CONSUMPTION AND LIQUIDITY CONSTRAINTS: AN EMPIRICAL INVESTIGATION

by

Stephen Zeldes

(24-85)

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

Consumption and Liquidity Constraints: An Empirical Investigation

Stephen Zeldes

Finance Department
The Wharton School
University of Pennsylvania

November 1985

I would like to thank Larry Ball, Olivier Blanchard, Stan Fischer, Whitney Newey, Stephen O'Connell, Julio Rotemberg, Larry Summers, members of the "Consumption Group" at the 1984 NBER Summer Institute, and members of the "Macro Lunch Group" at the University of Pennsylvania for helpful comments. This paper is a revised version of chapter 3 of my MIT Ph.D. dissertation.

I. Introduction

Most of the recent empirical work using aggregate time series data has rejected the restrictions on the data implied by the stochastic version of the permanent income hypothesis/life cycle hypothesis (PIH/LCH) and rational expectations. This includes work by Flavin (1981), Hansen and Singleton (1983), and Mankiw, Rotemberg, and Summers (1984), among others.[1] Some of these authors suggest that the rejections occur because some individuals are liquidity constrained.

Previous work has focused on deriving and testing implications for consumption under the null hypothesis that the PIH/LCH is the true model. In some cases, the alternative hypothesis is simply that the model does not fit the data, while in others the alternative hypothesis is that all or part of consumption is proportional to income (the "Keynesian" alternative). Little has yet been done, however, on deriving and testing implications under the specific alternative hypothesis that individuals maximize expected lifetime utility subject to a sequence of borrowing constraints.

In this paper, I derive testable implications for the behavior of consumption in the presence of borrowing constraints. The tests of these implications depend crucially on observing individual household behavior over time. I therefore use the PSID, a representative panel of U.S. families, to test these implications. My goal is to learn whether liquidity constraints are capable of explaining the rejections of the PIH discovered in the literature. This paper builds on panel data work on consumption by Hall and Mishkin (1982), Shapiro (1984), and Runkle (1983).[2]

I begin by describing three commonly mentioned divergences from the PIH/LCH model and examining the circumstances under which each will be implied by a model with the added restriction that net holdings of traded assets be non-negative in all periods. The first divergence is that the stochastic Euler equation for consumption between the current and next period is violated, the second is that the level of consumption is less and the sensitivity of consumption to transitory income is greater than the PIH would suggest, and the third is that consumption is equal to income ("Keynesian" behavior). In a two period model with a borrowing constraint, each occurs if and only if the others do. That is not true, however, in a multiperiod model with both present and future constraints. Only under rare circumstances will liquidity constraints imply "Keynesian" type behavior, and the Euler equation between t and t+1 will often be satisfied in cases where future constraints still lower consumption and lead to excess sensitivity.

I set up two formal tests, each of which depends on the first potential implication of liquidity constraints--the violation of the Euler equation. My tests involve splitting the observations into two groups on the basis of wealth to income ratios. Those observations with a low W/Y ratio (Group I) are most likely to be up against a binding liquidity constraint. If liquidity constraints are important, then the following should be true. First, the Euler equation should be satisfied for Group II but violated for Group I. These violations could take the form of implausible parameter estimates, a rejection of overidentifying restrictions, or both. Second, there should be a one-sided inequality in the Euler equation for Group I observations. If an individual would like to transfer additional resources from tomorrow to today, but is constrained from doing so, then the marginal utility of consumption must be higher today relative to tomorrow than would be predicted in a model with no constraints. In other words, the Lagrange multiplier associated with the borrowing constraint should be strictly positive. I derive an estimate of the Lagrange multiplier, equal to the part of consumption growth that is unexplained by the Euler equation. If borrowing constraints exist, this estimate should have a positive mean for Group I observations.

An advantage of these tests is that they do not require specifying either the closed form solution for consumption in the presence of borrowing constraints, or the particular income process. In addition, the technique that I use provides an estimate of the Lagrange multiplier associated with the constraint, and the technique could be useful to others who might like to estimate an approximation to the closed form model for the Lagrange multiplier.

If one were to hold all else constant (including future income), an increase in current income would relax a binding borrowing constraint and therefore would reduce the Lagrange multiplier. As a proxy for this partial correlation, I estimate the total correlation between current income and the Lagrange multiplier and test whether it is negative. This third test is suggestive, but not a formal test, because the sign of the overall correlation coefficient need not be the same as the sign of the partial correlation coefficient.

I test each of these implications for consumption in the presence of binding constraints on the transfer of resources between tomorrow and today. I use up to ten annual observations per family on food consumption and other variables from the Panel Study of Income Dynamics (PSID), a large panel of U.S. families.

The results are generally supportive of the view that liquidity constraints have important influences on consumption. The primary results come from a split of the sample based on liquid assets relative to income. For this sample split, the Euler equation is rejected for Group I observations, but not for Group II observations. For Group I, the average

Lagrange multiplier is positive and statistically significant, implying that there is the appropriate one-sided inequality in the Euler equation for constrained observations. The point estimate indicates that borrowing constraints caused annual food consumption growth for Group I to be 1.8 percentage points higher than it would have been in the absence of constraints. Finally, the total correlation between current income and the Lagrange multiplier was in fact negative, although statistically insignificant. In addition to the split based on liquid assets, an alternative split was tried that added together liquid assets and housing equity. These results are somewhat mixed, and are discussed below.

The remainder of the paper is structured as follows. In section II, I present the basic model with and without constraints. In section III, I distinguish between three possible effects of borrowing constraints: violation of the Euler equation, excess sensitivity, and "Keynesian" consumption. In section IV, I describe the tests that I perform to test for the importance of borrowing constraints. In section V, I describe the data and the techniques used to split the sample. (A detailed description of how I construct the variables used for the analysis is presented in the appendix.) In the sixth section, I examine the results, and concluding remarks are presented in section VII.

II. The Basic Model With and Without Borrowing Constraints

I begin by writing down the standard model without borrowing constraints and the corresponding set of Euler equations. The Euler equation approach to testing the permanent income/life cycle model under rational expectations was pioneered by Hall (1978) and extended by Mankiw (1981), Hansen and Singleton (1983), and others. Shapiro (1984) first applied this approach to time series/cross section data on consumption. A goal of this paper is to help determine whether the empirical rejections found by these authors are due to borrowing constraints rather than to a failure of other auxilliary assumptions. In the second part of this section, I write down the model that includes borrowing constraints and derive the corresponding set of Euler equations.

A. The Model Without Constraints

I assume that families maximize the expected value of a time separable lifetime utility function. In each period t, family i chooses consumption C_{it} and portfolio shares $\{w^j_{it}\}$ in order to:

$$\text{Max } E_t \sum_{k=0}^{T-t} (1/1 + \delta_i)^k \, U(C_{i,t+k}, H_{i,t+k}) \tag{1a}$$

$$\text{s.t. } A_{i,t+k} = (A_{i,t+k-1})*\Big(\sum_{j=1}^{M} w^j_{i,t+k-1}(1 + r^j_{i,t+k-1})\Big) + Y_{i,t+k} - C_{i,t+k} \tag{1b}$$

$$k = 0, \ldots, T-t$$

$$C_{i,t+k} \geq 0 \qquad k = 0, \ldots, T-t \tag{1c}$$

$$A_{iT} \geq 0 \tag{1d}$$

where:

A_{it} = real end of period financial wealth of family i in period t (after receiving income and consuming).

w^j_{it} = the share of end of period t wealth held in asset j.

r^j_{it} = the ex-post real after tax return on the j^{th} asset between periods t and t+1.

U = the one period utility function.

H_{it} = family i's tastes in period t.

C_{it} = real consumption of family i in period t.

Y_{it} = real disposable labor income of family i in period t.

E_t = the expectation operator, conditional on information available at time t.

As discussed in Zeldes (1984), analytic solutions to this problem when income is stochastic cannot in general be derived. Perturbation arguments can be used, however, to derive a set of first order conditions, or Euler equations, which are necessary for an optimum. An individual should be unable to increase expected lifetime utility by consuming one less unit today, increasing his holdings of any asset j between today and tomorrow, and consuming the extra gross returns tomorrow. Similarly, if the individual is not constrained from reducing his holdings of asset j below the current amount, then he should be unable to increase expected utility by consuming one more unit today, decreasing holdings of asset j, and reducing consumption by the corresponding amount tomorrow. In the unconstrained case, this leads to the following set of Euler equations:

$$U'(C_{it}, H_{it}) = E_t \frac{U'(C_{i,t+1}, H_{i,t+1})(1 + r^j_{it})}{(1 + \delta)} \qquad \begin{array}{l} i = 1, \ldots, N \\ t = 1, \ldots, T-1 \\ j = 1, \ldots, M \end{array} \quad . \tag{2}$$

where $U'(\cdot, \cdot)$ denotes the partial derivative of U with respect to C. If expectations are rational, this leads to:

$$\frac{U'(C_{i,t+1}, H_{i,t+1})(1 + r^j_{it})}{U'(C_{it}, H_{it})(1 + \delta_i)} = 1 + e^j_{i,t+1} \tag{3}$$

where $e^j_{i,t+1}$ is an expectational error uncorrelated with information known at time t. (For notational simplicity, I drop the superscript j in what follows.)

B. The Model With Borrowing Constraints

While no single definition of liquidity constraints has been agreed on in the literature, liquidity constraints must involve some price and/or quantity restrictions on the trading or holding of assets. The type of constraint considered here is a restriction on the total net stock of traded assets. Since A_{it} above is equal to total traded assets minus liabilities, this borrowing constraint adds the following set of restrictions to equations (1a)-(1d):

$$A_{i,t+k} \geq -B \qquad k = 0, \ldots, T-t-1 \tag{1e}$$

where B is the limit on net indebtedness. While consumer debt clearly exists in the U.S., it seems a reasonable working hypothesis that consumers cannot borrow on net against non-traded assets such as future labor income, in other words that debt cannot exceed total non-human assets. This would imply that B = 0. I therefore use this definition as the alternative hypothesis to the unconstrained model. Under the null hypothesis, agents can borrow and lend at the risk free rate. Under the alternative hypothesis, net wealth, excluding human capital and perhaps other thinly or non-traded assets, is restricted to be non-negative.[3] While I think of these "borrowing" constraints as being a subset of possible "liquidity" constraints, I use the terms interchangeably in this paper.[4]

While the alternative hypothesis is that all individuals face the set of constraints (1e), at any point in time constraint (1e) will be binding for some individuals and not for others. Those for whom the constraint is binding will be those who chose not to build up their wealth in earlier periods and/or those who received exceptionally bad draws of income or portfolio returns.

