The Reaction of Household Consumption to Predictable Changes in Payroll Tax Rates*

Jonathan A. Parker
University of Wisconsin

November 1997
First Draft: September 1994

Abstract

This paper exploits a natural experiment provided by the pattern of payroll taxation in the U.S. to test whether household consumption responds to predictable changes in after-tax income. Social Security taxes are withheld from individual paychecks until an annual contribution limit is reached, at which point take-home pay rises for the remainder of the calendar year. Using a set of statutory rate increases as well as this source of variation, the elasticity of nondurable consumption with respect to after-tax income is found to be 0.5 over three-month intervals. I find little evidence for or against the hypothesis that households with fewer liquid assets respond more to these income changes. However, consistent with some models of bounded rationality, the violations of consumption smoothing are most pronounced for semi-durable goods.

* Comments welcome. I thank the National Science Foundation and the Sloan Foundation for financial support. Further thanks are due to Olivier Blanchard, Ricardo Caballero, Charles Fleischman, Jonathan Gruber, Jerry Hausman, Steve Pischke, James Poterba, Joel Slemrod, Nicholas Souleles and seminar participants at M.I.T., Columbia, Princeton, Harvard, Yale and Stanford. Errors remain the responsibility of the author.

List of Tables

1.1 The Social Security Tax Structure, 1980-1993 2
4.1 Sample Statistics 10
6.1 The Reaction of Nondurable Consumption 13
6.2 Alternative Measures of the Social Security Tax Variable 15
6.3 High Consumption and Tax Cap Subsamples 16
6.4 Treatment vs. Control Group Regressions 18
7.1 Asset and Age Splits 21
7.2 The Reaction Across Categories of Consumption 23

1. Introduction

This paper evaluates the key implication of rational expectations and the **basic** Life Cycle/Permanent Income Hypothesis (LCH/PIH): that predictable changes in income should have no effect on the growth rate of consumption expenditures, and only new information should lead consumers to revise consumption behavior.[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. As is well known, to the extent that the basic LCH/PIH fails to hold, predictable changes in taxes can have real effects on consumption and investment, and government debt policy may affect national savings and possibly growth rates.

Using household-level consumption data from the Consumer Expenditure Survey (CEX), I examine whether expenditures on consumption increase contemporaneously with predictable changes in Social Security tax rates.[2] Individuals with wage and salary income earned in the United States are subject to Social Security tax withholding of around seven percent of their gross pay up to an annual maximum income level. Social Security tax rates provide two sources of variation. First, there have been a series of pre-announced tax rate increases in the 1980's which have affected all wage and salary earners. 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 1.1 shows the rates and caps for the sample period of 1980 to 1993.

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 this fluctuation in their after-tax wages annually, and, after 1982, the changes in the tax cap attempt to adjust the cap for average wage growth.[3] In the CEX, the average annual wage and salary income among heads of households who hit the tax cap is $52,385. In the middle year of my sample, 1986,

[1] As I discuss in Section 2, the theory predicts smoothing in expected marginal utility– often called Hall's Martingale hypothesis– which does not generally imply smoothing of expected consumption, which I refer to as the **basic** LCH/PIH.

[2] I am indebted to Joel Slemrod for originally pointing me to the Social Security tax cap as a nice test of the basic LCH/PIH.

[3] Begining in 1982, the maximum contribution was 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 this transition rule was accelerated because previous adjustments had ignored non-wage and deferred compensation which had been growing more rapidly than wage compensation.

Table 1.1: The Social Security Tax Structure, 1980-1993

Year	Tax Rate (percent)	Maximum Annual Contribution Per Earner	Maximum Annual Taxable Earnings
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*	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

* The tax rate in 1984 includes the tax credit.

Data refers to OASDI and HI except where noted. Source: Social Security Administration (1990) Tables 2.A3, 2.A4, and 2.A5 and The Social Security Bulletin, (January, 1993).

a person with this wage income would have a $312 temporary increase in monthly after-tax income from mid-October until the end of December. The Social Security tax system thus provides a pseudo-natural experiment to quantify consumption responses to predictable changes in income.

I find that households change their consumption expenditures in response to the predictable fluctuations in income induced by the Social Security tax system. An expected, one percent increase in after-tax income in a three month interval increases nondurable consumption expenditures by around half of a percent. While not statistically significant, similar (but more volatile) 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. Thus consumption smoothing is rejected.

What alternative model might account for this behavior? If some fraction of households either are liquidity constrained or behave as "buffer-stock" consumers, then households with few assets should show the strongest violations of consumption smoothing. Unfortunately, the test and data provide little evidence as to whether families with less cash-on-hand react more to expected changes in income, although there is evidence that young households, who are more likely to be constrained or buffer-stock agents, do react more strongly to these predictable income changes. I find that, consistent with some models of bounded-rationality, the strongest violations of consumption smoothing are found for semi-durable items– that is for subcategories of consumption in which households can substitute purchases across time with little utility loss and for which purchases generally do not employ credit. Finally, the point-estimate of the response of consumption is larger than what would be consistent with a typically parameterized buffer-stock model or with a model in which households are simply surprised by the income changes.

In part, this test uncovers strong violations of consumption smoothing because it improves on most previous household-level tests.[4] It does so in four main ways. First, the income changes caused by Social Security withholding are calculated at the household-level and are exogenous to the household. The predicable change in income is thus not highly correlated with labor supply effects or (group-level) family size movements. 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 nine 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 the change in income. Third, the Consumer Expenditure Survey provides comprehensive data on many different types of consumption expenditures at the household level. Previous microeconomic studies often have had to make do with food consumption as the only measure of nondurable 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 measures of expected change in income, demonstrated in Mariger and Shaw (1993).

In part, however, this test asks a different question than much of the previous literature. Most importantly, this test is asks whether expenditures are smoothed across 3-month periods. Most previous tests have employed annual data and often identified expected income changes from either lower-frequency life-cycle movements or cross-sectional differences in income growth. It is quite possible, and indeed consistent with the model of bounded rationality discussed subsequently, that consumption expenditures are poorly smoothed across income fluctuations at 3-month intervals but are well smoothed at lower-frequencies.

Two final characteristics of this test are worthy of note. First, if one is concerned with estimating the effect of fiscal policy, this test yields a key elasticity of interest: the extent to which current desired spending changes in response to an ex-

[4] Section 3 discusses these points and the previous literature in more detail.

pected, tax-driven change in income.[5] The distinction between tax-driven changes in income and other changes is important once the basic LCH/PIH is rejected, since some alternative theories of consumption behavior suggest that the source of the expected increase in part determines the magnitude of the consumption response. Second, while the dataset is representative of the entire U.S. population, the Social Security experiment relies more heavily on high-income households. Many previous studies have focussed on groups, such as union workers or the elderly, for which liquidity constraints are more likely to be important than for the general population. This test leans more heavily 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 of time.

The next section presents the basic LCH/PIH. Section 3 contains a brief critical review of previous empirical research on consumption smoothing. In section 4, the data is described, and in section 5 the test presented. Section 6 tests consumption smoothing and analyzes the robustness of the test. In section 7, I discuss and test alternative models of consumption behavior, while the final section concludes. Appendix A provides the details of the data and variable construction.

