

Consumption and Liquidity Constraints: An Empirical Investigation

Stephen P. Zeldes

University of Pennsylvania and National Bureau of Economic Research

Several recent studies have suggested that empirical rejections of the permanent income/life cycle model might be due to the existence of liquidity constraints. This paper tests the permanent income hypothesis against the alternative hypothesis that consumers optimize subject to a well-specified sequence of borrowing constraints. Implications for consumption in the presence of borrowing constraints are derived and then tested using time-series/cross-section data on families from the Panel Study of Income Dynamics. The results generally support the hypothesis that an inability to borrow against future labor income affects the consumption of a significant portion of the population.

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 (1985), among others. Some of these authors suggest

I would like to thank Andrew Abel, Joseph Altonji, Laurence Ball, Olivier Blanchard, Susan Collins, Robert Cumby, Wayne Ferson, Stanley Fischer, Robert Hall, James Heckman, Terry Marsh, Whitney Newey, Stephen O'Connell, Julio Rotemberg, Larry Summers, members of the Consumption Group at the 1984 NBER Summer Institute, members of the Macro Lunch Group at the University of Pennsylvania, and two anonymous referees for helpful comments. I am also grateful to Ananth Krishnamurthy, Caroline Naggjar, Ahmed Taha, and Andrea Wicks for research assistance, and to the Wharton Junior Faculty Research Fund for financial support. This paper is a revised version of chap. 3 of my MIT Ph.D. dissertation.

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. Because Keynesian consumers neither borrow nor save to smooth consumption while individuals subject to borrowing constraints are free to save, the presence of borrowing constraints will not in general lead to Keynesian behavior.

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 Panel Study of Income Dynamics (PSID), a representative panel of U.S. families, to test these implications. The 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).¹

I begin by presenting a model without borrowing constraints and examining the first-order condition that relates consumption in adjacent periods. Next, a set of exogenous quantity constraints on net indebtedness in each period are imposed. I set up two formal tests, each of which depends on the observation that the Euler equation derived from the unconstrained model will be violated when liquidity constraints exist. My tests involve splitting the observations into two groups on the basis of ratios of financial assets to income. Those observations with low assets (group 1) are most likely to be against a binding liquidity constraint. If liquidity constraints are an important source of departure from the PIH/LCH, then the following should be true. First, the Euler equation should be satisfied for group 2 but violated for group 1. These violations could take the form of implausible parameter estimates, a rejection of overidentifying restrictions,

¹ Hall and Mishkin directly examine the question of excess sensitivity of consumption to innovations of income in a certainty equivalent framework using seven years' worth of data from the PSID. Shapiro estimates the consumption Euler equation. He uses three years' worth of data on all the individuals in the PSID panel and rejects the overidentifying restrictions. 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 for the low-wealth sample but not for the high-wealth sample.

or both. Second, there should be a one-sided inequality in the Euler equation for group 1 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 1 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. This follows up on work on estimating Lagrange multipliers that has been done in the context of regulatory constraints on firms (e.g., Cowing 1978).

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 effect, I estimate the total derivative of the Lagrange multiplier with respect to current income and test whether it is negative. This third test is suggestive, but is not a formal test, because the sign of the total derivative need not be the same as the sign of the partial derivative.

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 10 annual observations per family on food consumption and other variables from the PSID, a large panel of U.S. families.

The results are generally, but not completely, 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 1 observations but not for group 2 observations. For group 1, the average Lagrange multiplier is positive, implying that there is the appropriate one-sided inequality in the Euler equation for constrained observations, but the coefficient is not statistically significant at the 5 percent level. The point estimate indicates that borrowing constraints caused annual food consumption growth for group 1 to be 1.7 percentage points higher than it would have

been in the absence of constraints (with rates of return held constant). In addition, the total correlation between current income and the Lagrange multiplier was in fact negative, although statistically insignificant. A second split based on liquid assets excluded some observations from both groups and yielded similar results, except that the average Lagrange multiplier was statistically significant (and positive). Finally, 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. This section also points out that individuals can be globally constrained even if the Euler equation that relates optimal consumption between the current and subsequent period is satisfied, and it describes the differences between Keynesian and borrowing-constrained behavior. In Section III, I describe the tests that I perform to test for the importance of borrowing constraints. In Section IV, 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 Section V, I examine the results, and concluding remarks are presented in Section VI.

II. The Basic Model with and without Borrowing Constraints

I begin by presenting the standard model, and the corresponding set of Euler equations, when no borrowing constraints are present. 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 auxiliary assumptions. In the second subsection, I present 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 $\{w_{it}^j\}$ in order to

$$\max E_t \sum_{k=0}^{T-t} \left(\frac{1}{1 + \delta_i} \right)^k U(C_{i,t+k}; \Theta_{i,t+k}) \quad (1a)$$

$$\text{subject to } A_{i,t+k} = (A_{i,t+k-1}) \times \left[\sum_{j=1}^M w_{i,t+k-1}^j (1 + r_{i,t+k-1}^j) \right] \quad (1b)$$

$$+ Y_{i,t+k} - C_{i,t+k}, \quad k = 0, \dots, T - t,$$

$$C_{i,t+k} \geq 0, \quad k = 0, \dots, T - t, \quad (1c)$$

$$A_{iT} \geq 0, \quad (1d)$$

where $U(\)$ = the one-period utility function, C_{it} = real consumption of family i in period t , Θ_{it} = family i 's tastes in period t , δ_i = the rate of time preference of family i , E_t = the expectation operator, conditional on information available at time t , T = the end of the family's horizon, A_{it} = real end-of-period financial (nonhuman) wealth of family i in period t (after receiving income and consuming), r_{it}^j = the ex post real after-tax return on the j th asset between periods t and $t + 1$, w_{it}^j = the share of end-of-period- t wealth held in asset j , M = the number of available assets, and Y_{it} = real disposable labor income of family i in period t .²

As discussed in Zeldes (1986), 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, that are necessary for an optimum. An individual should be unable to increase expected lifetime utility by consuming one fewer unit today, increasing her 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 her holdings of asset j below the current amount, then she 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:

$$U'(C_{it}; \Theta_{it}) = E_t \frac{U'(C_{i,t+1}; \Theta_{i,t+1})(1 + r_{it}^j)}{1 + \delta_i}, \quad (2)$$

$$i = 1, \dots, N, t = 1, \dots, T - 1, j = 1, \dots, M,$$

² It is assumed that individuals have access to a market in which they can borrow and lend at the riskless rate of interest. Other contingent claims markets may, but need not, exist. These markets may include, e.g., those for a set of securities sufficient to implement an Arrow-Debreu equilibrium.

where N is the number of families and $U'(\cdot; \cdot)$ denotes the partial derivative of U with respect to C .³ If expectations are rational, this leads to

$$\frac{U'(C_{i,t+1}; \Theta_{i,t+1})(1 + r_{it}^j)}{U'(C_{it}; \Theta_{it})(1 + \delta_i)} = 1 + e'_{i,t+1}{}^j, \quad (3)$$

where $e'_{i,t+1}{}^j$ is an expectational error uncorrelated with information known at time t . This Euler equation should hold with respect to any asset available to the consumer, including the riskless asset (with rate of return denoted simply r_{it}). The remainder of this paper focuses on the Euler equation that holds with respect to riskless borrowing and lending.

B. *The Model with Borrowing Constraints*

A number of recent theoretical papers have shown that liquidity constraints can have important effects on individual consumption behavior and on the behavior of aggregate consumption, output, and asset returns. For example, Scheinkman and Weiss (1986) examine a general equilibrium model with individual-specific uncertainty. They show that imposing the exogenous restriction that individuals cannot borrow against the future proceeds of their labor can induce price and output fluctuations that mimic actual business cycles but that are absent under a perfect markets assumption. In this paper, I empirically test for the presence of these borrowing constraints.

A variety of forms of liquidity constraints have been examined in the literature, each of which involves some price and/or quantity restrictions on the holding of assets. The type of constraint considered here is a simple quantity constraint: a floor on the total end-of-period net stock of traded assets. I do not attempt here to derive this constraint endogenously. However, this kind of borrowing constraint can arise in models in which individuals have private information about their future labor income.⁴ This type of exogenous constraint is similar to the constraints imposed by Bewley (1977) and Scheinkman and Weiss (1986), among others. The following set of restrictions on equations (1a)–(1d) constitutes what I refer to as a borrowing constraint

³ This Euler equation ignores the nonnegativity constraint on *consumption* (1c). As long as $U'(0) = \infty$ (as is assumed below), then this constraint will never be binding, and the Euler equation is correct as it is.

