

“The Reaction of Household Consumption to Predictable Changes in Social Security Taxes.” Parker, Jonathan A. *American Economic Review* Vol. 89, No. 4 (1999): 959-973. <http://doi.org/10.1257/aer.89.4.959>

Copyright & Permissions

Copyright © 1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017 by the American Economic Association.

Permission to make digital or hard copies of part or all of American Economic Association publications for personal or classroom use is granted without fee provided that copies are not distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation, including the name of the author. Copyrights for components of this work owned by others than AEA must be honored. Abstracting with credit is permitted.

The author has the right to republish, post on servers, redistribute to lists and use any component of this work in other works. For others to do so requires prior specific permission and/or a fee. Permissions may be requested from the American Economic Association Administrative Office by going to the Contact Us form and choosing "Copyright/Permissions Request" from the menu.

Copyright © 2017 AEA

The Reaction of Household Consumption to Predictable Changes in Social Security Taxes

By JONATHAN A. PARKER*

This paper evaluates the key implication of rational expectations and the *basic* life-cycle/permanent-income hypothesis (LCH/PIH): that predictable changes in income have no effect on the growth rate of consumption expenditures.¹ This implication is important for understanding the effectiveness and optimal timing of fiscal policy, the causes and propagation of business cycles, and the effects of income fluctuations on the growth rate of the economy.

Using household-level consumption data from the Consumer Expenditure Survey (CEX), this paper tests whether expenditures on nondurable goods increase contemporaneously with predictable changes in Social Security tax withholding.² Individuals with wage and salary income earned in the United States are subject to Social Security tax withholding of around 7 percent of their gross pay up to an annual maximum income level. The structure of the Social Security tax system provides two sources of variation. First, a series of preannounced tax

rate increases occurred in the 1980's. Since the share of after-tax labor income in total income differs across households and since some individuals are not subject to Social Security tax withholding, these changes produce different percentage changes in income for different households. Second, when an individual's income earned in a calendar year reaches the maximum taxable amount, that individual's take-home pay increases because Social Security taxes are no longer withheld from his or her paycheck. In January of the following year, when withholding begins again, take-home pay falls again. Table 1 shows the rates and caps for the time period used in this paper.

Under the null hypothesis, household consumption should not respond to these changes in tax rates since they are expected. The small changes in tax rates are legislated and announced well ahead of time. High-income individuals who hit the tax cap generally see fluctuations in their after-tax wages year after year and, after 1982, the tax cap is adjusted for average wage growth.³ In the CEX sample employed, the average annual wage and salary income among heads of households who hit the tax cap is \$55,424. In the middle year of the sample, 1986, a person with this wage income would have a \$330 temporary increase in monthly after-tax income from early October until the end of December. The Social Security tax system thus provides predictable variations in income, variations to which individuals should not respond if they are smoothing consumption.

This study finds that households do change their consumption expenditures in response to the

* Department of Economics and Woodrow Wilson School, Princeton University, Princeton, NJ 08544. My thanks to Olivier Jean Blanchard, Ricardo Caballero, Charles Fleischman, Jonathan Gruber, Jerry Hausman, Steve Pischke, James Poterba, Joel Slemrod, Nicholas Souleles, seminar participants at MIT, Columbia University, Princeton University, Harvard University, Yale University, Stanford University, and anonymous referees for helpful comments. I am indebted to the National Science Foundation and the Sloan Foundation for financial support. Any errors remain my responsibility.

¹ As subsequently discussed, the theory predicts smoothing of marginal utility which does not generally imply smoothing of consumption. Here, the *basic* LCH/PIH is used to refer to the versions of the theory that imply consumption smoothing, as in the certainty-equivalent version originally employed by Robert E. Hall (1978), and a version which assumes constant expected variance of consumption, as discussed in Section I. Consumption smoothing is also consistent with both Franco Modigliani and Richard Brumberg (1956) and Milton Friedman (1957), although they did not impose rational expectations.

² The author is indebted to Joel Slemrod for suggesting the Social Security tax cap as a means of testing consumption smoothing.

³ Beginning in 1982, the maximum contribution is adjusted upward automatically based on the average annual percent wage change and then rounded to the nearest figure divisible by \$300. From 1990 to 1992 the rate of increase is accelerated because previous adjustments ignored nonwage and deferred compensation which had been growing more rapidly than wage compensation.

TABLE 1—THE SOCIAL SECURITY TAX STRUCTURE, 1980–1993

Year	Individual tax rate (percent)	Maximum annual contribution per earner	Maximum annual taxable earnings per earner
1980	6.13	1,588	25,900
1981	6.65	1,975	29,700
1982	6.70	2,171	32,400
1983	6.70	2,392	35,700
1984 ^a	6.70	2,533	37,800
1985	7.05	2,792	39,600
1986	7.15	3,003	42,000
1987	7.15	3,132	43,800
1988	7.51	3,380	45,000
1989	7.51	3,605	48,000
1990	7.65	3,924	51,300
1991			
OASDI:	6.20	3,311	53,400
HI:	1.45	1,812	125,000
1992			
OASDI:	6.20	3,441	55,500
HI:	1.45	1,888	130,200
1993			
OASDI:	6.20	3,571	57,600
HI:	1.45	1,958	135,000

Sources: Social Security Administration (1990) Tables 2.A3, 2.A4, and 2.A5 and the *Social Security Bulletin* (January 1993). Data refer to both Old Age Survivor and Disability Insurance (OASDI) and to Health Insurance (HI), except where noted.