2. Consumption Smoothing

The starting point for most modern studies of consumption is the following canonical model of consumer optimization. Households choose consumption expenditures to maximize utility subject to a lifetime budget constraint:

$$\begin{aligned} Max \qquad & E_s\left[\sum_{t=s}^{T}\beta^{t-s}u(K_t)\nu(z_t)+\beta^{T+1-s}V(K_{T+1},A_{T+1},z_{T+1})\right] && (2.1)\\ S.T.\ A_{t+1} &= R_t(A_t+(1-\tau_t)Y_t-c_t) && (2.2)\\ K_t &= (1-\delta)K_{t-1}+c_t && (2.3)\\ A_{T+1} &\geq 0;\ A_s, K_s \text{ given} && (2.4) \end{aligned}$$

where E_s is the expectations operator conditional on all information available at time s; β is the discount factor; c_t represents real expenditure on consumption goods; K is the real stock of consumption goods which differs from c only to the extent that some consumption goods are durable, storable or provide lasting utility;[6] $\delta \in (0,1]$ is the depreciation rate of consumption goods; $u(\cdot)$ is an intertemporally-separable felicity function which is stable through time, increasing and concave; z_t contains variables such as family size or taste shocks

[5] Two caveats are in order. First, since this test is a reduced form test, this elasticity should only be taken to apply to similar tax changes. Most importantly, much larger income changes may well have much smaller elasticities. Secondly, the tax–cap-induced changes in income are effectively lump-sum. Changes in tax rates will have additional labor subsitution effects which will likely alter consumption behavior.

[6] When the depreciation rate is one, this is the cannonical model of nondurable consumption choice (Hall 1978); when the depreciation rate is less than one, the model is the standard model of optimal choice of durable expenditures (Mankiw 1982).

which change marginal utility (through the function $\nu(\cdot)$); A_t is household assets; R_t is the gross after-tax real interest rate; Y_t is stochastic labor income;[7] τ_t is the Social Security tax rate on labor income; and finally $V(\cdot,\cdot)$ captures the value of assets left at death in the form of bequests. Households are not allowed to die with negative assets. When $V \equiv 0$, equations (2.1)-(2.4) represent the life-cycle hypothesis. When individuals live S years and $V(K_{T+1}, A_{T+1}, z_{T+1}) \equiv \sum_{t=T+1}^{T+S} \beta^{t-T-1} u(K_t)\nu(z_t) + \beta^S V(K_{T+1}, A_{T+S+1}, z_{T+S+1})$, these equations represent the permanent income hypothesis (with occasional positive asset restrictions).

Prior to the last period of life, the Euler equation for the problem is:

$$u'(K_t)\nu(z_t) = \beta R_t E_t \left[u'(K_{t+1})\nu(z_{t+1})\right]. \tag{2.5}$$

This equation implies that households seek to equate marginal utility across time, with possible fluctuations due both to changes in the relative price of consumption across periods, R_t, and to changes in variables that shift marginal utility, such as family size, z_t.

Under rational expectations, if z_t is constant and the interest rate equals the discount rate, equation (2.5) implies that marginal utility follows a Martingale process. Since households plan to equate marginal utility today and tomorrow, any *ex post* difference between marginal utility today and tomorrow must be due to information not available today. That is, $u'(K_{t+1}) = u'(K_t) + \varepsilon_{t+1}$, where the innovation at time $t+1$ is orthogonal to all information available to the households at time t.[8] To derive an estimable Euler equation, it is typically assumed that: a) the consumption expenditures being examined are not durable or storable and do not provide lasting utility; b) the utility function is of the constant relative risk aversion (CRRA) form; c) that consumption is log-normally distributed; and d) z_{t+1} is known at time t. In this case the Euler equation becomes:

$$\Delta \ln(c_{t+1}) = \sigma \ln(\beta R_t) + \sigma \ln\left(\frac{\nu(z_{t+1})}{\nu(z_t)}\right) + \frac{\rho}{2} E_t \mathrm{Var}\left(\Delta \ln(c_{t+1})\right) + \epsilon_{t+1} \tag{2.6}$$

where σ is the intertemporal elasticity of substitution, ρ is the coefficient of relative risk aversion, and $\epsilon_{t+1} = \ln(c_{t+1}) - E_t[\ln(c_{t+1})]$.[9]

A large body of literature has proceeded under the assumption that the expected variance of the log of consumption is orthogonal to information available at time t. In this case, after conditioning on the real interest rate and variables which plausibly alter the marginal utility of consumption over the life cycle, equation (2.6) implies that the percent change in consumption should be uncorrelated with any variables dated t or earlier. This paper refers to this as the key implication

[7] Following convention, time t variables are assumed known when the household chooses consumption in t.

[8] In his seminal paper, Hall (1978) derives this result when $\delta = 1$ and shows that, if in addition the utility function is quadratic, then equation (2.5) implies that consumption follows a Martingale process and $c_{t+1} = c_t + \varepsilon_{t+1}$.

[9] Under these assumptions $\rho = \frac{1}{\sigma}$. Alternatively, a similar Euler equation can be derived as a Taylor approximation to (2.5). See Dynan (1993).

of the **basic** LCH/PIH and to the following equation as the **linear** consumption Euler equation:

$$\Delta \ln(c_{t+1}) = \sigma \ln(\beta R_t) + \sigma \ln \left(\frac{\nu(z_{t+1})}{\nu(z_t)} \right) + \epsilon'_{t+1}, \tag{2.7}$$

where the variance term has simply been added to the error term. Parsimony is the essence of modelling, and many insights can be drawn, especially at the aggregate level, from the assumption that the linear Euler equation is true at the household level. Further, to date the empirical evidence rejecting the basic LCH/PIH as a reasonable approximation to reality is not overwhelming. Given that the basic Euler equation is both highly useful and widely used, whether it is approximately true or not is a crucial issue in the study of consumer behavior.

If one estimates equation (2.7) and variables known as of time t do not enter significantly, then the coefficient on the real interest rate provides an estimate of the intertemporal elasticity of consumption, and this and the constant yield an estimate of the discount rate.

Most previous studies choose to test whether expected changes in household income have predictive power for changes in consumption in equation (2.7). One reason to test expected income is historical: prior to Hall (1978), economists generally modelled consumption as a function of current and past income and made no formal distinction between expected and unexpected changes in income. In addition, however, testing whether expected income growth is correlated with consumption growth allows one to evaluate the economic importance of the rejection and to assess what alternative hypothesis might be causing the failure. For this second reason, I remain in this tradition and test equation (2.7) by including a measure of the change in after-tax income due to individual-specific changes in Social Security tax rates.

3. A Brief History of Excess Sensitivity Tests

This section motivates the current study by clarifying its strengths relative to the previous literature. For more comprehensive reviews see Deaton (1992) and Browning and Lusardi (1996).

The main methodology for testing the basic LCH/PIH at the household level was pioneered by Hall and Mishkin (1982) and Shapiro (1984). The strategy consists of using the Panel Study of Income Dynamics (PSID) to examine whether consumption responds to expected changes in income.[10] There are two main weaknesses with this line of attack.

First, the only measure of consumption in the PSID is "usual weekly" food consumption, which is likely both to be a very noisy measure of consumption and to have a lower elasticity than many other types of consumption, since a certain

[10] Hall and Mishkin (1982) test more than just this, but this correlation was the main violation they uncovered.

amount of food consumption is necessary. The specific question used in the PSID survey does not clearly delineate whether the "usual week" is usual for the time of the interview or for the previous calendar year, to which many questions (and the previous survey question) refer. The methodology has been extended by Altonji and Siow (1987), Zeldes (1989a), Runkle (1991), Campbell (1987) and Japelli, Pischke and Souleles (1995), and there is no consensus among these papers as to whether the linear Euler equation is rejected by the data.

In response to the problems with food consumption, recent research has turned to using the Consumer Expenditure Survey (CEX) which contains much better information on household-level consumption. Using a similar technique, Attanasio and Weber (1995) find evidence that the linear Euler equation is violated and evidence that it is not.[11] Lusardi (1996) uses a two-sample instrumental variables procedure to match the higher-quality income data from the PSID with consumption data from the CEX. This yields a more powerful test than use of either dataset alone, and she rejects the Martingale hypothesis.