⁴ The resulting adverse selection and/or moral hazard problems can lead to credit rationing, a market failure that would not arise in a world of perfect information and enforcement (see, e.g., Stiglitz and Weiss 1981). Modeling this explicitly is beyond the scope of this paper, although by not doing so I run some risk of missing other implications of the model that might influence the empirical tests performed here.

throughout this paper:

$$A_{i,t+k} \geq 0, \quad k = 0, \dots, T - t - 1. \quad (1e)$$

While I think of these “borrowing” constraints as being a subset of possible “liquidity” constraints, I use the terms interchangeably in this paper.⁵ This set of constraints prohibits individuals from consuming today the proceeds from supplying labor in the future.⁶ While most consumers are able to borrow to purchase assets (e.g., mortgages or stock on margin), it seems a reasonable working (alternative) hypothesis that consumers cannot borrow, *on net*, against nontraded assets such as future labor income, in other words, that debt cannot exceed the total current value of traded assets.⁷

The empirical work will test two competing hypotheses. Under the null hypothesis, agents can borrow and lend at the same rate; under the alternative hypothesis, borrowing against future labor income (and other nontraded assets) is not allowed; that is, net wealth excluding nontraded assets is restricted to be nonnegative.

⁵ The type of constraint considered is clearly not the only form a liquidity constraint could take. A slightly more general form of the constraint would be $A_{i,t+k} \geq -B$ ($k = 0, \dots, T - t - 1$), where B is the limit on net indebtedness. Also, rather than considering exogenous quantity constraints on the stocks of particular assets, one might wish to consider quantity restrictions that depend on credit histories or credit market imperfections that introduce a spread between the borrowing and lending rates, imply transactions costs, or lead to imperfectly liquid assets (see, e.g., Pissarides 1978; Rotemberg 1984). Finally, n. 7 below describes another possible form that the constraints could take.

⁶ 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 equal zero in each period) 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 nonhuman 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 or if there is a positive floor on the distribution of future income. For a further discussion of this, see Zeldes (1987).

⁷ In addition to placing restrictions on borrowing risklessly (i.e., going short in the riskless asset), the constraints in the text also restrict short sales of risky assets. They do not rule out all short sales of risky assets, but the fact that the constraint must be satisfied next period with probability one forces individuals to rule out portfolio choices today that put at risk an amount greater than next period's labor income (i.e., there can be no asset return realizations such that the individual would end up owing more on short positions than the sum of the value of his long positions and labor income). One might want to consider an alternative set of constraints that is stronger than those in the text:

$$A_{i,t+k} \times \left[\sum_{j=1}^M w'_{j,t+k} (1 + r'_{j,t+k}) \right] \geq 0, \quad k = 0, \dots, T - t - 1.$$

In this case, the total *payoff* from the portfolio of traded assets has to be positive with probability one. This would prohibit individuals from even putting next period's income at risk. The difference between these constraints relates to restrictions on portfolio composition; if only a riskless asset is available, these constraints will be identical.

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. The constraint will be binding for those who chose not to build up their wealth in earlier periods or those who received exceptionally bad draws of income or portfolio returns.⁸

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

$$U'(C_{it}; \Theta_{it}) = E_t \left[\frac{U'(C_{i,t+1}; \Theta_{i,t+1})(1 + r_{it})}{1 + \delta_i} \right] + \lambda''_{it}, \tag{4}$$

where λ''_{it} is the Lagrange multiplier (known at time t) associated with constraint (1e) for time t .⁹ The term λ''_{it} is equal to the increase in

⁸ Bewley (1977) shows that in an infinite-horizon economy with uncertainty about income or tastes and with $\delta = r = 0$ (in the limit), individuals will tend to accumulate enough wealth over time so that they eventually act as if there is no uncertainty or borrowing constraint. Given many consumers with any finite initial wealth, however, there will be periods in which some individuals hit the borrowing constraint (i.e., have zero wealth) while others do not.

⁹ This is derived from an application of the Kuhn-Tucker first-order conditions to the standard Bellman equation, as follows. The i subscripts are dropped for simplicity. The following derivation assumes that only the riskless asset is available to the consumer. However, it can be shown that eq. (4) holds with respect to any asset, when many assets are available to the consumer. Also, it will be useful to define Z_t as *beginning-of-period* wealth (after Y_t received) (as opposed to A_t , which is end-of-period wealth), so $A_t = Z_t - C_t$ and $Z_{t+1} = A_t(1 + r_t) + Y_{t+1} = (Z_t - C_t)(1 + r_t) + Y_{t+1}$. In each period $t = 1, \dots, T$, the consumer solves the following problem:

$$V_t(Z_t) = \max_{C_t} \left[U(C_t) + \left(\frac{1}{1 + \delta} \right) E_t V_{t+1}(Z_{t+1}) + \lambda''_t(Z_t - C_t) \right]$$

(N1)

subject to $Z_{t+1} = (Z_t - C_t)(1 + r_t) + Y_{t+1}$ and $V_{T+1} \equiv 0$.

The first-order conditions for this problem are

$$U'(C_t) - \left(\frac{1}{1 + \delta} \right) E_t [V'_{t+1}(Z_{t+1}) \cdot (1 + r_t)] - \lambda''_t = 0,$$

$$Z_t - C_t \geq 0 \quad \text{with equality if } \lambda''_t > 0. \tag{N2}$$

Let $C_t^*(Z_t)$ be the solution to (N2). Then (N1) becomes

$$V_t(Z_t) = U(C_t^*(Z_t)) + \left(\frac{1}{1 + \delta} \right) E_t V_{t+1} \{ [Z_t - C_t^*(Z_t)](1 + r_t) + Y_{t+1} \}$$

$$+ \lambda''_t [Z_t - C_t^*(Z_t)]. \tag{N3}$$

Differentiating (N3) with respect to Z_t gives

$$V'_t(Z_t) = U'(C_t^*(Z_t)) \cdot C_t^{*'}(Z_t) + \left(\frac{1}{1 + \delta} \right) E_t \{ V'_{t+1}(Z_{t+1})$$

$$\cdot [1 - C_t^{*'}(Z_t)](1 + r_t) \} + \lambda''_t [1 - C_t^{*'}(Z_t)]. \tag{N4}$$

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 *greater than* or equal to the marginal utility from waiting until tomorrow to consume the extra amount. Also, if λ''_{it} is greater than zero, so that the current constraint is binding, then the end-of-period financial assets of the individual must be equal to zero. (If assets were positive, the individual could always transfer resources at the margin by consuming part of these assets.)

It will be convenient for future purposes to normalize λ''_{it} by a positive term that is nonstochastic as of time t :

$$\lambda'_{it} \equiv \frac{\lambda''_{it}}{E_t \left[\left(\frac{1+r_{it}}{1+\delta_i} \right) U'(C_{i,t+1}; \Theta_{i,t+1}) \right]}$$

As with λ''_{it} , λ'_{it} will be positive when the constraint is binding and zero when it is not binding. Rewriting (4), we get

$$E_t \left[\frac{U'(C_{i,t+1}; \Theta_{i,t+1})(1+r_{it})}{U'(C_{it}; \Theta_{it})(1+\delta_i)} \right] (1+\lambda'_{it}) = 1, \quad (5)$$

and, as before, rational expectations implies

$$\frac{U'(C_{i,t+1}; \Theta_{i,t+1})(1+r_{it})}{U'(C_{it}; \Theta_{it})(1+\delta_i)} (1+\lambda'_{it}) = 1 + e'_{it+1}, \quad (6)$$

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.

Plugging (N2) into (N4) yields

$$V'_i(Z_t) = [U'(C_t^*(Z_t)) \cdot C_t^*(Z_t)] + [1 - C_t^*(Z_t)][U'(C_t^*(Z_t))] = U'(C_t^*(Z_t)), \quad (N5)$$

which is the envelope condition. Note that at the optimum, the extra lifetime utility from being given an extra dollar in assets equals the extra utility from consuming it immediately, whether or not the current constraint is binding. Pushing (N5) forward one period and plugging into (N2) gives

$$U'(C_t) = E_t \left[\left(\frac{1}{1+\delta} \right) (1+r_t) U'(C_{t+1}) \right] + \lambda''_t, \quad (N6)$$

which is the equation given in the text.

