Union Contracts and the Life-Cycle/Permanent-Income Hypothesis

By JOHN SHEA*

This paper isolates households in the PSID whose heads can be matched to particular long-term union contracts with high confidence. I use published information on these contracts to construct a household-specific measure of expected wage growth. I find that predictable wage movements are significantly correlated with consumption changes, contrary to neoclassical consumption theory. I find that consumption responds more strongly to predictable income declines than to predictable income increases. This asymmetry is inconsistent with liquidity constraints and myopia but is qualitatively consistent with models in which preferences exhibit loss aversion. (*JEL* E21, D12)

Neoclassical consumption theory posits that households are forward-looking and that credit markets are perfect. A central implication of this life-cycle/permanent-income hypothesis (LCH/PIH) is that, aside from interest rates and taste shifts, future household consumption growth should be orthogonal to variables in the household's current information set. In particular, consumption should not respond to predictable movements in income. Intuitively, a household expecting income to change tomorrow should borrow or save today to smooth its consumption path.

Since the work of Robert E. Hall (1978), a large body of research has tested the dynamic implications of the LCH/PIH. This research has typically estimated equations of the following form:

$$(1)\quad \log(C_t) - \log(C_{t-1}) = \alpha + \varphi \mathbf{X}_{t-1} + \beta Z_{t-1} + \varepsilon_t$$

where C_t is real consumption at time t, $\mathbf{X}_{t-1}$ is a vector of controls such as the expected real after-tax rate of return between $t-1$ and t, and Z_{t-1} is a variable in the information set at $t-1$ correlated with income growth between $t-1$ and t.

Under the LCH/PIH, β should equal zero; predictable income movements should not cause consumption to change. Of course, β will differ from zero under plausible alternatives to the LCH/PIH, such as myopia or liquidity constraints, provided the instrument Z_{t-1} is truly correlated with future income growth.[1]

Recent tests of the LCH/PIH in aggregate data consistently reject the LCH/PIH. John Y. Campbell and N. Gregory Mankiw (1990), for instance, find that a predictable 1-percent income increase leads to a statistically significant consumption increase of

*Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706. I thank Joseph Altonji, John Kennan, Patrick Lampani, Steve Pischke, Ken West, Jeff Wooldridge, Steve Zeldes, anonymous referees, and seminar participants at MIT and the NBER Summer Institute for helpful discussions and suggestions. Support from the Graduate School at the University of Wisconsin and the National Science Foundation is acknowledged.

[1] I use the term "instrument" somewhat loosely, to refer to variables in the information set at $t-1$ used to test the Euler equation (1). One could, of course, also test the LCH/PIH by estimating (1) using instrumental variables, effectively setting Z equal to the projection of household income growth between $t-1$ and t on variables in the information set at $t-1$. This approach is difficult to implement using the Panel Study of Income Dynamics, because measured consumption at t refers to the flow of spending the time of the interview, while measured income at t refers to the entire previous calendar year.

between 0.351 percent and 0.713 percent. Recent tests using household data, however, have produced mixed results. Joseph G. Altonji and Aloysius Siow (1987) and David E. Runkle (1991) find little evidence against the LCH/PIH. Steven P. Zeldes (1989) rejects the LCH/PIH and attributes this rejection to liquidity constraints. Marjorie Flavin (1991) also rejects the LCH/PIH but attributes this rejection to myopia rather than liquidity constraints.

There are three reasons why tests of the LCH/PIH in household data may have low power. First, the Panel Study of Income Dynamics (PSID) data used in most household studies cover only food consumption.[2] In principle, one can test the LCH/PIH using food consumption alone, if food and other goods are separable in utility. In practice, however, the dynamic behavior of food appears to be atypical: Shea (1994) shows that standard aggregate time-series tests cannot reject the LCH/PIH for aggregate food consumption but can reject the LCH/PIH for other subcomponents of aggregate consumption. This suggests that estimates of (1) using food consumption may need an unusually powerful instrument to reject the LCH/PIH even if the true β is nonzero.

Second, consumption is poorly measured in the PSID. The PSID is primarily interested in tracking labor-market behavior, rather than consumption behavior. Households are not asked to keep a diary of food purchases; respondents are simply asked to report typical weekly food spending off the tops of their heads. As a result, a large fraction of consumption growth variability in the PSID appears to be noise (Matthew D. Shapiro, 1984). This noise will tend to make estimates of (1) imprecise.

Third, it is difficult to find variables in households' current information sets that are good predictors of future income growth.[3] For instance, Zeldes (1989) has several years of data for each household; he estimates (1) including household-specific constants and uses lagged household income as Z_{t-1}. Effectively, Zeldes's instrument is the deviation of lagged income from the household's mean income over the sample period. As Runkle (1991) points out, however, this instrument is biased against the LCH/PIH. Household mean income is estimated using data from before and after $t-1$. While lagged income per se is in the households' current information set, the deviation of lagged income from the sample mean is not; on average, a high level of this deviation at $t-1$ signals declining income to the econometrician, even if income movements are unpredictable to the household. Moreover, under the LCH/PIH, negative income surprises should reduce household consumption. Thus, Zeldes's instrument may reject the LCH/PIH even if the LCH/PIH is true.

Runkle (1991), meanwhile, estimates (1) without household-specific constants and uses instruments such as lagged income, lagged income growth, lagged consumption, and lagged consumption growth. Runkle finds that these variables are uncorrelated

[2] Micro studies of the LCH/PIH behavior using alternative data sources include Christopher D. Carroll (1994), which uses data from the Consumer Expenditure Survey (CES), and Flavin (1991), which uses data from the Michigan Survey of Consumer Finances (SCF). Unfortunately, the CES data is cross-sectional, while the SCF data have only two observations per family. The PSID is the only available data set that tracks a group of households for many periods.

[3] Shapiro (1984) also includes household-specific constants. Flavin (1991) has only two years of data per household. She uses survey evidence on households' expected income growth as her Z_{t-1}; this instrument is intuitively appealing, and Flavin's conclusions (the LCH/PIH is rejected, while liquidity constraints are not the chief culprit) are similar to mine. Altonji and Siow (1987) use a variety of instruments in their work, some of which (notably the household's lagged weeks spent unemployed) are relevant for future income growth; their inability to reject the LCH/PIH thus must be taken seriously as evidence contrary to my conclusions. I discuss Altonji and Siow's work further in what follows.

with consumption growth and concludes that the data support the LCH/PIH. Unfortunately, Runkle's instruments are not strongly correlated with future income growth, implying that his tests have little power to reject the LCH/PIH.