However, these tests still suffer from a second set of weaknesses related to the construction of "expected" income. First, it is difficult to construct quality measures of expected income growth that are not highly correlated with marginal utility shifters. Expected changes in income are usually measured by predicting the change in income using predetermined variables such as education, occupation, industry, and age.[12] Thus, in part, these studies are examining whether the age-profile of consumption tracks the age-profile of income.[13] Since the income profile has a hump-shape over the life cycle which is similar to the shapes of marginal utility shifters such as family size and labor supply, it is difficult to separate the relative roles of these variables in causing the hump-shape of consumption over the life cycle. Second, if there are permanent unexpected shocks to the age profile of earnings, then consumption should track these changes. The outlined procedure would include these changes in its measure of expected income change, and thus incorrectly reject the linear Euler equation. Third, most individual income movements are caused by idiosyncratic sources which are not predictable from information about an individual's age, occupation, education, etc. While the instruments employed in these studies do a reasonable job of predicting labor income, the studies employ a large set of instruments.[14] Instrumental variables

[11] Attanasio and Weber (1995) in Table 4 show that expected income growth predicts food consumption growth. The paper shows that when expected change in other nondurable consumption is included as an independent variable, it has a positive coefficient and the coefficient on expected labor income remains positive but becomes statistically insignificant. In contrast to the authors, I interpret this as demonstrating that predictable movements in consumption are correlated with predictable movements in income.

[12] Attanasio and Weber (1995) do not use industry, occupation, or education, but instead use age and lagged changes in consumption, family size, income and aggregate variables.

[13] These tests tests are thus low-frequency tests, and do not deal with the issue as to whether consumption tracks income over quarterly intervals.

[14] Lusardi (1996) reports an R^2 of around 0.01 when predicting income in the PSID. Attanasio and Weber (1995) have an R^2 of 0.24 in their first stage prediction of labor income using 21

estimation with weak instruments can have large finite-sample biases even in large finite samples.[15]

Fourth, preference parameters may in part determine which occupation, education and industry groups individuals select into, and these same parameters also determine the growth rate of consumption. This again could lead to spurious results.[16] Finally, as Mariger and Shaw (1993) have emphasized, the Euler equation provides a testable restriction on the time-series moments of consumption changes for each individual household. Since aggregate shocks may affect households differentially, cross-sectional moments cannot be used to consistently test the Euler equation.

In response to this set of critiques, several recent studies have turned to quasi natural experiments to identify households which experience idiosyncratic, expected fluctuations in income, and the evidence that expected income growth predicts consumption growth is building, but is not conclusive.[17] Shea (1995) exploits union contracts which cover several years and include publicly-known provisions for wage growth. Using PSID data, the paper finds an elasticity of food consumption with respect to these expected increases in wages around one.[18] Souleles (1995), in research contemporaneous with the present study and also using the CEX, finds that household consumption increases when and shortly after households receive their known Federal income tax refunds. Third, chapter 1 of Hayashi (1997) using surveys of household-level income expectations, estimates an elasticity of food consumption with respect to expected changes in income of 0.10 for Japanese households. The common weaknesses of these three papers is the short effective time horizon.[19] Finally, Paxson (1992) demonstrates that seasonal fluctuations in income due to different seasonal harvest times are well smoothed by agricultural households in Thailand.[20]

The current test avoids most of both sets of problems. I use detailed and relatively accurate consumption data from the Consumer Expenditure Survey. The experiment provided by the Social Security tax system yields a test with large, easily-measurable, changes in income which are exogenous to labor supply and family size. There is a large time horizon (168 months) and households which

instruments and with 288 observations on cohort data. Since the cohort techniques already effectively average the data, the true R^2 of the first stage regression is much smaller. The CEX is likely to have a sightly worse fit in the first stage due to the lower quality of the income data.

[15] See Bound, Jaeger and Baker (1995).

[16] For example, high education groups tend to have steeper income growth. If patient people get higher levels of education, then because patient people have higher consumption growth, consumption will seem to track expected changes in income.

[17] Deaton (1992), in his survey book on consumption, summarizes the micro evidence against the basic PIH as, "weaker and less transparent than in the aggregate data" (p.160).

[18] This finding is statistically significant at the 10% level for the entire sample, and at the 5% level for some subsamples.

[19] Souleles (1995) exploits a 11 year period of annual payments that vary across households.

[20] This finding is consistent with the idea that consumption is smoothed well when the utility costs to not doing so are large. That is, it is not inconsistent with lack of consumption smoothing when small utility losses are involved.

plausibly would be affected quite similarly by aggregate shocks hit the tax cap in different months and/or fall into different treatment/control groups. Section 6 deals in detail with the robustness of the findings to these and other critiques.

4. The Consumer Expenditure Survey

The data employed are from the Family, Member, and Detailed Expenditure files of the Consumer Expenditure Surveys(CEX) for the years 1980 to 1993. The CEX is structured as a rotating panel of households. Each household is interviewed 5 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. Demographic information is collected and a population weight assigned to the household. This weight is the only data released from the first interview. Each household is subsequently interviewed four more times, once every three months. In each of these interviews, detailed consumption expenditure information is collected on the past three months' expenditures. In each family's second and fifth interviews, a more detailed set of demographic and income information is collected. In these interviews, the family reports its pre-tax and post-tax income over the previous 12 months. In the fifth interview, each household is asked about its holdings of several categories of liquid assets and how much these holdings have changed over the past 12 months (i.e. since immediately prior to the start of their consumption reporting).

Families rotate into and out of the Survey, so that new households are being added every month. The Survey interviews about 1,500 households each month, and only about half of households contribute a complete one-year panel of four consumption observations each covering three months. Data on families and individual members are extracted and merged to make an unbalanced, overlapping panel of households covering January, 1980 to November, 1993. Any family which is missing the second interview reports of family size or age of reference person is dropped.[21] Households are also dropped if before-tax household income or after-tax household income in the second interview is topcoded, incomplete or missing. Differenced observations are dropped if nondurable consumption changes by more than 100%.[22]

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.

[21] The reference person in the CEX is defined as the person who is most responsible for paying the bills. This person is then interviewed, as he or she is probably the most knowledgable about household expenditures.

[22] I use other cuts when using other dependent variables. The results do not change much if outliers are not excluded. See section 6.

Table 4.1: Sample Statistics

Monthly Levels	Mean	Standard Deviation	Percent of Total Cons.
Total Consumption	1,436	1233	100
Nodurable Consumption	819	637	57
Non-Gift Consumption			
Total	1401	1205	98
Nondurable	794	612	55
Food+Alcohol	320	223	22
Apparel+Services	111	154	8
Entertain.+Personal Care	142	243	10
Family Size	2.64	1.54	
Age	46.4	17.5	
Head Labor Income	1,242	1,386	86
After-Tax Family Income	1,994	1,623	139
Before-Tax Family Income	2,191	1,808	153
First Differences			
Nondurable Consumption	−0.003	0.384	
Non-Gift Nondur. Cons.	−0.004	0.379	
$\Delta\tau^{ss}$	0.000	0.006	
Subsample: Head or Spouse Hits Soc Sec Cap; N=3522	Mean	S.D.	Max
Months Soc Sec Covers (Head)	9.38	1.87	12
Head Labor Income (Annual)	52,385	18,946	100,000
$\Delta\tau^{ss}$	−0.001	0.021	0.10

Based on sample for non-gift nondurable consumption regressions. Total consumption excludes expenditures on mortgages, health care, pensions, education, and cash contributions.

Average levels are across households except family size and age which are across observations.

The crucial variable for the test is the change in the share of after-tax income paid in Social Security taxes, $\Delta\tau^{ss}$. This is constructed as the first difference of:

$$\tau^{ss} = \sum_{i=1}^{2}\sum_{r=1}^{3}\left(\frac{\tau_{ss}Y_i}{(1-\tau_h)Y_h}\right)D_{ir}^{cap} \tag{4.1}$$

where i indexes individuals and r months in an interview period; τ_h is the average household tax rate; Y_i and Y_h are individual and household pre-tax income respectively; 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 equals 0 otherwise. All variables except D^{cap} are calculated from information given in the second interview pertaining to the 12 months prior to the second interview. To construct the best possible measure of when an individual hits the tax cap, D^{cap} is calculated from the retrospective labor income variable from the final interview.[23]

Finally, based on an individual's employment history, occupation, industry, reported Social Security contributions, and retirement plan payments, τ^{ss} is set to zero for any individual who might not be paying Social Security taxes. More details are provided in section 6 and in Appendix A.