C. *Euler Equation Violations, Global Constraints, and "Keynesian" Consumption*

Before I proceed to examine the Euler equation tests, two important observations need to be made. First, if the Euler equation between this period and next period 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.¹⁰ In any multiperiod model, equation (1e) constitutes a *set* of constraints, one for each time period. The Euler equation between t and $t + 1$ will be violated if, given all future constraints, the current constraint is binding, that is, if removing (1e) for $k = 0$ (but continuing to impose it for $k > 0$) would alter current consumption. However, even if the current constraint is not binding, so that the Euler equation between t and $t + 1$ is satisfied, the presence of constraints that will bind in the *future* with some positive probability will lower the current consumption of any risk-averse individual (Zeldes 1987). 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., $\lambda'_{it} = 0$).¹¹

The second point of this subsection is that borrowing constraints will not in general imply Keynesian behavior. The standard Keynesian consumption function is written $C_t = C + cY_t$. The only form of Keynesian behavior that could be explained by currently binding constraints would be the specific example $C_t = Y_t$, which would require 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 equaled income during the period.) Unconstrained PIH/LCH consumers can smooth out fluctuations in income by borrowing and saving. Keynesian behavior $C_t = Y_t$ involves neither borrowing nor saving to smooth consumption. Therefore, in order for constraints to generate Keynesian behavior, either there must be *both* borrowing and lending constraints (which might arise, e.g., in devel-

¹⁰ When referring to the unconstrained stochastic PIH/LCH model, I am referring to the model presented in subsection A. 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 discussion of this distinction, see Zeldes (1986).

¹¹ 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 assets from this period to the next. Given future constraints (i.e., for $k > 0$), the consumer is optimally allocating resources between today (t) and tomorrow ($t + 1$), and the current constraint (i.e., for $k = 0$) is not binding. The Euler equation that is commonly estimated, describing the relationship between consumption in t and $t + 1$, will be satisfied.

oping countries with poorly functioning capital markets) or, if there is only a borrowing constraint, individuals must be choosing never to save; that is, the borrowing constraints must be binding period after period. But individuals generally receive both good and bad draws of income, and when individuals receive high draws of income today relative to tomorrow, they will in general choose to save to avoid declines in consumption. 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 et al. 1985). 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 (1985a).¹² The discussion in this subsection 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.

III. Description of Euler Equation Tests

All the tests in the rest of the paper are based on the implication that a currently binding borrowing constraint will lead to a violation of the (unconstrained) Euler equation. 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 express a closed-form decision rule for consumption. This is important because a closed-form solution to the unconstrained problem has been derived only under very restrictive assumptions, and no one has derived a closed-form solution to the problem with constraints imposed.¹³ Hereinafter, when I refer to a

¹²Hall 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.

¹³See Zeldes (1986, 1987) for a further discussion.

liquidity constraint being binding this period, I mean that the Lagrange multiplier associated with transferring resources between tomorrow and today is positive, that is, 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 subsection *A*, and subsection *B* describes the econometric tests.

A. Assumptions about Preferences and Identification

In order to make equation (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 (Θ_{it}). I assume that the utility function is of the constant relative risk aversion form

$$U(C_{it}; \Theta_{it}) = \frac{C_{it}^{1-\alpha}}{1-\alpha} \times \exp(\Theta_{it}), \quad (7)$$

where α 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 . Substituting into equation (6), we get

$$\frac{C_{it+1}^{-\alpha} \times \exp(\Theta_{it+1} - \Theta_{it})(1 + r_{it})(1 + \lambda'_{it})}{C_{it}^{-\alpha}(1 + \delta_i)} = 1 + e'_{it+1}, \quad (8)$$

where e'_{it+1} is an expectational error about the left-hand side of equation (8) that has mean zero and is uncorrelated with any information available at time t .

Unlike some previous authors, I allow the family utility function (and thus consumption) to be influenced by tastes that may differ across families and shift across time.¹⁴ I allow for both observable (to the econometrician) and unobservable tastes. The observable factors, which vary across families and across time, are a measure of family size (AFN_{it}), age, and age squared. The unobservable part of tastes includes a fixed family component (ω_i), an aggregate component that is constant across families but varies across time (η_t), and a remaining component (u_{it}) that is orthogonal to the first two. I also allow each family to have a different rate of time preference, which is equivalent to including a fixed family component in the *change* in tastes. This gives

$$\Theta_{it} = b_0 \text{age}_{it} + b'_1 \text{age}_{it}^2 + b_2 \ln(\text{AFN}_{it}) + \omega_i + \eta_t + u_{it}. \quad (9)$$

¹⁴For discussions of taste shifters and identification in panel data, see, e.g., Heckman and Singer (1984), Ball (1986), and Ham (1986).

Substituting this into equation (8), taking logs of both sides, and rearranging yields

$$\begin{aligned}
 GC_{i,t+1} = & \frac{1}{\alpha} [b_0 - \ln(1 + \delta_i) + (\eta_{t+1} - \eta_t) + \ln(1 + r_{it}) \\
 & + b'_1(\text{age}_{i,t+1}^2 - \text{age}_{it}^2) + b_2\text{GAFN}_{i,t+1} \\
 & + (u_{i,t+1} - u_{it}) - \ln(1 + e'_{i,t+1}) + \ln(1 + \lambda_{it})],
 \end{aligned} \tag{10}$$

where $GC_{i,t+1} = \ln(C_{i,t+1}/C_{it})$ and $\text{GAFN}_{i,t+1} = \ln(\text{AFN}_{i,t+1}/\text{AFN}_{it})$.

Equation (10) is in log first differences, and thus the family-specific term ω_i drops out of the equation. The estimation procedure will account for the presence of the family-specific rate of time preference by including a fixed person effect and will account for the presence of the aggregate taste change by including a fixed time effect. The taste changes $(u_{it+1} - u_{it})$ are assumed to be stationary, with unconditional expectation equal to zero.

It is likely that the expectational error e'_{it+1} , while having mean zero, will be correlated across families because of macroeconomic shocks.¹⁵ I assume that the expectational error term $(1 + e'_{it+1})$ can be decomposed into the product of two independent components: a common aggregate component $(1 + e''_{t+1})$ and an idiosyncratic component $(1 + e_{it+1})$, where e''_{t+1} and e_{it+1} each have mean zero.

If e_{it+1} has mean zero, then $\ln(1 + e_{it+1})$ does not. Taking a second-order Taylor expansion gives $\ln(1 + e_{it+1}) \approx e_{it+1} - 1/2e_{it+1}^2$. This implies that, conditional on any information set Ω , $E[-\ln(1 + e_{it+1})|\Omega] \approx 1/2 \text{var}(e_{it+1}|\Omega)$, an approximation that holds exactly if e_{it+1} is lognormally distributed. Thus the *mean* of the error term in the Euler equation will be related to the *variance* of the expectation error. The conditional variance of the expectation error can include a time- and a person-specific component but is assumed to be unrelated to time t variables described below.

Incorporating all of the above gives the equation that is estimated:

$$\begin{aligned}
 GC_{i,t+1} = & k^1 + k_i^2 + k_i^3 \\
 & + \frac{1}{\alpha} [\ln(1 + r_{it}) + b_1\text{age}_{it} + b_2\text{GAFN}_{i,t+1}] \tag{11} \\
 & + v_{i,t+1} + \ln(1 + \lambda_{it}),
 \end{aligned}$$

where $v_{i,t+1} = (1/\alpha)[(u_{i,t+1} - u_{it}) - \ln(1 + e_{i,t+1}) - 1/2\sigma_{e_{i,t+1}}^2]$ has mean zero, and k^1 , k_i^2 , and k_i^3 are the constant, fixed family effect, and

¹⁵For example, if aggregate income were unexpectedly high in a given period, all individuals would tend to have higher than expected consumption this period and therefore unexpectedly higher consumption growth between last period and this period.

fixed time effect, respectively.¹⁶ For simplicity, λ'_{it} has been renormalized with a sign-preserving transformation: $1 + \lambda_{it} \equiv (1 + \lambda'_{it})^{1/\alpha}$. For use later on, I define $x_{it+1} \equiv [v_{it+1} + \ln(1 + \lambda_{it})]$; x_{it+1} is the growth in consumption above the amount that, all else held equal, 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 $t + 1$.

I assume that family composition and age of head at $t + 1$ are known to the family as of time t . However, the ex post after-tax return, $\ln(1 + r_{it})$, will in general not be known at time t and may be correlated with the expectation error on the growth of consumption because of correlations between time $t + 1$ consumption, income, and the marginal tax rate. I will therefore estimate the Euler equation with an instrumental variables procedure. The marginal tax rate and the log of income at time t are used directly as instruments, while the wealth to income ratio is used to split the sample.

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, which includes (approximately) an expectations error, the square of an expectations error, and the unobservable change in tastes $u_{it+1} - u_{it}$. By the assumption of rational expectations, any variable known at time t will be orthogonal to the expectations error.¹⁷ The key identifying assumptions made in this paper are that the conditional mean of the unobservable change in tastes and the conditional variance of the expectations error (after fixed time and family effects are accounted for) are unrelated to the time t marginal tax rate, log of real disposable income, and wealth to income ratio.