When the set of constraints (1e) are added to the model, the resulting Euler equation is:

$$U'(C_{it}, H_{it}) = E_t\left[\frac{U'(C_{it+1}, H_{it+1})(1 + r_{it})}{1 + \delta_i}\right] + \tilde{\lambda}_{it} \tag{4}$$

where $\tilde{\lambda}_{it}$ is the Lagrange multiplier (known at time t) associated with constraint (1e) for time t.[5] $\tilde{\lambda}_{it}$ is equal to the increase in expected lifetime utility that would result if the current constraint were relaxed by one unit; it is the extra utility that would result from borrowing an extra dollar, consuming the proceeds, and reducing consumption by the appropriate amount next period. Since agents are constrained from borrowing more, but not from saving more, λ_{it} enters equation (4) with a positive sign. At the constrained optimum, the marginal utility from consuming an extra unit today is always <u>greater than</u> or equal to the marginal utility from waiting until tomorrow to consumer the extra amount. As will be discussed below, $\tilde{\lambda}_{it}$ can be equal to zero even if consumption is constrained away from its global optimum, as long as the individual is optimally allocating resources between t and t+1.

It will be convenient for future purposes to normalize $\tilde{\lambda}_{it}$ by a positive term that is nonstochastic as of time t, as follows:

$\lambda_{it} \equiv \tilde{\lambda}_{it}/E_t\left[\left(\frac{1 + r_{it}}{1 + \delta_i}\right)U'(C_{it+1}, H_{it+1})\right]$. As before, λ_{it} will be positive when the constraint is binding and zero when it is not binding. Rewriting (4) we get:

$$E_t\left[\frac{U'(C_{it+1}, H_{it+1})(1 + r_{it})}{(1 + \delta_i)U'(C_{it}, H_{it})}\right](1 + \lambda_{it}) = 1 \tag{5}$$

and, as before, rational expectations implies:

$$\left[\frac{U'(C_{it+1}, H_{it+1})(1 + r_{it})}{(1 + \delta_i)U'(C_{it}, H_{it})} (1 + \lambda_{it})\right] = 1 + e_{it+1} \tag{6}$$

where e_{it+1} is $(1 + \lambda_{it})$ times an expectation error about the product of the interest rate and the marginal rate of substitution.

III. Euler Equation Violations, Excess Sensitivity, and "Keynesian" Consumption

Three descriptions of divergences from the unconstrained stochastic PIH/LCH model appear frequently in the literature.[6] In this section, I want to carefully distinguish each of these effects and describe the circumstances under which each will be implied by a model with borrowing constraints. The three are as follows.

(a) The Euler equation between t and t+1 is violated.

(b) The level of consumption is lower and there is greater sensitivity of consumption to current transitory income than the unconstrained PIH/LCH model would predict.

(c) Consumption is equal to income ("Keynesian" behavior).

In a two period model, all three of these occur whenever a liquidity constraint is binding. In a multi-period model with both present and future constraints, however, the direct correspondence disappears. In particular, two important observations need to be made. First, borrowing constraints will not in general imply "Keynesian" type consumption behavior. Second, if the Euler equation between this period and next is satisfied, this does not necessarily imply that the individual's current consumption behavior is identical to that of a totally unconstrained PIH/LCH consumer. In the following paragraphs, I describe why borrowing constraints will always imply (b), only sometimes imply (a), and almost never imply (c).

In a two period model, equation (1e) constitutes a single constraint. In a multiperiod model, however, (1e) is a set of constraints, one for each time period. Because of this, when considering a multiperiod world we must distinguish between two different concepts of constrained behavior. The first is that a consumer is constrained away from his global optimum; consumption is less than what it would be if the <u>entire set</u> of constraints (1e) were relaxed (i.e., C^* is less than what it would be if $\{A_{it+k} \geq 0 \ \forall k \geq 0\}$ were not imposed). The second notion of constrained behavior is that the current constraint is binding; consumption is less than what it would be if the <u>current</u> (time t) constraint were relaxed, but future constraints remained in place (i.e., C^* is less than what it would be if $A_{it} \geq 0$ were not imposed, but we continued to impose $\{A_{it+k} \geq 0 \ \forall k \geq 1\}$). If the current constraint is that net worth be non-negative and it is binding, then the end of period net worth of the individual must be equal to zero. (If net worth were positive, the individual could always transfer resources at the margin by consuming part of wealth.)

The first notion is clearly broader than the second. The time t consumption of a risk averse individual will be less than the global optimum any time there is a positive probability that for some $k \geq 0$ the constraint $A_{i,t+k} \geq 0$ will be binding.[7] While I have not seen a general proof in the literature, Zeldes (1984) provides numerical examples that show that a restriction on debt that has any probability of being binding in the future will also raise the sensitivity of consumption to innovations in transitory income. If we interpret this as a general result, this means that borrowing constraints lead to "excess sensitivity" of consumption to transitory income, even if the borrowing constraints are not binding today and will be binding in the future only with some positive probability. To understand the intuition behind the excess sensitivity result, consider a certainty model in which a constraint will be binding at some known future date. This will effectively shorten the horizon, and therefore the extra consumption from any transitory increases in income will be spread out over a shorter period, i.e., the MPC out of transitory income will be larger than without the constraint. This intuition carries through to an uncertainty model, in which constraints are binding in future periods only with some positive probability.

A lower level of consumption and higher sensitivity of consumption to transitory income will be implied any time there is a positive probability of at least one of the set of constraints (1e) being binding. Obviously, it need not be the current constraint that is the one that is binding. But only if the current (time t) constraint is binding will the Euler equation between t and t+1 be violated. Therefore it will often be the case that the consumer is globally constrained, but the Euler equation between t and t+1 is satisfied (i.e., $\lambda_{it} = 0$). In other words, for a given consumer at a given point in time, borrowing constraints can imply (b) without implying (a).

A concrete example might help in understanding this distinction. Consider an individual entering the period with positive financial wealth who expects income to be considerably higher in the distant future, but relatively constant in the short run. For many specifications of preferences, this individual will be constrained to consume less than his optimal amount, but will still carry over positive wealth from this period to next. Given future constraints (i.e., for $k > 0$), the consumer is optimally allocating resources between today and tomorrow, and the current constraint (i.e., for $k = 0$) is not binding. The Euler equation that is commonly estimated will be satisfied. The same phenomena can occur even if income is expected to be constant over time, because people may save today to reduce the chances of hitting future constraints.[8]

The second point of this section is that borrowing constraints will not in general imply Keynesian behavior. In what follows, I examine the conditions under which the consumption pattern of a PIH/LCH consumer subject to borrowing constraints would match that implied by a Keynesian consumption function. The standard Keynesian consumption function is written $C_t = C + cY_t$. This means that the expected growth of consumption between t and t+1 will in general be a function of Y_t. This constitutes a violation of the Euler equation, and if this behavior is due to borrowing constraints it must be that the time t constraint is binding. The only form of Keynesian behavior that could possibly be explained by currently binding constraints would be the specific example: $C_t = Y_t$. To get consumption equal to income in period t requires that the current borrowing constraint be binding in t-1 and t. (If the person enters and leaves the period with no wealth, then consumption must have equalled income during the period.) To get Keynesian behavior in both t and t+1 requires that the current borrowing constraint be binding three periods in a row (t-1, t and t+1). In order for an individual to continuously exhibit Keynesian behavior, the current constraint would have to be binding period after period.

Under what circumstances would we expect this to be the case? If in every period, regardless of the draw of income, individuals expected income to be higher in the next period, then the constraint could (for the right choice of parameters) be binding in every period. But in general, individuals receive good draws and bad draws of income. If a draw is especially good, the individual would want to save, and thus the current constraint would not be binding.

The general point here is the following. Unconstrained PIH/LCH consumers smooth out all fluctuations in income by borrowing and saving. Keynesian consumers never smooth out fluctuations in income. For borrowing constrained PIH/LCH consumers to act like Keynesians would require that they never borrow or save. But these consumers do have the option of saving, and will do so to smooth consumption in response to some fluctuations in income, in particular good current realizations of income. Therefore the consumption of individuals optimizing lifetime utility subject to a series of borrowing constraints will not in general match that implied by a Keynesian consumption function.

Most recent work on consumption has focused on deriving and testing implications for consumption under the null hypothesis that the PIH/LCH is the true model. In some cases, the alternative hypothesis is simply that the model does not fit the data (Hall (1978), Hansen and Singleton (1983), Mankiw, Rotemberg and Summers (1984)). In other cases, the Keynesian model is used as the benchmark alternative. The latter is the case, for example, in papers by Flavin (1981), Hall and Mishkin (1982), and Hayashi (1984).[9] The discussion in this section suggests that in order to test whether liquidity constraints are important, we need to improve on the Keynesian benchmark alternative. In the following sections, I construct tests that do just that: they test the null hypothesis that individuals are unconstrained PIH/LCH consumers against the specific alternative hypothesis that individuals are maximizing lifetime utility subject to a constraint on borrowing.

IV. Description of Euler Equation Tests

All of the tests in the rest of the paper are based on implication (a) above, the violation of the EE. The advantage of basing tests on the violation of the Euler equation rather than directly looking at the level of consumption or the sensitivity of consumption to current income is that to do the Euler equation tests one need not specify the exact income process or write down a closed form decision rule for consumption. This is important, because a closed form solution to the unconstrained problem has only been derived under very restrictive assumptions, and no one has derived a closed form solution to the problem with constraints imposed.[10] From herein, when I refer to a liquidity constraint being binding this period I mean that the Lagrange multiplier associated with transferring resources between tomorrow and today is positive, i.e., that the Euler equation between today and tomorrow is not satisfied.