Table 4.1 presents some summary statistics on the sample. There are $146,737$ total observations (NT) on $62,591$ total households (N) in the sample. $35,665$ households contribute a full 3 differenced observations; $12,816$ contribute only 2 differenced observations; and $14,110$ contribute only one.

5. The Test of the Basic LCH/PIH

As discussed in the introduction, I test the linear Euler equation by examining whether household consumption rises when any individual earner in a household undergoes a change in his or her Social Security tax rate. The following empirical specification of the linear Euler equation is employed:

$$\Delta \ln(c_{ht+1}) = \alpha_1 z_h + \alpha_2 \Delta\tau^{ss}_{ht} + \alpha_3 m_t + \alpha_4 y_t + \epsilon_{ht+1} \tag{5.1}$$

where h indexes households; m is a complete set of month dummies; y is a complete set of year dummies less one; z contains a second-order polynomial in family size and a fourth-order polynomial in age that capture the fact that household consumption is generally not flat over the life cycle.[24] Finally, $\Delta\tau^{ss}$ is the change in the fraction of household after-tax income which is paid in Social Security taxes. According to the basic LCH/PIH, α_2 should equal zero.

Three points about equation (5.1) are worth emphasizing. First, estimating the equation in first differences removes any household-specific effects. If these effects were not removed, a negative relationship between Social Security taxes and consumption might simply be due to the fact that only high-income households hit the cap and high-income households are more likely to have high levels of consumption.

Second, the regression includes a complete set of month effects.[25] Without the month dummies, it would be possible that the rise in consumption which occurs at the end of the calendar year would be attributed to the tax variable, which, on

[23] Consequently, D^{cap} uses information not available to the household when consumption decisions are made. Although it is unlikely that the month in which an individual hits the cap is systematically correlated with income shocks during the survey, it is technically possible. Thus I check the robustness of the results to calculating D^{cap} based on the individual salary report from the second interview instead.

[24] The age polynomial and the time dummies will also pick up cohort effects. This presents no problem for inference about α_2 but does imply that the year and time dummies should not be interpreted structurally.

[25] Since, as is discussed in the data section, individual observations are actually overlapping three-month periods, the 12 month dummies actually each represent a three-month period.

average, falls at the end of the calendar year. The month dummies however should completely absorb this 'Christmas' effect.

Third, in addition to the month dummies, time dummies for years are included in the regression. This allows the possibility that interest rates and other macro shocks affect families' consumption growth rates.

Equation (5.1) then has three sources of variation which identify the effect of the changing tax rates on consumption. First, high-income households hit the Social Security tax cap in different months, and τ^{ss} falls and after-tax income rises.[26] There is no variation along this dimension between December and January when all households who hit the cap begin to pay Social Security payroll taxes again. Second, for different households, Social Security taxes represent different amounts 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.[27] 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 their becoming affective. Thus all households sometimes experience small changes in after-tax income between December and January.[28]

6. Estimation and Results

Due largely to measurement error in the level of consumption, observations are serially correlated within households. The first-order serial correlation is -0.4 while higher orders are insignificantly different from zero.[29]

The first result in the first column of Table 6.1 reports the results of ordinary least squares (OLS) estimation of equation (5.1). The point estimate implies that when Social Security taxes paid fall by 10% of income (a typical amount when the tax cap is hit), nondurable consumption rises by 5.4%. Reported standard errors are consistent (as $N \to \infty$) in the face of arbitrary heteroskedasticity and

[26] After 1992, the tax caps for OASDI and HI differ and so an individual can hit each cap at a different times during a year.

[27] As discussed in more detail in the next section, the payroll tax rate is adjusted from the percent of gross labor income to the percent of after-tax family income using information reported in the first interview about the retrospective 12 months.

[28] The next setion explores the consequences of focussing on different sources of identification.

[29] With no noise in the consumption data, time aggregation (aliasing) implies that change in consumption should follow an $MA(1)$ with a positive coefficient of 0.25. Under the assumptions that consumption is truly a random walk and the measurement error is classical, observed first-order serial correlation of -0.4 implies that 87 percent of the variance in observed change in consumption is due to measurement error. Given that these are the best consumption data available, it is not surprising that conclusive tests of the linear Euler equation are rare.

Table 6.1: The Reaction of Nondurable Consumption

Dependent Variable	(1) Month & Year Dummies OLS	(2) Month & Year Dummies FGLS	(3) Time Dummies OLS	(4) Time Dummies FGLS
%ΔNondurable Cons.	−0.544	−0.494	−0.550	−0.515
NT= 146,650	(0.171)	(0.158)	(0.171)	(0.159)
%ΔNondurable Non-Gift Cons.	-0.461	−0.420	−0.468	−0.442
NT= 146,737	(0.173)	(0.158)	(0.173)	(0.158)

within-household serial correlation.[30] The coefficient is statistically significant at the 99% level.

One can also estimate equation (5.1) efficiently using a Feasible Generalized Least Squares ($FGLS$) estimator.[31] Assuming that the temporal variables capture any cross-household correlation, the covariance matrix, Ω, is specified as block diagonal and the same across households. Each block Ω^h, the within-household covariance matrix, is estimated nonparametrically as:

$$\widehat{\Omega}^h_{ij} \equiv N_{ij}^{-1} \sum_{n=1}^{N} e_{ni} e_{nj}$$

where $i, j \epsilon (1, 2, 3)$ denote the elements of Ω^h, N_{ij} is the number of households for which both the i^{th} and j^{th} change in consumption is observed, and e_{ni} is household n's i^{th} OLS residual. As shown in the first row and the second column of Table 6.1, $FGLS$ estimation gives a slightly lower but still highly significant point estimate of -0.494.

One could however remain skeptical on the grounds that the month and year dummies might not fully capture the seasonal in consumption, caused in part by Christmas. That is, consumption rises in November and December and declines in January. If this "Christmas effect" were larger in years in which the Social Security tax rate increased, this could lead to spuriously large coefficients. Put

[30] For OLS, standard errors are

$$(X'X)^{-1}(\sum_{n=1}^{N} X_n' e_n e_n' X_n)(X'X)^{-1}$$

where X is the full $NT \times K$ matrix of data, $X_n$is the $T_n \times K$ matrix of data for a household, and e_n is the vector of OLS residuals for that household. These errors are also consistent in the presence of household-specific random effects in first-differences (equation (5.1)).

[31] The FGLS procedure is efficient if the covariance matrix is correctly specified. The point estimates remain consistent regardless.

slightly differently, perhaps there are macroeconomic effects contemporaneous with the January increases in the Social Security tax rate which cause consumption to grow less in the six years of the increases than in the seven years without. However, as reported in the last two columns of the first row of Table 6.1, when a complete set of time dummies (164) are included in the regression instead of month and year dummies the estimated coefficients actually rise. Further, one can take steps to eliminate the effect of Christmas itself on the estimates. In the CEX expenditure files, every expenditure which is a gift for someone not living in the household is flagged. The second row of Table 6.1 reports the coefficient on the Social Security tax variable when all such gift purchases are excluded from the definition of consumption. The point estimates decline slightly and remain highly significant, implying a 0.45 percent decline in nondurable consumption for every one percent decline in after-tax income.[32] This may not be a better estimate of the true elasticity, however. It is possible that one of the more responsive categories of consumption expenditure is the purchase of gifts for people outside the immediate household. Nevertheless, the remainder of the paper excludes gift purchases from all measures of consumption so that any possible bias is bias in favor of the null hypothesis.[33]

Before exploring the sources of identification and robustness of this finding, I pursue the possibility that measurement error in the Social Security tax variable could be *attenuating* the coefficients in Table 6.1. The income data are not of the quality of the income data in the PSID, for example.[34] However, in addition to the income data employed so far, which are the individual's reported earnings over the past 12 months, each individual is also asked the amount of their last paycheck and the length of the pay period.[35] From this alternative measure of annual income, another measure of $\Delta\tau^{ss}$ is constructed. The first column of Table 6.2 shows the results of FGLS estimation when the average of the two measures of $\Delta\tau^{ss}$ is used. The estimated coefficient does indeed rise slightly, consistent with there being measurement error in both income reports.[36]

As mentioned in section 4, the construction of $\Delta\tau^{ss}$ uses information from the end of each household's tenure in the survey. This technically violates the information structure of the test, although it seems unlikely to have generated the results

[32] There are not enough households with only one member to accurately identify the coefficient of interest in such a subsample.