To review, I have accounted for heterogeneity in tastes by estimating an equation in the first difference of consumption (thus eliminating any family-specific effects on the level of consumption), including observable variables that should influence the change in tastes, including a fixed time effect to account for aggregate taste changes over

¹⁶Note that $\text{age}_{i,t+1}^2 - \text{age}_{i,t}^2 = 2 \times \text{age}_{i,t} + 1$.

¹⁷Since the panel has a large cross-section dimension (N) but a relatively small time-series dimension (T), the required orthogonality condition is that the plim as N goes to infinity of $(\mathbf{W}'e)/N$ equals zero (where \mathbf{W} is the matrix of instruments). As Chamberlain (1984) points out, this would not hold if e_{it} contains aggregate expectations errors. Under the assumption made here that the aggregate shock hits all households equally, the orthogonality condition holds once time dummies are included in the regression. If the aggregate shock hits different groups of households in a way that is related to permanent characteristics of the household, then the orthogonality condition could fail even with fixed time effects included (see Hayashi 1985*b*, p. 1094, n. 13). However, the inclusion of fixed family effects in this case will again make the orthogonality condition valid.

time, and including a fixed family effect to account for a family-specific (non-time-varying) *change* in tastes. Despite this, there may be circumstances under which the identifying assumptions above will be violated. First, the unobservable changes in tastes for consumption could be related to the level of income in a period because of exogenous shifts in preferences between consumption and leisure.¹⁸ Second, the conditional variance of the forecast error could be a function of wealth or disposable income. For example, when household assets are especially low, uncertainty about the growth rate of consumption could be higher. The conditional mean of the residual will then also be a function of assets (because of the nonlinearity in the residual term). In each of these cases the estimation scheme presented below will be inconsistent.

As described below, the empirical tests use data on food consumption rather than on total consumption. Also, throughout the paper I abstract from modeling the decision rule for labor supply or the purchase of durable goods. Therefore, in order for the Euler equation (11) to be valid for food consumption, I need to assume that the utility function is additively separable in food, leisure, and other consumption goods. If the utility function is not separable in leisure and (food) consumption, then leisure in periods t and $t + 1$ will enter the Euler equation.¹⁹ If the utility function is nonseparable in (food) consumption and the service flows from durables, then the service flows from durables in t and $t + 1$ would enter the Euler equation. Standard assumptions about functional forms would imply that the Euler equation in that case would include an extra additive term equal to a coefficient times the growth rate of the stock of durables.²⁰ Unfortunately, this model cannot be tested with the PSID data set because of a lack of data on consumer durables. For an (aggregate time-series) analysis that includes this type of nonseparability, see Bernanke (1985).

¹⁸ Recall, however, that I have accounted for a component of tastes for leisure vs. consumption that varies across households but not time.

¹⁹ For analyses of this case, see Heckman (1974), Mankiw et al. (1985), and Altonji (1986).

²⁰ Let D_t denote the stock of durables in period t , let the flow of services from durables in period t be proportional to the stock in period t ($= m \cdot D_t$), and let the utility function be

$$U(C_{it}, mD_{it}; \Theta_{it}) = \frac{(C_{it})^{1-\alpha}}{1-\alpha} \frac{(mD_{it})^{1-\beta}}{1-\beta} \cdot \exp(\Theta_{it}).$$

Substituting into eq. (6), taking logs, and rearranging yields eq. (11) with a different constant term and with $[(1-\beta)/\alpha]GD_{i,t+1}$ added to the right-hand side.

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

Splitting the Sample

The approach of splitting a sample of households into two sets of observations—those likely to be constrained and those not likely to be constrained—has been used in a number of studies, including an early study by Juster and Shay (1964), as well as recent papers by Runkle (1983), Bernanke (1984), and Hayashi (1985*a*).²¹ I follow this approach, dividing the sample on a priori grounds on the basis of the theory presented above. For one group of observations (group 1), the current constraint is binding ($\lambda_{it} > 0$), and for the other (group 2), it is not ($\lambda_{it} = 0$).²² The sample is split by observation, so that a given family can sometimes fall in group 1 and other times fall in group 2. If the form of the borrowing constraint is that nonhuman wealth must be nonnegative, then λ_{it} will be positive only if end-of-period nonhuman wealth is equal to zero. It therefore would be appropriate to split the sample into observations for which end-of-period nonhuman 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 1) not only observations in which wealth equals zero but also observations in which the wealth to (average past) income ratio is small but nonzero. Each of the tests performed relies on the consistency of the group 2 parameter estimates under both the null and the alternative hypotheses; it is therefore crucial that group 2 contain only observations for which the current constraint is not binding. However, the tests are robust to the inclusion of some unconstrained observations in group 1 (although this will reduce the power of the tests and the size of the estimate of the average Lagrange multiplier under the alternative hypothesis).

²¹ My paper is similar in a number of ways to that of Juster and Shay. They define consumers as “rationed” if they would like to borrow more at the current finance rate than they are able to, split the observations of families into those with low and high liquid assets, and test whether individuals were constrained in their borrowing against nonfinancial assets (including equity in automobiles). My approach differs from theirs in that they do not examine the implications for consumption, but rather directly examine the notional demand for loans by using survey data in which households were asked to preference-rank hypothetical (auto) loan packages. Their findings support the prediction that rationed families should be willing to pay a higher interest rate in order to extend the maturity of their loans (while unrationed families should not) and that the loan demand of rationed families should be less responsive to changes in interest rates than that of unrationed families.

²² Recall that even if the observation falls into group 2, neither the current level of consumption nor the sensitivity of consumption to income need be equal to what it would be in the absence of all future constraints.

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.

Equation (11) derived above includes λ_{it} , a measure of how severely the borrowing constraint is affecting consumption growth. The approach of estimating a Lagrange multiplier associated with an external constraint has been used previously in other contexts. For example, Cowing (1978) estimates a static Averch-Johnson model of a profit-maximizing firm subject to a regulatory constraint on the rate of return on capital. He presents closed-form decision rules for a firm that is subject to the constraint, calculates an estimate of the unobservable Lagrange multiplier for each observation, and tests whether it is zero. Unfortunately, I cannot derive a closed-form expression for λ_{it} . It will in general be a nonlinear 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. I therefore use the technique of splitting the sample in order to estimate λ and derive testable implications for consumption with borrowing constraints *without* having to utilize such a closed-form expression. The following are three such implications.

Implication and Test i: Euler Equation Estimation on the Two Subgroups

This first test follows Runkle (1983) and tests the implication that the Euler equation should be satisfied for group 2 but not for group 1.²³ 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.²⁴ This is done here by testing whether an additional time t variable, the log of real disposable income (y_{it}), enters significantly when equation (11) is estimated. I estimate equation (11) 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

²³ The differences between test i and Runkle's test are that (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.

²⁴ For a more detailed description of the test of overidentifying restrictions, see Hansen and Singleton (1983).

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 2 (the constraint is not binding for this group). For group 1, however, λ_{it} will not equal zero and will be correlated with y_{it} and presumably the other variables in equation (11). Since λ_{it} is in the error term, we would expect to get a significant coefficient on income (i.e., reject the over-identifying restriction) or get implausible parameter estimates for group 1, but not for group 2.

An advantage of this test is that it can help distinguish between the failure of auxiliary assumptions and the presence of liquidity constraints. While the failure of any of the auxiliary assumptions described above could lead to a rejection of the Euler equation, it seems unlikely that such a failure would lead to a rejection for low-asset consumers but not for high-asset consumers. However, we would expect to see this pair of results if the rejection of the Euler equation is due to liquidity constraints.

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 2 but not for group 1. 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 today 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.²⁵ I derive a consistent estimate of the true group 1 population average of $\ln(1 + \lambda_{it})$ and test whether it is strictly positive. In order to learn

²⁵ Note that $\ln(1 + \lambda_{it})$ is positive if and only if λ_{it} is positive.

something useful about the importance of liquidity constraints from the sign and magnitude of this estimate, it is crucial that it be consistent under both the null *and* the alternative hypotheses. To arrive at such an estimate, I first estimate equation (11) (without λ_{it} included) for group 2. The resulting parameter estimates are consistent estimates of group 1 parameters whether or not liquidity constraints exist, as long as the constraint is not binding for group 2 and the underlying parameters are the same for the two groups.²⁶ I then use the group 2 parameter estimates to construct estimates of group 1 residuals $x_{i,t+1}$ (call the estimated residuals $\hat{x}_{i,t+1}$).

