

Princeton University

Household Expenditure and the Income Tax Rebates of 2001

David S. Johnson, Jonathan A. Parker and Nicholas S. Souleles

Discussion Paper #231

August 2004

Discussion Papers in Economics

Woodrow Wilson School
of Public and International Affairs



Household Expenditure and the Income Tax Rebates of 2001

David S. Johnson, Jonathan A. Parker and Nicholas S. Souleles*

August 2004

Abstract

Under the Economic Growth and Tax Relief Reconciliation Act of 2001, most U.S. taxpayers received a tax rebate between July and September, 2001. The week in which the rebate was mailed was based on the second-to-last digit of the taxpayer's Social Security number, a digit that is effectively randomly assigned. Using special questions about the rebates added to the Consumer Expenditure Survey, we exploit this historically unique experiment to measure the change in consumption expenditures caused by receipt of the rebate and to test the Permanent Income Hypothesis and related models. We find that households spent about 20-40 percent of their rebates on non-durable goods during the three-month period in which their rebates were received, and roughly another third of their rebates during the subsequent three-month period. The implied effects on aggregate consumption demand are significant. The estimated responses are largest for households with relatively low liquid wealth and low income, consistent with liquidity constraints.

Keywords: consumption, saving, Life-Cycle model, Permanent-Income Hypothesis, liquidity constraints; fiscal policy, tax cuts, tax rebates, windfalls.

JEL classification: E21, E62, H31

* Author affiliations: Johnson: Division of Price and Index Number Research, Bureau of Labor Statistics, Washington, DC 20212; Parker: Department of Economics, Bendheim Center for Finance, and Woodrow Wilson School, Princeton University, Princeton, NJ 08544. Souleles: Finance Department, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6367. We thank Orazio Attanasio, Alan Blinder, Christopher Carroll, Angus Deaton, Karen Dynan, Mary McCarthy, John Sabelhaus, and Matthew Shapiro for their help in getting the 2001 tax rebate questions added to the CE Survey, and the staff of the BLS, especially Mary McCarthy, for expeditiously adding the questions. We also thank for their comments Miles Kimball, Jeffrey Kling, David Laibson, Matthew Shapiro; discussants Nicola Fuchs-Schuendeln, Julia Coronado, and Norbert Michel; and participants in seminars at the July 2004 NBER Economic Fluctuations and Growth meeting, the Bureau of Labor Statistics, Princeton, Michigan, Harvard, the American Enterprise Institute, the Federal Reserve Bank of New York, the Board of Governors of the Federal Reserve System, the July 2003 NBER Summer Institute, the Federal Reserve Bank of Philadelphia, and the University of Pennsylvania. Parker acknowledges support from National Science Foundation grant SES-0096076, the Sloan Foundation, and a National Bureau of Economic Research Aging and Health Economics Fellowship through National Institute on Aging grant number T32 AG00186. The analyses and conclusions expressed in this paper are those of the authors alone, and do not necessarily coincide with those of the Bureau of Labor Statistics.

I. Introduction

Policymakers often try to use tax policy to reduce the magnitude of economic fluctuations. They cut income taxes in recessions, assuming that the resulting increase in disposable income increases household spending, thereby reducing the severity of recessions. Academic economists, however, tend to be more skeptical about the use of tax policy to stabilize economic fluctuations, in large part because the canonical theory of the consumer suggests that consumption should not respond much to temporary changes in taxes, such as a one-time tax rebate. Instead, any increase in consumption in response to a tax cut should in theory be spread out over consumers' entire lifetimes (or forever), which implies only a small change in spending in the short-term after a tax cut. Moreover, consumption should increase as soon as consumers learn of an upcoming tax cut, rather than wait to increase only once disposable income actually rises.