^a The tax rate in 1984 includes the tax credit.

predictable fluctuations in income induced by the Social Security tax system. A predictable, 1-percent increase in after-tax income in a three-month interval contemporaneously increases expenditures on nondurable consumption by around a half of a percent. To put this in perspective, since nondurable consumption averages about 40 percent of income, expenditures on nondurable goods rise 20 cents for each dollar of predictable increase in income. Several steps are taken to reduce the possibility that the results are spurious or that differential seasonal patterns of consumption across wealth levels are driving the results. Similar but less statistically significant results are found among more homogeneous subsamples. Additionally, while not statistically significant, even larger point estimates are obtained when identification is derived from the differences in behavior between a treatment group of earners who have Social Security taxes withheld from their paychecks and a control group of earners who do not. Finally, motivated by models in which consumption is not smoothed across predictable fluctuations, the paper examines the consumption response at different levels of assets and for different categories of consumption.

In part, this work provides evidence against consumption smoothing by asking a different question than much of the previous literature. This work asks if households smooth expenditures across three-month periods. Most previous tests at the household level use annual data and identify expected income changes from either lower-frequency life-cycle movements or cross-sectional differences in income growth. It is possible that consumers smooth expenditures poorly across predictable income fluctuations at three-month intervals while smoothing expenditures well at lower frequencies. This test also leans more than previous works on high-income individuals who are continuously employed, a group that is less likely than average to be liquidity constrained, but also a group that has a high value for time.

In part, however, this test uncovers evidence against consumption smoothing by improving on most previous household-level tests.⁴ It does so in four main ways. First, the income changes caused by Social Security

⁴ The main methodology is exemplified by such papers as Matthew D. Shapiro (1984), Joseph Altonji and Aloyisius

withholding are calculated at the household level and are exogenous to the household. The predictable change in income is thus not highly correlated with labor supply or family size. Such a correlation reduces the power of many previous tests. Second, the predictable changes in income are large and easily identifiable. Income changes due to the tax cap are around 9 percent of net income. Measurement error in the income change is relatively small because the income change is calculated from the *level* of individual labor income, not its growth rate. Third, the CEX provides comprehensive data on many different types of consumption expenditures at the household level. Many previous microeconomic studies have employed only a (noisy) measure of food consumption. Finally, because this test has a large time-series dimension (167 three-month periods) and individual-level measures of change in income (including a control group), it avoids the poor performance of tests using short panels with common sources of expected change in income, as critiqued, for example, in Randall P. Mariger and Kathryn Shaw (1993).

These benefits are shared by some previous studies of consumption smoothing.⁵ James M. Poterba (1988), David W. Wilcox (1989), and Alec R. Levenson (1996) use fiscal experiments and aggregate consumption data and reach similar conclusions to those of this paper. At the household level, Ronald Bodkin (1959), Shapiro and Joel Slemrod (1995), John Shea (1995), Fumio Hayashi (1997 Ch. 1), and Nicholas S. Souleles (1999) all identify institutional features that deliver significant, plausibly exogenous, predictable variations in household income. The findings of these papers also are in general agreement with the findings of the present paper: that there are small but economically significant deviations from consumption smoothing. These findings stand in contrast to

the literature that employs lagged information to construct measures of predictable income movements, which often fails to reject consumption smoothing (e.g., Orazio P. Attanasio and Guglielmo Weber, 1995).

I. Consumption Smoothing

Consider a canonical consumption Euler equation:

$$(1) \quad u'(c_t)\nu(\mathbf{z}_t) = \beta R_t E_t[u'(c_{t+1})\nu(\mathbf{z}_{t+1})]$$

where E_t is the expectations operator; $u'(\cdot)$ is a marginal utility function, assumed decreasing; c_t is nondurable consumption; \mathbf{z}_t is a vector of deterministic variables that alter marginal utility through the function $\nu(\cdot)$; β is the discount factor; and R_t is the gross after-tax real interest rate between t and $t + 1$. One can derive a testable relationship from equation (1) by assuming that consumption is approximately log-normally distributed and the utility function exhibits constant relative risk aversion leading to the linear Euler equation:

$$(2) \quad \Delta \ln(c_{t+1}) = \sigma \ln(\beta R_t) + \sigma \ln\left(\frac{\nu(\mathbf{z}_{t+1})}{\nu(\mathbf{z}_t)}\right) \\ + \frac{1}{2\sigma} E_t[\text{Var}(\Delta \ln(c_{t+1}))] \\ + \varepsilon_{t+1}$$

where $\varepsilon_{t+1} \equiv \ln(c_{t+1}) - E_t[\ln(c_{t+1})]$ and σ is the intertemporal elasticity of substitution.⁶ If expectations are formed rationally and variations in the conditional variance term are unpredictable, then anything known to the household at time t beyond βR_t and $\nu(\mathbf{z}_{t+1})/\nu(\mathbf{z}_t)$, such as income fluctuations generated by hitting the Social Security tax cap, should not alter the growth rate of consumption in equation (2).

Equation (2) highlights the reasons that predictable changes in income might affect

Siow (1987), Stephen P. Zeldes (1989), David E. Runkle (1991), and Annamaria Lusardi (1996).

⁵ A full characterization of the empirical evidence on consumption smoothing is not possible here, and many excellent papers are not discussed. The reader interested in further discussion is referred to Angus Deaton (1992), Martin Browning and Lusardi (1996), and Parker (1997).

⁶ The parameter σ is also one over the coefficient of relative risk aversion. The same substantive implications also can be derived from a Taylor approximation to equation (1) for a general utility function, as in Karen Dynan (1993).

consumption. First, expectations might not be formed using all available information, so that ε_{t+1} is not uncorrelated with all things known at time t . Second, there might be predictable variations in the conditional variance term (the precautionary saving motive) so that marginal-utility smoothing does not imply consumption smoothing. Finally the Euler equation might not hold for some households, for example due to difficulties in borrowing at R_t .

