analyzing dynamic relationships.

²Keynes (1936) noted this relationship. He claimed that higher levels of income would lead to "a greater *proportion* of income being saved" (pp. 96-97, emphasis original)

- As the economy grows, APC should decline, resulting in an unstable steady state. If the economy consumes a smaller and smaller average of total income, then as $Y \rightarrow \infty$, savings approaches infinity and households consume nothing.
- This would also suggest that rich countries should have very high savings rates and that countries reduce consumption (relative to output) as they grow.
- Using aggregate U.S. time series data, Kuznets (1946) found that the savings rate was stable over time, even though income increased over this period.

- **Cross section implications**

- Across households, this consumption function implies that households with higher levels of income should save a proportionately smaller proportion of their current incomes.
- Cross section evidence by Brumberg and Modigliani (1954) and Ando and Modigliani (1963) suggests that households with higher income levels tend to save a *higher* proportion of their current income.

In addition to the apparent internal inconsistencies associated with the Keynesian consumption function, empirical evidence in the 1950s and 1960s lead to a revision modeling household behavior.

Most macroeconomists believe this view of household decisions is simplistic. For example, households probably look at not just current disposable income, but the path of expected income over their lifetimes. We saw this behavior in earlier models (Diamond model and the RBC approach). Fisher (1930) is an early attempt at formalizing how consumption and savings depend on lifetime income and interest rates. One important strategy used by Fisher (1930) used here is that we can express the household's consumption-savings decision in terms of consumption today and consumption tomorrow.³

As we will see in this chapter, modeling household's behavior in terms of lifetime income/utility dramatically changes their response to economic fluctuations. Specifically, households are said to behave according to the permanent income hypothesis (PIH) associated with work by Friedman (1957) and Hall (1978). The PIH says that households only respond to changes in their permanent income. Therefore, the changes in their temporary income (associated with business cycles) have little effect on their behavior. Another key component of the PIH is that households use savings to buffer against unexpected changes to their income. So, when times are exceptionally good, households will save their extra earnings. When times are bad, they use this accumulated wealth in order to enjoy the same level of consumption.

³This requires that we assume that utility is time separable. This assumption is used in the majority of dynamic stochastic general equilibrium models.

Understanding how households make consumption-savings decisions is critical for understanding the social costs of business cycles (Lucas, 1987), but also for understanding the pricing of assets (section 7.5 of this chapter).

7.1 Consumption Under Certainty: The Permanent-Income Hypothesis

7.1.1 Assumptions

Each household lives for T periods. Lifetime utility is defined as:

$$U = \sum_{t=1}^T u(C_t)$$

where $u'(\cdot) > 0$ and $u''(\cdot) < 0$, so that the marginal utility of consumption each period is positive and diminishing.

Each household has initial wealth, A_0 , that it enters the period with. Each period, the household earns income, Y_t , from working. Since there is no uncertainty here, we are assuming that lifetime income and initial wealth are known and taken as given. If the household knows its lifetime income, it can deduce what level of consumption each period will maximize its lifetime utility. The lifetime budget constraint is:

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

so that lifetime consumption cannot exceed lifetime income plus initial wealth. The household must payoff all debts and the end of life in period T . For simplicity, we are assuming no interest and that the household does not discount future consumption. We will relax these assumptions later on in the chapter.

7.1.2 Household Behavior

The household's problem can be set up using the following Lagrangian:

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

This is virtually identical to the Diamond model we solved before, only we have T first order conditions for consumption (instead of only two in the two-period Diamond model). Expanding the above Lagrangian, we have:

$$\begin{aligned} \mathcal{L} &= [u(C_1) + u(C_2) + \cdots + u(C_T)] + \\ &\quad \lambda (A_0 + Y_1 + Y_2 + \cdots + Y_T - [C_1 + C_2 + \cdots + C_T]) \end{aligned}$$

The first order condition for C_1 is:

$$\frac{\partial \mathcal{L}}{\partial C_1} = 0 : u'(C_1) - \lambda = 0$$

Similarly, for any given C_t :

$$\frac{\partial \mathcal{L}}{\partial C_t} = 0 : u'(C_t) - \lambda = 0$$

Since this first order condition is the same for each period, the marginal utility of consumption is constant and equal to λ every period. This implies that the level of consumption must also be constant each period (otherwise the marginal utility would vary over time). Therefore, the household chooses a consumption path such that it consumes an equal amount each period.