This paper attempts to address the third problem—finding valid, yet powerful, instruments—with brute force. I isolate a subset of households in the PSID whose heads can be matched to particular labor-union contracts with high confidence. I use published information on these contracts to construct a household-specific measure of expected wage growth. To the extent that my matches are accurate, this measure should be strongly correlated with future income growth and should thus be a powerful instrument to test the LCH/PIH in equation (1).

The intuition for this approach is simple. Suppose worker A is covered by a union contract that specifies no wage increase between $t-1$ and t, while worker B's contract specifies a large wage increase between $t-1$ and t. According to the LCH/PIH, B's consumption growth between $t-1$ and t should not systematically exceed A's, since B's expected raise should be fully reflected in his consumption at $t-1$. Of course, B's consumption may grow faster under plausible alternatives to the LCH/PIH.

Obviously, matching households to particular union contracts is the hardest part of this strategy. The key variables available in the PSID that make such assignments possible are the household's state and county of residence, the household head's industry and occupation, and the head's union status. Successful matches of households to contracts fall into two categories. First, some heads work in industries subject to national or regional pattern bargaining, such as railroads, trucking, and construction. Second, some heads work in industries characterized by dominant local employers. For instance, a Seattle household whose head works in aircraft manufacturing can be assigned to Boeing with reasonably high confidence. This second type of assignment relies both on PSID data and on outside information about the location of establishments.

The empirical results can be summarized as follows. I find that expected wage growth is significantly correlated with consumption growth, contrary to the LCH/PIH. I cannot reject the LCH/PIH using other commonly employed instruments such as lagged consumption and income; this suggests that my instrument is relatively powerful. I find that consumption is more sensitive to predictable real wage declines than to predictable wage increases. This finding is inconsistent with both myopia and liquidity constraints but is qualitatively consistent with recent theoretical work (David Bowman et al., 1993) incorporating loss aversion into intertemporal preferences.

While this paper focuses on the consumption Euler equation, my method of matching households to particular observable shocks may be useful in other settings. For example, recent research (e.g., John H. Cochrane, 1991) tests for the existence of complete contingent markets by studying the response of household consumption to exogenous idiosyncratic shocks. If company-level shocks are more exogenous to households than other variables typically found in panel data sets, linking households to their employers may permit a cleaner test of full insurance. This paper shows that one can often make such matches given detailed information on household location and industry affiliation.

I. Sample Selection

This section briefly describes my sample-selection methods. More thorough descriptions of the selection techniques and the composition of the sample are contained in Shea (1992), which is available from the author upon request.

This paper uses data from the Michigan Panel Study of Income Dynamics (PSID). The PSID has gathered annual data on food consumption and other variables for a group of households and their "split-offs" since 1968. However, the industry and occupation of the household head are available at the three-digit level only since 1981. Since I need such detailed information to assign households to particular union contracts, I use only data from 1981–1987.[4] For my purposes, an *observation* consists of two consecutive years of data (which I call $t-1$ and t) for a single household. There are 7,061 households on the 1987 PSID tape, implying 42,366 potential sample observations (six per household, ending 1982–1987).[5]

Before proceeding, I must take a stand on the timing of the consumption data. Following Zeldes (1989), I assume that the PSID's questions measure the flow of consumption *at the time of the interview*.[6] For the period 1981–1987, over 90 percent of PSID interviews were conducted in March, April, or May. For the sake of uniformity, I assume that consumption data refer to March of the interview year. Thus, for each observation I have consumption data for March in year $t-1$ and March in year t.

I begin by narrowing my sample to observations satisfying the following criteria: (i) the identities of the household head and spouse (if any) do not change in year $t-1$ or year t; (ii) the household consists only of immediate family in year $t-1$ and year t; (iii) consumption data are reported by the household, rather than imputed by the PSID, in years $t-1$ and t; (iv) the household head is employed at the time of the $t-1$ interview; and (v) the head's job at $t-1$ is covered by a union contract. These criteria narrow the sample to 4,381 observations.

I next attempt to assign as many of these observations as possible to particular union contracts. Successful matches fall into two categories. First, some heads work in industries subject to national or regional pattern bargaining. National pattern industries include trucking, the postal service, and railroads. Regional pattern industries include lumber in the Pacific Northwest, shipping on the East Coast, and construction. I identify pattern-bargaining industries using the monthly Bureau of Labor Statistics publication *Current Wage Developments*.

Second, some household heads work in industries with a dominant local employer. For instance, an automobile worker living in Genesee County, Michigan, can be assigned to GM in Flint with reasonably high confidence; other examples include public employees living in metropolitan areas, such as a policeman living in the Bronx. Table 1 lists additional examples of sample observations linked to dominant local employers. The existence of dominant local employers is verified using the Census Bureau's annual *County Business Patterns*, which reports employment and the number of establishments by four-digit industry for every county in the United States. The actual identity of dominant local employers is determined using *Current Wage Developments*, *Ward's Manufacturing Directory*, the *United States Manufacturers Directory*, and various *Moody's* manuals, all of which have information on the location of particular establishments. Further details are provided in Shea (1992).

Once households have been matched to particular unions, I use contract information to construct a measure of the head's expected wage growth. Given the timing of the consumption data, I measure the head's expected wage growth between the end of March in year $t-1$ and the end of March in year t. I assume that expectations use all information available at the end of February of year $t-1$; to be included in the final sample, an observation ending in year t must be assigned to a contract that is negotiated before the end of February in year

[4]The PSID did not ask food-consumption questions in 1988.

[5]Some of these observations are duplicates. The 1987 PSID tape consists of retrospective histories of all households interviewed by the PSID in 1987, including children or spouses who "split off" from another "source" household at some time t between 1981 and 1987. Data for such split-off households before t duplicates the data for the source household. I eliminate such duplicate observations when selecting my sample and when computing summary statistics for the complete PSID random subsample.

[6]The PSID gathers consumption data in the following manner. First, the household is asked about the bonus value of food stamps received *last month*. Immediately thereafter, the household is asked about its typical weekly grocery and restaurant expenditure *beyond the value of these food stamps*. See Zeldes (1989) for more discussion.