In order to estimate the Euler equations, we first need to make some assumptions about preferences and some identifying assumptions. These are spelled out in the next section, and in the section after that I describe the econometric tests that I carry out.

A. Assumptions about Preferences and Identification

In order to make equations (3) or (6) operational, assumptions need to be made both about the general form of the utility function, and also about the factors that shift tastes (H_{it}). I assume that the utility function is of the constant relative risk aversion form:

$$U(C_{it}, H_{it}) = \frac{C_{it}^{1-A}}{1 - A} * \exp(H_{it})$$

where A is the coefficient of relative risk aversion, assumed equal across households. I allow each family to have a different rate of time preference δ_i. Plugging this into equation (6), we get:

$$\frac{C_{it+1}^{-A} * \exp(H_{it+1} - H_{it})(1 + r_{it})(1 + \lambda_{it})}{C_{it}^{-A}(1 + \delta_i)} = 1 + e_{it+1} \qquad (7)$$

where e_{it+1} is an expectational error about the left hand side of equation (7) that has mean zero and is uncorrelated with information available at time t.

Unlike some previous authors, I allow changes in tastes to shift the utility function. It may be reasonable to assume that there are no aggregate changes in tastes, and thus ignore taste shifts in aggregate time series work. However, changes in tastes seem much more important when studying changes in consumption of individual families. As Ball (1985) points out, ignoring these changes in tastes when using micro data can lead to identification problems and inconsistent estimates.

I decompose the taste shifter (H_{it}) into a part related to observable (to the econometrician) factors and an (orthogonal) unobservable part, where the observable factors are age, age squared, and a measure of family composition, i.e.:

$$H_{it} = b_0 age_{it} + b_1' age_{it}^2 + b_2 \ln(AFN_{it}) + u_{it}$$

where AFN_{it} = a measure of family composition. Plugging this into equation (7), taking logs of both sides, and rearranging yields:

$$GC_{it+1} = 1/A\Big[-\ln(1+\delta_i) + b_0 + \ln(1+r_{it}) + b_1'(age_{it+1}^2 - age_{it}^2) + b_2 GAFN_{it+1} + (u_{it+1} - u_{it}) - \ln(1+e_{it+1}) + \ln(1+\lambda_{it})\Big] \quad (8)$$

where $GC_{it+1} = \ln(C_{it+1}/C_{it})$

and $GAFN_{it+1} = \ln(AFN_{it+1}/AFN_{it})$.

I assume that family composition and age of head at t+1 are known to the family as of time t and that the unobservable taste changes ($u_{it+1} - u_{it}$) are i.i.d. across time and families.

If e_{it+1} has mean zero, then $\ln(1 + e_{it+1})$ does not. To arrive at an error term with mean approximately zero, it is standard to take a second order Taylor expansion of $\ln(1 + e_{it+1})$ around $e_{it+1} = 0$: $\ln(1 + e_{it+1}) \approx e_{it+1} - \frac{1}{2}e_{it+1}^2$. This implies $E_t[\ln(1 + e_{it+1})] \approx -\frac{1}{2}\sigma_{e,it+1}^2$. Adding and subtracting from equation (8) gives:

$$GC_{it+1} = \frac{1}{A}\Big[\tfrac{1}{2}\sigma^2_{e,i,t+1} - \ln(1+\delta_i) + b_0' + \ln(1+r_{it})$$

$$+ b_1'(age^2_{i,t+1} - age^2_{it}) + b_2 GAFN_{it+1} + \tilde{v}_{it+1} + \ln(1+\lambda_{it})\Big] \qquad (9)$$

where $\tilde{v}_{it+1} \equiv (u_{it+1} - u_{it}) - (e_{it+1} - \tfrac{1}{2}e^2_{it+1} + \tfrac{1}{2}\sigma^2_{e,it+1})$ has expectation zero. The ex-post after tax return $[\ln(1 + r_{it})]$ may be correlated with the expectation error on the growth of consumption, due to correlations between the time t+1 consumption, income, and marginal tax rate. Therefore, I will estimate the Euler equations with two stage least squares. In order to qualify as a valid instrument, a variable must be correlated with the variables included in the regression, but uncorrelated with the error term $\tilde{v}_{it+1}$ which includes both an expectations error and the unobservable change in tastes. By the assumption of rational expectations, any variable known at time t will be orthogonal to the expectations error.[11] I assume that the marginal tax rate in t and the log of real disposable income in t are uncorrelated with the unobservable changes in tastes between t and t+1, and use these as the instruments. I also assume that the variable used below to split the sample is uncorrelated with the unobservable change in tastes. These assumptions are likely to be better than the assumptions used by others that there are no taste changes or that the instruments are orthogonal to <u>all</u> taste changes.[12]

It is likely that the expectational error e_{it+1}, while having mean zero, will be correlated across families due to macro shocks. For example, if aggregate income were unexpectedly high in a given period, this would tend to cause all individuals to have higher than expected consumption this period and therefore to have unexpectedly high consumption growth between last period and this period. In other words, there would be an aggregate expectational error. I assume that the expectational error e_{it+1} can be decomposed into the sum of two orthogonal mean zero components: a common (aggregate) component e^a_{t+1} and an idiosyncratic component e^b_{it+1} that is independently distributed across i and t. I subtract off the aggregate error component to define:

$$v_{it+1} \equiv \tilde{v}_{it+1} - (-e^a_{t+1})$$

where v_{it+1} is independently distributed across i and t, has mean zero, and is orthogonal to the instruments used. The variances of v_{it+1} and e_{it+1} are assumed to be constant across i and t.[13] This yields the equation that I estimate:[14]

$$GC_{it+1} = \frac{1}{A}\left[e^a_{t+1} + K + \ln(1+\delta_i) + \ln(1+r_{it}) + b_1 age_{it} + b_2 GAFN_{it+1} + v_{it+1} + \ln(1+\lambda_{it})\right] . \qquad (10)$$

For use later on I define $X_{it+1} \equiv \frac{1}{A}\left[v_{it+1} + \ln(1+\lambda_{it})\right]$. X_{it+1} is the growth in consumption above the amount that would be predicted as of time t in a model with no constraints. Note that, all else equal, a higher λ_{it} corresponds to a faster expected growth of consumption between t and t+1.

B. Implications of Borrowing Constraints and Corresponding Tests Based on Euler Equation Estimation

Splitting the Sample

Bernanke (1984), Hayashi (1985), and Runkle (1983), have used the approach of splitting a sample of households into two sets of observations--those likely to be constrained and those likely not to be constrained. I follow this approach, dividing the sample on a priori grounds based on the theory presented above. For one group of observations (Group I), the current constraint is binding ($\lambda_{it} > 0$), and for the other (Group II) it is not ($\lambda_{it} = 0$). The sample is split by observation, so that a given family can sometimes fall in Group I and other times fall in Group II. If the form of the borrowing constraint is that non-human wealth must be non-negative, then λ_{it} will be positive only if end of period non-human wealth is equal to zero. It therefore would be appropriate to split the sample into observations for which end of period non-human wealth is equal to zero and those for which it is positive. Unfortunately, however, wealth is imperfectly measured and is measured as an annual average. Therefore, I include in the "liquidity constrained" group (Group I) not only observations where wealth equals zero, but also observations where the wealth to (average past) income ratio is small but non-zero. Consistency of these tests requires that Group II contain only observations for which the current constraint is binding. It does not require, however, that Group I contain only constrained observations. The possibility that some extra observations have been put in Group I does not, therefore, pose a problem.

I try two different measures of wealth. One includes housing equity while the other does not, because of the possibility that housing wealth was not liquid and could not be easily borrowed against. I split the sample by variables known at time t to avoid the sample selection bias that would occur if the expectation errors in the Euler equation were correlated with the splitting variable.

Before describing the tests, two points are worth emphasizing. First, recall from above that even if the constraint is not currently binding in the sense that an observation falls into Group II, neither the level nor sensitivity of consumption need be equal to what they would be in the absence of any constraints. In other words, the Euler equation should be satisfied for Group II, but the consumption patterns of these observations need not otherwise mimic those of pure life cycle/PIH consumers.

Second, Equation (10) derived above includes λ_{it}, a measure of how severely the borrowing constraint is affecting consumption growth. Unfortunately, we cannot derive a closed form expression for λ_{it}. It will in general be a non-linear function of variables known at t such as income at time t and wealth carried between t-1 and t, as well as moments of future variables such as future income, interest rates and tastes. A goal of this paper is to derive testable implications for consumption with borrowing constraints without having to write down such a closed form expression. The following are three such implications.

Implication and Test (i): Euler equation estimation on the two sub-groups.

This first test follows Runkle (1983), and tests the implication that the Euler equation should be satisfied for Group II but not for Group I.[15] The standard orthogonality test of this type of model involves testing the overidentifying restrictions of the model that information known at time t is orthogonal to the error term.[16] This is done here by testing whether additional time t variables enter significantly when equation (10) is estimated. The variable that I use is the log of time t real disposable income (y_{it}). I estimate equation (10) separately for each group of observations, in each case including disposable income in the regression. Under the null hypothesis that borrowing constraints do not exist λ_{it} will equal zero for both groups, which means that the parameter estimates should be plausible and similar across groups, and income should be insignificant in both cases. Under the alternative hypothesis that borrowing constraints exist, λ_{it} will still be equal to zero for Group II (the constraint is not binding for this group). For Group I, however, λ_{it} will not equal zero, and will be correlated with y_{it} and presumably the other variables in equation

(10). Since λ_{it} is in the error term, we would expect to get a significant coefficient on income (i.e., reject the overidentifying restriction) and/or get implausible parameter estimates for Group I, but not for Group II.

Implication and Test (ii): The one-sided inequality of the Euler equation.