We can use this fact to solve for any given C_t . Note that because consumption is equal each period: $C_t = C_1 = C_2 = \dots = C_T$, and $\sum_{t=1}^T C_t = T \times C_t$. Plugging this into the budget constraint yields:

$$\begin{aligned} \sum_{t=1}^T C_t &= A_0 + \sum_{t=1}^T Y_t \\ TC_t &= A_0 + \sum_{t=1}^T Y_t \\ C_t &= \frac{1}{T} \left[A_0 + \sum_{t=1}^T Y_t \right] \quad \forall t \end{aligned}$$

7.1.3 Implications and Analysis

The key implication from the model above is that current consumption depends not only on current income, but on lifetime income. This is a significant departure from the Keynesian consumption function we used in the IS/MP/IA model. Friedman (1957) used a version of this model to study consumption behavior, distinguishing between transitory income versus permanent income. Transitory income refers to one-time windfalls or losses that have a small effect on lifetime income (e.g., permanent income). The model above implies that only lifetime, or what Friedman (1957) calls, permanent income matters.

To see why this is the case, suppose the household receives a one-time payment of Z . While Z may be large relative to current income, this gain is spread over the household's lifetime, increasing permanent income by only Z/T . As long as the household has a relatively long time horizon, this temporary change will have a small effect on consumption. This approach can be used in fiscal policy analysis. Those changes that households perceive as temporary have little effect on consumption patterns, whereas those that are perceived to be permanent changes have a larger effect on consumption.

We can also use this model to understand savings behavior. While the model demonstrates that household maximize utility by consuming a constant amount

each period, this does not imply a constant savings rate. Current savings, S_t is equal to current income less current consumption:

$$S_t = Y_t - C_t$$

Substituting the solution for consumption from above:

$$S_t = Y_t - \frac{1}{T} \left[A_0 + \sum_{t=1}^T Y_t \right]$$

$$S_t = Y_t - \frac{1}{T} \sum_{t=1}^T Y_t - \frac{1}{T} A_0$$

From the above expression, if income is higher than average, savings will be higher. Similarly, when income is lower than average, households dissave ($S_t < 0$). Friedman (1957) used a similar framework to study how households will borrow and save to smooth consumption over their lifetimes. Brumberg and Modigliani (1954) apply this same concept to the savings/borrowing behavior of working/retired households.

There are some key assumptions that we've made in this analysis. Each will be addressed in later parts of the chapter:

- *Households do not discount future consumption*
This assumption does not significantly change the analysis, when coupled with the introduction of the interest rate earned on savings.
- *Households do not earn interest on savings.*
See previous bullet point. One can imagine that different levels of interest (even if they are known and constant) would change the consumption path households choose to maximize utility.
- *Households know their future income. (non-stochastic model)*
When households forecast future income, this could significantly affect their consumption decisions each period. For example, if new information about lifetime income arrives between t and $t + 1$, the household may change the consumption path each period. One can build in risk behavior and allow for precautionary savings by households.
- *Households have unlimited access to resources in the financial system.*
That is, we assumed that households are no liquidity constrained. When a household has temporarily low income, it can access funds from a lender without restrictions. The key here is that one can argue that some households have limited access to credit, perhaps because lenders are unable to observe their permanent income.

7.1.4 What is Saving?

The permanent-income hypothesis (PIH) suggests that savings is nothing more than future consumption. Unless an individual household simply gets added

utility from the idea of saving,⁴ savings is a way to transfer wealth over time for future consumption (either for the individual or heirs).