TABLE 1—EXAMPLES OF OBSERVATIONS ASSIGNED TO DOMINANT LOCAL EMPLOYERS

Industry	County and state	Assigned employer
Steel mills	Baltimore City, MD	Bethlehem Steel, Sparrows Point
Radio, TV, and phone equipment	Marion, IN	Western Electric, Indianapolis
Aircraft	St. Louis City, MO	McDonnell Douglas, St. Louis
Shipbuilding	Los Angeles, CA	Todd Shipyards, San Pedro
Paper products	Chatham, GA	Union Camp, Savannah
Meat packing	Minehaha, SD	John Morrell, Sioux Falls
Automobiles	Onondaga, NY	Chrysler, East Syracuse
Electric utilities	Harris, TX	Houston Lighting and Power
Street rails and bus lines	Fulton, GA	Metropolitan Atlanta Rapid Transit Authority
Local government (policeman)	Lucas, OH	Police and Fire, Toledo

$t-1$ and is scheduled to expire after the end of March in year t.[7]

For example, most workers in the U.S. Postal Service reached agreements in July 1981 which extended until July 1984. Postal-worker observations ending in 1982 or 1985 are excluded from my sample; in these cases, a new contract was scheduled to be negotiated between $t-1$ and t, making it difficult to infer the head's expected wage growth at $t-1$. However, observations ending in 1983 or 1984 are eligible for sample inclusion; in these cases, information from the ongoing contract can be used to measure expected wage growth.

I take dates and wage provisions of particular contracts from *Current Wage Developments* and the Bureau of National Affairs' *Government Employee Relations Reporter*. The former typically reports only settlements covering more than 1,000 workers, while the latter reports public-employee settlements for states and large cities. Unfortunately, some observations were dropped from the sample due to lack of contract information.

To summarize, my final sample consists of observations satisfying criteria (i)–(v) above, as well as the following: (vi) the household head can be assigned to a particular union contract with high confidence; (vii) contract information is readily available from published sources, and (viii) the contract is negotiated by the end of February in year $t-1$ and is scheduled to expire after the end of March in year t. I also eliminate observations in which any element of consumption at $t-1$ or t is top-coded (at 9,999) in the PSID; in which any element of consumption grows by more than 5,000 in nominal terms between $t-1$ and t; or in which real consumption grows or shrinks by more than a factor of three between $t-1$ and t. These restrictions are designed to eliminate obvious cases of measurement error from the sample.

The final sample consists of 647 observations, drawn from 285 households. Pattern bargains account for 374 observations, while dominant local employers account for the remaining 273 observations. Additional information about the sample is provided in Shea (1992).

[7] I require only that the contract be *scheduled* to expire after March 31 of year t; I do not require that the contract actually be carried out. Some contracts in my samples are unexpectedly renegotiated between March of year $t-1$ and March of year t; similarly, some household heads in my sample retire, are laid off, are disabled, or switch industries between years $t-1$ and t. Including such observations reduces the correlation between expected and actual wage growth and thus reduces power. However, including these observations should not bias my results under the LCH/PIH null, while excluding these observations might subject my study to sample-selection bias by censoring the distribution of bad news arriving between $t-1$ and t.

II. Empirical Results

A. *Specification and Variable Definition*

This section presents results from testing the LCH/PIH using the sample described above. I employ the following specification, for which theoretical justification can be found in Zeldes (1989) and elsewhere:

$$(2)\quad \mathrm{GC}_t = \alpha + \gamma\,\mathrm{GAFN}_t + \sigma \log(1 + r_t) + \beta Z_{t-1} + \varepsilon_t.$$

GC represents the growth rate of real food consumption between $t-1$ and t. Consumption equals the value of food consumed at home and in restaurants, plus the bonus value of food stamps.[8] The PSID reports nominal data; I convert food at home and food stamps to 1982 dollars using the March CPI-W for food consumed at home, and I convert restaurant food using the March CPI-W for food consumed away from home.

Following Zeldes (1989) and others, I control for predictable taste shifts induced by changes in family composition. GAFN represents the growth rate of the household's annual food needs; the index of annual food needs is constructed by the PSID using information on household composition along with information on nutritional needs by age and sex. I also control for the expected real interest rate r_t. Because inflation is unpredictable, the real return is not known to the household at $t-1$. Accordingly, most researchers include the *ex post* return in (2) and instrument using lagged *ex post* returns. Instead, I measure the expected March-to-March real return directly:

$$r_t = (1 - \mathrm{MTR}) \times (i_{t-1}) - E_{t-1}(\Pi_t)$$

where MTR is the household's marginal tax rate reported at $t-1$, i_{t-1} is the average nominal secondary-market one-year Treasury bill yield in February of year $t-1$, and $E_{t-1}(\Pi_t)$ is a measure of expected March-to-March inflation generated from naive rolling AR(12) forecasts of the CPI-W inflation rate; this forecast is described further in the Data Appendix.

The chief instrument Z_{t-1} is EDWAGE, the expected growth rate of the head's real straight-time wage or salary between April 1 of year $t-1$ and March 31 of year t. Expected nominal wage growth includes guaranteed raises and expected cost-of-living adjustments (COLA's), computed using contract COLA provisions and the inflation forecasts described above (see the Data Appendix for further details). Expected real wage growth equals expected nominal wage growth minus forecast March-to-March inflation.

Provided that my matches of households to contracts are accurate, EDWAGE should be a good predictor of actual wage growth between $t-1$ and t. Of course, instrument relevance should never be accepted on faith alone. The head's nominal wage is available at $t-1$ and t for 401 sample observations. For these observations, I regress actual real wage growth on EDWAGE and the other right-hand-side variables in (2). The coefficient on EDWAGE is 0.857, with a t statistic of 3.826.[9] Thus, EDWAGE seems relevant enough to provide a powerful test of the LCH/PIH.

Notice that estimates of (2) measure the elasticity of consumption with respect to predictable growth in the head's wage. Other papers in the literature instead estimate the elasticity of consumption with respect to predictable household income growth. Expected wage and income growth may differ for two reasons: first, expected income growth includes expected movements in hours as well as wages; and second, labor income of the head is only part

[8] My measure of food stamps is the reported bonus value of stamps received last month. This variable must be multiplied by 12 to be comparable to other PSID consumption data.