This paper uses unique data and features of the 2001 income tax rebates to estimate the causal effect of these rebates on household expenditure. The Economic Growth and Tax Relief Reconciliation Act of 2001 sent tax rebates, typically \$300 or \$600 in value, to about two-thirds of U.S. households over a ten-week period from late July to the end of September, 2001. The unique feature of these rebates is that the timing of the mailing of each rebate was based on the second-to-last digit of the Social Security number of the tax filer who received it, a digit that is effectively randomly assigned.¹

The unique data that we use is part of the Consumer Expenditure (CE) Survey, which, among large household surveys in the U.S., contains the most comprehensive measures of households' expenditures. The regular CE data does not contain sufficient information to study the 2001 tax rebates. In particular, the ongoing CE survey does not record the timing of taxes and transfers, nor the Social Security numbers of households' tax filers. However, shortly after the passage of the 2001 Tax Act, the authors worked with the staff of the Bureau of Labor Statistics (BLS) and other government agencies to add a special module of questions about the tax rebates to the Survey. This module asked households about the timing and amount of each rebate check they received, and was included in the survey from shortly

¹The last four digits of a Social Security number (SSN) are assigned sequentially to applicants within geographic areas (which determine the first three digits of the SSN) and a "group" (the middle two digits of the SSN). The main reason for this staggered disbursement schedule is that it was difficult in practice to print and mail the rebate checks all at once. Accordingly, the disbursement schedule was keyed to the randomized social security digit for purposes of fairness.

after the rebate mailing began until the end of 2001. This is the first paper to use the new tax rebate module and exploit the randomized timing of the rebates in the CE.

We estimate the change in household expenditures due to rebate receipt by comparing the expenditures of households who received rebates at different times. The natural experiment provided by the randomized mailing dates allows us to directly identify the causal effect of the rebate. While our results can be interpreted as a test of the Permanent Income Hypothesis (PIH) and related models, the experiment provides identification under much weaker assumptions than required in previous tests of the PIH that rely on the time-series properties of the consumption Euler equation.

We begin our analysis using all of the available information about the rebates, including the magnitudes of the rebates, and all available CE households. We then progressively reduce the variation that we utilize, until we are left with only variation in the timing of when households received their rebates, conditional on receiving a rebate. Given the structure of the data, this leads to progressively smaller samples and less power, with large standard errors in our final specifications. Nonetheless, all of the results suggest that the rebates caused an economically significant increase in spending.

Summarizing the main results, the average household spent about 20-40% of its 2001 tax rebate on nondurable goods during the three-month period in which the rebate was received, depending on the specification. We also find evidence of additional, smaller but still substantial, lagged effects on spending. Roughly two-thirds of the rebate was spent during the quarter of receipt and subsequent three-month period. Any additional lagged effects on spending cannot be estimated with precision.

To shed light on the reasons behind the estimated increase in spending, we contrast the spending responses across different types of households and different subcategories of nondurable goods. Households with low levels of liquid assets and low income spent significantly more of the rebate than typical, consistent with their facing liquidity constraints. While not statistically significant, the point estimates also suggest somewhat larger responses among those with high levels of liquid assets and high income (relative to households with intermediate levels of assets and income). Finally, we also find some evidence that expenditures on food away from home, apparel, and personal care and miscellaneous items responded disproportionately strongly to the rebate.

Given that the Treasury distributed 38 billion dollars in tax rebates, our estimates imply that the rebates directly increased aggregate consumption expenditures by about 0.8 percent in the third quarter of 2001 and 0.6 percent in the fourth quarter of 2001. The ultimate effects of the rebate on the economy also depend on other factors beyond the scope of this paper, such as the extent to which the increased demand for consumption goods caused the relative price of current goods to increase and/or had a multiplier effect.

The paper is organized as follows. The next section relates our paper to the prior literature. The third section describes the relevant tax law changes and the fourth our use of the CE Survey data. The fifth section discusses our empirical methodology. The sixth section presents our main results regarding the short-run response to the rebate, while the seventh section examines the lagged response. The eighth section examines differences in the response across different types of households and consumption goods. A final section discusses the aggregate impact of the rebates and concludes. Appendixes contain additional information about the data.

II. The Literature

Previous tests of consumption smoothing have often had trouble identifying predictable (or transitory versus permanent) changes in income, and separating the effect of a change in income from other factors concurrently impacting the consumption decision (e.g., changes in monetary policy or the stock market). They also usually required the assumption that the characteristics of a household that determine the size or timing of its income change be uncorrelated with other reasons for differential consumption growth rates.² By contrast, the random variation in the timing of the 2001 tax rebates avoids these recurrent problems.