With this interpretation of saving, the cross-sectional differences in consumption-savings behavior may not be attributed to different levels of income. Romer gives the example of poor versus rich households. Poorer households have relatively low savings not because they earn only a little above subsistence income, but rather because their lifetime income is low. In other words, if a poor household (say a 18-year old college student) expects to receive significantly higher income in the future (upon graduation), then this household should borrow now while its income is relatively lower than lifetime income. Savings behavior is dictated by current income relative to lifetime income, for rich and for poor.

Another common belief is that households suffer from a "keeping up with the Joneses" mentality. That is, they care not only about the level of consumption, but also on their level of consumption relative to others. Consider what would happen if a household consumes more in an attempt to "keep up" with their neighbors' living standards. Consuming more today means saving less today. So, over time, the household's ability to keep up will decrease.

7.1.5 Empirical Application: Estimating Consumption Functions

As discussed in detail in the introduction above, the empirical evidence on the Keynesian consumption function is at best mixed. Even though Kuznets' (1946) work supports the assertion that the *APC* should decline over time, this implication would lead not allow the economy to be on a balanced growth path in the long run. Moreover, in cross section work, the *APC* varies not only by income level, but also by demographic groups. This suggests that the Keynesian consumption function is an oversimplification that misses some key aspects of consumption behavior.

Friedman (1957) claims the PIH can explain this apparent consumption puzzle. Let consumption be equal to permanent income, Y^P : $C = Y^P$. Total income, Y , is equal to the sum of transitory income, Y^T , and permanent income: $Y = Y^P + Y^T$. The Keynesian consumption function can be estimated using the following regression equation:

$$C_i = a + bY_i + e_i$$

Using ordinary least squares (OLS), the estimated coefficient for income is, \hat{b} is:

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

⁴The character Trina from *McTeague: A Story of San Francisco* (1899) is an example of an individual who receives utility from saving. This book was adapted into Eric Von Stronheim's *Greed* (1924).

Transitory income is, by definition, mean zero ($\bar{Y}^T = 0$) and uncorrelated with permanent income ($E(Y^T Y^P) = 0$). Therefore, $Cov(Y^P + Y^T, Y^P) = Var(Y^P)$:⁵

$$\hat{b} = \frac{Var(Y^P)}{Var(Y^P) + Var(Y^T)}$$

Notice that \hat{b} is bounded by 0 and 1, just as we assumed in the Keynesian consumption function. As long as there are some transitory shocks to income, $\hat{b} > 0$. Because the variance is by definition positive, $\hat{b} > 0$.

From the expression, the slope of the consumption function (marginal propensity to consume) depends on how much permanent income varies relative to total income. For example, in an economy with stable permanent income, but highly variable business cycles (e.g., transitory income shocks), the \hat{b} will be smaller. That is, households in this economy will be relatively less responsive to changes in their current income because they expect that these shocks are largely transitory.

The estimate of the intercept, \hat{a} , is:

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

Over the long run, as the economy grows, the fluctuations in income are associated with changes in the economy's capacity to produce. Therefore, in the long run, the term $Var(Y^T)$ is relatively small, compared to $Var(Y^P)$, so $\hat{b} \rightarrow 1$ and $\hat{a} \rightarrow 0$. Friedman (1957) applied this model to describe differences in consumption-savings behavior across different races.

7.2 Consumption Under Uncertainty: The Random-Walk Hypothesis

7.2.1 Individual Behavior

This section takes the PIH model above and allows for uncertainty with respect to the level of income. In order to solve the model, we will assume a specific functional form for utility. We use the quadratic utility function employed by

⁵ $Cov(Y^P + Y^T, Y^P) = E[(Y^P + Y^T - (\bar{Y}^P + \bar{Y}^T))(Y^P - \bar{Y}^P)]$

Note, $\bar{Y}^T = 0$ by definition.

$E[(Y^P + Y^T - (\bar{Y}^P)) (Y^P - \bar{Y}^P)] =$

$E[(Y^P Y^P + Y^T Y^P - (\bar{Y}^P Y^P)) - (Y^P \bar{Y}^P + Y^T \bar{Y}^P - (\bar{Y}^P \bar{Y}^P))]$

Note, $E(Y^T Y^P) = 0$ because transitory and permanent income are by definition uncorrelated.