The following specification is employed to test the linear Euler equation:

$$(3) \quad \Delta \ln(c_{h,t+1}) = \alpha'_1 \mathbf{z}_h + \alpha_2 E_t[\Delta y_{h,t+1}^{sst}] \\ + \alpha'_3 \mathbf{m}_t + \alpha'_4 \mathbf{d}_t + \varepsilon_{h,t+1}$$

where h indexes households; \mathbf{m} is a complete set of month dummies;⁷ \mathbf{d} is a complete set of year dummies less one; \mathbf{z} contains a second-order polynomial in family size in the second interview, a second-order polynomial in family size in the fifth interview, and a fourth-order polynomial in age; and $E_t[\Delta y_{h,t+1}^{sst}]$ is the percent increase in income due to Social Security taxes. According to the basic LCH/PIH, α_2 should equal zero. Note that estimating the equation in first differences removes any household-specific effects in the level of consumption. Further, the regression includes a complete set of month effects. Without the month dummies, the seasonal rise in consumption that occurs at the end of the calendar year could incorrectly be attributed to the tax variable which, on average, falls at the end of the calendar year.

Equation (3) then has three sources of variation which identify the effect of the changing tax rates on consumption. First, high-income individuals hit the Social Security tax cap and their after-tax income rises in different months depending on each individual's income.⁸ Second, for different households, Social Security taxes represent different percentages of after-tax income. For example, a household with two

earners in which one earner hits the tax cap undergoes a smaller change in its after-tax income than a household with only one earner in which that earner hits the cap. Finally, there are small changes in the Social Security tax rate across calendar years which, like the tax caps, are public knowledge well in advance of becoming effective.

The residual from equation (3) is serially correlated within households due to time averaging of consumption over three-month periods, measurement error, and possible random effects in growth rates. Thus, when estimating equation (3), standard errors are calculated to allow for arbitrary heteroskedasticity and within-household serial correlation.⁹

II. The Consumer Expenditure Survey

The data used to estimate equation (3) are constructed from the Family, Member, and Detailed Expenditure files of the CEX for the years 1980 to 1993. The CEX is a rotating panel of households, with new households entering every month. Data on families and individual members are extracted and merged to make an unbalanced, overlapping panel of households covering January 1980 to November 1993. Each household is interviewed five times. In the first interview, the CEX procedures are explained to the members of the household and they are asked to keep track of their expenditures for future interviews. Each household is subsequently interviewed four more times, once every three months. In each of these interviews, detailed information is collected on the past three months' consumption expenditures. In each family's second and fifth interviews, demographic and income data is collected, including income and earnings information about the previous 12 months. The CEX reports information on more

⁹ Let $\mathbf{X}_{h,t}$ represent the row vector of regressors for household h in period t and \mathbf{X}_h represent the $\mathbf{X}_{h,t}$ vectors vertically stacked. The variance-covariance matrix is estimated as $(\sum_h \mathbf{X}_h' \mathbf{X}_h)^{-1} (\sum_h \mathbf{X}_h' \varepsilon_h) (\sum_h \mathbf{X}_h \mathbf{X}_h')^{-1}$ where ε_h is the column vector of residuals for household h from estimation of equation (3). For two-stage least-squares regressions, the same formula is employed with \mathbf{X}_h replaced by $\tilde{\mathbf{X}}_h = \mathbf{Z}_h \hat{\gamma}$ where $\hat{\gamma} = (\sum_h \mathbf{Z}_h' \mathbf{Z}_h)^{-1} (\sum_h \mathbf{Z}_h' \mathbf{X}_{h,k})$ where k indexes the regressors and the columns of $\hat{\gamma}$, and \mathbf{Z}_h is the matrix of exogenous instruments. In practice, the first-order intrahousehold serial correlation is -0.4 while higher orders are insignificantly different from zero.

⁷ Since individual observations are actually overlapping three-month periods, each "month" dummy represents a three-month period.

⁸ As shown in Table 1, after 1992, an individual can hit two caps at a different times during a year because the tax caps for Old Age Survivor and Disability Insurance (OASDI) and Health Insurance (HI) differ.

than 1,500 households each month, and just over half of households contribute a complete one-year panel of four consumption observations.

The crucial independent variable for the test is the percent change in after-tax income caused by fluctuations in Social Security taxes, denoted $\Delta y_{t+1}^{ss\tau}$. This variable is $T_t^{ss} - T_{t+1}^{ss}$, the negative first difference of the percent of after-tax income paid in Social Security taxes, which in turn is:

$$(4) \quad T^{ss} = \sum_{i=1}^2 \sum_{r=1}^3 \left(\frac{\tau_{ss} Y_i}{(1 - \tau_h) Y_h} \right) D_{ir}^{cap}$$

where i indexes individuals and r months in an interview period; τ_h is the average household tax rate; Y_i is individual labor income subject to Social Security taxation; Y_h is total household pretax income; and τ_{ss} is the statutory Social Security tax rate in that month. D^{cap} is a variable which equals 1 if the individual has not hit the tax cap, a fraction representing the fraction of the month the individual pays Social Security if the individual hits the tax cap during the month, and 0 otherwise. All variables used to construct $\Delta y_{t+1}^{ss\tau}$ are calculated from retrospective information reported in the second interview, except D^{cap} , which can be, and is, calculated in several different ways.

First, D^{cap} can be calculated using each measure of individual income that the CEX provides. In addition to the main income datum in the CEX, which is the individual's earnings over the past 12 months, each individual also is asked the amount of his or her last paycheck and the length of the pay period, from which annual income can be calculated. D^{cap} (and thus $\Delta y^{ss\tau}$) is calculated using both measures of income and the measure used throughout the paper is the average of these two constructs. The second choice one faces is whether to construct D_{ir}^{cap} using the incomes reported in the fifth interview or those reported in the second.¹⁰ The use of the final interview technically violates the information restriction—that is, one would be using information to calculate a

change in income which is not available to the household when it made its initial consumption decisions. Thus, the paper focuses on results that employ solely second-interview information to construct $E_t[\Delta y_{t+1}^{ss\tau}]$ and on results that employ second-interview information to predict (instrument for) a measure of $\Delta y_{t+1}^{ss\tau}$ which uses both the fifth- and second-interview reports to construct D^{cap} . Note that this “fifth-interview” measure of the expected change in income is only using fifth-interview information to calculate when the individual is most likely to have hit the Social Security tax cap.