As just pointed out, under the hypothesis that borrowing constraints exist, λ_{it} is equal to zero for Group II but not for Group I. We can say more, however. Since individuals can be constrained from borrowing more, but not from borrowing less (saving more), λ_{it} must be strictly positive. Even though individuals cannot smooth consumption by borrowing in response to bad draws of current income or high expected future income, they can save and thereby smooth away an exceptionally good draw of income or an expected drop in future income. This is the reason that consumption in the face of borrowing constraints need not look at all like the consumption of a "Keynesian" consumer who sets consumption equal to income. In a world with borrowing constraints, the marginal utility of consumption can be too high relative to what is expected tomorrow, but never too low. Because U" < 0, the growth in marginal utility is inversely related to the growth in consumption. Therefore, consumption can be expected to grow faster than if it were unconstrained, but can never be expected to grow more slowly than if it were unconstrained.

Recall that $\ln(1 + \lambda_{it})$ is equal to the increase in expected consumption growth that is due to the presence of the borrowing constraint.[17] X_{it+1} $(= \frac{1}{A} [\ln(1 + \lambda_{it}) + v_{it+1}])$ is an unbiased estimate of $\ln(1 + \lambda_{it})$ since v_{it+1} is a mean zero error term. I cannot directly measure $\ln(1 + \lambda_{it})$ or X_{it}, but I can derive a consistent estimate $(\hat{X}_{it})$ of X_{it}. The sample average $(\bar{\hat{X}}_{it})$ of $\hat{X}_{it}$ for Group I is a consistent estimate of the sample average of $\ln(1 + \lambda_{it})$ for that group.[18] In words, $\bar{\hat{X}}_{it+1}$ is an estimate of the average expected excess growth in consumption for Group I due to liquidity constraints. If liquidity constraints are important for consumption behavior, this estimate should be positive, statistically significant, and quantitatively large.

The goals of this test are to come up with an estimate of the average excess growth of consumption due to borrowing constraints, and to test whether it is strictly positive. In order to achieve these goals, it is crucial to have an estimate of $\overline{\ln(1 + \lambda_{it})}$ that is consistent both under the null <u>and</u> the alternative hypothesis. To arrive at such an estimate, I first estimate equation (10) (without λ_{it} included) for Group II. The resulting parameter estimates are consistent estimates of Group I parameters whether or not liquidity constraints exist, as long as the constraint is not binding for Group II and the underlying parameters are the same for the two groups. I then use these parameter estimates to construct estimates for Group I of X_{it+1}, the difference between actual consumption growth and that predicted by the right hand side variables. These estimates are consistent estimates of X_{it+1} for Group I, <u>whether or not Group I is liquidity constrained</u>. If I had instead used parameter values estimated from Group I observations (or the entire sample), they would <u>not</u> have been consistent if Group I were in fact liquidity constrained.

<u>Implication and Test (iii): The relationship between λ_{it} and y_{it}.</u>

For an individual facing a binding borrowing constraint, if disposable income at time t increases, and nothing else in the model changes, this will directly relax the constraint and therefore λ_{it} will fall. Consumption will rise today relative to tomorrow, lowering the expected growth in consumption. In other words, if borrowing constraints exist, they should imply a negative partial correlation between λ_{it} and y_{it}. Since y_{it} is assumed to be uncorrelated with v_{it+1}, there should be a negative correlation between X_{it+1} and y_{it}, holding all else equal. As a test of this, I regress the consistent estimate of X_{it+1} calculated above on y_{it} and test whether the sign is negative. Unfortunately, however, this is an estimate of the total correlation between λ_{it} and y_{it} which need not be of the same sign as the partial correlation coefficient. For example, if an increase in current income signals an even larger increase in future income, it is possible that such an increase will worsen rather than relax the binding constraint. This third test, therefore, should be taken only as suggestive evidence. The more transitory are changes in income and the less correlated income is with the other factors that influence λ_{it}, the closer the partial correlation coefficient will be to the total correlation coefficient.

V. Data

The data were collected by the Survey Research Center at the University of Michigan for the Panel Study of Income Dynamics (PSID).[19] The Survey has been conducted annually each spring and has followed the same families and their "splitoffs" over time. The first survey, "wave" 1, was conducted in the spring of 1968, and I use data through wave 15 (1982 survey). The questions in the survey refer primarily to the prior calendar year. Some questions, however, are aimed at capturing economic data as of the survey date. Throughout this paper, when I refer to the value of a variable in year or wave t, I mean the value of the variable as reported in the survey taken in year t, even though for some variables (such as family composition) this value is for the spring of the survey year, while for others (such as total income) it is for the previous calendar year.

In what follows, I briefly describe the variables used for my analysis. An exact description of the selection procedure for the sample, and the techniques used to construct each of the variables, is presented in the appendix.

Consumption. The PSID unfortunately does not include questions about total consumption. The survey does ask about food consumption, however. I use food consumption data for this study, as have others (such as Hall and Mishkin (1978)) who use the PSID to study consumption behavior. Questions were asked about the amount spent on food consumed at home and at restaurants. I deflate the responses by the CPI for food consumed at home and away from home, respectively, and sum the two real components to arrive at total real food consumption.

There are both advantages and disadvantages in using food consumption data. In order for the Euler equations presented in Section II to be rigorously justified, the utility function must be additively separable in food consumption and other consumption.[20] However, the assumption of separability between food and non-food consumption seems more plausible than the assumptions typically made in aggregate time series studies: separability between non-durables/services and durables, and perfect substitutability between non-durables and services. In addition, the consumption terms that belong in the Euler equation are the service flows in the given periods, not measured expenditures in those periods. Food expenditures are far less likely to include durable components that given service flows over subsequent periods,[21] and therefore the correspondence between current expenditures and consumption flows is certainly better for food expenditures than for total consumption expenditures.

It is not exactly clear what period the survey questions about consumption refer to. The food consumption question in the PSID is: "How much do you spend on the food that you use at home in an average week?" Since

1977, that question has been asked immediately following the question "Did you (or anyone else now living in your family) receive or buy food stamps, last month?" If the answer is yes, the subsequent question was: "In addition to what you spent on food stamps, did you spend any money on food that you use at home," and if so, "how much?" It seems clear that the objective of the surveyors was to measure current consumption flow. I interpret the responses (which are collected in the spring), therefore, as pertaining to the first quarter of the interview year, and time my price and interest rates accordingly.

Food consumption is undoubtedly measured with error.[22] The standard deviation of the growth rate of consumption is equal to 36%, which seems implausibly high to be explained by differences in expected growth rates, expectational errors, and/or changes in tastes. The estimates presented in the next section are consistent as long as the first difference of the log of the measurement error of consumption is unrelated to the instruments used.

The Real After Tax Interest Rate. The nominal interest rate used is the one year T-Bill rate (1st quarter average). I use the marginal tax rates (MTR) on unearned income for the head and wife for the year t+1. For waves 9-15, the MTR was estimated for each observation by the Survey Research Center, based on detailed income data and the relevant tax tables. For waves 3-8, only estimates of total taxes paid were provided on the data tapes. I use these, together with the tax tables each year, including surcharges where relevant, to estimate the MTR of the head and wife.

The measure of the inflation rate used is the growth of the overall food CPI between the first quarters of years t and t+1.

Real Disposable Income. Detailed questions were asked about income in the calendar year immediately preceding the interview. I calculate real disposable income as total family income, minus taxes and plus transfers, deflated by the yearly average of the NIPA personal consumption expenditures deflator.

As discussed above, I use two different measures of wealth--one includes housing equity while the other does not.

Housing Equity. One of the questions asked directly was the current value of the family's house. In most of the waves, the amount of outstanding mortgage principal was ascertained. For the remaining waves, I estimate the mortgage principal using answers from the surrounding waves, as long as no move took place and the mortgage principal declined over the interval. The difference between house value and outstanding mortgage principal equals house equity.[23]

Non-Housing Equity Wealth (NHEW). Unfortunately, no data is available on asset holdings other than housing equity. There are questions, however, about asset income received over the year. To estimate NHEW, I divide the first $250 of interest and dividend income by the passbook rate at commercial banks over the year, and all additional interest and dividend income by the average yield on three month T-bills over the year. This is an approximation that attempts to account for the amount of wealth typically held in savings accounts. Because of the difficulty of "scaling up" asset income other than interest and dividends, I exclude observations with substantial "other asset income."

Annual Foods Needs (AFN). I assume above that tastes for food consumption depend on family composition. AFN is a measure of the low budget food needs of the family based on composition at the time of the interview and serves as a measure of family composition. It is a weighted sum of the current ages of family members, adjusted for total family size. The weights are the same in each period.

Sample Selection

The original 1968 data included a special "poverty sample," which means that low income families are oversampled. I eliminate this poverty group from my sample, so that in 1968 the remaining sample was representative of the U.S. economy.[24] I include splitoff families as separate families, but eliminate them from the sample during their first year as a new family. Observations on a family currently living with another unrelated family are discarded because of the difficulty of accurately measuring family consumption. Finally, the interviewer occasionally could not elicit a response about the amount of a family's food consumption, and estimated the amount from past family data or from tables. I eliminate these observations from my study.

The food consumption questions were not asked in wave 6, and waves 1 and 2 lacked some other necessary questions. This, together with the fact that growth rates are used in the regressions, means that a maximum of 10 observations per family are used in the estimation.

Criteria for Splitting the Sample

As discussed in Section IV, the sample is split on the basis of two different wealth to income ratios. One is based on total wealth, while the other excludes housing equity. In each case, the sample is split by observation, not by person, so that a given family could be placed in one group some years and in the other group in other years.

In six of the relevant waves, there is a question in the survey that asks whether the family has savings, such as checking or savings accounts or government bonds, that amount to two months worth of income or more. I base the first split both on the response to this question and on the level of a constructed measure of liquid wealth that is available for all waves. In order to avoid incorrectly placing observations in Group II, I place an observation on the growth rate of consumption between t and t+1 in the liquidity constrained group (Group I) if either of the following held:

a) Their response to the above question in year t was no.

b) Their non-housing equity wealth (as described above) in period t was less than 2/12 of their average income during t and t-1. This second criterion is chosen to correspond most closely to the survey question.

For the second split, I ignore the direct question on savings, and place an observation in Group I if estimated total wealth, including housing equity, was less than two months worth of average income.