[33] An alternative specification one might consider is to perform the analysis in levels. The results of this experiment lead to even greater statistical significance, but much smaller economic effects. Another question one might ask is whether the results in the logarithmic specification change when outliers are not dropped. Including outliers led to standard errors of around 0.19 and a slight decline in coefficients so that the eight reported in Table 6.1 averaged -0.404.

[34] See Lusardi (1996).

[35] This income measure is similar to that in the PSID and is likely to be superior to the main variable which is most commonly used.

[36] I do not instrument one measure with the other since classical measurement error in income generates non-classical measurement error in the Social Security tax variable which in turn causes two-stage least squares estimates to be inconsistent.

Table 6.2: Alternative Measures of the Social Security Tax Variable

Independent Variable:	Avg. of Measures	Interview 2 Measure	Average of Int. 2 Meas.	Avg. Int. 2 Meas. Dropping Int.2
Coefficient:	−0.514	−0.407	−0.508	−0.684
Std. Error:	(0.180)	(0.165)	(0.183)	(0.222)
NT:	134,512	146,737	134,512	79,659

FGLS estimation on non-gift, nondurable consumption.

so far. To examine this, the Social Security tax rate is recalculated using only information from each household's second interview.[37] The results are reported in the second column in Table 6.2. The third column reports the results of using the average of both Social Security tax measures from the second interview information to estimate equation (5.1). The estimated coefficients are again consistent with the presence of measurement error and provide further confirmation that consumption falls by half a percent for every percent increase in the effective net Social Security tax rate. Finally, since 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, equation (5.1) is re-estimated dropping all the changes in consumption between the second and third interviews from the sample. As the final column of Table 6.2 shows, the coefficient actually rises nearly one standard deviation to 0.68. Thus, these finding are not driven by the use of information from household's final interviews to construct $\Delta\tau^{ss}$.

While the previous results would seem a strong 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 are just less than 1%. This is due, in part, to the fact that the Social Security tax changes are small relative to the swings in consumption (true and due to measurement error).[38] Second, the sample size is very large.[39] Under these conditions, only a small correlation between the error term and the independent variable could perhaps cause the significant results. Note, however, that such a story is likely to be at odds with the fact that a better measure of $\Delta\tau^{ss}$ increases the estimated coefficient. Further, one would have to argue that it is pure coincidence that the coefficient estimate is negative and of a reasonable

[37] This is the *first* interview in which consumption data are collected.

[38] There is no problem of the sort encountered in weak intrumental variables regressions. The point estimates are unbiased and the standard errors are correct for making statistical inference.

[39] Actually, only about twelve percent of observations have nonzero $\Delta\tau^{ss}$.

Table 6.3: High Consumption and Tax Cap Subsamples

Dependent Variable:	Main Measure	Average of Measures	Average of Int. 2 Meas.
High Consumption Sample:			
Coefficient:	−0.475	−0.267	−0.432
Std. Error:	(0.248)	(0.293)	(0.290)
NT:	17,756	15,792	15,792
Households Hitting Cap:			
Coefficient:	−0.393	−0.196	−0.413
Std. Error:	(0.228)	(0.246)	(0.259)
NT:	9,358	11,011	9,688

FGLS estimation on non-gift, nondurable consumption.

magnitude.

Nevertheless, consider the following scenario. Suppose that households for which the small increases in Social Security tax rates represent the largest fraction of after-tax income are also those households with the largest consumption declines between December and January. Alternatively, suppose that larger families have larger increases in consumption around Christmas and larger incomes so are more likely to hit the tax cap.[40] If either of these stories were true, the change in the Social Security tax variable would be spuriously correlated with the change in consumption. Two tacks are taken to eliminate such possible interpretations.

First, the analysis is performed on two subsamples: households in which average consumption exceeds $30,000 (1987) dollars a year; 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. The standard deviation of $\Delta\tau^{ss}$ in the high consumption subsample is 1.0%, and in the cap-hitting subsample, 2.1%. These numbers are roughly double and quadruple that in the entire sample, respectively. Table 6.3 shows that the coefficients decline slightly and the standard errors increase from those garnered from the full sample. Identification can rest almost exclusively on when households hit the Social Security tax cap and the estimated coefficients still tell the same story.

The second tack is more elaborate. The institutional structure provides a control group of identical households who do not pay Social Security taxes and one

[40] These stories are simply more elaborate versions of the critique raised in Mariger and Shaw (1993).

can identify the coefficient of interest using only the variation across otherwise identical households that do and do not pay Social Security taxes. This should eliminate any spurious correlation between change in consumption and $\Delta\tau^{ss}$.

Each earner in the dataset is assigned to one of three groups: a treatment group comprised of those individuals who are almost certainly 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.[41] Assignment is made on the basis of five sets of information from individuals' fifth interviews.[42] First, the CEX contains information on both occupation and industry of the reference person and spouse. Individuals who are self-employed or, prior to 1984, Federal government employees, are assigned to the control group.[43] Second, individuals are asked to estimate their contributions to Social Security over the past 12 months. Starting in 1982, individuals also are asked whether Social Security is normally deducted from their paycheck, and, starting in 1986 whether this deduction covers Medicare only. These variables are used to assign a second set of individuals to groups. Third, individuals who change jobs or work more than one job may overpay Social Security taxes and thus the time at which they hit the tax cap is uncertain. The CEX asks whether households overpaid Social Security in the last 12 months, and any individual in such a household is assigned to the neither group. Fourth, if an individual spends time unemployed they may not be employed during any tax change and it is not possible to accurately calculate when they will hit the tax cap. Thus any individual reporting fewer than 50 weeks worked in the past year is put into the neither group. Finally, the family reporting of contributions to Railroad Retirement, which was merged with the Social Security system in 1985, is used to make a further set of assignments. Appendix A give a complete description of the allocation to treatment, control, and neither groups. In sum, anyone in the treatment group is almost certainly employed and paying Social Security taxes and the calculated timing of when they hit the tax cap is quite accurate.

From a hypothetical Social Security tax variable for every individual, two additional variables are constructed: $\Delta\tau^{ssE}$, the change in the Social Security tax rate that every household would have experienced had it been paying Social Security Taxes; and $\Delta\tau^{ssN}$, the hypothetical change in the Social Security tax rate for households that are in the "neither" group. The new estimating equation is:

$$\Delta\ln(c_{it+1}) = \alpha_1 z_i + \alpha_2\Delta\tau_t^{ss} + \alpha_3\Delta\tau_t^{ssN} + \alpha_4\Delta\tau_t^{ssE} + \alpha_5 m_t + \alpha_6 y_t + \epsilon_{t+1}. \quad (6.1)$$

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

[41] All previous results only employ the treatment group.

[42] If the final interview is missing, information from the second interview is used.

[43] Self-employed individuals pay Social Security taxes, but they do so as part of their paying quarterly estimated tax payments. Thus there is no withholding pattern which directly affects after-tax income. After 1984, some Federal government employees were covered by Social Security instead of Federal government retirement plans.