The residual $x_{i,t+1}$ in equation (11) ($= \ln[1 + \lambda_{it}] + v_{i,t+1}$) includes both the Lagrange multiplier term and the mean zero disturbance term $v_{i,t+1}$. For each observation, $\hat{x}_{i,t+1}$ is therefore the sum of the Lagrange multiplier term, the mean zero disturbance term, and an estimation error term. When $\hat{x}_{i,t+1}$ is averaged across all observations in group 1 (call this average $\bar{\hat{x}}_{i,t+1}$), the estimation error component and the sample average of $v_{i,t+1}$ (as $N \rightarrow \infty$) will approach zero, and the sample average of $(1 + \lambda_{it})$ will approach the true population average of $(1 + \lambda_{it})$. In other words, $\bar{\hat{x}}_{i,t+1}$ will be a consistent estimate of the group 1 population average of $\ln(1 + \lambda_{it})$, which is equal to the average excess growth in consumption in group 1 due to liquidity constraints. This estimate is consistent *whether or not group 1 is liquidity constrained*. If I had instead used parameter values estimated from group 1 observations or the entire sample, they would *not* have been consistent if group 1 were in fact liquidity constrained. If liquidity constraints are important for consumption behavior, this estimate should be positive, statistically significant, and quantitatively large.

Implication and Test iii: The Relationship between λ_{it} and y_{it}

For an individual facing a binding borrowing constraint, if disposable income at time t increases and nothing else in the model changes, the constraint will be directly relaxed 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 $v_{i,t+1}$, there should be a negative partial correlation between $x_{i,t+1}$ and y_{it} , all else equal. As a test of this, I regress the estimate $\hat{x}_{i,t+1}$ calculated above on y_{it} and test

²⁶ The estimate of each individual fixed effect is not consistent as $N \rightarrow \infty$ (consistency would require $T \rightarrow \infty$). Nevertheless, the estimate of the population average of $\ln(1 + \lambda_{it})$ is still consistent as $N \rightarrow \infty$.

whether the sign is negative. Unfortunately, however, this is an estimate of the total derivative of λ_{it} with respect to y_{it} , which need not be of the same sign as the partial derivative. 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 changes in income are and the less correlated income is with the other factors that influence λ_{it} , the closer the partial derivative will be to the total derivative.²⁷

IV. Data

The data were collected by the Survey Research Center at the University of Michigan for the Panel Study of Income Dynamics (PSID).²⁸ The survey has been conducted annually each spring and has followed the same families and their "split-offs" over time. The first survey, "wave" 1, was conducted in the spring of 1968, and I use data through wave 15 (1982 survey). Most questions in the survey (such as total income) ask for values of the prior calendar year's economic variables. Some questions (such as family composition), 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 as reported in the survey taken in year t .

In what follows, I briefly describe the variables used for the 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.

²⁷ There are similarities between these tests and tests for constraints in the micro labor supply literature. Two different types of "corner" solutions are considered in that literature. First, individuals who choose not to participate in the labor force are presumably not at an interior solution with respect to their choice of leisure (at an observed labor supply of zero, the marginal utility of leisure is "too high"). This problem is successfully dealt with in Heckman and MaCurdy (1980), through the use of Tobit estimation (a procedure that might usefully be applied to the consumption equation but is not done so here). Second, workers who respond that they are unemployed or underemployed may be constrained from working as much as they would like at the going wage (at the observed labor supply, the marginal utility of leisure is "too low"). These workers will be off of their unconstrained labor supply function, and thus for these workers the Euler equation relating leisure in t and $t + 1$ will not be satisfied. Tests analogous to my test i are performed in Ball (1986) and Ham (1986), and (under certain assumptions about unobservable preferences for labor supply) the results indicate that many workers are constrained in their choice of hours. Note that as long as consumption and leisure are additively separable in the utility function, the consumption Euler equation (11) will be satisfied even if labor supply constraints are present.

²⁸ The data utilized in this paper were made available by the Inter-University Consortium for Political and Social Research. Neither the original source or collectors of the data nor the consortium bears 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.

A. *Variables Used*

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 others have (such as Hall and Mishkin [1982]) 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 consumer price index (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.²⁹ However, the assumption of separability between food and nonfood consumption seems more plausible than the assumptions typically made in aggregate time-series studies: separability between nondurables/services and durables, and perfect substitutability between nondurables 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 give service flows over subsequent periods, 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 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, a subsequent question is, "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.³⁰ The

²⁹ An additional problem with food consumption as constructed is that the value of labor for preparation is included in one component (food out) but excluded from another (food at home).

³⁰ For a description of the measurement error problems with these data, see Shapiro (1982) and Altonji and Siow (1987).

standard deviation of the growth rate of consumption is equal to 32 percent, which seems implausibly high to be explained by differences in expected growth rates, expectational errors, 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 1-year Treasury bill rate (first 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, on the basis of 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 personal consumption expenditures deflator from the *National Income and Product Accounts* (NIPA).

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.³¹

Nonhousing wealth.—Unfortunately, no data are available on asset holdings other than housing equity. There are questions, however, about asset income received over the year. To estimate nonhousing wealth, 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 3-month Treasury bills over the year. This is an approximation that attempts to account for the amount of wealth typically held in savings accounts.

³¹ 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 because of changes in market interest rates.

Because of the difficulty of “scaling up” asset income other than interest and dividends, I exclude observations with substantial “other asset income.”

Annual food needs (AFN).—I assume above that tastes for food consumption depend on family composition. This 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.

B. *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.³² I include split-off 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.

C. *Criteria for Splitting the Sample*

As discussed above, the sample is split on the basis of two different wealth to income ratios: one ratio is based on nonhuman wealth excluding housing equity, while the other includes housing equity.

There are two sources of information about nonhousing wealth. First, as described above, this wealth is estimated from the numbers on asset income. Second, in five of the relevant waves, there is a question that asks whether the family has any savings such as checking or savings accounts or government bonds, and if so whether they amount to two or more months’ worth of income.

For the first split, I use both sources of information and place an observation on the growth rate of consumption between t and $t + 1$ in

³² 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.

the liquidity-constrained group (group 1) if *either* they responded directly that they did not have at least two months' worth of savings *or* their estimated nonhousing wealth was less than two-twelfths of their average disposable income during t and $t - 1$. This second criterion is chosen to correspond most closely to the direct survey question.

I also try a more extreme split based on nonhousing wealth that excludes from either group the middle 30 percent of the sample. Under this split, group 1 contains individuals who responded that they had no savings or whose estimated nonhousing wealth was zero.³³ Group 2 contains individuals with estimated nonhousing wealth of at least six months' worth of average income.

Finally, for the split that is based on *total* wealth (including housing equity), I ignore the direct question on savings and place an observation in group 1 if estimated total wealth was less than two months' worth of average income.

Table 1 contains a summary of the responses to the direct question on savings and the fractions of observations falling into groups 1 and 2 for each split. For the first nonhousing wealth split, about two-thirds of the population fall into group 1 and one-third into group 2. For the more extreme split, about 45 percent fall into group 1 and about 25 percent into group 2, with the middle 30 percent excluded. For the total wealth split, about one-third of the sample are in group 1 and two-thirds in group 2.

There is some discrepancy between the direct question on liquid wealth and my calculations. For the years that the question was asked directly, only 46 percent 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 nonconstrained 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 nonhousing wealth. As mentioned above, the tests rely on the consistency of the group 2 parameter estimates and are therefore valid even if some non-liquidity-constrained observations are included in group 1.

V. Empirical Analysis

A. Estimation

As discussed above, the equations are estimated with an instrumental variables procedure. The basic Euler equation that I estimate follows

³³ I exclude individuals who had estimated nonhousing wealth equal to zero but responded to the direct question that they had greater than two months' worth of liquid assets.

TABLE 1
FRACTIONS OF SAMPLE IN DIFFERENT WEALTH POSITIONS

Liquid Wealth Position*		
1a.	Between 0 and 2 months' worth of income now; > 2 months' worth of income sometime in the last 5 years	.16
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 anytime in the last 5 years	.10
2.	> 2 months' worth of income now	.54
Nonhousing Wealth Splits [†]		
Basic split:		
Group 1:	≤ 2 months' worth of (average) income or liquid wealth position = 1 ($a-d$)	.67
Group 2:	> 2 months' worth of (average) income and liquid wealth position = 2	.33
Extreme split:		
Group 1:	0 current nonhousing wealth or liquid wealth position = 1c or 1d	.46
Group 2:	> 6 months' worth of (average) income and liquid wealth position = 2	.24
Total Wealth Split [‡]		
Group 1:	≤ 2 months' worth of (average) income	.29
Group 2:	> 2 months' worth of (average) income	.71

* Direct survey questions. Five waves; 9,613 observations.

† Also based on liquid wealth question. Entire sample for which nonhousing wealth could be calculated. Ten waves; 18,181 observations.