Research using aggregate data to measure whether tax cuts increase consumption expenditures has had difficulty distinguishing the effects of the tax cuts themselves from the economic changes that led to the tax cuts, as well as other concurrent macroeconomic factors. Due also in part to the limited number of significant changes in tax policy, there is a lack of consensus about the effects of tax rebates and other tax changes on consumption

²For example, in his paper discussed below, Ronald G. Bodkin was aware that his insurance dividend variable might have been picking up the correlation of the dividend with omitted variables in turn correlated with permanent income. On adding such control variables to Bodkin's original regression, Roger C. Bird and Bodkin (1965) found smaller spending responses.

expenditures (see Franco Modigliani and Charles Steindel (1977), Alan S. Blinder (1981), Blinder and Angus Deaton (1985), James M. Poterba (1988), and David W. Wilcox (1990)).³

Our paper builds more directly on the literature using micro data to test whether expected or transitory changes in household income affect household consumption expenditures (see the surveys by Deaton (1992) and Martin Browning and Annamaria Lusardi (1996)). The seminal studies of Ronald G. Bodkin (1959) and Mordechai E. Kreinin (1961) examined windfalls like insurance dividends to WWII veterans and German restitution payments. More recently, Jonathan A. Parker (1999), Nicholas S. Souleles (1999, 2002), and Chang-Tai Hsieh (2003) study more directly changes in fiscal policy and use larger, more representative samples.⁴

Two other papers study the impact of the 2001 tax rebates on household spending. Using innovative questions added to the Michigan Survey of Consumers, Matthew D. Shapiro and Joel B. Slemrod (2003) found that only 21.8% of respondents who received (or expected to receive) a rebate report that they will mostly spend their rebate. They calculate that this result is consistent with an average marginal propensity to consume of about one third, very close to the present paper's estimate of the short-run response of expenditures. However, they find no evidence that liquidity constraints play a role in this response and no evidence of a lagged effect on expenditures.⁵

³Modigliani and Steindel, Blinder, and Poterba studied the 1975 tax rebate. They found that consumption expenditures responded too much to the rebate, though they came to somewhat different conclusions regarding the relative magnitude of the initial versus lagged response. Blinder and Deaton (1985) found smaller consumption responses when they considered jointly the 1975 rebate along with the 1968-70 tax surcharge. Nonetheless they found consumption to be too sensitive to the pre-announced changes in taxes in the later phases of the Reagan tax cuts. They note that their mixed results are "probably not precise enough to persuade anyone to abandon strongly held a priori views".

⁴Souleles (1999) found that spending responds significantly to the federal income tax refunds that most taxpayers receive each spring. Hsieh found smaller responses by Alaskans to their oil rebates than to their federal income tax refunds. Parker found that household spending responded significantly to changes in take-home pay that occurred for high-income households that hit the Social Security tax cap. Other related studies include Wilcox (1989), Shapiro and Slemrod (1995), John Shea (1995), Souleles (2000), Browning and M. Delores Collado (2001), David Gross and Souleles (2002), and Melvin Stephens Jr. (2003), among others.

⁵Of the 78% of respondents who report they will mostly save their rebate, the majority (about three-fifths) report that they will mostly pay down debt (as opposed to accumulate assets). Shapiro and Slemrod (2002, forthcoming) used a novel follow-up survey in 2002 to try to determine whether there was a lagged response to the rebate. They found that, of respondents who said they initially mostly used the rebate to pay down debt, most report that they will "try to keep [down their] lower debt for at least a year." They found similar results for those who report they will save by accumulating assets.