The final step in constructing $E_t[\Delta y_{t+1}^{ss\tau}]$ is to set it to zero for any individual who might not be paying Social Security taxes. This calculation is made on the basis of an individual's employment history, occupation, industry, reported Social Security contributions, and retirement plan payments.

Turning to the construction of the consumption data, the main task of the statistical analysis is to separate the effects of differential seasonal patterns of consumption from the effect of hitting the Social Security tax cap. Since Christmas gift giving is such a seasonal variation, all expenditures on gifts for someone outside the household are excluded from the consumption data. While this helps to minimize the possibility of contamination of the regressions by seasonal factors, if this category belongs in the regressions and is one of the more responsive categories of consumption expenditure, then the omission will bias the coefficients in favor of the null hypothesis of consumption smoothing.

Table 2 presents some summary statistics on the sample. There are 133,820 observations on 57,051 households in the sample. 32,554 households contribute a full three differenced observations; 11,661 contribute only two differenced observations; and 12,836 contribute only one. The Appendix discusses additional details of data construction such as the components of nondurable consumption and the dropping of missing, topcoded, extreme, or incomplete observations.

III. Estimation and Results

The first entry in Table 3 reports the estimated response of consumption to predictable changes in income (α_2) from ordinary least-squares (OLS)

¹⁰ The correlation between the measures from the same interview is 0.75. The correlation between the averages from different interviews is 0.90. The two measures from the same interview are entered separately when used as instruments, as subsequently discussed.

TABLE 2—SAMPLE STATISTICS

Monthly rates	Mean	Standard deviation	Percent of total consumption
Panel A: Full sample, 133,820 observations			
Total expenditures	1,449	960	100
Nondurable consumption	823	545	57
Expenditures excluding gifts to people outside the household:			
Total	1,414	934	98
Nondurable	797	523	55
Food and alcohol	325	191	22
Apparel and services	84	96	6
Entertainment and personal care	136	149	9
Family size	2.65	1.54	
Age	46.7	17.7	
Before-tax family income	2,241	1,812	155
After-tax family income	2,022	1,607	140
Head labor income	1,271	1,403	88
Fifth-interview Δy^{sst}	0.000	0.005	
Second-interview Δy^{sst}	0.000	0.006	
Growth in nondurable consumption excluding gifts to people outside the household	-0.004	0.380	
Panel B: Sample of households hitting cap, 11,828 observations			
Nondurable consumption excluding gifts to people outside the household	1,438	661	
Before-tax family income	4,922	2,183	
Head labor income	3,939	1,680	
Months Social Security covers head	9.8	2.2	
Fifth-interview Δy^{sst}	0.000	0.017	
Second-interview Δy^{sst}	0.000	0.017	
Growth in nondurable consumption excluding gifts to people outside the household	-0.002	0.366	

Notes: Based on samples for regressions on nondurable consumption excluding gifts to people outside the household. Total consumption excludes expenditures on mortgages, health care, pensions, education, and cash contributions. See the Appendix and text for additional definitions. Averages are across observations. All variables are current dollars.

estimation of equation (3) on the entire sample using the “fifth-interview” measure of the change in income caused by the Social Security taxes. The response of consumption is highly statistically significant and implies that when a household’s Social Security payments fall so that income rises by 10 percent, nondurable consumption rises by 5.4 percent. Since, as just described, the construction of $E_t[\Delta y_{t+1}^{sst}]$ uses some information from the end of each household’s tenure in the survey, the final two columns report results for a version of both $E_t[\Delta y_{t+1}^{sst}]$ constructed solely from second-interview information and for two-stage least-squares (TSLS) estimation of the equation containing the fifth-interview measure of $E_t[\Delta y_{t+1}^{sst}]$. The instrument set consists of the

dummy variables, family size and age regressors; both second-interview measures of $E_t[\Delta y_{t+1}^{sst}]$; and, to capture some of the nonlinearity in the transformation of income, dummy variables indicating for each measure if it is positive, negative, or zero. As reported in the second and third columns of results in the first row of Table 3, this possible endogeneity is not driving the results.¹¹

¹¹ The following alternative specifications are examined and found not to eliminate the statistical or substantive significance of the finding in the first row of Table 3: feasible generalized least-squares estimation, a complete set of time dummies rather than month and year dummies, not dropping outliers, including gift expenditures for people outside the household, and various measures of y^{sst} which

TABLE 3—THE REACTION OF NONDURABLE CONSUMPTION TO PREDICTABLE CHANGES IN INCOME

	Measure of income growth due to Social Security		
	Fifth-interview $\Delta y_{t+1}^{SS\tau}$ OLS regression	Second-interview $\Delta y_{t+1}^{SS\tau}$ OLS regression	Fifth-interview $\Delta y_{t+1}^{SS\tau}$ TSLS regression
Panel A: Full sample, 133,820 observations			
Coefficient:	0.538	0.617	0.661
Standard error:	(0.197)	(0.202)	(0.220)
Panel B: Sample with high consumption, 13,895 observations			
Coefficient:	0.615	0.718	0.784
Standard error:	(0.344)	(0.351)	(0.380)
Panel C: Sample of households hitting cap, 11,828 observations			
Coefficient:	0.446	0.563	0.608
Standard error:	(0.272)	(0.279)	(0.313)

Notes: The dependent variable is the change in the log of nondurable consumption excluding gifts to those outside the household. In addition to the measure of the change in income caused by the Social Security tax, all regressions also include a fourth-order polynomial in age, a second-order polynomial in family size in the second interview and in the last interview, and complete sets (less one) of month and year dummies. The instrument set includes all these additional regressors and both measures of the income change due to the Social Security tax calculated from second-interview information, and dummy variables for whether the changes thus calculated are strictly positive for that observation and dummy variables for whether the changes are strictly negative.