$E[(Y^P Y^P - (\bar{Y}^P Y^P)) - Y^P \bar{Y}^P + (\bar{Y}^P \bar{Y}^P)] =$

$E[(Y^P - \bar{Y}^P)(Y^P - \bar{Y}^P)] = Var(Y^P)$

Hall (1978) in his test of the permanent income hypothesis:

$$E(U) = E \left[\sum_{t=1}^T \left(C_t - \frac{a}{2} C_t^2 \right) \right]$$

subject to the lifetime budget constraint:

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

This can be set up as a Lagrangian:

$$\mathcal{L} = E \left[\sum_{t=1}^T \left(C_t - \frac{a}{2} C_t^2 \right) \right] + \lambda \left(A_0 + \sum_{t=1}^T Y_t - \sum_{t=1}^T C_t \right)$$

The first order condition with respect to first-period consumption is:⁶

$$\frac{\partial \mathcal{L}}{\partial C_1} = 0 : E_1(1 - aC_1) - \lambda = 0$$

Since first-period consumption is known at time 1 :

$$(1 - aC_1) = \lambda$$

At time 1, the household's choice of consumption for any period t :

$$\frac{\partial \mathcal{L}}{\partial C_t} = 0 : E_1(1 - aC_t) - \lambda = 0$$

Combining these two expressions we have the choice of first-period consumption, C_1 :

$$\begin{aligned} (1 - aC_1) &= E_1(1 - aC_t) \\ C_1 &= E_1(C_t) \end{aligned}$$

If the lifetime budget constraint is binding, then it is also binding in expectations at time 1 (or any time t) :

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

From the utility-maximization problem, we know that $C_1 = E_t(C_t)$. Therefore, $\sum_{t=1}^T E[C_t] = \sum_{t=1}^T C_1 = TC_1$. Plugging this into the expression above:

$$\begin{aligned} TC_1 &= A_0 + \sum_{t=1}^T E[Y_t] \\ C_1 &= \frac{1}{T} \left(A_0 + \sum_{t=1}^T E_1[Y_t] \right) \end{aligned}$$

⁶The notation $E_1(x_t)$ is the same as $E(x_t | I_1)$ where I_t is the information set at time 1. That is, the household is making a decision based on information available at that point in time. So, if we had $E_2(C_3)$, this would denote the household's choice of consumption for period 3, conditioning on information available at time 2.

Therefore, at time 1, the household consumes a fraction $1/T$ of its expected lifetime resources.

7.2.2 Implications and Analysis

Recall we had the first order condition for consumption today, relative to some future period t :

$$C_1 = E_1(C_t)$$

This condition holds for each consumption value chosen by the household at time 1. So, consumption today is simply equal to expected consumption next period, C_2 :

$$C_1 = E_1(C_2)$$

Or, more generally:

$$C_{t-1} = E_{t-1}(C_t)$$

This expression implies that changes in consumption are unpredictable. Suppose that e_2 is the error between forecasted and actual consumption in period 2, $e_2 = C_2 - E_1(C_2)$. More generally, we can let the error, e_t , denote the difference between actual and forecasted consumption for a given period t :

$$\begin{aligned} e_t &= C_t - E_{t-1}(C_t) \\ C_t &= E_{t-1}(C_t) + e_t \end{aligned}$$

From the expression for C_{t-1} , we know that $E_{t-1}(C_t) = C_{t-1}$:

$$\begin{aligned} C_t &= C_{t-1} + e_t \\ C_t - C_{t-1} &= e_t \\ \Delta C_t &= e_t \end{aligned}$$

This implies that consumption follows a random walk. Recall, a random walk is an AR(1) process where the autoregressive parameter is equal to 1. This is also referred to an integrated of order 1, and I(1) process. We also used this empirical specification for describing stochastic trends.