[9] Except where noted, all standard errors and t statistics reported in this paper are robust to heteroscedasticity of unknown form and to arbitrary serial correlation of errors within households; see footnote 10.

TABLE 2—SUMMARY STATISTICS

Variable	Mean	Standard deviation
GC	0.015	0.310
GAFN	−0.002	0.120
$\log(1+r_t)$	0.024	0.018
EDWAGE	0.006	0.025
INCOME	27,960	10,200
WEALTH	8,289	24,705

Notes: Number of observations = 647. This table presents sample means and standard deviations for all variables in my analysis. The first four variables are explained in the text. INCOME is disposable income in 1982 dollars for calendar year $t-2$, as reported in March $t-1$. WEALTH is estimated liquid wealth in 1982 dollars for calendar year $t-2$.

of household income. In principle, I could estimate the income elasticity of consumption by including income growth in (2) and instrumenting using EDWAGE. Unfortunately, data timing in the PSID makes this strategy impossible. While GC and EDWAGE are March-to-March, PSID income data reported in March of year t refers to the entire previous calendar year $t-1$. In my sample, regressing household disposable income growth between years $t-2$ and $t-1$ on EDWAGE and the other variables included in (2) yields a coefficient on EDWAGE of 0.304 with a t statistic of only 0.864; results are similar if I use income growth between years $t-1$ and t. Thus, due to timing mismatch, EDWAGE is insufficiently correlated with income growth to estimate reliably the sensitivity of consumption with respect to predictable income.

Finally, I discuss the constant term. Gary Chamberlain (1984) notes that estimates of (2) are not consistent for short time periods if ε_t contains an aggregate component. I therefore include time dummies in all regressions, although results are similar if only one constant is included. These time dummies should remove any systematic mismeasurement of expected real wage growth or the expected after-tax return caused by incorrect specification of inflation expectations. Thus, I primarily test whether *cross-sectional* variation in expected wage growth leads to cross-sectional variation in consumption growth.

Table 2 presents descriptive statistics for my sample. Real household wealth and disposable income in 1982 dollars are constructed using data for calendar year $t-2$, as reported in March of year $t-1$; the construction is described in the Data Appendix. In the working version of the paper, I discussed differences between the composition of my sample and of the U.S. population. While my sample is somewhat younger, less educated, less wealthy, and has more blacks than the overall population, I showed in the working version that these differences are unlikely to have a large impact on my empirical results.

B. *Tests of the LCH/PIH*

I report the central result of this paper in column (i) of Table 3, which estimates equation (2) setting Z_{t-1} equal to expected wage growth. Standard errors are robust to heteroscedasticity of unknown form and to arbitrary serial correlation of disturbances within households.[10] The LCH/PIH is rejected; the coefficient on expected wage growth is positive and significant at the 10-percent level (and nearly significant at the 5-percent level). Changes in household

[10]These standard errors are computed as follows. Suppose there are I households in the sample, indexed by i. Let $\mathbf{X}_i$ denote the matrix of right-hand-side variables for household i; this matrix has dimension $T_i \times k$, where T_i is the number of observations for household i and k is the number of right-hand-side variables. Finally, let $\hat{\varepsilon}_i$ denote the vector of estimated disturbance terms for household i. Then, the estimated variance-covariance matrix equals

$$\left[\sum_{i=1}^{I}\mathbf{X}_i'\mathbf{X}_i\right]^{-1}\left[\sum_{i=1}^{I}\mathbf{X}_i'\hat{\varepsilon}_i\hat{\varepsilon}_i'\mathbf{X}_i\right]\left[\sum_{i=1}^{I}\mathbf{X}_i'\mathbf{X}_i\right]^{-1}.$$

I thank Joe Altonji for suggesting this method. Results are similar if standard errors are robust to heteroscedasticity alone.

TABLE 3—TESTS OF THE LIFE-CYCLE/PERMANENT-INCOME HYPOTHESIS

GC_t = Time Dummies + $\gamma GAFN_t + \sigma \log(1+r_t) + \beta Z_{t-1} + \varepsilon_t$

Variable	(i)	(ii)	(iii)	(iv)	(v)
GAFN	0.290	0.294	0.289	0.294	0.288
	(0.141)	(0.146)	(0.143)	(0.145)	(0.142)
	[2.057]	[2.014]	[2.021]	[2.028]	[2.028]
$\log(1+r_t)$	1.556	1.847	1.139	1.780	1.301
	(1.418)	(2.358)	(1.810)	(1.436)	(1.489)
	[1.104]	[0.783]	[0.629]	[1.195]	[0.865]
Z_{t-1}	0.888	0.008	−0.013	0.030	−0.021
	(0.455)	(0.045)	(0.034)	(0.040)	(0.031)
	[1.952]	[0.178]	[0.382]	[0.750]	[−0.677]
R^2:	0.023	0.019	0.019	0.019	0.019
Number of observations:	647	647	647	647	647

Notes: This table presents estimates of equation (2). The dependent variable is real consumption growth between March of year $t-1$ and March of year t. Each column reports results using a different instrument Z_{t-1}:

(i) expected real wage growth (EDWAGE),
(ii) log real disposable income for year $t-2$,
(iii) log real disposable income for year $t-3$,
(iv) growth rate of disposable income from $t-3$ to $t-2$,
(v) log real consumption in March of year $t-2$.

Each regression includes six year dummies. Standard errors (in parentheses) and t statistics (in brackets) are robust to heteroscedasticity of unknown form and to arbitrary serial correlation of disturbances within households. Further details can be found in the text and Data Appendix.

food needs have a significant effect on consumption, as in Zeldes (1989). The interest elasticity of consumption is positive but insignificant. The failure to estimate the interest elasticity precisely is not surprising given that I identify this elasticity using only cross-sectional variation in marginal tax rates, and given that cross-sectional variation in marginal tax rates was reduced by the tax cuts of 1981.

In fairness, I must point out two unattractive features of the results in column (i). First, the coefficient on EDWAGE is imprecise, in part because EDWAGE itself has low variation; from Table 2, the sample standard deviation of EDWAGE is only 2.5 percent.