Table 6.4: TREATMENT VS. CONTROL GROUP REGRESSIONS

	MAIN MEASURE	AVERAGE OF MEASURES	AVERAGE OF INT. 2 MEASURES
ENTIRE SAMPLE:	−0.020	−0.696	−0.725
	(0.437)	(0.538)	(0.541)
NT:	146,737	134,512	134,512
INDIVIDUALS:			
% IN TREAT:	65	66	66
% IN CONTROL:	5	4	4
% IN NEITHER:	31	29	29
HOUSHOLDS HITTING CAP:	−0.183	−0.711	−0.899
	(0.484)	(0.583)	(0.591)
NT:	9,358	11,011	9,688
INDIVIDUALS:			
% IN TREAT:	73	71	71
% IN CONTROL:	8	7	7
% IN NEITHER:	19	21	22

FGLS estimation on non-gift, nondurable consumption.

NT reports the number of oservations used in the regressions, The percent of invdividuals is the number of heads or spouses with positive earnings in each group divided by the total number with positive earnings.

the control group and that of the treatment group.

The top half of Table 6.4 displays the share of earners in each group and the results of estimating equation (6.1) for the entire sample. The point estimates are noisier than earlier results and are statistically insignificant. However, the average across the columns still suggests an elasticity just under 0.5. The large standard errors are consistent with the fact that the control group is quite small. If it were the case that consumption happened to decline across households in Januaries in the same way in which Social Security taxes caused after-tax incomes to decline, then one would expect these point estimates to return to zero. In fact consumption falls for those households who are paying Social Security taxes and not for those that are not. As the second half of Table 6.4 shows, estimation on the much smaller subsample of those households who hit and would have hit the Social Security tax cap yields the same conclusion with perhaps even more bite.

The stability of the point estimates across alternative specifications, subsamples, and identification strategies strongly suggests that the significant coefficients found for the previous regressions are not driven by odd cross-household seasonal patterns in consumption.

7. Alternative Models of Consumption Behavior

This section discusses and evaluates three alternative theories of consumption behavior which imply that consumption should to some respond to expected income changes.

First, the linear Euler equation might fail because the expected variance term in equation (2.5) is correlated with the included time t variable.[44] Consider as an example, the realistic specification in which income is lognormally distributed and, as above, the utility function is CRRA. First, note that in this case households are *unwilling* to borrow against future income. Why? If a household were to arrive in its last period of life with negative assets and it were to receive an income of less than this amount, then it would have to consume nothing. Since this yields infinite marginal utility, the household would never allow this to be a possibility. By iterating backward though life, it follows that households will never borrow. Thus a household with few assets and low current income consumes roughly its income – rather than some fraction of its permanent or life-cycle income. If income is expected to grow, as it does so, consumption will rise. Consequently, expected income growth and asset levels can be used to predict consumption growth. The mechanism for this relation is exactly the expected variance of consumption: because a household with few assets is unable to smooth consumption in response to negative shocks to income, its expected variance of consumption is high and its expected consumption growth is high.[45]

This point is more general than the example above. Carroll and Kimball (1996) demonstrate that a wide range of utility functions (the HARA class) and income processes generate consumption functions which depend upon the share of the expected present discounted value of total wealth that is comprised of currently available assets. To summarize, this alternative hypothesis predicts that the variance term in the consumption Euler equation is correlated with household assets and thus expected income growth should have predictive power for consumption growth. The relationship should be strongest for households with few assets.

A second reason for failure of the linear Euler equation is that some households might face liquidity constraints.[46] When a household is liquidity constrained, its Euler equation ceases to hold and it is *unable* to smooth consumption. With

[44] The role of uncertainty in linear Euler equation violations was first explored by Zeldes (1989b) and then Deaton (1991). Carroll (1997) demonstrates that within this framework, assets are closely correlated with the variance term. See also Gourinchas and Parker (1996), Dynan (1993), and Caballero (1990) for studies that evaluate the importance of this mechanism.

[45] This requires prudence of the utility function, a feature which CRRA exhibits. See Kimball (1990).

[46] This setup is motivated by that considered by Hubbard, Skinner and Zeldes (1994).

liquidity constraints, predictable rises in income, as would occur following a temporary layoff, for example, can lead to corresponding rises in consumption. For concreteness, consider the following additional two constraints on the household optimization problem (equations (2.1) to (2.4)):

$$A_t \geq 0 \; \forall t \tag{7.1}$$

$$c_t = \widehat{c} \text{ if } A_{t+1} = 0 \text{ and } c_t \leq \widehat{c} \tag{7.2}$$

where $\widehat{c}$ is a government-provided strictly positive floor on consumption. Equation (7.1) implies that due to liquidity constraints, households are unable to borrow against future income, while equation (7.2) implies that if the household consumes all its income and assets, the government provides a consumption floor of $\widehat{c}$. In this scenario, a household with no assets that has high expected future income is unable to consume more than $\widehat{c}$ due to liquidity constraints. When the income of such a household rises, consumption rises also. Thus, as in the first alternative considered, expected income growth should have predictive power for consumption growth for low-asset households.[47]

To evaluate whether the linear Euler equation fails due to liquidity constraints or a correlation between the expected variance of consumption and the growth rate of consumption, I now test whether the relationship between expected income growth and consumption growth is strongest for those households with few liquid assets. Unfortunately, asset information in the CEX is quite limited and is often topcoded or missing. Further, the CEX only asks asset information in the final interview. However, this interview contains information about 4 categories of liquid assets and about how much the wealth in each category has changed over the past 12 months. From these sets of information, one can construct a measure 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.[48] The ratio of this variable to average annual household 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.

There is little evidence that the Euler equation failure is concentrated among households with the fewest assets. The first row of Table 7.1 reports the results of FGLS estimation of the linear Euler equation with the Social Security tax variable and an interaction between this variable and a dummy variable which is set to one if the household has more than enough assets to finance 0.4 years of consumption or if any asset category is topcoded.[49] Households with the ability to finance more

[47] Note that in both of the scenarios discussed, the possibility of having low assets (or being liquidity constrained) in the future changes the consumption decision even for households with many assets. That is, the basic LCH/PIH is not strictly true even for households with large wealth holdings. However, for these households, the linear Euler equation should provide a close approximation to the non-linear Euler equation.

[48] Details are provided in Appendix *A*.

[49] Typically, studies estimate that, for those consumers who are estimated to be buffer stock

Table 7.1: Asset and Age Splits

Interaction Variable	NT	$\Delta\tau^{ss}$	$\Delta\tau^{ss}$ Interaction
High Asset Ratio	78,530	−0.322	−0.527
		(0.242)	(0.400)
Young Age	134,512	−0.298	−0.285
		(0.227)	(0.308)
Results for Treatment vs. Control			
High Asset Ratio	78,530	−1.872	1.101
		(0.784)	(1.463)
Young Age	134,512	−0.482	−0.464
		(0.645)	(0.938)

FGLS estimation using average of measures and non-gift, nondurable consumption.

than 0.4 years of consumption from assets seem to react more to the changes in Social Security tax rates than do low-asset individuals, although the result is not statistically significant. Row 3 of Table 7.1 shows the same regression but identifying the coefficients using the difference between treatment and control group reactions; the results are reversed, although again with large standard errors. The size of these estimates and the reversal of sign makes it difficult to take much from this experiment.[50] It may well be that since the asset data is 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 which the coefficients reflect.

As another tack, age can be used to proxy for the probability of being liquidity constrained or having few assets.[51] Young households have larger expected income growth and fewer assets than older households. Due to either liquidity constraints or the optimal choice not to borrow, young households may be more likely to violate the linear Euler equation. Rows 2 and 4 of Table 7.1 report a similar pair of regressions to those for the asset classification but based on whether a household is younger than 43.[52] Gourinchas and Parker (1996) 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. While again the coefficients are

consumers, the target level of cash-on-hand to permanent income around 0.4. See Gourinchas and Parker (1996) or Carroll (1997)

[50] Interacting the ratio itself leads to similar inconclusions.