Intuitively, why are changes in consumption unpredictable? If consumption is expected to rise in the future, then the household will respond by consuming more today, in order to smooth out this increase over its lifetime. We can see this by solving for C_2 using the same approach we used to solve for C_1 to show that consumption in period 2. Romer covers this derivation in detail. The end result is:

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

Notice the term in parentheses on the right is the difference between the forecasted lifetime income between periods 1 and 2. So, consumption in period 2 will differ from consumption in period 1, only if there was a shock that caused lifetime income to deviate from the forecast in period 1 $\sum_{t=2}^T E_1[Y_t]$.

This expression is consistent with certainty equivalence. Certainty equivalence implies that the individual consumes the same amount if lifetime income were certain to be equal to their average values. That is, $C_2 = C_1$ if $\sum_{t=2}^T E_2 [Y_t] = \sum_{t=2}^T E_1 [Y_t]$. Certainty equivalence is a result of our assumption about the utility function - it does not hold for a more general set of utility functions. This feature of models (in consumption or in explaining other variables) results from linear marginal utility.

7.2.3 Hall (1978)

Hall (1978) uses the above specification to construct an empirical test of the permanent-income hypothesis. In short, he runs a series of regressions to see whether the data support the PIH. It is worth noting that some of the authors assumptions are critical to his empirical specification: linear marginal utility and constant interest rates.

1. Does consumption follow a random walk? Yes

Data

Hall (1978) uses quarterly data on real consumption per capita of non-durables and services, 1948-1977. Durable consumption is excluded because household purchases of these items are problematic in terms of the model. These are large purchases that the household receives utility from over time. In the model, consumption is time separable, so that the household receives utility from consumption in period t and no utility from that consumption purchase thereafter.

Methodology

Using the quadratic utility function allows for linear marginal utility. In the context of the model, Hall (1978) regresses current marginal utility on its past value:⁷

$$C_t = \beta C_{t-1} + \varepsilon_t$$

The model tells us that $\beta = 1$, so that consumption follows a random walk. However, random walk tests are notoriously difficult because by failing to reject the null hypothesis that $\beta = 1$ we know very little. What we know is that the data are not inconsistent with the random walk, but statistically, this isn't the same as saying that $\beta = 1$. As Hall (1978) points out, this test is severely limited because it does not allow us to distinguish this theory from other theories of consumption.

Results

The data fail to reject the random walk specification. For the reasons mentioned above, this is weak evidence of the PIH.

2. Can consumption be predicted based on past values of consumption? No

Methodology

⁷He does a few different specifications, including one with a constant term.

According to the model, consumption should be based solely on one lag of consumption.⁸ All of the past values/information from further lags of consumption are already contained in the first lag. If households observe a change in their income, they immediately adjust current consumption. Therefore, if we include additional $AR(p)$ terms, they should be statistically insignificant.

Results

An F-test of the coefficients on the lags of consumption beyond the first lag fail to reject the null hypothesis that they are jointly equal to zero. This is evidence in support of the PIH. Again, because we fail to reject a null hypothesis, this evidence is at best weakly in support of the PIH.

3. Can consumption be predicted based on past values of disposable income? Mixed

Data

Hall (1978) uses quarterly data on real disposable income per capita (1948-1977) to see whether past values of income matter for current consumption. He uses the same implicit price deflator and population values used in constructing his consumption variables.

Methodology

According to the model, consumption should be based solely on one lag of consumption. All of the past values/information from further lags of disposable income are already contained in the first lag of consumption. Lags of disposable income should be statistically insignificant.

Results

In most specifications, the evidence fail to reject the joint statistical significance of lags of disposable income. However, there is some evidence that more recent lags of disposable income affect current consumption. For those lags that are significant, they are relatively small in magnitude, and therefore explain a relatively small portion of the variation in current consumption. However, joint tests longer lags of disposable income are not statistically significant.

4. Can consumption be predicted based on past values of wealth?

Mixed

Data

Hall (1978) uses Standard and Poor's 500 stock indices deflated by the implicit price deflator to measure wealth. While a broader measure of wealth would be ideal (because this excludes large portions of the households who do not actively participate in the stock market), these other measures are difficult to work with. For example, property values can be affected by tax treatment relative to other types of wealth.