Second, the point estimate of β is implausibly large. An extreme Keynesian alternative to the LCH/PIH holds that households myopically consume current income. Under this alternative, consumption should rise one-for-one with predictable income. However, *food* consumption should rise less than one-for-one, due to Engel's law. Furthermore, the elasticity of income with respect to the wage may also be less than 1, since the head's wages are only one component of household income.[11]

Under a Keynesian alternative to the LCH/PIH, one would expect estimated β to be comparable to the cross-sectional elasticity of consumption with respect to wages. I estimate this elasticity in three steps. First, there are 464 observations for which I observe the head's wage in March of year $t-1$. For these observations, I regress the log of the head's real labor income for calendar year $t-1$ on six year dummies and the log real wage. The wage coefficient is 1.032 and is highly significant. This suggests

[11] In my sample, labor income of the head comprises 80.7 percent of total family income on average.

that labor income is a good proxy for wages in my sample. Second, I estimate the cross-sectional elasticity of household disposable income with respect to the head's labor income. I regress log disposable income in year $t-1$ on six year dummies and log labor income in year $t-1$; in order to reduce bias from measurement error, I instrument for labor income using labor income in year $t-3$.[12] The labor-income coefficient is 0.907 and is highly significant. Finally, I estimate the cross-sectional elasticity of food consumption with respect to disposable income. I regress log real food consumption in March of year $t-1$ on six year dummies, the log of the annual-food-needs index for year $t-1$, and log real disposable income for year $t-1$. In order to isolate the "permanent" component of income, I instrument for income using real income in year $t-3$; this should correct for transitory income fluctuations due to measurement error or temporary shocks.[13] The income coefficient is 0.363 and is highly significant. Putting these results together, the cross-sectional elasticity of consumption with respect to wages in my sample is roughly 0.363 times 0.907, or 0.329. Thus, even under an extreme Keynesian alternative to the LCH/PIH, estimated β should be no higher than about one-third, rather than almost 1.

For the sake of comparison, I also estimate equation (2) using Runkle's (1991) instruments. In column (ii) of Table 3, I set Z_{t-1} equal to log real disposable income for calendar year $t-2$. In column (iii), I use income for year $t-3$; in column (iv) I use income growth between $t-3$ and $t-2$; in column (v), I use log real consumption in March of year $t-2$.[14] The LCH/PIH is not rejected for any of Runkle's instruments; other parameter estimates are similar to those reported in column (i).

The contrast between column (i) and columns (ii)–(v) reveals the advantages and disadvantages of this paper's approach to testing the LCH/PIH. EDWAGE has a low variance relative to measured consumption growth; its coefficient is thus estimated imprecisely. However, EDWAGE is strongly correlated with actual wage growth; this relevance enables me to reject the LCH/PIH despite low precision. In contrast, Runkle's instruments are highly variable and thus have precisely estimated coefficients. However, these instruments are weakly correlated with future income growth, and thus have little power to reject the LCH/PIH.[15] Ideally, one should study household consumption using instruments that are both highly relevant and variable enough to produce precise estimates. My results suggest that, in the absence of ideal instruments, one may learn more by sacrificing precision to obtain a highly relevant instrument than by using highly variable yet irrelevant instruments.

C. *Tests for Liquidity Constraints*

In Table 4, I examine whether liquidity constraints can explain the failure of the

[12] Results are similar if I use labor income at $t-2$ or labor income at both $t-2$ and $t-3$ as instruments.

[13] Results are similar if I use disposable income at $t-2$ or disposable income at both $t-2$ and $t-3$ as instruments; results are also similar if I estimate "permanent income" as a two- or three-year average of income instead of as a projection of $t-1$ income on past income.

[14] I do reject the LCH/PIH when I use log $t-1$ consumption as an instrument ($\hat{\beta}=-0.285$; t statistic $=-7.500$). As Runkle (1991) points out, however, this instrument is invalid if there is serially uncorrelated measurement error in the level of consumption.

[15] To check the Runkle instruments for relevance, I regress the growth rate of disposable income between $t-2$ and $t-1$ on six year dummies, the Runkle instruments (one at a time), and the other right-hand-side variables in (2). Neither income at $t-3$ nor consumption at $t-2$ is a significant predictor of income growth between $t-2$ and $t-1$. I do find that $t-2$ income (coefficient $=-0.182$; t statistic $=2.984$) and income growth between $t-3$ and $t-2$ (coefficient $=-0.195$; t statistic $=3.250$) are significant predictors of income growth between $t-2$ and $t-1$; however, this could be due to serially uncorrelated measurement error in the level of income, rather than true negative first-order autocorrelation in income.

TABLE 4—TESTS FOR LIQUIDITY CONSTRAINTS

GC_t = Time Dummies + $\gamma GAFN_t + \sigma \log(1+r_t) + \beta(EDWAGE) + \varepsilon_t$

Variable	(i)	(ii)	(iii)	(iv)	(v)
GAFN	0.154	0.486	0.226	0.388	0.294
	(0.192)	(0.121)	(0.187)	(0.093)	(0.143)
	[0.802]	[4.017]	[1.209]	[4.172]	[2.056]
$\log(1+r_t)$	0.631	2.702	0.756	4.748	1.720
	(2.098)	(2.007)	(1.675)	(2.876)	(1.407)
	[0.301]	[1.346]	[0.451]	[1.651]	[1.222]
EDWAGE	0.997	0.765	0.961	0.867	—
	(0.703)	(0.539)	(0.614)	(0.593)	
	[1.418]	[1.419]	[1.565]	[1.462]	
Positive EDWAGE	—	—	—	—	0.063
					(0.785)
					[0.080]
Negative EDWAGE	—	—	—	—	2.242
					(0.951)
					[2.358]
R^2:	0.010	0.067	0.014	0.121	0.011
Number of observations:	372	275	475	172	647
Split:	Zero wealth	Positive wealth	Low ratio	High ratio	Full sample

Notes: This table presents estimates of equation (2) designed to test for liquidity constraints. In columns (i) and (ii) the sample is split into households reporting zero and positive wealth at $t-1$. In columns (iii) and (iv) the sample is split into households reporting a ratio of wealth to disposable income less than and greater than two months at $t-1$. Column (v) uses the full sample but splits expected wage changes into increases and decreases. In all columns, the dependent variable is real consumption growth between March of year $t-1$ and March of year t. Each regression includes six year dummies. Standard errors (in parentheses) and t statistics (in brackets) are robust to heteroscedasticity of unknown form and to arbitrary serial correlation of disturbances within households. Further details can be found in the text and Data Appendix.