[51] Japelli et al. (1995) find that age is a significant predictor of whether a household reports that it is liquidity constrained.

[52] I also drop households older than 70.

not statistically significant, they are consistent with the hypothesis that younger households react more strongly to changes in expected income.

While there is at best weak, suggestive evidence for the relevance of liquidity constraints of precautionary savings, it should be noted that, for everyone except those who hit the tax cap, the changes in Social Security taxes are always increases. In the face of liquidity constraints or buffer-stock behavior, consumers do a much better job of smoothing consumption across expected declines in income. Further, since I discard individuals who are not fully employed during the previous 12 months, the sample has fewer candidates to be liquidity constrained than the population which is typically studied.

As a third alternative hypothesis which can be evaluated in the current context, 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 level of consumption which a fully-rational agent would choose.[53] For example, household may face constraints on time which make it optimal for them to ignore the complex tax code when calculating $E_t[\Delta Y_{t+1}]$, and simply to choose a marginal propensity to consume out of income innovations between occasional state of time dependent optimizations. 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. Since small deviations from optimal behavior have small costs, only small deviations from full rationality or small costs of shifting income across time are required for this story to be plausible.

The testable implications of this alternative theory are twofold. First, expenditures should track income more closely for goods that are purchased using current income– that is, that are not purchased using credit. Second, expenditures should track income more closely for goods with high intertemporal elasticities of substitution of expenditures– that is, for which swings in consumption expenditure provide little utility loss. Many goods typically considered nondurable are in fact semi-durable or storable, or provide lasting utility over three-month intervals. The effective intertemporal substitutability of these goods is much higher.[54] 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. Thus expenditures should respond to expected income most strongly for goods that are generally financed out of current cash on hand, and that are semi-durable, can be stored, or provide some lasting utility flow.

[53] This hypothesis is similar in spirit to that proposed by Hall and Mishkin (1982). Here "boundedly rational" and "fully-rational" apply to the behavior interpretted within the context of the model of Section 2. That is, boundedly rational behavior is not meant to imply that the behavior is not prefectly rational within the context of a more detailed model in which the contraints on calculation, information gathering, or time are made explicit.

[54] Heaton (1993) and Parker (1996) both explore the implications of higher intertemporal substitution due to lasting utility/durability.

Table 7.2: The Reaction Across Categories of Consumption

Dependent Variable Consumption Category	NT	Average of Measures	Average of Int. 2 Measures
Total	129,165	−0.294	−0.326
		(0.183)	(0.187)
Food and Alcohol	131,856	0.119	0.002
		(0.181)	(0.184)
Entertainment and Personal Care	129,178	−0.802	−0.704
		(0.333)	(0.340)
Apparel and Services	105,249	−2.729	−2.459
		(0.424)	(0.433)

In regressions on total and food consumption, as for nondurable consumption, observations are discarded if consumption changes more than 100% between quarters. For the other categories, the cutoff is 200%. These cutoffs are all around two standard deviations.

More generally, however, durability or lasting utility can alter the specification of tests of the basic LCH/PIH. Specifically, if there are adjustment costs associated with upgrading durable goods, then shocks to income feed through to purchases with delays. Caballero (1993) finds that expenditures on durable goods do not follow the theoretically predicted first-order moving-average process and argues that this is due to nonconvex costs of adjusting the stocks of these goods. If this is true, then the timing of these purchases is influenced by the cumulation of innovations to household wealth since the previous adjustment (which took all information known at that point into account). This can invalidate some tests because shocks to income cause delayed increases in income and consumption expenditures. The current test is robust to this specification. If the announcement of Social Security taxes contains news to the household about its permanent income, there is no reason for the expenditure which adjusts the household's stock of durables to occur at the same time as the expected income fluctuation; it is just as likely to occur in the months before or after.

Table 7.2 shows the results of estimating equation (5.1) across subcategories of goods. The elasticity of total expenditures to expected changes in take-home pay is estimated as around 0.3, suggesting that nondurable consumption reacts more that durable consumption.[55] One interpretation of this finding is that there are adjustment costs in the sale and purchase of durables and therefore the purchase

[55] While the standard error is similar to those found in the nondurables regressions, the lower point estimate implies that the response is statistically insignificant.

of durables does not react as quickly or as much as the purchase of nondurables. An alternative explanation is that durable goods are often purchased with credit, as in an automotive loan, so that there is no reason for expenditures on durable goods to be linked to fluctuations in income as expenditure on nondurables are.

Rows 2 through 4 of Table 7.2 show the results for estimation involving three subcategories of nondurable consumption. First, note that food consumption does not seem to respond at all to expected changes in income. This is not entirely consistent with the alternative theory at hand, since dinners out and alcohol purchases can both be easily timed to track income.[56] However, food consumption also consists of purchases of necessities which are relatively inelastic. Apparel and services consumption reacts the most of all categories, with an elasticity around 2.5. Since apparel and services make up roughly 5.7% of after-tax income, these estimates suggest that for every extra expected dollar of income which a Social Security tax change induces, 17 cents are spent on apparel and services.[57] Finally Entertainment and Personal Care, which also includes expenditures on reading materials and tobacco and smoking supplies, shows a just slightly stronger reaction than total nondurable consumption.[58]

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, it is precisely such households who are most likely 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.

8. Conclusion

This paper finds that consumers do not perfectly smooth consumption at high frequencies across expected income changes. Contrary to the basic LCH/PIH, consumption reacts to predictable changes in tax rates. The findings are economically significant: consumption falls roughly half as much as the expected decline in income.

I evaluate three alternative theories to the linear Euler equation, each of which can explain why the basic PIH/LCH fails. Tests of models of precautionary saving

[56] It is interesting to note that this is the only subcategory of consumption which can be used by PSID-based studies.

[57] This number is quite sensitive to outliers. The Table reports results when observations in which apparel purchases change more than 200% between quarters are dropped. If one instead makes the cut at 100%, one finds an elasticity around unity.

[58] The pattern and magnitude of coefficients is similar in treatment vs.control group regressions, although the coefficients are more volatile and the standard errors increase on a scale almost exactly as that for nondurable consumption.

or liquidity constraints, which evaluate whether low-asset households have stronger violations of the linear Euler equation, yield inconclusive results. Some scant evidence is found that young households, which generally have lower assets and higher income growth, are more likely to have consumption track these expected changes in income. 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 and in which purchases are generally made with cash or short-term credit.

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, these tax changes may come largely as surprises to households who do not take the time and effort to understand the Social Security tax code. Thus, consumption rises because the analyzed changes in after-tax income surprise many households. Note however that if this theory is to rationalize the magnitude of the violation, households must expect the income shock caused by hitting the tax cap to be quite persistent.

Whatever interpretation one chooses to take, the mounting evidence against consumption smoothing has far-reaching implications. First, if expected changes in tax rates influence contemporaneous consumption behavior, then fiscal stabilization, such as that undertaken by President Bush in 1992 or that provided by automatic fiscal stabilizers, may have large and important effects on consumption. Second, while the results of this paper do not imply that consumption smoothing is a poor approximation over long horizons, they do imply that when studying shorter horizons, such as those addressed in business cycle models, the linear Euler equation should not be employed as a close approximation to the correct structural equation.

References

Altonji, Joseph and Aloyisius Siow, "Testing the Response of Consumption to Income Change with (Noisy) Panel Data," *Quarterly Journal of Economics*, 1987, *103*, 293–328.

Attanasio, Orazio and Guglielmo Weber, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, Dec 1995, *103* (6), 1121–57.

Bound, John, David A. Jaeger, and Regina Baker, "Problems with Instrumental Variables Estimation when the Correlation between the Instruments and the Endogenous Explanatory Variable is Weak," *Journal of the American Statistical Association*, Jun 1995, *90*, 443–50.