A concurrent paper by Sumit Agarwal, Chunlin Liu, and Souleles (2004) also exploits the random timing of the rebate mailing, to identify the dynamic response of credit-card payments, spending, and debt to the rebates. They find that households initially used some of their rebates to increase credit card payments and thereby pay down debt, but soon afterwards credit card spending rose such that debt returned back near its pre-rebate levels. These dynamics of spending are consistent with the dynamics we find in this paper. This paper complements Agarwal, Liu and Souleles (2004) by estimating the magnitude of the effect of the rebate on nondurable expenditure, not just credit card spending.⁶

III. The 2001 Tax Rebates

The Economic Growth and Tax Relief Reconciliation Act of 2001 enacted substantial reductions in personal and estate tax rates, which were forecasted to reduce revenues by around 10 trillion dollars over ten years. The Tax Act reduced the income tax rate applied to income in the lowest income tax bracket from 15 percent to 10 percent, with this change applied retroactively to income earned from the start of 2001. The tax rebates represented an advance payment of this tax cut. The first income tax bracket applied to the first \$6,000 of income for a single individual filing a return, and to the first \$12,000 of income for a married couple filing jointly, so that, of the approximately two-thirds of U.S. households that received a rebate, most received rebates of \$300 or \$600. The Internal Revenue Service determined the rebate amounts for each tax filer based on his or her year 2000 tax return.

We exploit two key features of the rebate disbursement. First, and more importantly, the rebate checks were not mailed all at once, but rather in different weeks that were randomly assigned to households, as described in the introduction. Thus, the date at which each household received its rebate is independent of other household characteristics.⁷ Second,

⁶That is, their credit card dataset does not record spending via cash or checks, nor spending on other credit cards held by the account-holders in their data. Also, it does not record whether a card-holder actually received a rebate. Agarwal, Liu and Souleles identify the timing of rebate receipt based on an indicator of the corresponding randomized digit of the card-holder's SSN. However they note that the card-holder might not be the tax filer (whose SSN determined the actual timing of receipt), and that some card-holding households did not receive a rebate at all. Agarwal, Liu and Souleles attribute the slight delay before credit card spending rises to a couple of features particular to credit cards; e.g., the need for constrained card-holders (those whose balances start near their credit limits) to first make payments before they can spend using their cards, and to the delay before payments and spending register on the credit card statement.

⁷ One potential exception is that a household that filed its year 2000 tax return late may have been mailed its rebate after the ten-week period of randomized disbursement ended in September. Since 92 percent of

Congress passed the Tax Act in May, 2001, and of course expectations of some tax cut arose even earlier.⁸ Given this, our empirical methodology allows the rebates to be treated as pre-announced. This distinction matters for interpreting the results as a test of the PIH, but not for measuring per se the effect of the rebates on expenditure.⁹

In aggregate, the 2001 tax rebates totaled 38 billion dollars, and so represent about 1.5 percent of GDP, and 2.2 percent of aggregate personal consumption expenditures, in the third quarter of 2001. The rebates were the dominant component (about 84%) of the tax cuts implemented in the first year of the Tax Act. The timing of the remaining, smaller components in 2001 is independent of the randomized timing of the rebates analyzed here. For more details about the Tax Act, see Alan J. Auerbach (2002), Donald Kiefer et al. (2002), and Shapiro and Slemrod (2003; 2002, forthcoming).

IV. The Consumer Expenditure Survey

The CE interview survey contains detailed measures of the expenditures of a large, stratified random sample of U.S. households. CE households are interviewed four times, three months apart. In each interview the households report their expenditures during the preceding three months. New households are added to the survey every month so that the data are effectively monthly in frequency. In addition to surveying households about their expenditures, the CE also gathers some information about their demographic characteristics, income, and wealth.

taxpayers typically file at or before the normal April 15th deadline (Slemrod et al. (1997)), this non-randomized source of variation from the previous year is small and likely to be exogenous to the rebate. We present results below that exclude rebates received late in 2001.

⁸ Indeed, tax cuts were a central element of George W. Bush's platform in the 2000 election. Moreover, the Treasury sent taxpayers a letter shortly in advance of the rebate informing them of the size of their upcoming rebate and the particular week in which it would be disbursed: "We are pleased to inform you ... you will be receiving a check in the amount of \$[amount] during the week of [mm/dd/yy]."