LCH/PIH in my sample.[16] As Zeldes (1989) points out, households with liquid wealth can dissave to smooth consumption when income is temporarily low, even if borrowing against future income is difficult. Thus, if the LCH/PIH fails because of liquidity constraints, wealthy households should obey the LCH/PIH, while low-wealth households should not. Following Zeldes (1989) and Runkle (1991), I reestimate (2) splitting the sample according to the household's wealth at $t-1$; the measure of wealth is similar to that employed by Zeldes (1989) and is discussed in the Data Appendix. In columns (i) and (ii) I examine households with zero and positive liquid wealth, respectively; in columns (iii) and (iv) I examine households whose ratio of wealth to disposable income is less than and greater than two months, respectively. The results are

[16] In the working version of this paper, I also examined whether near-rationality (as discussed by Cochrane [1989]) could explain the failure of the LCH/PIH in my sample. I found that the utility losses implied by my estimate of β were small when evaluated at the sample standard deviation of EDWAGE, consistent with near-rationality. However, I also found that consumption is more sensitive to large absolute expected wage changes than to small expected wage changes, inconsistent with near-rationality or with small fixed costs of consumption adjustment.

mildly supportive of liquidity constraints. The sensitivity of consumption to expected wage growth is larger for low-wealth than for high-wealth households, although the difference between subsamples is small.

It is not clear that such sample splits can discriminate between liquidity constraints and other alternatives to the LCH/PIH. Suppose, for instance, that the LCH/PIH fails because some households are myopic. Suppose, too, that intelligence is positively correlated with the tendency to be forward-looking. If intelligence is also positively correlated with success, then rich households may be more likely to obey the LCH/PIH than poor households even if there are no liquidity constraints; in this case, wealth is merely a proxy for forward-looking behavior.

Fortunately, liquidity constraints have an additional testable implication, as pointed out by Altonji and Siow (1987): the response of consumption to predictable income should be asymmetric. If the LCH/PIH fails because of liquidity constraints, then households will be more likely to violate the LCH/PIH when income is expected to grow than when income is expected to fall, since liquidity constraints inhibit borrowing but not saving; more precisely, liquidity constraints will not cause the Euler equation between adjacent periods to fail as long as optimal frictionless consumption growth exceeds expected income growth.[17] Accordingly, I reestimate equation (2) for the full sample, allowing the coefficients on positive and negative expected wage changes to differ.[18]

The results, presented in column (v) of Table 4, indicate that consumption is significantly correlated with predictable wage *declines*, but not with predictable wage increases. This result is inconsistent with liquidity constraints, which predict the opposite asymmetry. This result is also inconsistent with a simple myopic alternative, which predicts a symmetric relationship between consumption and predictable income. Moreover, a similar asymmetry can be found in aggregate U.S. time-series data (Shea, 1995): aggregate consumption responds more strongly to predictable declines in aggregate income than to predictable increases in income. Thus, the findings in column (v) do not seem to be a fluke. On the other hand, my results do contradict Altonji and Siow (1987), who find that households expecting income to rise exhibit a higher sensitivity of consumption to predictable income than households expecting income to fall (although neither sensitivity is significantly different from zero).

Instrument relevance may be the key to explaining this contradiction. Altonji and Siow (1987) estimate (2) using as their instrument the projection of *ex post* household income growth on a variety of household-specific variables dated $t-1$ or earlier. Unfortunately, Altonji and Siow's most powerful instruments—lagged hours of unemployment and disability, and lagged wages interacted with dummy variables indicating past quits or layoffs—are far better predictors of income *increases* than declines. Thus, Altonji and Siow's (1987) tests may have little power to detect asymmetry. On the other hand, a substantial number of contracts in my sample call for long-term

[17] Angus Deaton (1991) formally investigates optimal consumption by a liquidity-constrained consumer under various stochastic income processes. For stationary income processes, he finds that the response of consumption to income is asymmetric; temporarily high income draws are smoothed by saving, but low income draws are not smoothed unless the household has wealth. For difference-stationary income processes with positive serial correlation in growth rates, he finds (counterfactually) that savings should actually rise at the beginning of recessions, since negative income growth today implies predictably declining income in the near future.

[18] For the full sample, expected real wage growth is positive for 412 observations, and negative for 235 observations. Empirical results are qualitatively similar if only a low-wealth subsample is used, rather than the full sample, or if all coefficients (not just β) are allowed to depend on whether expected wage growth is positive or negative.

nominal-wage freezes, which imply predictable real income declines for the affected households. Thus, I can arguably conduct a more powerful test for asymmetry than Altonji and Siow (1987).

Bowman et al. (1993) propose a model that is potentially consistent with the "perverse asymmetry" found here and in Shea (1995). In their model, preferences exhibit *loss aversion*: utility is concave when consumption increases above a reference level of consumption but is *convex* when consumption declines below reference consumption. This asymmetry captures the intuition that households suffer relatively large psychic losses when forced to cut living standards by even a small amount; such loss aversion has been documented in numerous experimental studies, surveyed in Daniel Kahneman et al. (1991). Bowman et al. show that loss-averse households may optimally refuse to reduce consumption today in the face of expected but uncertain declines in future income. Intuitively, households resist cutting back today as long as there is a chance that income will not fall tomorrow; the cost of not smoothing losses is smaller than the potential benefit of never having to reduce consumption below reference levels. Thus, loss aversion can potentially explain why consumption is more sensitive to predictable income decreases than increases.

I must emphasize that, at this stage, loss aversion is consistent with my results only in a broad qualitative sense. Bowman et al. analyze only a two-period model; results on the dynamic behavior of consumption under loss aversion with infinite horizons and general income processes are not yet available. Moreover, it seems unlikely that even a model with loss aversion could explain my *quantitative* results. I find an elasticity of consumption with respect to expected income declines in excess of 2, while loss aversion would seem to imply an elasticity of 1 for overall consumption (and, again, less than 1 for food consumption due to Engel's law) since, under loss aversion, consumption should fall one-for-one with income once the decline in income actually occurs. Once again, it appears that my elasticity estimates are implausibly high, although my estimates are, of course, imprecise.