Table 1 contains the fractions of observations falling into Groups I and II for each split, and also a more detailed breakdown of liquid wealth positions. I expected that both methods of splitting the sample would lead to a fairly small Group I sample, but this is not the case. For the second way of splitting (based on total wealth), one third of the sample falls into Group I, and for the first way of splitting (based on liquid wealth), two thirds of the sample falls into Group I. There is some discrepancy between the direct question on liquid wealth and my calculations. For the years that the question was asked directly, only 45% responded that they did not have two months' worth of savings. It is therefore likely that I put in the liquidity constrained group some observations that belong in the non-constrained group. The reason for this is probably that some forms of wealth, such as cash or checking accounts, generated no interest payments and therefore are excluded from the "scaling up" procedure that I use to calculate non-housing wealth. As mentioned above, the tests rely on the consistency of the Group II

TABLE 1

Fractions of Population in Different Wealth Positions

Liquid Wealth Position

(Direct Survey Question. 7 waves. Entire Sample. 14,209 Observations)

1a.	Between 0 and 2 months worth of income now. > 2 months worth of income sometime in the last 5 years.	.15
1b.	Between 0 and 2 months worth of income now. Not > 2 months worth of income anytime in the last 5 years.	.15
1c.	0 now, > 2 months worth of income sometime in the last 5 years.	.05
1d.	0 now, not > 2 months worth of income in the last 5 years.	.10
2.	> 2 months worth of income now.	.55

Non-Housing Equity Wealth Split

(Also based on liquid wealth question; 13 waves, Entire sample for which NHEW could be calculated. 23,752 observations.)

Group I.	⩽ 2 months worth of (average income or liquid wealth position = 1 (a-d)	.67
Group II.	> 2 months worth of (average) income and liquid wealth position = 2	.33

Total Wealth Split

(13 waves. Entire sample for which total wealth could be calculated. 21,713 observations.)

Group I.	⩽ 2 months worth of (average) income.	.30
Group II.	> 2 months worth of (average) income.	.70

parameter estimates, and are therefore valid even if some non-liquidity constrained observations are included in Group I.

VI. Empirical Analysis

A. Estimation

As discussed above, the equations are estimated with two stage least squares. The basic Euler equation that I estimate follows directly from equation (10):

$$GC_{it+1} = \sum_{j=1}^{N} c^j \, FD^j_{it} + \Sigma_k d^k \, WD^k_{it} + a_1 age_{it} + a_2 GAFN_{it+1} + a_3 \ln(1 + r_{it}) + \varepsilon_{it+1} \qquad (11)$$

and the instruments for $\ln(1 + r_{it})$ are: y_{it} and MTR_{it}

where GC_{it+1} = the growth rate of real food consumption between t and t+1.

$\{FD^j_{it}\}$ = N family dummies, one for each family in the sample (= 1 if i = j, 0 otherwise).

$\{WD^k_{it}\}$ = nine wave k dummy variables, one for each included wave except the last (= 1 if t = k, 0 otherwise).

age_{it} = the age of the head in year t.

$GAFN_{it+1}$ = the growth in annual food needs between t and t+1.

$\ln(1 + r_{it})$ = the ln of one plus the real after tax riskless rate between t and t+1.

y_{it} = the log of disposable real income in year t.

MTR_{it} = the marginal tax rate of family i in year t.

The wave dummies are included to capture the aggregate component of expectations errors described earlier. Even though these aggregate errors have expected value zero, if left in the error term of the regression the resulting errors would not be i.i.d. and asymptotic inference would not be valid.[25]

It is likely that different families have different rates of time preference and it is possible that the rate of time preference is correlated with the instruments or the variable used to split the sample. In particular, families with a low rate of time preference are more likely to accumulate wealth, but will also tend to have a high growth rate of consumption. Ignoring this problem would lead to a sample selection bias, and for test (ii) the bias would be against the alternative hypothesis that families are liquidity constrained.[26] To avoid this problem, family dummies are included in the regressions. The family dummies were first "partialled out," by subtracting off family means from each variable, i.e., fixed effects estimation was performed.

For the first test (estimating the Euler equation on each subsample) all observations without missing relevant data are used. The second and third tests involve estimating the Euler equation for one group and imposing these parameter estimates on the other group in order to derive an estimate of λ_{it}. These parameters include a set of family dummies, so this means that each family dummy is estimated on family observations in Group II and then imposed on the family observations in Group I. Therefore, observations can be used for these tests only if the family has observations both in Group I and in Group II at some time during the sample. This reduces the number of Group II observations available for these tests, sometimes dramatically.

The standard errors given for tests (ii) and (iii) by a standard regression package are correct only under the assumption that estimated Group II parameters that are imposed on Group I are exactly equal to the true underlying parameters. Since this assumption is not true, the standard errors needed to be corrected. The standard errors reported here for tests (ii) and (iii) take into account the sampling error in the construction of the X_{it}, i.e., they account for the fact that the parameter estimates imposed on Group I, while consistent, are not necessarily equal to the true underlying parameters.[27]

B. Results

I first present the results based on the ratio of liquid (non-housing equity) wealth because they are more clear cut than the others.

Test (i): Estimation on Each of the Two Sub-Groups

The results of separate estimation of equation (11) for Groups I and II (all possible observations), are reported on the top of Table 2.

Group II: Let me first discuss the results for the high asset "unconstrained" group (Group II). For this group of observations, the Euler equation is not violated. The coefficient on the interest rate is positive, and implies a coefficient of relative risk aversion of 2, an estimate in line with previous empirical estimates.[28] The standard error on the coefficient is very large, however, and the coefficient is not statistically different from zero. The coefficient on age is negative, but not statistically different from zero. The coefficient on the growth in annual food needs is positive, as expected, and very significant. Finally, as would be predicted by a model with no currently binding liquidity constraints, the coefficient on the log of income is not significantly different from zero at the 5% level.

Group I: The Group I coefficient on GAFN is almost the same as for Group II. The coefficient on the interest rate implies a coefficient of relative risk aversion of .7 but again has a large associated standard error. However, the coefficient on income is negative, larger in magnitude than the Group II

TABLE 2

Split based on liquid (non-housing) wealth relative to average income.

Test (i): Euler Equation Estimates for Two Subsamples:
Dependent Variable: Log $C_{i,t+1}/C_{it}$

Independent Variable	Group I (Low W/Y)	Group II (High W/Y)
Age of Head	-.00851 (-1.27)	-.00442 (-.91)
Growth in Annual Food Needs t, t+1	.25 (9.85)*	.24 (6.59)*
Real After Tax T-bill rate t, t+1	1.42 (.85)	.49 (.34)
Log of Real Disposable Income t	-.0686 (-4.50)*	-.0378 (-1.75)
DOF	9164 (2724 families)	4433 (1566 families)

Tests on the Group I residuals constructed using Group II parameter estimates.

Test (ii) Estimate of average excess consumption growth for Group I due to binding constraint:

$\hat{\bar{X}}_{it}$	.018
S.E.:	.0107
(one sided test)	(1.72)*

Test (iii) Regression of estimate of excess consumption growth for Group I on the log of real disposable income.

Y_{it}:	-.0213 (-1.39)

DOF 3053
(1098 families)

t-statistics are in parentheses.
*indicates significance at 5% level.

Equations estimated with two stage least squares, and include time and family fixed effects.

coefficient, and statistically significant. As discussed in Section II, this is inconsistent with the permanent income hypothesis, and is what would be predicted for individuals subject to a currently binding liquidity constraint.

The results therefore indicate no violation of the Euler equation for observations with high non-housing assets relative to income, but a violation of the Euler equation for observations with less than two months' worth of income in non housing assets--exactly those observations we expect to be liquidity constrained.

Test (ii): The one-sided inequality in the Euler equation.

In the bottom of Table 2, I present a consistent estimate of the average Lagrange multiplier for Group I observations, equal to the average unexplained consumption growth for this group. Recall that if Group I observations face a binding borrowing constraint, this term should be positive. Since only a positive (and not a negative) estimate would cause us to reject the null hypothesis in favor of the alternative that constraints are important, a one-sided test is clearly appropriate. The estimate of the average Lagrange multiplier is positive, and statistically significant at the 5% level. The point estimate is equal to .018, indicating that the growth of consumption is on average 1.8 percentage points higher for Group I than would be predicted by a model with no binding constraints. (Raw average consumption growth for the entire sample was approximately zero.) Recall that this is the excess growth of food consumption. It seems reasonable that the excess growth of other forms of consumption would be even greater.

In summary, the results of both of these formal tests support the alternative hypothesis that borrowing constraints exist and affect consumption in important ways.

Test (iii): The relationship between unexplained consumption growth (t, t+1) and income (t).

Table 2 also presents the results of a regression (for Group I observations) in which the left hand side variable is a consistent estimate of X_{it} (the consumption growth unexplained by the Euler equation) and the right hand side variable is the log of real disposable income in time t. The coefficient is a consistent estimate of the relationship between the Lagrange multiplier and current income. As explained above, if liquidity constraints are important, the partial regression coefficient on income should be negative. This, however, is an estimate of the total regression coefficient. I find that the coefficient on income is in fact negative, but not statistically significantly different from zero.

Alternative Specifications

In Table 3 I present results based on splitting the panel by the total wealth to income ratio. I now include an observation in Group I if total estimated wealth (including housing equity) in period t is less than two months' worth of the average income in t and t-1. If, over the sample period, low liquid wealth individuals were able to borrow against housing equity, then this is the right way to split the sample.

The results are similar, but not as strong or clear cut as before. Lagged income enters significantly (with a negative sign) for both Group I and Group II, although the coefficient for Group I is closer to zero. The estimate of the average Lagrange multiplier is positive (.8 percent), but smaller than for the first split and not statistically significant. Finally, the estimate of the Lagrange multiplier is negatively related to the level of income (and statistically different from zero). One possible interpretation of this set of results is that borrowing against housing equity was difficult over the sample period and that therefore this second split is not the appropriate one.

Because there is some ambiguity as to exactly where to split the sample, I tried one last way of splitting the sample. The sample is split into three sets of observations, but the middle set of observations is not used, leaving only the two extreme sets. The observations in Group I did not own a house, and either responded that they had no liquid assets or reported zero asset income for the year. The observations in Group II possessed non-housing wealth of at least six months' worth of average income. With this split, Group I has about 20% and Group II 25% of the total observations, and the middle 55% is eliminated from the sample. The results are in Table 4.