⁹ We focus on the behavior of expenditure after rebate receipt. A not-mutually-exclusive alternative empirical approach would be to try to estimate the expected effect of the tax cuts on permanent income, even before actual rebate receipt. However, such an approach would be much harder to implement, as discussed below. First, the passage of the Tax Act per se cannot be separated from aggregate effects captured by time dummies, such as seasonality. It does not provide randomized cross-sectional variation that one can exploit, unlike the timing of rebate receipt. Second, it is unclear how large and permanent consumers actually expected the tax cuts to be (see the survey evidence in Shapiro and Slemrod, 2003). Third, consumers' expectations of tax-cut induced changes in their permanent income would have evolved over time, starting at least as early as the 2000 election campaign. (E.g., there is no single point in time at which the tax cut went from being entirely unexpected to being entirely expected.)

The special module of questions about the 2001 rebates covers the crucial period during which and after the rebates were mailed: the module went into the field in the second week of August, and remained there through the end of December. The new questions asked households whether they received a rebate, how many rebate checks they received, and then the month and amount of each check received. These questions were asked at the end of the CE interview, after households completed their usual reporting of consumption expenditures and other information. The questions were written so as to be consistent with the style of other CE questions. Appendix A contains the survey instrument. Appendix B describes how we construct from the raw data the measures used below of the rebates received in each three-month expenditure reference-period. The response rate to the new module was rather good. Only about 3% of the rebate amounts were flagged as invalidly missing (e.g., ‘don’t know’ or refusals), and only about another 4% of the months-of-receipt were flagged as invalidly missing.

We focus on three measures of consumption expenditures. First, we study expenditures on food, which include food consumed away from home, food consumed at home, and purchases of alcoholic beverages. Much previous research has studied expenditures on food, largely because of its availability in the Panel Study of Income Dynamics, but it is a narrow measure of expenditures. Our second and main measure of consumption expenditures is nondurable expenditures (broadly defined). However, the NIPA definition of nondurable goods includes some semi-durables like apparel. Hence we also consider a third, intermediate definition of consumption expenditures, nondurable goods strictly defined, following Lusardi (1996). In preliminary analysis we also considered a fourth definition, BLS total expenditures, including durable expenditures like auto and truck purchases. However the response of total expenditures to the rebates was never statistically significant. This is not surprising. The rebates are small relative to the cost of autos and trucks and, more importantly, including expenditures on durable goods dramatically increases the variability of the dependent variable and decreases precision in estimation.¹⁰ Thus, in keeping with previous research, we focus on nondurable expenditures.¹¹

¹⁰Also, the introduction of zero-percent auto financing during our sample period significantly altered expenditures on autos.

¹¹Generalizing across specifications, estimates of the effect of the tax rebate on total expenditures are measured with a standard error about four times the size of that on nondurable goods, and the point

Our baseline sample uses the 2000 and 2001 waves of the CE survey, with the sample period starting with interviews in January 2001 (when period $t+1$ in equation (2) below covers expenditures in October 2000 to December 2000) and running through interviews in March 2002 (when period $t+1$ covers December 2001 to February of 2002). Where mentioned, we extend this baseline sample period by adding data from the recently released 2002 wave in order to allow for additional lags of the rebate in the analysis. The sample includes only households that had at least one interview during the period in which the tax rebate module was in the field. Also, we drop from our sample any households with implausibly low expenditures (the bottom 1% of nondurable expenditures), unusually large changes in age or family size, and uncertain tax rebate status. Appendix B describes our data and sample in more detail.

Table 1 presents summary statistics for our dataset. For each household-reference quarter, we sum all rebate checks received by the household in that quarter to create our main rebate variable, *Rebate*. The pattern of reported rebates appears consistent with the limited information about the rebates that we have from other sources. The average value of *Rebate*, conditional on receiving at least one rebate check in the reference quarter, is \$480. Of households receiving rebates, 27 percent report receiving \$300 in rebates and 54 percent report receiving \$600 in rebates.¹² The three-month reference period (July-September) for households interviewed in October 2001 covers the entire ten-week period during which the rebate checks were mailed. Of these households, 57 percent report receiving a rebate during this period.¹³

V. Economic Theory and Empirical Methodology

The recent literature testing the PIH and related models has typically relied on the time-series properties of the consumption Euler equation. Formally, according to the Euler

estimates often (and implausibly) imply that somewhat less money is spent on total expenditures than on nondurable goods.