D. *Can Bad Luck Explain My Results?*

One possible objection to my results is that my sample has too short a time frame to exploit orthogonality between expected wage growth and *ex post* shocks to permanent income. Suppose, for instance, that certain unionized industries or occupations experienced a string of consecutive negative surprises in the early 1980's. Then a household head with low expected wage growth (due to a concessionary agreement negotiated by her union following bad shocks in the past) might also tend to receive bad news *ex post*, causing expected wage growth and consumption growth to be correlated even if the LCH/PIH is true. In a long time series, of course, bad luck should even out; negative surprises should not be systematically serially correlated. But bad luck need not even out within six years.

I perform two tests to check whether serially correlated surprises are important in my sample. First, there are 401 sample observations for which wage rates are available at $t-1$ and t. For these observations, I regress the difference between actual real wage growth and EDWAGE on year dummies and EDWAGE. The estimated coefficient on EDWAGE is -0.143, with a robust standard error of 0.224.[19] Thus, low expected wage growth is not systematically associated with negative wage surprises in my sample. Second, there are 54 sample observations which experience observable nonwage negative shocks between $t-1$ and t. Probit estimates on the full sample (including year dummies) indicate that EDWAGE has an insignificant negative effect on the probability of being unemployed or laid off at t (t statistic $= -0.242$), an insignificant positive effect on the probability of switching industries between $t-1$ and t

[19] Of course, this finding is not surprising given that the regression of *ex post* wage growth on EDWAGE and the other right-hand-side variables in (2) produced a coefficient on EDWAGE near 1!

(t statistic = 0.103), and a marginally significant positive effect on the probability of experiencing a contract renegotiation between $t-1$ and t (t statistic = 1.828); since all renegotiations in my sample involve union concessions, this last effect is of the wrong sign to generate spurious rejection of the LCH/PIH.[20] Thus, low expected wage growth does not systematically signal negative nonwage shocks in my sample. It does not appear that bad luck can explain my findings.[21]

E. *Can Good Luck Explain My Results?*

Another potential objection to my results is that my data give me an unusually high chance of being "lucky." Recall that EDWAGE has a low standard deviation, while consumption growth has a high standard deviation, due in large part to measurement error. Intuitively, it seems that false rejections of the LCH/PIH might be likely in this setting; a few households with large consumption measurement errors whose EDWAGE happens to have the same sign might be enough to generate a type-I error. More formally, perhaps the usual t distribution understates the actual size of my test in finite samples, given the unusual properties of my data.

To investigate this possibility, I perform the following Monte Carlo experiment. For each trial, I assign each of the 647 observations in my sample a randomly generated rank-order number. I then sort the variable EDWAGE according to this rank-ordering. The resulting sorted variable, which I denote EDWAGEX, is a "scrambled" version of EDWAGE; EDWAGEX has the same distribution as EDWAGE, but the actual value assigned to a particular observation will not equal EDWAGE except by sheer luck. I then estimate equation (2) for the full sample using EDWAGEX as the instrument, and save the t statistic for $\hat{\beta}$. I repeat this procedure 1,000 times and examine the resulting empirical distribution of $t(\hat{\beta})$. Since EDWAGEX is irrelevant to future income, its true coefficient in (2) is zero whether or not the LCH/PIH is true. On the other hand, if the probability of a type-I error is unusually high in my sample, then the empirical distribution of $t(\hat{\beta})$ should exhibit a high number of false rejections.

Results are reported in Table 5. Column (i) shows the percentiles of the asymptotic t distribution, while column (ii) shows the percentiles of the empirical distribution of $t(\hat{\beta})$. The results suggest that my data are not especially prone to type-I errors. The empirical distribution is slightly fatter in the tails than the asymptotic distribution, but the difference is very small. The fraction of trials in which t is greater than 1.645 is only 5.4 percent, and the fraction of trials in which t exceeds 1.960 is only 3 percent. It does not appear, then, that "good luck" is a particularly compelling explanation of my results.

TABLE 5—MONTE CARLO RESULTS

Percentile	Asymptotic t distribution (i)	Empirical t distribution (ii)
1	−2.326	−2.552
5	−1.645	−1.706
10	−1.282	−1.399
50	0.000	0.055
90	1.282	1.330
95	1.645	1.708
97.5	1.960	2.021
99	2.326	2.330

Notes: This table presents results from the Monte Carlo experiment described in the text. Column (i) presents the percentiles of the asymptotic t distribution. Column (ii) presents the percentiles of the empirical t distribution generated by estimating equation (2) using the scrambled variable EDWAGEX as the instrument. The empirical distribution is based on 1,000 trials.

[20] These t statistics are not robust to heteroscedasticity or serial correlation within households.

[21] It is possible that bad news received between $t-1$ and t could affect observable outcomes only with a long lag, so that serially correlated negative shocks would not be detected by my tests. To investigate this possibility, I regressed leads 1 through 3 of the head's labor-income growth on year dummies and EDWAGE; since I have data only through 1987, this required leaving out observations taken from later years in the sample. EDWAGE was not a significant predictor of future labor-income growth at any lead. Details are available from the author upon request.

III. Conclusion

This paper isolates PSID households that can be matched to particular long-term union contracts with high confidence and tests the life-cycle/permanent-income hypothesis by using published contract provisions to construct a household-specific measure of predictable wage growth. While my estimates of the elasticity of consumption with respect to expected wage growth are imprecise and implausibly high, my results still advance our knowledge of household consumption behavior. First, expected wage growth is significantly correlated with consumption growth, contrary to the LCH/PIH. Second, other instruments commonly employed in the literature are not correlated with consumption growth in my sample, suggesting that these instruments' inability to predict future income may have caused previous tests to have low power. Third, and most strikingly, consumption is more sensitive to predictable income declines than increases. This asymmetry is inconsistent with myopia and liquidity constraints but is potentially consistent with models involving loss aversion. Further research should investigate the implications of loss aversion for the dynamic behavior of consumption in more general settings and should attempt to derive additional testable implications of loss aversion beyond asymmetric rejection of the LCH/PIH.

Data Appendix

Inflation Expectations.—The construction of expected real wage growth and the expected after-tax real return requires a measure of agents' expected inflation between March of year $t-1$ and March of year t, conditional on information available at the end of February in year $t-1$. I assume that agents form their inflation expectations using an AR(12) forecasting model of the monthly percentage change in the BLS CPI-W. Expectations as of February in year $t-1$ are formed using data available through December of year $t-2$, and coefficients are estimated using monthly data from February 1967 through December of year $t-2$. My forecasting equation does not include deterministic terms. Experiments using deterministic terms (seasonal dummies and a time trend) led to large overpredictions of inflation for every year in the sample period.