‡ Entire sample for which total wealth could be calculated. Ten waves; 16,628 observations.

directly from equation (11):

$$\begin{aligned}
 GC_{it+1} = & \sum_{j=1}^N c^j FD_{it}^j + \sum_k d^k WD_{it}^k + a_1 \text{age}_{it} + a_2 \text{GAFN}_{it+1} \\
 & + a_3 \ln(1 + r_{it}) + \epsilon_{it+1},
 \end{aligned}
 \tag{12}$$

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 $t + 1$; $\{FD_{it}^j\}$ = N family dummies, one for each family in the sample (= 1 if $i = j$, 0 otherwise); $\{WD_{it}^k\}$ = nine wave 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 family i 's annual food needs between t and $t + 1$; $\ln(1 + r_{it})$ = the log of one plus family i 's real after-tax riskless rate between t and $t + 1$; y_{it} = the log of family i 's disposable real income in year t ; and MTR_{it} = the marginal tax rate of family i in year t .

The wave dummies are included to capture the aggregate compo-

ment of expectations errors, and the family dummies are included to capture the family-specific effects, each described earlier. The family dummies were first "partialled out" by subtracting off family means from each variable; that is, fixed effects estimation was performed.³⁴

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 2 and then imposed on the family observations in group 1. Therefore, observations can be used for these tests only if the family has observations *both* in group 1 and in group 2 at some time during the sample. This reduces the number of group 2 observations available for these tests, sometimes dramatically.

The residuals in the estimated equations are unlikely to be independent and identically distributed for three reasons. First, as discussed previously, the PSID consumption data contain substantial measurement error. The Euler equation error will therefore include a term in the growth rate of this measurement error, which is unlikely to be serially uncorrelated. Second, the taste shifter may not follow a random walk; that is, the change in tastes may not be independent and identically distributed. Finally, the variance of the expectations error may differ across families. For these reasons, I correct the standard errors to allow for general serial correlation and heteroscedasticity. The procedure allows each family to have a different and unrestricted covariance structure but assumes that the errors are uncorrelated across families.³⁵

³⁴ Since the level of income and the marginal tax rate are potentially correlated with lagged innovations in the Euler equation, the use of fixed effects introduces potential biases that are analogous to those in dynamic regression models with fixed effects (see Hsiao 1986, chap. 4).

³⁵ The standard errors are calculated along the lines of White (1984) and Altonji and Siow (1987) as

$$\hat{\mathbf{V}}(\hat{\beta}) = \left[\sum_{i=1}^N (\hat{\mathbf{X}}_i' \hat{\mathbf{X}}_i) \right]^{-1} \left[\sum_{i=1}^N (\hat{\mathbf{X}}_i' \hat{\boldsymbol{\epsilon}}_i \hat{\mathbf{X}}_i) \right] \left[\sum_{i=1}^N (\hat{\mathbf{X}}_i' \hat{\mathbf{X}}_i) \right]^{-1},$$

where $\hat{\mathbf{X}}_i$ is the matrix of observations for family i of the right-hand-side variables in the second-stage estimation of the Euler equation, and $\hat{\boldsymbol{\epsilon}}_i$ is the vector of estimated Euler equation residuals for family i . Note that these standard errors account for the within-family correlation of the errors induced by removing fixed effects, so that no additional degree of freedom correction is necessary. I thank Joe Altonji and Chris Paxson for their help with calculating these standard errors. For tests ii and iii, these standard errors are correct only under the assumption that estimated group 2 parameters that are imposed on group 1 are exactly equal to the true underlying parameters. I calculated a second set of t -statistics for tests ii and iii, which take into account the

B. Results

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

Test i: Estimation on Each of the Two Subgroups

The results of separate estimation of equation (12), with y included, for groups 1 and 2, are reported on the top of table 2.

Group 2.—For this high-asset “unconstrained” 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 2.3, an estimate in line with previous empirical estimates.³⁶ 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 statistically 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 percent level.

Group 1.—The group 1 coefficient on GAFN is almost the same as that for group 2. The coefficient on the interest rate implies a coefficient of relative risk aversion of 2.7, about the same as that for group 2, but again has a large associated standard error. However, the coefficient on income is negative, has almost twice the magnitude as the group 2 coefficient, and is statistically significant. As discussed in Section III, this is inconsistent with the permanent income hy-

sampling error in the group 2 parameter estimates but do not correct for serial correlation or heteroscedasticity. This variance-covariance matrix is

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

where \mathbf{Z} is 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), $\hat{\gamma}$ is the corresponding estimated parameters, \mathbf{X} is the matrix of right-hand-side variables in the Euler equation, $\hat{\beta}$ is the corresponding two-stage least-squares parameter estimates for the unconstrained group, and $\hat{V}(\hat{\beta})$ is the estimated variance matrix of these parameters. The first term is what would be printed by a standard regression package. The second term captures the fact that the estimated $\hat{\beta}$ for the unconstrained group is a consistent but noisy estimate of the true β . The resulting t -statistics were approximately 30 percent smaller than those without any corrections. I thank Whitney Newey and Andy Lo for their help in figuring out the appropriate formulas for these standard errors. It is likely that the standard errors that correct jointly for both of these problems would be somewhat larger than those reported in the text. Note, however, that the standard errors in the text do correct for the estimation error in the fixed person effects (which generates a family-specific component of the error) and for any heteroscedasticity arising from the estimation error of the other coefficients.

³⁶ Estimates of the coefficient of relative risk aversion include Friend and Blume (1975) (2), Mankiw (1981) (4), and Hansen and Singleton (1983) (1).

TABLE 2

SPLIT BASED ON LIQUID (Nonhousing) 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 1 (Low W/Y)	Group 2 (High W/Y)
Age of head	-.0084 (-.88)	-.0044 (-1.97)*
Growth in annual food needs $t, t + 1$.25 (8.26)*	.23 (3.97)*
Real after-tax Treasury bill rate $t, t + 1$.37 (.24)	.43 (.31)
Log of real disposable income t	-.071 (-4.40)*	-.039 (-1.49)
Degrees of freedom	9,362 (2,731 families)	4,477 (1,583 families)
TESTS ON THE GROUP 1 RESIDUALS CONSTRUCTED USING GROUP 2 PARAMETER ESTIMATES		
Test ii: Estimate of Average Excess Consumption Growth for Group 1 Due to Binding Constraint		
\bar{x}_{it}	.017 (1.63)	
Test iii: Regression of Estimate of Excess Consumption Growth for Group 1 on the Log of Real Disposable Income		
y_{it}	-.024 (-1.31)	
Degrees of freedom	4,267 (1,114 families)	

NOTE.—Equations are estimated with instrumental variables and include time and family fixed effects. t -statistics are in parentheses.

* Significant at the 5 percent level.

pothesis 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 nonhousing assets relative to income, but a violation of the Euler equation for observations with less than two months' worth of income in nonhousing assets—exactly those observations we expect to be liquidity constrained.

Test ii: The One-Sided Inequality in the Euler Equation

At the bottom of table 2, I present a consistent estimate of the average Lagrange multiplier for group 1 observations, equal to the average

unexplained consumption growth for this group. Recall that if group 1 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 but is not statistically significant at the 5 percent level (the coefficient is, however, significant at the 10 percent level). The point estimate is equal to .017, indicating that the growth of food consumption is on average 1.7 percentage points higher for group 1 than would be predicted by a model with no binding constraints.^{37,38} (Raw average consumption growth for the entire sample was approximately zero.)³⁹

In summary, the results of test i lead to a rejection of the Euler equation for group 1 but not for group 2. Test ii indicates that the sign and magnitude of the one-sided inequality are consistent with the view that borrowing constraints exist and affect consumption in important ways, but this conclusion is clearly weakened by the low significance level of this coefficient.

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 1 observations) in which the left-hand-side variable is an 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 at time t . The coefficient is a consistent estimate of the relationship between the Lagrange multiplier and current income. As explained above, if li-

³⁷ An earlier draft of this paper reported a coefficient of .018 and a t -statistic of 1.72, which was significant at the 5 percent level. The slight difference in the coefficient estimate occurs because the current estimates use the 1984 data tape to replicate the 1982 sample and data, whereas the previous draft used the 1982 data tape. In addition, the standard error differs because the previous draft did not adjust the standard errors for general serial correlation and heteroscedasticity.

³⁸ The estimate of .017 does not necessarily imply that consumption growth for this group would be 1.7 percentage points lower if borrowing constraints were eliminated for the economy as a whole because there could be mitigating effects due to resulting changes in equilibrium asset returns.

³⁹ One might expect both the fraction of the population that is constrained and the severity of the constraint for a given individual to be related to the stage of the business cycle, although it is not clear as a matter of theory exactly what the timing of this relationship should be. I examined the average of the Lagrange multiplier for each year. The averages were positive for eight of the ten years, but I saw no obvious relationship to growth in gross national product. In addition, the fraction of the sample in group 1 was relatively constant across years. For an examination with aggregate data of differences in consumption Euler equation estimates across stage of the business cycle, see Ferson and Merrick (1987).

quidity 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.