The results for the first test support the liquidity constraint hypothesis. Lagged income enters significantly (with a negative sign) in the Group I Euler equation, while the point estimate on lagged income is much smaller and insignificant for Group II. Unfortunately, very few families had observations in both Groups I and II under this split, leaving only 325 observations (vs. 4152 under the first split) and extremely imprecise estimates for tests (ii) and (iii).

Each of the tests above was also performed under the assumption that all families had the same rate of time preference (i.e., no family dummies were included). The results for tests (i) and (iii) were similar, but the estimate of the average Lagrange multiplier (test (ii)) was close to zero and statistically insignificant. This suggests to me that the sample selection problem that arises when the dummies are omitted is probably important.

VII. Concluding Remarks

In those studies where the permanent income hypothesis is rejected empirically, liquidity/borrowing constraints are often suggested as a possible

TABLE 3

Split based on total wealth relative to average income.

Test (i): Euler Equation Estimates for Two Subsamples:
Dependent Variable: Log $C_{i,t+1}/C_{it}$

Independent Variable	Group I (Low W/Y)	Group II (High W/Y)
Age of Head	-.00850	-.00485
	(.78)	(.70)
Growth in Annual Food Needs t, t+1	.30	.22
	(6.87)*	(8.79)*
Real After Tax T-bill rate t, t+1	1.56	1.55
	(.34)	(1.42)
Log of Real Disposable Income t	-.0699	-.0458
	(-2.46)*	(-3.14)*
DOF	3198	9334
	(1532 families)	(2509 families)

Tests on the Group I residuals constructed using Group II parameter estimates.

Test (ii) Estimate of average excess consumption growth for Group I due to binding constraint:

$\bar{\hat{X}}_{it}$	.0083
S.E.:	.0127
(one sided test)	(0.66)

Test (iii) Regression of estimate of excess consumption growth for Group I on the log of real disposable income.

Y_{it}:	-.0573
	(-2.78)*

DOF 1405
(860 families)

t-statistics are in parentheses.
*indicates significance at 5% level.

Equations estimated with two stage least squares, and include time and family fixed effects.

TABLE 4

Extreme Split

Test (i): Euler Equation Estimates for Two Subsamples:
Dependent Variable: Log $C_{i,t+1}/C_{it}$

Independent Variable	Group I (Low W/Y)	Group II (High W/Y)
Age of Head	.02328	-.00435
	(1.77)	(-.87)
Growth in Annual Food Needs t, t+1	.28	.27
	(5.35)*	(5.91)*
Real After Tax T-bill rate t, t+1	1.09	.85
	(.17)	(.47)
Log of Real Disposable Income t	-.0632	-.0194
	(-1.98)*	(-.79)
DOF	2278	3056
	(1220 families)	(1188 families)

Tests on the Group I residuals constructed using Group II parameter estimates.

Test (ii) Estimate of average excess consumption growth for Group I due to binding constraint:

$\bar{\hat{X}}_{it}$	-.0032
S.E.:	.0339
(one sided test)	(-.09)

Test (iii) Regression of estimate of excess consumption growth for Group I on the log of real disposable income.

y_{it}:	-.0313
	(-.61)

DOF 172 (152 families)

t-statistics are in parentheses.
*indicates significance at 5% level.

Equations estimated with two stage least squares, and include time and family fixed effects.

explanation. Most tests of the PIH/LCH, however, either do not specify an alternative hypothesis or specify the Keynesian alternative that consumption is proportional to income.

In a model with constraints on borrowing (but not on saving), borrowing constraints will not in general imply Keynesian behavior. In this paper, therefore, I examine some of the properties of consumption under the specific alternative hypothesis that individuals optimize subject to a set of borrowing constraints and derive tests that shed light on whether or not borrowing constraints are important empirically. Each of the tests involves splitting panel data observations into two groups according to wealth-income ratios, and examining the behavior of the groups through Euler equation estimation. My method does not require specifying a closed form solution for consumption with borrowing constraints, yet yields an estimate of the Lagrange multiplier associated with a borrowing constraint. This technique might prove useful to others who would like to specify and estimate a model for this Lagrange multiplier in order to study what determines how constrained people are.

The empirical results presented should be interpreted with some caution. The data are considerably less than ideal, in that some variables need to be constructed and others are extremely noisy indicators of the true measure. With that said, these results generally support the view that borrowing constraints affect consumption in the U.S. At least for the split based on liquid assets, the Euler equation is violated for observations for which a constraint is likely to be binding (the low wealth/income ratio group), and is not violated for the remaining observations. In addition, an estimate of the average Lagrange multiplier associated with the borrowing constraint is significantly greater than zero for the low wealth/income ratio group, indicating that there is a one-sided inequality of the Euler equation in the appropriate direction. The results using a split based on total assets were similar, but less clear cut.

The results presented here suggest that borrowing constraints are important. Further research is clearly needed, however, developing and testing this and other specific alternative hypotheses to the unconstrained life cycle/permanent income model.

APPENDIX:

A Description of the Constructed Variables and the Sample Selection and Splitting Procedures.

Fifteen years or "waves" of PSID data was available at the time of this study. The surveys are conducted each spring and most of the questions refer to the preceding calendar year. The first survey used (wave 1) was conducted in the spring of 1968. The most recent survey used (wave 15) was conducted in the spring of 1982.

When I use the term "observation," I mean the data for a particular family in a particular year. For some observations, interviewers did not get a response for some questions. When data was missing for certain questions, the answers were estimated by the interviewer or the PSID staff. It was considered a minor assignment if the answer was approximated using other responses from current or past surveys. It was considered a major assignment if the answer was approximated using statistical tables.

Many of the questions asked about the "head of the household." In virtually all of the cases where the family included a married couple, the male was automatically considered the "head" and the female the "wife." For families with only one adult, this adult was referred to as head regardless of gender.

I. CONSTRUCTED VARIABLES

I describe below how each of the variables used either for estimation or splitting the sample is constructed. These include growth in food consumption (GC), the real after tax interest rate (r), growth in annual food needs (GAFN), the log of real disposable income (Y), non-housing wealth, and housing equity.

A. Food Consumption

In waves 10-15, the question on the amount spent on food at home was explicitly designed to exclude the amount saved on food stamps. The PSID researchers were trying to measure "out of pocket" expenditures on food. In waves 7-9, food purchased with food stamps was not excluded explicitly, however a follow-up question asked whether or not the value of food stamps received was included in their food expenditure answer. In waves 1-5, the surveyor attempted to exclude the value of food stamps received, although this was not done systematically. No food consumption questions were asked in wave 6.

For the Euler equation, I want to measure total food consumption, and I therefore include food purchased with food stamps in my measure of food consumption. This requires adding, in the appropriate waves, the net value of food stamps (amount of food stamps received - amount paid for food stamps) to the out of pocket expenditures on food. This is done for waves 1-5 and waves 10-15. For waves 7-9 it was done if the follow-up question indicated that food stamps had not been included.

I calculate the net value of food stamps as follows. For waves 8-15, I scale up the response to a question on the net value of food stamps in the previous month. This question was not asked in waves 1-7. For these waves, I use the answer to the question on the net value of food stamps received in the previous calendar year.

As discussed in the text, I interpret the survey questions on consumption as referring to the first quarter of the interview year. The questions on food consumption were designed to capture flow consumption at the time of the interview, although it is not obvious how they were interpreted by each respondent. The 1972 PSID book on study design and survey procedures stated (referring to the 1972 survey year): "All the income questions refer to the year 1971, while the food expenditure questions refer to the flow existing at the time of the interview."

I deflate the adjusted nominal amount spent on food at home and the nominal amount spent on food away from home by the CPI for food at home and the CPI for food away from home, respectively, in the first quarter of each year. The two real quantities are summed to arrive at total real food consumption (C_{it}). The growth rate of consumption (GC_{it+1}) is equal to $\ln(C_{i,t+1}/C_{it})$.

B. The Real After Tax Interest Rate

The log of one plus the real after tax interest rate is equal to $\ln\left(1 + RN_t(1 - MTR_{i,t+1})\right) - \ln\left((1 + P_{t+1})/(1 + P_t)\right)$, where RN_t is the nominal interest rate between t and t+1, $MTR_{i,t+1}$ is the marginal tax rate in t+1 and P_t is the price level in t.

The Nominal Interest Rate: The interest rate is timed to cover the period between the first quarter of the survey year and the first quarter of the subsequent survey year, in order to coincide with the consumption data. The nominal rate used is the quarterly average of the market yield on one year Treasury bills.

The Price Level: The price level is the quarterly average of the overall CPI for food.

The Marginal Tax Rate: For waves 3-15, total federal taxes paid by head and wife are estimated by the PSID staff (by computer since wave 13) based on data on income, filing status, number of dependents, and the appropriate tax tables. For waves 13-15, mortgage payments and property taxes are taken into account when estimating the amount of deductions.

For waves 9-15, the corresponding marginal tax rate is reported on the data tapes, but for the earlier waves it is not. For these years, I therefore take the PSID estimate of taxes paid and work backwards on the relevant tax table (single, married, or head of household) for the appropriate year to get the marginal tax rate. It appears that the PSID made no correction for the ceiling on taxes on earned income, and therefore my technique of reading directly off of the tables yields the correct marginal tax rate on unearned income. The marginal tax rates are multiplied by 1.075 in 1968, 1.10 in 1969, and 1.025 in 1970, to account for surcharges in those years.

The after tax nominal rate is calculated using the MTR in year t+1.

C. Annual Food Needs (AFN)

This variable is included as a measure of family composition. The level of AFN is based on the 1967 U.S. Department of Agriculture's Low Cost Plan estimates of weekly food costs. The variable is calculated by the PSID as follows. First, each family member is given a value, based on Table A1 below. Then the values are summed up for the family as it existed at the time of the interview. Finally, the sums are adjusted for food economies of scale based on Table A2 below. Note that the same tables are used for all 15 waves. The growth in annual food needs ($GAFN_{i,t+1}$) is equal to $\ln(AFN_{i,t+1}/AFN_{it})$.