While the first row of Table 3 constitutes a rejection of the linear Euler equation, several steps can still be taken to eliminate possible alternative interpretations. Why might one be concerned with the results so far? First, the fit of the regressions is small: the R^2 's are just less than 1 percent. This is due, in part, to the fact that the Social Security tax changes are small relative to the swings in consumption (true movements and those from measurement error). Second, the sample size is very large.¹² Under these conditions, only a small spurious correlation between the error term and the independent variable could cause the significant results. Consider, for example, the following source of such a correlation. Suppose that larger or

wealthier families have larger increases in consumption around Christmas and larger incomes and so are more likely to hit the tax cap. Then the change in the Social Security tax variable would be spuriously correlated with the change in consumption through different seasonal preferences for consumption by wealth or family size. Two tacks are taken to reduce this possible spurious correlation.¹³

First, the analysis is performed on two subsamples: households in which average nondurable consumption exceeds 1,500 (1987) dollars

are not averages of measures. Analysis in levels (rather than percents) uncovers even greater statistical significance, but smaller economic effects. Finally, even the income measure taken from the second interview includes some information not available to the household when it made its consumption decisions at the beginning of the three-month recall period of the second interview. Dropping all changes in consumption from interview 2, both OLS estimation using only interview-two information and TSLS estimation consistently yield coefficients of around or above 0.5.

¹² Actually, only about 13 percent of the sample have nonzero $\Delta \tau^{SS}$.

¹³ A third and fourth tack are also tried. As suggested by Christina H. Paxson, since all gifts to members outside the household are excluded, by restricting the sample to households of size one, all gifts can be eliminated from consumption. Doing so yields coefficients greater than one which, while much less precisely estimated, remain statistically different from zero. As suggested by a referee, one can drop observations in which one member of a household hits the Social Security tax cap late in the year, when seasonal spending causes the largest fluctuations in consumption. Using the average of the second-interview reports of income to construct the tax-induced change in income uncovers an elasticity of 0.502 with a standard error of 0.207 when dropping any household in which either member hits the cap in November or December. A slightly higher elasticity of 0.535 with a standard error of 0.225 is found when dropping any household in which either member hits the cap in October, November, or December.

a month; and households in which either the head or spouse hit the tax cap. These samples are more homogeneous in terms of income and consumption levels, and are smaller by an order of magnitude. They also contain more individuals who hit the tax cap and therefore more variation in expected income changes. The standard deviation of $E_t[\Delta y_{t+1}^{ss\tau}]$ in the high consumption subsample is 1.1 percent, and in the cap-hitting subsample, 1.6 percent. These numbers are roughly double and triple that in the entire sample, respectively. The bottom two panels of Table 3 show that similar coefficient estimates are obtained from these samples while the standard errors increase from those garnered from the full sample.

As a second tack, the response of consumption to expected changes in income is identified using only the variation across otherwise identical individuals that do and do not pay Social Security taxes. If there is a spurious correlation between consumption growth and $E_t[\Delta y_{t+1}^{ss\tau}]$, such as from differential seasonal patterns of consumption, then the correlation will be present for all households regardless of whether they actually have Social Security taxes withheld. Comparing the consumption responses of those individuals known to be covered by Social Security to the responses of a control group of individuals who are not covered by Social Security should eliminate the effect of any such spurious correlation.

Using the fact that the Social Security tax withholding rules do not cover the self-employed, many government workers, and some smaller groups like clergy, each earner in the data set is assigned to one of three groups: a treatment group comprised of those individuals who almost certainly are subject to Social Security tax withholding; a control group comprised of those individuals who almost certainly are not; and a "neither" group comprised of the remaining individuals. Assignment is made on the basis of five sets of information: an individual's occupation and industry of employment; the amount of Social Security taxes paid by an individual in the past year and whether Social Security or Medicare are usually deducted from his or her paycheck; whether a household overpaid Social Security in the last 12 months; the number of weeks employed in the last year;

and finally, the household's contributions to a Railroad Retirement pension.

For every household, a hypothetical variable is constructed that represents the change in income due to the Social Security tax rate that this household would have experienced had it been paying Social Security taxes. This variable is denoted $E_t[\Delta y_{t+1}^{ssE}]$ and its coefficient captures the effect of any spurious correlation, such as the pattern of seasonal variation discussed above. Similarly, $E_t[\Delta y_{t+1}^{ssN}]$ is the hypothetical change for households that are in the "neither" group. The response of interest is then estimated by adding these two variables to the estimating equation:

$$(5) \quad \Delta \ln(c_{h,t+1}) = \alpha'_1 \mathbf{z}_h + \alpha_2 E_t[\Delta y_{h,t+1}^{ss\tau}] \\ + \alpha_3 E_t[\Delta y_{h,t+1}^{ssN}] \\ + \alpha_4 E_t[\Delta y_{h,t+1}^{ssE}] \\ + \alpha'_5 \mathbf{m}_t + \alpha'_6 \mathbf{d}_t + \varepsilon_{h,t+1}.$$

The significance of α_2 again provides a test of the basic LCH/PIH; however the identification comes only from the difference between the consumption response of the control group and that of the treatment group.

The left panel of Table 4 displays the coefficients of interest from estimation of equation (5) on the entire sample. Results are reported for both OLS using the second-interview measures of expected income change for each group and TSLS estimation. Relative to correlation present in the control group, the response of consumption for households in the treatment group is estimated to be one—a much larger estimate than the previous estimate of one-half. However, this estimated elasticity of one from the relative response is not statistically different either from the estimate from the absolute response (that is, one-half) or from zero. The large standard errors are consistent with the fact that the control group is quite small.