⁸Hall (1978) does include a constant term in this test. If he excluded it, the results could be biased (if further lags of consumption do matter). By including the constant, he sidesteps this potential criticism. The basic rule is that even if you believe the constant is equal to zero, including it in your regression causes a loss of efficiency, but avoids the potential bias associated with omitted variables.

Methodology

According to the model, consumption should be based solely on one lag of consumption. Again, all of the past values/information from further lags of other economic variables (like stock prices) are already contained in the first lag of consumption. Lags of stock prices should be statistically insignificant.

Results

In most specifications, the evidence fail to reject the joint statistical significance of lags of stock prices. However, there is some evidence that more recent lags of disposable income affect current consumption. For those lags that are significant, they are still relatively small in magnitude, although larger than the magnitudes from #3 above. Joint tests longer lags of stock prices are not statistically significant. Still, if there is a variable that helps to predict current consumption, besides one lag of consumption, it appears to be changes in stock prices.

General implications of Hall's (1978) test

The data support the PIH. When households face an unexpected decline in income, consumption declines only by the amount of the fall in permanent income. Therefore, there is no reason to expect that consumption would rebound from a decline in permanent income.

Hall's (1978) test has sparked a number of empirical studies attempting to determine whether predictable changes in income lead to predictable changes in consumption.

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

Before moving ahead to these specific tests, it is worth noting that these tests are mainly a response to Hall (1978). Recall, that his test made specific assumptions about the PIH model to arrive at the random walk conclusion. Therefore, evidence against the random walk is not necessarily evidence against a more general PIH model. This is a relatively straightforward application of the two-stage least squares (2SLS) method.□

7.3.1 Campbell and Mankiw (1989)

Methodology

Testing against alternatives

Campbell and Mankiw (1989) use instrumental variables to test the PIH against a specific alternative. Specifically, they test against the alternative that consumers spend a constant fraction of their current income:

$$C_t = \alpha + \beta Y_t$$

That is, we are looking to see whether Hall's (1978) specification outperforms the Keynesian consumption function discussed in the introduction. According to the Keynesian consumption function, the change in consumption should be equal to the change in income over the same period:

$$\Delta C_t = \Delta Y_t$$

In contrast, the random walk hypothesis implies that the change in consumption is random:

$$\Delta C_t = e_t$$

where e_t is the unexpected change to permanent income. The change in consumption is not a function of other lagged variables.

Campbell and Mankiw (1989) assume that some fraction λ of consumers follow the constant fraction rule, while $1 - \lambda$ behave according to the PIH. Therefore, the aggregate change in income should be:

$$\Delta C_t = \lambda Y_t + (1 - \lambda)\varepsilon_t$$

In regression notation, this would be:

$$\Delta C_t = \lambda Z_t + v_t$$

The problem with this specification is that Z_t and v_t correlated. When there are unexpected increases to permanent income, there are changes in total income, Y_t . Therefore, we cannot estimate λ using OLS.

In this specification, if $\lambda = 1$, this would imply that households behave according to the Keynesian consumption function. If $\lambda = 0$, then households behave according to the PIH/random walk specification from Hall (1978).

2SLS

The solution is to run two regressions (two-stage least squares, 2SLS) so that we can identify the variation in Z_t independent of v_t . To do this, we need to find variables that are correlated with the explanatory variables, but uncorrelated with the dependent variable. This can easily be checked by looking at the correlation between the instrument and the dependent variable. If the instruments were correlated with ΔC_t , then we should include them in the regression above.

The first stage is to run a regression of Z_t on the instruments. The second stage is to take the residuals from this regression, \hat{Z}_t - this is the variation in Z_t that is uncorrelated with v_t - and use this as a measure of the change in income:

$$\Delta C_t = \lambda \hat{Z}_t + \hat{v}_t$$

In this regression, \hat{Z}_t and \hat{v}_t are uncorrelated, yielding an unbiased estimate of λ .