Browning, Martin and Annamaria Lusardi, "Household Saving: Micro Theories and Macro Facts," *Journal of Economic Literature*, Dec 1996, pp. 1797–1855.

Bureau Of Labor Statistics, *Consumer Expenditure Survey, 1993: Interview Survey*, Washington, DC: US Department of Commerce, Bureau of the Census, 1980-1993. Ann Arbor, MI: Inter-university Consortium of Political and Social Research.

Caballero, Ricardo J., "Consumption Puzzles and Precautionary Saving," *Journal of Monetary Economics*, 1990, *25*, 113–36.

__ , "Durable Goods: An Explanation for their Slow Adjustment," *Journal of Political Economy*, April 1993, *101*, 351–84.

Campbell, John Y., "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis," *Econometrica*, Nov 1987, *55* (6), 1249–73.

Carroll, Christopher D., "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, Feb 1997, *107* (1), 1–56.

__ **and Miles S. Kimball**, "On the Concavity of the Consumption Function," *Econometrica*, Jul 1996, *64* (4), 981–92.

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, 1991, *59*, 1121–1142.

__ , *Understanding Consumption*, Oxford: Clarendon Press, 1992.

Dynan, Karen, "How Prudent Are Consumers?," *Journal of Political Economy*, Dec 1993, *101*, 1104–1113.

Friedman, Milton, *A Theory of the Consumption Function*, Princeton: Princeton University Press, 1957.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption Over the Lifecycle," 1996. MIT mimeo.

Hall, Robert E., "The Stochastic Implications of the Life Cycle Permanent Income Hypothesis," *Journal of Political Economy*, 1978, *86*, 971–87.

__ **and Frederic S. Mishkin**, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, 1982, *50*, 461–481.

Hayashi, Fumio, *Understanding Saving: Evidence from the U.S. and Japan*, Cambridge, MA: MIT Press, 1997.

Heaton, John, "The Interaction Between Time-nonseparable Preferences and Time Aggregation," *Econometrica*, Mar 1993, *61* (2), 353–85.

Hubbard, Glenn, Jonathan S. Skinner, and Stephen P. Zeldes, "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," in Allan H. Meltzer and Charles I. Plosser, eds., *Carnegie-Rochester Conference Series on Public Policy*, Vol. 40 Jun 1994, pp. 59–125.

Japelli, Tulio, Jörn-Steffen Pischke, and Nicholas Souleles, "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," 1995. MIT mimeo.

Kimball, Miles S., "Precautionary Saving in the Small and in the Large," *Econometrica*, 1990, *58*, 53–73.

Lusardi, Annamaria, "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets," *Journal of Business and Economic Statistics*, 1996, *14* (1), 81–90.

Mankiw, N. Gregory, "Hall's Consumption Hypothesis and Durable Goods," *Journal of Monetary Economics*, Nov 1982.

Mariger, Randall P. and Kathryn Shaw, "Unanticipated Aggregate Disturbances and Tests of the Life-Cycle Consumption Model Using Panel Data," *The Review of Economics and Statistics*, Feb 1993.

Modigliani, Franco and Richard Brumberg, "Utility Analysis and the Consumption Function: an Interpretation of Cross-Section Data," in Kenneth K. Kurihara, ed., *Post-Keynesian Economics*, New Brunswick, N.J. Rutgers University Press 1956, pp. 388–436.

Parker, Jonathan A., "The Timing of Purchases, Market Power, and Economic Fluctuations," 1996. MIT mimeo.

Paxson, Christina H., "Consumption and Income Seasonality in Thailand," *Journal of Political Economy*, Feb 1992, *101* (1), 39–72.

Runkle, David, "Liquidity Constraints and the Permanent Income Hypothesis," *Journal of Monetary Economics*, 1991, *27*, 73–98.

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

__ **and Joel Slemrod**, "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding," *American Economic Review*, 1995, *85*, 274–83.

Shea, John, "Union Contracts and the Life Cycle-Permanent Income Hypothesis," *American Economic Review*, 1995, *85*, 186–200.

Social Security Administration, *Social Security Bulletin Annual Statistical Supplement*, United States Department of Health and Human Services, Government Printing Office, 1990.

Souleles, Nicholas, "The Response of Household Consumption to Income Tax

Refunds," 1995. M.I.T. Mimeo.

Zeldes, Stephen P., "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, Apr 1989, *97*, 305–346.

__ , "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, May 1989, *104*, 275–98.

Appendices

A. The Consumer Expenditure Survey.

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 is obtained from Bureau Of Labor Statistics (1980-1993) and conversations with Bureau of Labor Statistics (BLS) statisticians. Households should not be matched across 1985 to 1986, and are not. Care is taken to assure consistency in the data despite variable classification changes through time, and across reference person and spouse. Information was kindly provided by the Division of the CEX in the BLS about various issues including the matching of occupation codes from $1980-81$ to later years.

As discussed in the text, households are discarded if they are missing any of the information necessary for the regressions, or if they are classified as incomplete income reporters or if family income is topcoded or missing in the second interview. 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, I reset reported tax rates above 60% to 60%, and those below 0 are reset to 0. The results are quite insensitive to alternatives to this correction.

Individuals are assigned to treatment, control and neither groups as described in the following paragraphs. At the end of the procedure, any individual assigned to both the treatment and the control group is moved into the neither group.[59] Individuals without labor income are not assigned to any group.

An individual is assigned to the control group if he or she reports his or her industry as Federal government employee prior to 1984 or occupation as self-employed in any year. All remaining individuals who report their industry as any level of government employee are placed in a government category. Further, 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 in the government group which reports paying Social Security in their 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.

Next, any individual who is missing industry or occupation data is put into the treatment group unless they report not paying Social Security taxes in either of these variables. Households which report overpaying Social Security in the past 12 months are put into neither. Individuals who report working less than 50 weeks in the past 12 months are put into neither. Finally, if a household reports paying into a Railroad retirement account, if there is only one earner, that earner is put into the control group prior to 1985 and into the treatment group after 1985. If the family has multiple earners, all earners are assigned to the neither group.

Consumption data is compiled from the detailed expenditure files. Monthly expenditures are averaged to get consumption at a quarterly rate. The denominator for the

[59] The procedure never assigns an individual into the neither group and one of the other groups.

average is the number of nonzero consumption months for nondurable consumption. Thus for example, if no apparel expenditures are reported but nondurable consumption is positive, this is considered a valid month of data on apparel purchases. If nondurable purchases are zero, then there is assumed to be no data for this month. Some consumption observations are reported to have occurred prior to the three months recall period for an interview and some in a later month. BLS statisticians recommend treating these expenditures as if they occurred in the reported month. Households with only 1 or 2 months of consumption data (one percent of the sample) were dropped. Households with more than 12 months of data have the last few observations dropped unless the first observation is less than the 13th in which case the first observation was dropped and the next 12 used.

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. Apparel and Services is simply this CEX category. Entertainment and Personal Care is the sum of Entertainment, Personal Care, Reading, and Tobacco and Smoking expenditures. Nondurable consumption is defined in the main text.

Observations with large changes in consumption between adjacent 3-month periods are dropped. For total, food, and nondurable consumption, observations which had changes in excess of 100% were dropped. For both Apparel and Services and for Entertainment and Personal Care observations were dropped only if the change in consumption exceeded 200%. All cuts are just around two standard deviations in the ex ante data.

The main measure of income is constructed by extrapolating using the second and fifth interview reports to make a best guess about income for any month. If income changes by more than 25% between the second and fifth reports, or if the second interview measure is missing, I simply use the fifth interview measure. If the fifth measure is missing, I use the second.

Assets are considered missing if both savings and checking account information is missing. When this account information is not missing, bonds and stocks amounts were added to the amounts in the accounts to create total asset measures. Households with annual total non-gift consumption less than $\$1,000$ are dropped from the regressions with asset interactions.