C. 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 1 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.

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 1 (Low W/Y)	Group 2 (High W/Y)
Age of head	.0098 (.84)	-.0050 (-.63)
Growth in annual food needs $t, t + 1$.30 (5.95)*	.22 (6.07)*
Real after-tax Treasury bill rate $t, t + 1$	-1.46 (-.36)	1.44 (1.34)
Log of real disposable income t	-.081 (-2.98)*	-.047 (-2.59)*
Degrees of freedom	3,203 (1,533 families)	9,353 (2,511 families)
TESTS ON THE GROUP 1 RESIDUALS CONSTRUCTED USING GROUP 2 PARAMETER ESTIMATES		
Test ii: Estimate of Average Excess Consumption Growth for Group 1 Due to Binding Constraint		
\bar{x}_{it}	.0084 (.66)	
Test iii: Regression of Estimate of Excess Consumption Growth for Group 1 on the Log of Real Disposable Income		
y_{it}	-.058 (-2.38)*	
Degrees of freedom	2,268 (862 families)	

NOTE.—Equations are estimated with instrumental variables and include time and family fixed effects. t -statistics are in parentheses.

* Significant at the 5 percent level.

The results are similar but not as strong or clear-cut as before. Lagged income enters significantly (with a negative sign) for both group 1 and group 2, although the coefficient for group 1 is closer to zero. The estimated coefficient of relative risk aversion is positive for group 2 but negative (and thus implausible) for group 1, although again neither is significantly different from zero. The estimate of the average Lagrange multiplier is positive (0.8 percent) but smaller than that 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 about exactly how to split the sample, I tried one last way of splitting it. The sample is split into three sets of observations, but the middle set of observations is not used, leaving only the two extreme sets. As described above, the observations in group 1 either responded that they had no liquid assets or reported zero asset income for the year. The observations in group 2 possessed nonhousing wealth of at least six months' worth of average income. 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 1 Euler equation, while the point estimate on lagged income is much smaller and insignificant for group 2. The estimate of the average Lagrange multiplier for group 1 is 4.3 percent and is statistically significant at the 5 percent level using a one-sided test. The Lagrange multiplier is negatively related to lagged income, but the coefficient is not statistically significant.

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.⁴⁰ This suggests that a sample selection bias that arises in test ii when the dummies are omitted may be important.⁴¹

⁴⁰ For test i (for each split) the coefficient on income was negative and statistically significant for group 1 and was smaller in magnitude and not significant for group 2. Also, the estimated coefficient of relative risk aversion was negative (and in some cases significantly so) for group 1 and positive (but not significantly so) for group 2.

⁴¹ If the family-specific effect is correlated with the instruments or the variable used to split the sample, then excluding it might lead to inconsistent estimates. In particular, those families with high rates of time preference, all else equal, will 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 1. Measured

TABLE 4
EXTREME SPLIT

Test i: Euler Equation Estimates for Two Subsamples: Dependent Variable: $\log C_{i,t+1}/C_{it}$		
Independent Variable	Group 1 (Low W/Y)	Group 2 (High W/Y)
Age of head	.0031 (.31)	-.0044 (-2.00)*
Growth in annual food needs $t, t + 1$.26 (6.70)*	.25 (4.09)*
Real after-tax Treasury bill rate $t, t + 1$	1.92 (.81)	.58 (.32)
Log of real disposable income t	-.063 (-3.22)*	-.021 (-.67)
Degrees of freedom	6,066 (2,318 families)	3,079 (1,197 families)
TESTS ON THE GROUP 1 RESIDUALS CONSTRUCTED USING GROUP 2 PARAMETER ESTIMATES		
Test ii: Estimate of Average Excess Consumption Growth for Group 1 Due to Binding Constraint		
\bar{x}_{it}	.043 (2.27)*	
Test iii: Regression of Estimate of Excess Consumption Growth for Group 1 on the Log of Real Disposable Income		
y_{it}	-.015 (-.52)	
Degrees of freedom	1,315 (537 families)	

NOTE.—Equations are estimated with instrumental variables and include time and family fixed effects. t -statistics are in parentheses.

* Significant at the 5 percent level.

Finally, it is possible that since wealth itself is an endogenous variable, the error term is correlated with the variable used to split the sample. To guard against this possibility, the sample was split on the basis of predicted rather than actual wealth as follows. A dummy variable was set equal to one if the observation fit into group 1 and zero if it fit into group 2. A logistic regression procedure was then used to predict this dummy variable on the basis of age, age squared, family composition, change in family composition, race, religion, sex,

excess growth in consumption for group 1 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 (test ii) *against* the hypothesis that individuals are liquidity constrained.

marital status, and education. An observation was considered constrained if the predicted probability was greater than .6 and unconstrained if the predicted probability was less than .4.⁴² For the basic split based on nonhousing wealth, the results for test i were very similar to those reported in table 2 above. The coefficient on income for (predicted) group 1 was $-.07$ with a t -statistic of -4.27 , while the coefficient on income for (predicted) group 2 was only $-.02$ with a t -statistic of -0.56 . Since most of the variables used to predict the assets/income ratio do not vary across time for a given family, only 150 observations were available to perform tests ii and iii. The excess consumption growth was positive, and the coefficient on income was negative (and each was statistically significant), but these should be interpreted with caution given the small number of observations on which they are based.⁴³

VI. Concluding Remarks

In those studies in which the permanent income hypothesis is rejected empirically, liquidity/borrowing constraints are often suggested as a possible 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 I 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 the extent to which individuals are constrained.

The empirical results presented should be interpreted with caution. The data are considerably less than ideal in that some variables need

⁴² This procedure was used, rather than predicting wealth directly, because the original group 1/group 2 split was based in part on a zero/one response to the question whether the household had assets of more than two months' income.

⁴³ The results for the other splits were mixed but generally less supportive of the liquidity constraints hypothesis.

to be constructed and others are extremely noisy indicators of the true measure. In addition, the results are not always consistent across variations in the testing and sample selection procedures. With that said, however, these results generally support the view that borrowing constraints affect consumption in the United States. For the basic and extreme splits 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 positive for the low wealth/income ratio group corresponding to a one-sided inequality of the Euler equation in the direction consistent with borrowing constraints. The estimates of the excess consumption growth are 1.7 percent (but not statistically significant) and 4.3 percent (statistically significant) for the basic and extreme splits, respectively. The results using a split based on total assets are less conclusive.

The results presented here suggest that borrowing constraints are important. Further research is clearly needed, however, to develop and test 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 were available at the time of the beginning 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 to some questions. When data were 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 the cases in which 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.

A. *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), nonhousing wealth, and housing equity. Additional information is available in Institute for Social Research (1972–84, 1984).

1. 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 the 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 minus 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 is 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 (referring to the 1972 survey year) stated that “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” (Institute for Social Research 1972, p. 302).

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 ($GC_{i,t+1}$) is equal to $\ln(C_{i,t+1}/C_{it})$.

2. The Real after-Tax Interest Rate

The log of one plus the real after-tax interest rate is equal to $\ln\{[1 + RN_t(1 - MTR_{i,t+1})]/[(1 + P_{t+1})/(1 + P_t)]\}$, 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 1-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) on the basis of 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 backward 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.

3. Annual Food Needs

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 on the basis of table A1. 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 on the basis of table A2. Note that the same tables are used for all 15 waves. The growth in annual food needs ($GAFN_{i,t+1}$) is equal to $\ln(AFN_{i,t+1}/AFN_{it})$.

4. 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 plus taxable income of others plus total transfers of head and wife plus total transfers of others plus bonus value of food stamps minus federal income taxes of head and wife minus federal income taxes of others minus social security taxes of head and wife, each for the entire calendar year prior to the survey year. (The PSID did not include the bonus value of food stamps as part of transfers.) I

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 OF SCALE

Size	Adjustment
One-person family	+20%
Two-person family	+10%
Three-person family	+5%
Four-person family	0
Five-person family	-5%
Family with six or more persons	-10%

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 the waves on income, transfers, and food stamps, as well as the data for waves 9, 10, and 12–15 for property taxes, come from direct survey questions. Total income taxes are estimated by the PSID (see the description of the marginal tax rate above). I add the surcharges for 1968–70 and subtract the low-income tax credit in the years in which the PSID did not do this (waves 9–12). For waves 1–8, property taxes were imputed by the PSID on the basis of house value and distance from the nearest city. There are no data on the tapes about wave 11 property taxes. I estimate these on the basis of 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 or 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.

5. Nonhousing Wealth

I construct the wealth variable from questions on asset income, and to do so requires some rather bold assumptions. Since this variable is used only 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 [i.e., as being in one of 10 brackets] 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 were reported only in bracketed form. No asset data on others were 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 nonzero, I have no way of estimating their wealth. I therefore record wealth as missing in waves 9–15 if others' asset income is nonzero. For waves 1–8, I record wealth as missing if others' asset income in wave 8 is nonzero.