The coefficient on the hypothetical income change for the entire sample, $E_t[\Delta y_{h,t+1}^{ssE}]$, would be zero if there were no spurious correlation between the constructed measure of expected income change and the residual in equation (5). In fact this coefficient is nega-

TABLE 4—THE RELATIVE REACTION OF NONDURABLE CONSUMPTION FOR THE TREATMENT GROUP

Δy^{ss} Measures employed: Method of estimation:	Panel A: Full sample, 133,820 observations		Panel B: Sample of households hitting cap, 10,361 observations	
	Second-interview OLS regression	Fifth-interview TSLS regression	Second-interview OLS regression	Fifth-interview TSLS regression
Treatment group coefficient ($\hat{\alpha}_2$):	1.009	1.043	1.213	1.265
Standard error:	(0.637)	(0.676)	(0.694)	(0.742)
Neither group coefficient ($\hat{\alpha}_3$):	0.295	0.288	0.711	0.779
Standard error:	(0.828)	(0.941)	(0.902)	(1.045)
Everyone group coefficient ($\hat{\alpha}_4$):	-0.394	-0.385	-0.705	-0.710
Standard error:	(0.606)	(0.643)	(0.686)	(0.736)
Percent of individuals in				
Treatment group:	70	70	75	75
Control group:	5	5	8	8

Notes: The dependent variable is the change in the log of nondurable consumption excluding gifts to those outside the household. In addition to the measure of the change in income caused by the Social Security tax, all regressions also include a constant, fourth-order polynomial in age, a second-order polynomial in family size in the second interview and in the last interview, and complete sets (less one) of month and year dummies. The instrument set includes all these additional regressors and both measures of the income change due to the Social Security tax calculated from second-interview information, and dummy variables for whether the changes thus calculated are strictly positive for that observation and dummy variables for whether the changes are strictly negative. The percentages reported are the number of heads or spouses with positive earnings in each group divided by the total number of heads and spouses with positive earnings.

tive, suggesting that the spurious correlation is in fact negative, although the estimate is far from statistically significant. The negative sign is also consistent with the rise in the coefficient for the treatment group. The coefficient on the hypothetical income change for the neither group, $\hat{\alpha}_3$, lies between the other coefficients, as it should since it is a mixture of households that actually have Social Security taxes withheld and those that actually do not.

In sum, if it were the case that preference-driven increases in consumption around the end of the calendar year occurred in the same way in which Social Security taxes caused after-tax incomes to rise, then one would expect the point estimates of the relative response of the treatment group to be zero. In fact, consumption responds more for households that are covered by Social Security taxes relative to those that are not. As the right panel of Table 4 shows, estimation on the much smaller subsample of only those households who hit (and would have hit) the Social Security tax cap yields a similar conclusion.

So far, the evidence suggests significant failures of consumption smoothing. The next

section uses the data and test to evaluate two possible explanations for this rejection.

IV. Evidence on Alternative Theories

A. Liquidity Constraints and Precautionary Saving

If the linear Euler equation fails due to liquidity constraints or a correlation between the expected variance of consumption and the growth rate of consumption, then the relationship between expected income growth and consumption growth should be strongest for those households with few liquid assets. In each household's fifth interview, the CEX collects information on the current level of liquid assets and how this level has changed over the past 12 months. Measures of the value of checking accounts, savings accounts, stock and mutual-fund holdings, and bonds as of immediately before each household's first consumption observation are constructed. The ratio of this variable to average monthly nondurable consumption yields a measure of how likely a household is to be liquidity constrained or,

nearly equivalently, whether consumption is likely to track income due to the changes in the expected variance of consumption.¹⁴

The first two columns of Table 5 report the results of estimation of the linear Euler equation on two different subsets of the data. The first subset includes only households with asset ratios below 1, that is without enough asset wealth to finance one month of nondurable consumption. The second subset includes only households with asset ratios above 6. There is little evidence that the Euler equation failure is concentrated among households with the fewest assets. It may well be that since the asset data are retrospective, the significant amount of noise in the data is correlated in some way with the growth rate of consumption—or with households' increasing or decreasing fortunes—and it is this aspect of the data that the coefficients reflect.¹⁵

As another tack, age can be used to proxy for the probability of being liquidity constrained or having few assets. Young households typically have larger expected income growth and fewer assets than do older households. Due to either liquidity constraints or to the optimal choice not to borrow, young households may be more likely to violate the linear Euler equation. The final two columns of Table 5 report a similar pair of regressions to those for the asset classification but based on a sample of household 43 or younger and a sample of households 44 through 70.¹⁶ While the coefficients are consistent with the hypothesis that younger households react more to expected changes in income, this difference is neither economically nor statistically significant.

While almost no evidence for liquidity constraints or precautionary savings is found, it should be noted that, since individuals who are

not fully employed during the previous 12 months are discarded, the sample has fewer candidates to be liquidity constrained than the population that is typically studied.

B. *Near Rationality*¹⁷

Consider a model of boundedly rational consumers who allow consumption to track income provided that this strategy does not take them too far from the utility level that a fully rational strategy would imply.¹⁸ Thus, households that get a few hundred dollars extra in take-home pay for a few months simply spend some fraction of this money when they get it, rather than completely smoothing consumption expenditures. In this case, one might expect that purchases would track expected income changes more closely for goods with high intertemporal elasticities of substitution of expenditures—that is, for which swings in consumption expenditure provide little utility loss. The more a good is durable or storable, or provides lasting utility, the higher the effective intertemporal substitutability of such a good. For example, relative to the swing in expenditures, there is a small gain in utility for a household which takes its monthly trip to the movies a week early. Buying lunch a few hours later may cause a large utility loss.

Table 6 displays the results of estimating the response of expenditures for different categories of goods. The elasticity of total expenditures to expected changes in take-home pay is estimated as similar to that of nondurables. This finding suggests that nondurable and durable expenditures react similarly to predictable changes in income.¹⁹

¹⁷ The idea of these “rules of thumb” is similar in spirit to those proposed by Hall and Frederic S. Mishkin (1982).