D. Disposable Income

The log of real disposable income is included in some of the regressions. I compute disposable income as taxable income of head and wife + taxable income of others + total transfers of head and wife + total transfers of others + bonus value of food stamps - Federal income taxes of head and wife - Federal income taxes of others - Social Security taxes of head and wife, each for the entire calendar year prior to the survey year. (The PSID did not

TABLE A1

Individual Food Standard

Age	Male	Female
Under 4	$3.90	$3.90
4-6	4.60	4.60
7-9	5.50	5.50
10-12	6.40	6.30
13-15	7.40	6.90
16-20	8.70	7.20
21-35	7.50	6.50
36-55	6.90	6.30
56 and Older	6.30	5.40

TABLE A2

Adjustment for Economies to Scale

+20% for one-person families
+10% for two-person families
+ 5% for three-person families
No adjustment for four-person families
- 5% for five-person families
-10% for families with six or more persons

include the bonus value of food stamps as part of transfers.) I deflate the nominal value by the annual average of the NIPA Personal Consumption Expenditures (PCE) deflator (1972 = 100).

Each of the components of disposable income was included in the data tapes except for Social Security taxes. The data for all of the waves on income, transfers, and food stamps, as well as the data for waves 9, 10, 12-15 for property taxes, come from direct survey questions. Total income taxes are estimated by the PSID (see by description of the marginal tax rate above). I add on the surcharges for 1968-1970 and subtract off the low income tax credit in the years when the PSID did not do this (waves 9-12). For waves 1-8, property taxes were imputed by the PSID based on house value and distance from the nearest city. There is no data on the tapes about wave 11 property taxes. I estimate these based on the wave 10 or wave 12 property tax rate (if no move took place) and the wave 11 house value.

I impute Social Security taxes by multiplying the appropriate Social Security tax rate by the lesser of annual wages and the ceiling on wages taxable by Social Security. I use the self employed tax rate for the head if he/she is self-employed. The regular rate is used for the spouse's wages because there is no information available in many waves on whether the spouse is self employed.

E. Non-Housing Wealth

I construct the wealth variable from questions on asset income, and to do so requires some rather bold assumptions. Since this variable is only used to split the sample, however, I do not think that the results presented in the paper are sensitive to those assumptions.

I estimate the stock of wealth using data on the flow of asset income and an assumed rate of return on wealth. Because of the difficulty of approximating the rate of return on family owned businesses, I use the response to a question on "dividends, interest, rent, trust funds, or royalties" for the head, and a similar question for the wife. I then treat the wealth data as missing for anyone with home business income of more than $100 in absolute value. (Actually, the dividends, interest, etc. question is used in waves 9-15, but because this was reported only as a bracketed variable for waves 1-8, I use total asset income for waves 1-8. In these waves, I record missing data if total asset income lay outside of the range of the sum of the brackets (for head and wife), for dividends, interest, etc. income.)

In waves 9-15, asset income data on other family members was reported only in bracketed form, i.e., as being in one of ten brackets. No asset data on others was available in waves 1-7. (Actual dollar values were reported for wave 8.) I want to capture total family wealth, but if asset income of others is non-zero, I have no way of estimating their wealth. I therefore record wealth as missing in waves 9-15 if others' asset income is non-zero. For waves 1-8, I record wealth as missing if others' asset income in wave 8 is non-zero.

To estimate the stock of assets, I divide the first $250 of interest, dividends, etc. income by the annual average of the passbook rate at commercial banks, and the rest of such income by the annual average of the yield on three month Treasury Bills. I do this to attempt to account for the amount of wealth typically held in low interest accounts.

Real non-housing wealth is equal to the nominal stock deflated by the PCE deflator.

F. <u>Housing Equity</u>

Real housing equity is equal to house value minus outstanding mortgage principal, deflated by the PCE deflator. House value was a direct survey question for all 15 waves. Unfortunately, the question about outstanding mortgage principal was not asked in waves 7 and 8. (It was also not asked in waves 1, 6 and 15, but observations from these waves were not included in the regressions.) I estimate the outstanding mortgage principal in waves 7 and 8 by interpolating waves 5 and 9, or by extrapolating the changes in waves 4-5 or 9-10. This was done only if the family did not move between the relevant years, and the reported mortgage principal declined over time.

The survey asked whether the family had a second mortgage, but did not ask the corresponding principal. Therefore, I record missing data for housing equity for observations where a second mortgage exists, (and for observations based on interpolations or extrapolations of observations with a second mortgage).

II. Sample Selection

A. Families Excluded

Families are followed through time, and new families that are formed from the original ones (through "splitoffs" such as children leaving home or parents separating) are also included in the sample.

The 1968 sample contained a subsample that was representative of the U.S. population, and a subsample of poverty families. I use only the initial representative subsample of families (and their "splitoffs"), and exclude the poverty subsample.

I follow a family by keeping track of the head of a household. In the representative subsample there were 4130 individuals who were "heads" for at least one wave of the fifteen wave panel.

B. Observations Excluded due to Family Change

Observations dated at t include data on growth rates between t and t+1. I exclude observations if either of the following is true in waves t or t+1.

If the family is living with another family I exclude these observations because of the difficulty of separating out food expenditures. If there is a change (from t-1 to t or from t to t-1) in either the head of the family or the wife, I exclude these observations for two reasons. First, I wanted to allow new families time to adjust. Second, when there is a major family change between surveys, it is not obvious to which family (old or new) the questions for the preceding calendar year refer.

After these exclusions, and the exclusion of waves 1, 2, 5, 6, and 15 (for reasons discussed below), there were 23,337 observations (based on multiple time series observations per family).

C. Extreme Outliers

I exclude an observation if the natural log of the ratio of consumption in t+1 to consumption in t is greater in absolute value than 1.1. This excludes observations where the level of food consumption rose or fell by more than a factor of 3 (= $e^{1.1}$) in a year (i.e., it eliminates observations where food consumption fell to less than a third of its prior year's value or rose to more than three times its prior year's value).

D. Missing Data

The question regarding amount spent on food at restaurants was not asked in wave 1, and neither food question was asked in wave 6. Neither total taxes for the head and wife nor the marginal tax rate were estimated for wave 2. Because I use growth rates, I need to eliminate waves 5 and 7. This leaves me with a maximum of ten observations per family.

Some data were missing for particular observations for the variables that I use. When this occurs (except as noted above), I record missing data for the constructed variable for the relevant observation. I exclude from the sample, for a given year, any observation with missing data on any of the variables included in the basic regression (including disposable income). I also exclude a year t observation if either a major or minor assignment was made for food consumption in t or t+1, or if a major assignment was made for asset or labor income of the head or wife in year t. After steps C and D, there are 17,915 total observations for the liquid asset split, and 16,601 total observations for the total asset split.

III. Splitting into Sub-Groups

A. Split Based on Two Months' Worth of Income in Liquid Assets (Table 2)

In waves 1-5, 8, and 13 the following questions were asked: 1) "Do you have any savings such as checking or savings accounts or government bonds?" and, if so, 2) "Would they amount to as much as two months' income or more?"

In order to replicate those questions for the other waves, I constructed a variable equal to the ratio of non-housing wealth in t to the average disposable income in t and t-1. If the data for this ratio is missing, I exclude the observation from both groups.

I put observations into Group I if either: the answer to either questions (1) or (2) above is "no," or the above ratio is less than 2/12 (corresponding to 2 months worth of income.) If neither of the above was true, I put the observation in Group II. This results in 11,902 observations in Group I and 6013 observations in Group II.

For tests (ii) and (iii) only Group I observations corresponding to families that also had observations in Group II could be used. This leaves 4152 observations for Group I.

B. Split Based on Total Wealth (Table 3)

For this split, I construct a variable equal to the ratio of total wealth (including housing) in t to the average of disposable income in t and t-1. If data on this variable is missing, I exclude the observation from both groups.

For this split, I put observations into Group I if the ratio is less than 2/12; otherwise the observation goes into Group II. This split puts 4744 observations in Group I and 11,857 in Group II. For tests (ii) and (iii), 2266 observations are included in Group I.

C. Most Stringent Split (Table 4)

For this split I put the two extremes into Groups I and II and I exclude the middle observations. If the ratio of non-housing wealth to income is missing, the observation is excluded from both groups.

An observation is placed in Group I if the family did not own a house in year t, and at least one of the following is true: a) the answer to question (1) (see III.A. above) is no (i.e., they say they have no current savings), or b) the ratio of non-housing wealth to income is equal to zero.

An observation is placed in Group II only if both of the following are true: a) they report that they have current savings, and b) the ratio of non-housing wealth to income is greater than .5.

This results in 3512 observations in Group I and 4258 observations in Group II. For tests (ii) and (iii), only 325 observations are available in Group I.

FOOTNOTES

[1]An exception is Miron (1984), who uses seasonally unadjusted quarterly data, and cannot reject the PIH.

[2]Hall and Mishkin directly examine the question of excess sensitivity in a certainty equivalence framework using seven years' worth of data from the same data set that I use (PSID). Shapiro was the first to use panel data to estimate the consumption Euler equation. He used three years' worth of data from the PSID and rejects the overidentifying restrictions using the entire panel. Runkle estimates the consumption Euler equation using four years worth of data collected for the Denver Income Maintenance experiment. He splits his sample by net worth and finds a violation of the Euler equation of the low wealth sample but not for the high wealth sample. He does not, however, examine the direction of the rejection or test the one sided inequality for the low net worth group.

[3]In some cases, such a ban on debt will never be binding. For example, if all future income is risky (i.e., it is possible for future labor income to be zero), and the marginal utility at zero consumption equals infinity, then a ban on riskless borrowing will never influence consumption. In this case, individuals always choose to carry positive non-human wealth to insure against the possibility of receiving zero income in a future period. In most cases, however, a borrowing constraint can be binding, either currently or in future periods. This would be the case if U'(0) were less than infinity, if there is a positive floor on the distribution of future income, or if the restriction is on certain forms of risky borrowing. An example of the latter would be a ban on borrowing with default allowed, contingent on income. These ideas are discussed in more detail in Zeldes (1984).