Data

Campbell and Mankiw (1989) use quarterly real consumption purchases of nondurables and services per capita and real disposable income per capita, 1953-1986. They consider two instruments. First, they examine lagged changes

on income and find that lagged changes in income have almost no predictive power for future changes in income. Thus, they would be a poor instrument. Ultimately, they use lagged values of changes in consumption as an instrument (as they do not help to predict future changes, and are therefore uncorrelated with ΔC_t).

Results

The authors find that λ is around 0.5 and statistically significant. This, a \$1 change in income is associated with a \$0.50 change in consumption. This is evidence against the random walk hypothesis. However, because the $\lambda < 1$, this suggests that the Hall's (1978) test explain a sizable fraction of consumption behavior.

7.3.2 Shea (1995)

Data

Shea (1995) takes a different approach by using cross section data on households. Recall, much of the debate surrounding the Keynesian consumption function stemmed from an inability to resolve cross section and time series evidence. The disadvantage of using time series data include: small number of observations, difficulty in testing the random walk because we cannot observe expectations, and the application of individual behavior to aggregate data.

Shea (1995) uses PSID data on wage-earners (647 sample size) to study how union contracts serve as a predictor of future expected income.

Methodology

Shea's (1995) sample includes wage-earners on long-term union contracts, allowing him to proxy for lifetime income. He regresses actual real wage growth on and estimate of permanent income (constructed from the union contract terms) and other control variables. The coefficient on union contract terms is large and significant.

In the second step, he regresses consumption growth on the measure of expected wage growth. The PIH predicts the coefficient should be zero - all of the information contained in expected future income is already contained in the lag of consumption. That is, any change in consumption should be a reflection of only unexpected changes in future income (information arriving between $t - 1$ and t).

Results

Shea (1995) finds that expected future wages are statistically significant and large in magnitude. This is strong evidence against the random walk specification from Hall (1978). He goes on to test why the random walk test fails by looking at whether households are liquidity constrained. He finds no evidence of liquidity constraints in his sample - it appears that households with lower than average income are able to borrow. Shea (1995) finds that predictable wage decreases do not lead to predictable consumption decreases, as one would expect if households are liquidity constrained.

7.4 The Interest Rate and Saving

In this section, we relax the assumption of no interest rate and allow for households' discounting of future utility/consumption. We expect that changes in interest rates affect consumption decisions because they affect the relative price of consumption today versus tomorrow. For simplicity, we will deal with non-stochastic environments. One could extend the analysis to allow for uncertainty, but it would not greatly change the implications of the model with interest rates.

7.4.1 Interest Rates and Consumption Growth

Incorporating interest rates into the PIH model we saw before implies the following lifetime budget constraint (this should look familiar from the two-period Diamond model):

$$\sum_{t=1}^T \frac{1}{(1+r)^t} C_t \leq A_0 + \sum_{t=1}^T \frac{1}{(1+r)^t} Y_t$$

This lifetime budget constraint expresses future consumption and income in terms of their present values.

Next, we allow for a more flexible utility function. Specifically, we will use the CRRA utility function commonly used in RBC models and later in asset pricing models:⁹

$$U = \sum_{t=1}^T \beta^t \frac{C_t^{1-\theta}}{1-\theta}$$

We can set up this problem as a Lagrangian:

$$\mathcal{L} = \sum_{t=1}^T \beta^t \frac{C_t^{1-\theta}}{1-\theta} + \lambda \left(A_0 + \sum_{t=1}^T \frac{1}{(1+r)^t} Y_t - \sum_{t=1}^T \frac{1}{(1+r)^t} C_t \right)$$

For the household making a choice of consumption in period t , the first order condition is:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial C_t} &= 0 : \beta^t C_t^{-\theta} - \lambda \frac{1}{(1+r)^t} = 0 \\ \lambda &= \beta^t (1+r)^t C_t^{-\theta} \end{aligned}$$

This first-order condition holds for each period. So, for period $t+1$:

$$\frac{\partial \mathcal{L}}{\partial C_{t+1}} = 0 : \lambda = \beta^{t+1} (1+r)^{t+1} C_{t+1}^{-\theta}$$

⁹Here, $\beta = 1/(1+\rho)$ in Romer.