¹⁸ Here “boundedly rational” and “fully rational” apply to the behavior interpreted within the context of the model of Section I. That is, boundedly rational behavior is not meant to imply that the behavior is not perfectly rational within the context of a more detailed model in which the constraints on calculation, information gathering, or time are made explicit. An interesting and relevant test is found in Paxson (1993) which finds that Thai farmers—a group likely to have large costs associated with letting consumption track income—smooth their consumption quite well over predictable seasonal variations in harvests.

¹⁹ The Social Security tax changes are known far enough in advance that this is a valid test when expenditures on durable goods are added to nondurable expenditures as studied in N. Gregory Mankiw (1992) and Ricardo J. Caballero (1993).

¹⁴ See Deaton (1992) and Christopher D. Carroll (1997) for discussions of the relationship between liquidity constraints and precautionary saving.

¹⁵ Interacting the ratio itself or picking a single cutoff near the middle of the distribution leads to similar inconclusions.

¹⁶ Tullio Jappelli et al. (1998) find that age is a significant predictor of whether a household reports that it is liquidity constrained and Pierre-Olivier Gourinchas and Parker (1997) estimate that the typical household moves from “buffer-stock” type behavior to behavior more consistent with that of the basic LCH/PIH around age 43.

TABLE 5—THE REACTION OF NONDURABLE CONSUMPTION BY AGE-GROUP AND ASSET LEVEL

	Low asset ratio	High asset ratio	Young age	High age
Panel A: Full sample				
Coefficient:	0.513	0.828	0.693	0.466
Standard error:	(0.638)	(0.349)	(0.299)	(0.320)
Number of observations:	33,795	29,460	67,276	49,626
Panel B: Sample of households hitting cap				
Coefficient:	1.727	0.570	0.614	0.570
Standard error:	(1.099)	(0.488)	(0.428)	(0.452)
Number of observations:	1,284	4,358	6,938	4,867

Notes: The dependent variable is the change in the log of nondurable consumption excluding gifts to those outside the household. High age is 44 to 70; low age is 43 or younger. High asset ratio is more than enough assets to finance six months of nondurable consumption; low asset ratio is less than enough assets to finance one month of nondurable consumption; the remaining households are dropped. All regressions employ TLSLS on the fifth-interview Δy^{sst} . In addition to the fifth-interview measure of the change in income caused by the Social Security tax all regressions also include a constant, a fourth-order polynomial in age, a second-order polynomial in family size in the second interview and in the last interview, and complete sets (less one) of month and year dummies. The instrument set replaces the Social Security variable with the following variables: both measures of the income change due to the Social Security tax calculated from second-interview information, and dummy variables for whether the changes thus calculated are strictly positive for that observation and dummy variables for whether the changes are strictly negative.

Consistent with this version of near rationality, food consumption (Table 6, column 2) responds less than nondurable consumption to expected changes in income, while expenditures on entertainment and personal care, which also includes expenditures on reading materials and tobacco and smoking supplies, shows a slightly stronger reaction in the entire sample, and no response in the subsample of households in which one member hits the Social Security tax cap. Finally, apparel and services consumption reacts the most of all categories, with point estimates of 1 to 2. Since expenditures on apparel and services make up roughly 5.6 percent of after-tax income, this estimate suggests that for every extra expected dollar of income which a Social Security tax change induces, 6 to 12 cents are spent on apparel and services.

Two final points about this alternative hypothesis are worth mentioning. First, this test of consumption smoothing leans most heavily on high-income households who are likely to have the highest dollar value of time. If optimizing takes time, such households are more likely than the typical household to imperfectly smooth consumption over these tax-induced swings in income. Second, the low fit of the regression suggests that the Social Security tax code does not cause most fluctuations in consumption. Thus

there would not be a large percentage reduction in the variance of consumption if households did completely smooth consumption across Social Security tax changes.

V. Conclusion

Consumers do not perfectly smooth their demand for goods at quarterly frequencies across expected income changes. Contrary to the basic LCH/PIH, consumption reacts in an economically significant manner to predictable changes in tax rates: the elasticity of expenditures on nondurable goods with respect to the predictable declines in income that are studied is around one-half. Consistent with some stories of bounded rationality, the strongest violations of consumption smoothing occur in subcategories of consumption in which households can easily substitute purchases across time. No evidence is found that precautionary saving or a constraint on borrowing is causing the failure of consumption smoothing.

This paper tests the joint hypothesis of rational expectations and the basic LCH/PIH. Consumer behavior may be in accord with the basic LCH/PIH but expectations may not be formed as predicted by rational expectations theory. That is, some subset of households may face constraints on time which make it optimal for them to ignore

TABLE 6—THE REACTION OF DIFFERENT CATEGORIES OF CONSUMPTION

Dependent variable consumption category:	Total	Food and alcohol	Entertainment and personal care	Apparel and services
Panel A: Full sample				
Coefficient:	0.564	0.133	0.835	2.145
Standard error:	(0.241)	(0.206)	(0.407)	(0.515)
Number of observations:	128,437	131,076	128,709	103,799
Panel B: Sample of households hitting cap				
Coefficient:	0.631	0.192	-0.281	1.010
Standard error:	(0.347)	(0.285)	(0.574)	(0.714)
Number of observations:	11,089	11,745	11,784	10,671

Notes: In regressions on total and food consumption observations are discarded if consumption changes more than 100 percent between quarters. For the other categories, the cutoff is 200 percent. These cutoffs are all around two standard deviations. See the Appendix for exact definitions of the categories. All regressions employ TSLS on the fifth-interview Δy^{sst} . In addition to the fifth-interview measure of the change in income caused by the Social Security tax all regressions also include a constant, a fourth-order polynomial in age, a second-order polynomial in family size in the second interview and in the last interview, and complete sets (less one) of month and year dummies. The instrument set replaces the Social Security variable with the following variables: both measures of the income change due to the Social Security tax calculated from second-interview information, and dummy variables for whether the changes thus calculated are strictly positive for that observation and dummy variables for whether the changes are strictly negative.

the complex Social Security tax code when forming expectations about their future income. In this case, consumption for these households rises because the analyzed changes in after-tax income come as a surprise. Note, however, that if this theory is to rationalize the magnitude of the point estimate, households must expect the income shock caused by hitting the tax cap to be more persistent than it actually is.