To estimate the stock of assets, I divide the first \$250 of interest and dividend 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 3-month Treasury bills. I do this to attempt to account for the amount of wealth typically held in low-interest accounts.

Real nonhousing wealth is equal to the nominal amount deflated by the PCE deflator.

6. Housing Equity

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 in which a second mortgage exists (and for observations based on interpolations or extrapolations of observations with a second mortgage).

B. Sample Selection

1. Families Excluded

Families are followed through time by keeping track of the head of the household, and new families that are formed from the original ones (through split-offs 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 split-offs) and exclude the poverty subsample.

2. Observations Excluded because of Family Change

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

3. 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 in which 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 in which 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).

4. 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 was estimated for wave 2. Because I use growth rates, I need to eliminate waves 5 and 7. This leaves me with a maximum of 10 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 either if 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 3 and 4, there are 18,181 total observations for the liquid asset splits and 16,628 observations for the total asset split.

C. *Splitting into Subgroups*

1. 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 nonhousing wealth in t to the average disposable income in t and $t - 1$. If the data for this ratio are missing, I exclude the observation from both groups.

I put observations into group 1 if either the answer to either question 1 or 2 above is no or the ratio above is less than $2/12$ (corresponding to two months' worth of income). If neither of the above was true, I put the observation in group 2. This results in 12,107 observations in group 1 and 6,074 observations in group 2.

For tests ii and iii, only group 1 observations corresponding to families that also had observations in group 2 could be used. This leaves 4,269 observations for group 1.

2. 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 are missing, I exclude the observation from both groups.

For this split, I put observations into group 1 if the ratio is less than 2/12; otherwise the observation goes into group 2. This split puts 4,750 observations in group 1 and 11,878 in group 2. For tests ii and iii, 2,270 observations are included in group 1.

3. Most Stringent Split (Table 4)

For this split I put the two extremes into groups 1 and 2 and I exclude the middle observations. If the ratio of nonhousing wealth to income is missing, the observation is excluded from both groups.

An observation is placed in group 1 if at least one of the following is true: (a) the answer to question 1 (see Sec. C1 above) is no (i.e., they say they have no current savings) or (b) the ratio of nonhousing wealth to income is equal to zero and the answer to question 2, if asked, is no (i.e., they do not have at least two months' worth of savings).

An observation is placed in group 2 only if both of the following are true: (a) they report that they have current savings (if the direct question is asked in that wave), and (b) the ratio of nonhousing wealth to income is greater than or equal to .5.

This results in 8,398 observations in group 1 and 4,290 observations in group 2. For tests ii and iii, only 1,317 observations are available in group 1.

References

- Altonji, Joseph G. "Intertemporal Substitution in Labor Supply: Evidence from Micro Data." *J.P.E.* 94, no. 3, pt. 2 (June 1986): S176-S215.
- Altonji, Joseph G., and Siow, Aloysius. "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data." *Q.J.E.* 102 (May 1987): 293-328.
- Ball, Laurence. "Intertemporal Substitution and Constraints on Labor Supply: Evidence from Panel Data." Manuscript. New York: New York Univ., Grad. School Bus., 1986.
- Bernanke, Ben S. "Permanent Income, Liquidity, and Expenditure on Automobiles: Evidence from Panel Data." *Q.J.E.* 99 (August 1984): 587-614.
- . "Adjustment Costs, Durables, and Aggregate Consumption." *J. Monetary Econ.* 15 (January 1985): 41-68.
- Bewley, Truman F. "The Permanent Income Hypothesis: A Theoretical Formulation." *J. Econ. Theory* 16 (December 1977): 252-92.
- Chamberlain, Gary. "Panel Data." In *The Handbook of Econometrics*, vol. 2, edited by Zvi Griliches and Michael D. Intriligator. New York: North-Holland, 1984.
- Cowing, Thomas G. "The Effectiveness of Rate-of-Return Regulation: An Empirical Test Using Profit Functions." In *Production Economics: A Dual Approach to Theory and Applications*, vol. 2, edited by Daniel McFadden and Melvyn Fuss. New York: North-Holland, 1978.
- Ferson, Wayne E., and Merrick, John J., Jr. "Non-Stationarity and Stage-of-the-Business-Cycle Effects in Consumption-based Asset Pricing Relations." *J. Financial Econ.* 18 (March 1987): 127-46.
- Flavin, Marjorie A. "The Adjustment of Consumption to Changing Expectations about Future Income." *J.P.E.* 89 (October 1981): 974-1009.

- Friend, Irwin, and Blume, Marshall E. "The Demand for Risky Assets." *A.E.R.* 65 (December 1975): 900–922.
- Hall, Robert E. "Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence." *J.P.E.* 86 (December 1978): 971–87.
- Hall, Robert E., and Mishkin, Frederic S. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica* 50 (March 1982): 461–81.
- Ham, John C. "Testing Whether Unemployment Represents Intertemporal Labour Supply Behaviour." *Rev. Econ. Studies* 53 (August 1986): 559–78.
- Hansen, Lars Peter, and Singleton, Kenneth J. "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *J.P.E.* 91 (April 1983): 249–65.
- Hayashi, Fumio. "The Effect of Liquidity Constraints on Consumption: A Cross-sectional Analysis." *Q.J.E.* 100 (February 1985): 183–206. (a)
- . "The Permanent Income Hypothesis and Consumption Durability: Analysis Based on Japanese Panel Data." *Q.J.E.* 100 (November 1985): 1083–1113. (b)
- Heckman, James J. "Life Cycle Consumption and Labor Supply: An Explanation of the Relationship between Income and Consumption over the Life Cycle." *A.E.R.* 64 (March 1974): 188–94.
- Heckman, James J., and MaCurdy, Thomas E. "A Life Cycle Model of Female Labour Supply." *Rev. Econ. Studies* 47 (January 1980): 47–74.
- Heckman, James J., and Singer, Burton. "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data." *Econometrica* 52 (March 1984): 271–320.
- Hsiao, Cheng. *Analysis of Panel Data*. Econometric Society Monograph no. 11. New York: Cambridge Univ. Press, 1986.
- Institute for Social Research. *A Panel Study of Income Dynamics: Procedures and Tape Codes*. Waves 1–15. Ann Arbor: Univ. Michigan, Survey Res. Center, 1972–84.
- . *User Guide to the Panel Study of Income Dynamics*. Ann Arbor: Univ. Michigan, Survey Res. Center, 1984.
- Juster, F. Thomas, and Shay, Robert P. *Consumer Sensitivity to Finance Rates: An Empirical and Analytical Investigation*. Occasional Paper no. 88. New York: NBER, 1964.
- Mankiw, N. Gregory. "The Permanent Income Hypothesis and the Real Interest Rate." *Econ. Letters* 7, no. 4 (1981): 307–11.
- Mankiw, N. Gregory; Rotemberg, Julio J.; and Summers, Lawrence H. "Intertemporal Substitution in Macroeconomics." *Q.J.E.* 100 (February 1985): 225–51.
- Pissarides, Christopher A. "Liquidity Considerations in the Theory of Consumption." *Q.J.E.* 92 (May 1978): 279–96.
- Rotemberg, Julio J. "Consumption and Liquidity Constraints." Manuscript. Cambridge: Massachusetts Inst. Tech., 1984.
- Runkle, David E. "Liquidity Constraints and the Permanent Income Hypothesis: Evidence from Panel Data." Ph.D. dissertation, Massachusetts Inst. Tech., 1983.
- Scheinkman, José A., and Weiss, Laurence. "Borrowing Constraints and Aggregate Economic Activity." *Econometrica* 54 (January 1986): 23–45.
- Shapiro, Matthew D. "A Note on Tests of the Permanent Income Hypothesis in Panel Data." Manuscript. New Haven, Conn.: Yale Univ., 1982.

- . “The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data.” *Econ. Letters* 14, no. 1 (1984): 93–100.
- Stiglitz, Joseph E., and Weiss, Andrew. “Credit Rationing in Markets with Imperfect Information.” *A.E.R.* 71 (June 1981): 393–410.
- White, Halbert. *Asymptotic Theory for Econometricians*. Orlando, Fla.: Academic Press, 1984.
- Zeldes, Stephen P. “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence.” Working Paper no. 20-86. Philadelphia: Univ. Pennsylvania, Wharton School, Rodney L. White Center, August 1986.
- . “The Effects of Borrowing Constraints on Consumption in a Multiperiod Model with Income Uncertainty.” Manuscript. Philadelphia: Univ. Pennsylvania, Wharton School, 1987.