Combining these conditions we obtain the Euler equation:

$$\begin{aligned}\beta^t (1+r)^t C_t^{-\theta} &= \beta^{t+1} (1+r)^{t+1} C_{t+1}^{-\theta} \\ \frac{C_t^{-\theta}}{C_{t+1}^{-\theta}} &= \beta(1+r) \\ \frac{C_{t+1}}{C_t} &= [\beta(1+r)]^\theta\end{aligned}$$

In this case, consumption will follow a random walk only if $\theta = 0$ or if $\beta(1+r) = 1$.

We can see from the above expression that the interest rate affects the relative value of consumption today versus tomorrow in the utility function. As the interest rate rises, the household saves more, and consumption tomorrow will rise relative to consumption today. Using this more general utility function, we can see that the household seeks to equate not consumption, but the marginal utility of consumption, discounted to the present value. Hall's (1978) test is a subset of this, but we can see from above, that this test is not inclusive of more general specification for consumption.

7.4.2 Two-Period Case

This section looks at the basic implications of the model above using familiar indifference curve graphs. The idea is that income each period (Y_1, Y_2) and the interest rate affect the consumption bundles that households choose.

Thinking of consumption as two goods that households consume, we can look at the budget constraint to see how these two goods are priced. Changes in the interest rate affect the slope of the budget constraint. We can rewrite the budget constraint for two periods:

$$C_1 + \frac{1}{(1+r)}C_2 = A_0 + Y_1 + \frac{1}{(1+r)}Y_2$$

Solving for C_2 (on the vertical axes of the graphs from Romer):

$$C_2 = (A_0 + Y_1 - C_1)(1+r) + Y_2$$

Therefore, the slope of the budget constraint, $\frac{\partial C_2}{\partial C_1} = (1+r)$. So, an increase in the interest rate will cause the budget constraint to pivot to a steeper slope (for any given initial wealth and income levels). Romer provides case studies for a number of scenarios where the interest rate changes to study how households will respond when interest rates increase while income levels remain unchanged. As expected, there is a substitution effect and an income effect. The size of each effect depends on our assumptions about the parameter θ in the CRRA utility function.

We could use this same framework to study how changes in initial wealth, and changes in income levels would affect consumption decisions. However, this graphical analysis and framework are somewhat limiting. First, interest rates are

assumed to be constant, when they could actually be a function of government policy. Second, households can accumulate wealth over several periods (in a more general model), but we are unable to consider this in the simple two-period case. To look at these issues in more detail, one needs to consider the savings behavior of households over long horizons.

Bibliography

- [1] Ando, A. & Modigliani, F. (1963). The life cycle hypothesis of saving: Aggregated implications and tests. *American Economic Review*, 53, 55-84.
- [2] Brumberg, R. & Modigliani, F. (1954). Utility analysis and the consumption function: An interpretation of cross-section data. In Kenneth K. Kurihara, ed., *Post-Keynesian Economics*, 388-436. New Brunswick, NJ: Rutgers University Press.
- [3] Campbell, J. & Mankiw, N.G. (1989). Consumption, income, and interest rates: Reinterpreting the time series evidence. *NBER Macroeconomics Annual*, 4, 185-216.
- [4] Fisher, I. (1930). *The theory of interest*. New York: MacMillan.
- [5] Friedman, M. (1957). *A theory of the consumption function*. Princeton, NJ: Princeton University Press.
- [6] Hall, R. (1978). Stochastic implications of the life-cycle permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86, 971-987.
- [7] Keynes, J.M. (1936). *The general theory of employment, interest, and money*. London: Macmillan.
- [8] Kuznets, S. (1946). *National product since 1869*. New York: National Bureau of Economic Research, Princeton University Press.
- [9] Lucas, R. (1987). *Models of business cycles*. New York: Basil Blackwell.
- [10] Shea, J. (1995). Union contracts and the life-cycle/permanent-income hypothesis. *American Economic Review*, 85, 186-200.