This paper and other recent evidence against pure consumption smoothing have potentially far-reaching implications. If expected changes in taxes influence contemporaneous consumption behavior, then fiscal stabilization, such as that undertaken by President Bush in 1992 or that provided by automatic tax stabilizers, is likely to have important effects on consumption. While one must keep in mind that the current experiment is a partial-equilibrium result and involves small swings in income, the findings of this paper suggest that when studying horizons such as those addressed in business-cycle models, the linear Euler equation is not a close approximation to the correct structural equation.

APPENDIX

The CEX family, member, and detailed expenditure files for years 1980 to 1993 were

provided by the National Bureau of Economic Research. Most information about the CEX was obtained from the U.S. Bureau of Labor Statistics (BLS) (1980–1993) and conversations with statisticians at Division of the CEX in the Bureau of Labor Statistics.

Any family that is missing the second-interview reports of family size or age of reference person is dropped. Households also are dropped if before-tax household income or after-tax household income in the second interview is topcoded, incomplete, or missing. Age is the average of both head and spouse if there is a spouse, otherwise it is the head's. Due to some extreme reports, reported tax rates above 60 percent are reset to 60 percent, and those below 0 to 0. The results are quite insensitive to alternatives to this correction. Households should not be matched across 1985 to 1986, and are not.

Consumption data are compiled from the detailed expenditure files. Monthly expenditures are averaged over the number of months with nonzero nondurable consumption to get consumption at a quarterly rate except total expenditures, which are averaged over the number of months with positive total expenditures. Any three-month period with only one month of expenditures is dropped. Households with non-

durable consumption less than 1,000 1987 dollars are dropped.

Nondurable consumption is the sum of expenditures on the following categories of goods: food, excluding food as pay and school meals; alcohol; house furnishings and equipment, excluding furniture, major appliances, and floor coverings; apparel and services; transportation, excluding new and used vehicle spending and financing; entertainment; personal care; reading; and tobacco and smoking. Total consumption is defined as total expenditures less outlays for mortgage payments, education, health care, pensions, and cash contributions. Food expenditures are all expenditures on food and alcohol less food as pay and school meals. Entertainment and personal care is the sum of entertainment, personal care, reading, and tobacco and smoking expenditures. Nondurable consumption observations which have changes in consumption across three-month periods in excess of 100 percent are dropped.

The measure of income using fifth-interview information is constructed by interpolating using the second- and fifth-interview reports. If income changes by more than 25 percent between the second and fifth reports, the fifth measure is used. Assets are considered missing if both savings and checking account information is missing. When this account information is not missing, bonds and stock accounts are added to the amounts in the accounts to create total asset measures.

Individuals are assigned to treatment, control, and neither groups as follows. Individuals who are federal government employees prior to 1984 or self-employed are assigned to the control group. All individuals who are not federal government employees prior to 1984 and who are government employees are placed in a government category. All individuals in any family that reports paying into a government retirement account and that has no members already assigned to the government group are moved into the government group. Next any individual who reports paying Social Security in his or her normal paycheck or during the last 12 months is assigned to the treatment group. Those who report not paying Social Security in either of these questions are put into the control group. The remaining government workers are placed in the neither group. Next, any individual who is

missing industry or occupation data is put into the treatment group unless he reports not paying Social Security taxes. Any individual who reports overpaying Social Security (which generally occurs only because individuals switch jobs) or working less than 50 weeks during the past year is placed into the neither group. If a family reports paying into a Railroad Retirement (RRR) account and both adults work, the household is moved into the neither group. If a household reporting paying RRR consists of only one worker, or if the other worker is already assigned to the treatment group, the individual is put into the control group pre-1985 and into the treatment group from 1985 onwards. At the end of the procedure, any individual assigned to both the treatment and the control group is moved into the neither group. Individuals without labor income are not assigned to any group.

REFERENCES

- Altonji, Joseph and Siow, Aloysius.** "Testing the Response of Consumption to Income Change with (Noisy) Panel Data." *Quarterly Journal of Economics*, May 1987, 102 (2), pp. 293–328.
- Attanasio, Orazio P. and Weber, Guglielmo.** "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *Journal of Political Economy*, December 1995, 103 (6), pp. 1121–57.
- Bodkin, Ronald.** "Windfall Income and Consumption." *American Economic Review*, September 1959, 49 (4), pp. 602–14.
- Browning, Martin and Lusardi, Annamaria.** "Household Saving: Micro Theories and Macro Facts." *Journal of Economic Literature*, December 1996, 34 (4), pp. 1797–855.
- Caballero, Ricardo J.** "Durable Goods: An Explanation for Slow Adjustment." *Journal of Political Economy*, April 1993, 101 (2), pp. 351–84.
- Carroll, Christopher D.** "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics*, February 1997, 112 (1), pp. 1–55.
- Deaton, Angus.** *Understanding consumption*. Oxford: Oxford University Press, 1992.

Ann Arbor, MI: Inter-university Consortium of Political and Social Research.

- Wilcox, David W.** "Social Security Benefits, Consumption Expenditure, and the Life Cycle Hypothesis." *Journal of Political Economy*, April 1989, 97 (2), pp. 288–304.
- Zeldes, Stephen P.** "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy*, April 1989, 97 (2), pp. 305–46.