Expected Real Wage Growth.—I first compute each household's expected nominal straight-time wage change between April 1 ($t-1$) and March 31 (t). Expected nominal wage changes equal guaranteed raises plus expected cost-of-living adjustments. In some cases, *Current Wage Developments* reports a range of guaranteed raises (e.g., a raise of 20–52 cents based on experience). I assume that all workers with 60 months or more tenure at their current employers at $t-1$ get the maximum raise; that workers with 24 months or less tenure at $t-1$ get the minimum; and that workers with intermediate tenure get the midpoint. This scheme is in line with the practice at AT&T, the only contract for which I could actually determine the pattern of raises paid by experience. Results were not sensitive to alternative rules in any event. Expected COLA's are computed using the formulas reported for each contract in *Current Wage Developments*, along with expected movements in the CPI-W generated using the inflation forecasts described above. I assume a two-month lag between inflation and COLA payments; for instance, a quarterly COLA payment scheduled for July 1 of year $t-1$ is computed on the basis of expected inflation between February of year $t-1$ and May of year $t-1$. Some contracts specify a price index other than the BLS CPI-W; in most cases, the alternative index is the BLS CPI-U. I use only CPI-W forecasts in computing expected COLA's; between March 1979 and December 1986, the correlation between the CPI-U and CPI-W monthly inflation rates was 0.98.

Next, I must convert expected nominal wage changes into percentage terms. In some cases, all expected raises are expressed in percentage terms in the contract, so that no conversion is necessary. In other cases I need a measure of the head's wage at $t-1$. Where available, I use the head's reported wage at $t-1$. If the $t-1$ wage is not available, I use wages reported at (in order of priority) $t-2$, t, $t-3$ or $t+1$, adjusted for actual wage changes in the interim as reported in *Current Wage Developments*, provided that the head has not changed jobs in the interim. If no such adjacent wages are available, I estimate the head's wage at $t-1$ as the average of all reported $t-1$ wages in the head's industry in the sample. Details of such imputations are available from the author upon request. Note that I deliberately avoid estimating $t-1$ wages using reported income at $t-1$ divided by reported hours at $t-1$. My aversion is based on two facts: (i) reported hours and earnings at $t-1$ refer to the entire year $t-2$, whereas I want an estimate of the wage at the time of interview at $t-1$, and (ii) reported hours information in the PSID is notoriously noisy.

Finally, I convert expected nominal wage growth to expected real wage growth using the forecast of March-to-March inflation described above.

Disposable Income.—Following Zeldes (1989), nominal household disposable income reported in March of year $t-1$ consists of total household money income received in year $t-2$, plus the bonus value of food stamps received in year $t-2$, minus federal income taxes paid by the head, wife, and other family members in year $t-2$, minus Social Security taxes paid by the head and wife in year $t-2$. All components except Social Security taxes are available in the PSID. Social

Security taxes for, say, the head are computed by multiplying the Social Security tax rate by the maximum of the head's labor income and the Social Security ceiling income. Ceilings and tax rates were taken from annual supplements to the *Social Security Bulletin*. Nominal disposable income is converted to 1982 dollars using the annual average of the CPI-W.

Wealth.—As in Zeldes (1989), household liquid wealth is estimated by dividing asset income reported in March of year $t-1$ (which refers to calendar year $t-2$) by an estimate of the rate of return for the year $t-2$. Asset income consists of dividends, interest, rent, trust funds, and royalties for the head and wife. To determine nominal wealth, the first $250 of asset income is divided through by the maximum passbook savings rate at commercial banks; additional asset income is divided by the annual average three-month Treasury bill secondary-market yield. Nominal wealth is converted to 1982 dollars using the annual average of the CPI-W.

REFERENCES

Altonji, Joseph G. and Siow, Aloysius. "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data." *Quarterly Journal of Economics*, May 1987, *102*(2), pp. 293–328.

Bowman, David; Minehart, Debby and Rabin, Matthew. "Loss Aversion in a Savings Model." Mimeo, University of California, Berkeley, 1993.

Campbell, John Y. and Mankiw, N. Gregory. "Permanent Income, Current Income, and Consumption." *Journal of Business and Economic Statistics*, July 1990, *8*(3), pp. 265–79.

Carroll, Christopher D. "How Does Future Income Affect Current Consumption?" *Quarterly Journal of Economics*, February 1994, *109*(1), pp. 111–47.

Chamberlain, Gary. "Panel Data," in Zvi Griliches and Michael Intriligator, eds., *Handbook of econometrics*, Vol. 2. Amsterdam: Elsevier, 1984, pp. 1247–1313.

Cochrane, John H. "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives." *American Economic Review*, June 1989, *79*(3), pp. 319–37.

______. "A Simple Test of Consumption Insurance." *Journal of Political Economy*, October 1991, *99*(5), pp. 957–76.

Deaton, Angus. "Saving and Liquidity Constraints." *Econometrica*, September 1991, *59*(5), pp. 1221–48.

Flavin, Marjorie. "The Joint Consumption/Asset Demand Decision: A Case Study in Robust Estimation," National Bureau of Economic Research (Cambridge, MA) Working Paper No. 3802, August 1991.

Hall, Robert E. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, December 1978, *86*(6), pp. 971–87.

Kahneman, Daniel; Knestch, Jack L. and Thaler, Richard H. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives*, Winter 1991, *5*(1), pp. 193–206.

Runkle, David E. "Liquidity Constraints and the Permanent-Income Hypothesis: Evidence from Panel Data." *Journal of Monetary Economics*, February 1991, *27*(1), pp. 73–98.

Shapiro, Matthew D. "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data." *Economics Letters*, 1984, *14*(1), pp. 93–100.

Shea, John. "Union Contracts and the Life Cycle–Permanent Income Hypothesis, Appendix: Sample Selection and Description." Mimeo, University of Wisconsin, 1992.

______. "Should We Test the Life Cycle–Permanent Income Hypothesis with Food Consumption Data?" *Economics Letters*, May 1994, *45*(1), pp. 63–68.

______. "Myopia, Liquidity Constraints, and Aggregate Consumption: A Simple Test." *Journal of Money, Credit, and Banking*, 1995 (forthcoming).

Zeldes, Steven P. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy*, April 1989, *97*(2), pp. 305–46.