[4]In this paper, I discuss exogenous stock constraints or credit rationing - restrictions on the level of a particular asset or portfolio. Alternatively, one might wish to consider credit market imperfections that introduce a spread between the borrowing and lending rates, imply transactions costs, or lead to imperfectly liquid assets. See, for example, Rotemberg (1984), and Pissarides (1978).

[5]This is derived from an application of the Kuhn-Tucker first order conditions to the standard Bellman equation.

[6]When referring to the unconstrained stochastic PIH/LCH model, I am referring to the model presented in section IIA. I am not necessarily referring to the specific example of that model (the "certainty equivalence" version) in which optimal consumption is equal to permanent income, defined as the annuity value of the sum of financial wealth and the expected present discounted value of future income. For a detailed discussion of this distinction, see Zeldes (1984).

[7]For a further discussion, see Zeldes (1984).

[8]See Zeldes (1984).

[9]Hall and Mishkin's first test for excess sensitivity does not use a Keynesian alternative, but rather tests whether the response of consumption to transitory innovations in income is larger than would be predicted by the certainty equivalence version of the PIH. Their second test tests against a Keynesian alternative of sorts: they allow a fraction of consumption to be proportional to income. Note that I interpret the Keynesian consumption function as consumption equal to income; I can see no way that borrowing constraints could ever imply that consumption is proportional but not equal to income, or that a fraction of household consumption is equal to income while the rest is set according to the life cycle model.

[10]For a further discussion, see Zeldes (1984).

[11]For $\ln(1 + e_{it+1})$ to be uncorrelated with the instruments requires the stronger assumption that e_{it+1} be independent of the instruments and the variables with which I split the sample.

[12]Ignoring changes in tastes entirely (i.e., leaving out age and change in family composition from the estimated equation) would lead to an omitted variables bias if these variables are related to the instruments or the sample splitting variable.

[13]Consistency of the estimates only requires the weaker assumption that the variance of e_{it+1} is equal to the sum of a component that varies over i, a component that varies over t, and an i.i.d. component. For consistency of the standard errors, the variance of v_{it+1} must be constant across i and t. Zeldes (1984) discusses the relationship between this variance term and optimal consumption plans.

[14]Note that $\text{age}^2_{i,t+1} - \text{age}^2_{i,t} = 2 * \text{age}_{it} + 1$.

[15]The differences between the test I perform and Runkle's are: (1) I allow a given family to have observations in both groups and Runkle does not, (2) I split by net worth relative to average income instead of just net worth, and (3) I test for the significance of lagged income in the Euler equation rather than net worth.

[16]For a more detailed description of the test of overidentifying restrictions, see Hansen and Singleton (1983).

[17]Note that the log of $1 + \lambda_{it}$ is positive if and only if λ_{it} is positive.

[18]This can be formally shown as follows. $X_{it+1} \equiv \frac{1}{A}\,[\ln(1 + \lambda_{it}) + v_{it+1}]$. Let $\hat{X}_{it+1}$ be a consistent estimate of X_{it+1}, so $\hat{X}_{it+1} = X_{it+1} + \hat{\eta}_{it+1}$, where $\hat{\eta}_{it+1}$ is the estimation error and $\text{plim}(\hat{\eta}_{it+1}) = 0$. Let a bar over a variable denote the sample average of that variable. Then $\overline{X}_{it+1} = \frac{1}{A}\,[\overline{\ln(1 + \lambda_{it})} + \overline{v_{it+1}}]$. By the law of large numbers, $\text{plim}\ \overline{v_{it+1}} = 0$, so $\text{plim}\ \overline{X}_{it+1} = \frac{1}{A}\,[\overline{\ln(1 + \lambda_{it})}]$. $\overline{\hat{X}}_{it+1} = \overline{X_{it+1}} + \overline{\hat{\eta}}_{it+1}$. Since $\text{plim}\ \hat{\eta}_{it+1} = 0$,

plim $\bar{\hat{\eta}}_{it+1}$ also equals 0, so plim $\bar{\hat{X}}_{it+1} = \text{plim}(\overline{X_{it+1}}) + \text{plim}(\bar{\hat{\eta}}_{it+1}) = \frac{1}{A}\left[\overline{\ln(1 + \lambda_{it})}\right]$.

[19]The data utilized in this paper were made available by the Inter-University Consortium for Political and Social Research. The data were originally collected by the Survey Research Center at the University of Michigan. Neither the original source or collectors of the data nor the Consortium bear any responsibility for the analyses or interpretations presented here. I thank Debbie Laren of the University of Michigan for her help with some of the trickier data problems.

[20]An additional problem with food consumption as I have constructed it is that the value of labor for preparation is included in one component (food out) but excluded from another (food at home).

[21]Hayashi (1984), however, finds evidence of durability of food consumption.

[22]For a description of the measurement error problems with this data see Shapiro (1982).

[23]A problem with this measure of housing equity is that it does not capture any differences between the outstanding mortgage principal and the market value of the mortgage that may have arisen due to changes in market interest rates.

[24]Even though these observations may provide interesting information on consumers likely to be subject to borrowing constraints, they were excluded in order to arrive at a representative sample.

[25]The panel used has a large cross section dimension (N), but a relatively small time series dimension (T). In order for the instrumental variables procedure to yield consistent estimates, it must be the case that the plim as N goes to infinity of $\frac{W_t'\varepsilon_t}{N}$ equals zero, where W_t is the matrix of

predetermined variables. If the aggregate expectations errors were included in ε_t, then $\operatorname{plim}_{N\to\infty} \frac{W_t'\varepsilon_t}{N}$ would not equal zero. For a further discussion of this point, see Chamberlain (1984).

[26]Those families with high rates of time preference will, all else equal, have a lower growth rate of consumption. They will also, however, tend to accumulate less wealth and therefore tend to be associated with observations that fall into Group I. Measured excess growth in consumption for group I would therefore tend to be lower than the true amount and failure to account for differences in rates of time preference would bias the one sided inequality test <u>against</u> the hypothesis that individuals are liquidity constrained. See the "alternative specifications" part of Section VI for a brief description of the results when all families were restricted to have the same rate of time preference.

[27]Define $\tilde{X}$ as the matrix of right hand side variables in the Euler equation. Define $\tilde{\beta}$ as the corresponding two stage least squares parameter estimates for the unconstrained group, and $\hat{V}(\tilde{\beta})$ as the estimated variances of these parameters. Define Z as the matrix of right hand side variables for tests (ii) and (iii) (i.e., a constant for test (ii) and a constant and y for test (iii)). Define $\hat{\gamma}$ as the corresponding estimated parameters, and $\hat{V}(\hat{\gamma})$ as the estimated variance matrix. Then:

$$\hat{V}(\hat{\gamma}) = (Z'Z)^{-1}\hat{\sigma}^2 + (Z'Z)^{-1}Z'X\hat{V}(\tilde{\beta})(X'Z)(Z'Z)^{-1}$$

is the consistent estimate of the variance matrix (reported in Tables 2-4). The first term is what would be printed by a standard regression package. The second term captures the fact that the estimated β for the unconstrained group is a consistent but noisy estimate of the true β. I thank Whitney Newey and Andy Lo for their help in figuring out the appropriate formulas for these standard errors.

[28]Estimates of the coefficient of relative risk aversion include: Friend and Blume (1975) (2), Mankiw (1981) (4) and Hansen and Singleton (1983) (1).

REFERENCES

Ball, Laurence, "Intertemporal Substitution and Constraints on Labor Supply: Evidence from Panel Data," manuscript, MIT, 1985.

Bernanke, Ben S., "Permanent Income, Liquidity, and Expenditures on Automobiles: Evidence from Panel Data," Quarterly Journal of Economics, August 1984, 589-614.

Chamberlain, Gary, "Panel Data" in: The Handbood of Econometrics: Volume 2, M. Intrilligator and Z. Grilliches ed., North-Holland, 1984.

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

Friend, Irwin and Marshall Blume, "The Demand for Risky Assets," American Economic Review, December 1975, 900-923.

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

Hall, Robert E. and Frederic S. Mishkin, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," Econometrica, 50, March 1982.

Hansen, Lars Peter and Kenneth J. Singleton, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," Journal of Political Economy, April 1983, 249-266.

Hayashi, Fumio, "Permanent Income Hypothesis and Consumption Durability: Analysis Based on Japanese Panel Data," NBER Working Paper, No. 1305, March 1984.

Hayashi, Fumio, "The Effect of Liquidity Constraints on Consumption: A Cross Section Analysis," Quarterly Journal of Economics, February 1985.

Institute For Social Research, A Panel Study of Income Dynamics, Procedures and Tape Codes, Waves I-XV, Survey Research Center, University of Michigan, 1972-1984.

Institute For Social Research, User Guide to the Panel Study of Income Dynamics, Survey Research Center, University of Michigan, 1984.

King, Mervyn, "The Economics of Saving," NBER Working Paper No. 1247, December 1983.

Mankiw, N. G., "The Permanent Income Hypothesis and the Real Interest Rate," Economics Letters, 7, 1981.

Mankiw, N. Gregory, Julio J. Rotemberg and Lawrence Summers, "Intertemporal Substitution in Macroeconomics," Quarterly Journal of Economics, 1984.

Miron, Jeffrey A., "Seasonality and the Life Cycle-Permanent Income Model of Consumption," manuscript, University of Michigan, November 1984.

Pissarides, Christopher A., "Liquidity Considerations in the Theory of Consumption," Quarterly Journal of Economics, 1978.

Rotemberg, Julio J., "Consumption and Liquidity Constraints," manuscript, MIT, February 1984.

Runkle, David E., "Liquidity Constraints and the Permanent Income Hypothesis: Evidence from Panel Data," manuscript, MIT, May 1983.

Shapiro, Matthew D., "A Note on Tests of the Permanent Income Hypothesis in Panel Data," manuscript, MIT, May 1982.

Shapiro, Matthew D., "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data," Economics Letters, 14, 1984.

Zeldes, Stephen, "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," manuscript, MIT, July 1984.