¹²The household rebate value need not be equal to \$300 or \$600. Households with 2000 tax liabilities smaller than \$300 (or \$600) could receive smaller rebates; households with multiple tax filers could receive multiple checks; taxpayers filing as heads of households could receive a \$500 check.

¹³Despite the potential for measurement error, this result is close to estimates of rebate receipt based on (unpublished) Treasury estimates: about 89.5m tax returns received a rebate while 23.5m did not receive a rebate, and about 22.9m households did not file and so also did not receive rebates (Office of Tax Analysis).

equation, households equate the marginal utility of consumption in the current period with the expected marginal utility of consumption in the next period:

$$u'(C_{i,t}) = E_t [(1-\delta)^{-1}(1+r_t)u'(C_{i,t+1})], \quad (1)$$

where $C_{i,t}$ is nondurable consumption by household i in period t ; E_t is the expectations operator; $u'(\cdot)$ is the marginal utility function, decreasing because of diminishing marginal benefits of consumption; $1+r_t$ is the price of consumption in t relative to $t+1$; and δ_t is the discount rate. Starting with Robert E. Hall (1978), many papers assume a specific functional form for the utility function, or linearize equation (1), and exploit the time-series properties of the expectation error to estimate preference parameters and test the Euler equation. Motivated by the alternative hypothesis (which actually predates the null hypothesis) that households to some extent consume income when it arrives, the tests often focus on whether predictable changes in income are statistically significant when added to the Euler equation.

While this approach to the Euler equation can estimate the model and test the null hypothesis of the PIH, it is not suitable for our purposes because it cannot estimate outside of the null hypothesis the causal impact of a predictable change in income on consumption growth.¹⁴ Moreover, in the present context, the traditional approach to Euler equation estimation is inappropriate because there is insufficient time-series variation across our sample period to effectively exploit the usual time-series properties of the expectation error (Gary Chamberlain (1984), Souleles (2004)). By contrast, our approach does not rely on time-series asymptotics. We can directly identify and estimate the impact of the rebate on consumption growth using the fact that the randomized rebate receipt is uncorrelated with households' expectation errors and any other unobserved heterogeneity.

Consistent with specifications in the previous literature, our main estimating equation is

$$C_{i,t+1} - C_{i,t} = \sum_s \beta_{0s} * month_{s,i} + \beta_1' X_{i,t} + \beta_2 R_{i,t+1} + u_{i,t+1}, \quad (2)$$

where C is either consumption expenditures or their log; $month$ is a set of indicator variables for every period in the sample, used to absorb the seasonal variation in consumption expenditures as well as other concurrent macro factors; and X are control variables (here age and changes in family composition) included to absorb some of the preference-driven

¹⁴That is, this estimation method is based on the time-series properties of the expectation error under the assumption that the model is true and there is no effect of predictable income changes on consumption growth. Finding a statistically significant effect rejects the model, so a significant coefficient on the predictable change in income cannot be interpreted causally.

differences in the growth rate of consumption expenditures across households. $R_{i,t+1}$ represents our key rebate variables, which take one of three forms: i) the total dollar amount of rebates received by household i in period $t+1$ ($Rebate_{i,t+1}$); ii) a dummy variable indicating whether any rebate was received in $t+1$ ($I(Rebate_{i,t+1} > 0)$); and iii) a distributed lag of $Rebate$ or $I(Rebate > 0)$.¹⁵ If consumption expenditures are smoothed across rebate receipt, the null hypothesis under the PIH treating the rebates as pre-announced, then β_2 should equal zero.¹⁶ We correct the standard errors to allow for arbitrary heteroskedasticity and within-household serial correlation.

In light of potential measurement error and sample-size limitations, in working with data on household expenditure it is generally important to use the largest possible sample and as much variation as possible in the independent variables. Hence we begin by utilizing all of the available information about the rebates received by each household, using $Rebate$ as the key regressor. This variable includes variation in the magnitudes of the rebates received, which is not randomized. While this variation is analogous to that used in most previous tests of consumption smoothing, we can further investigate its validity. We progressively limit the variation that we utilize, until we are left with variation in just the timing of rebate receipt, conditional on receipt. This limited variation is guaranteed to be exogenous because it is randomized. However, given the structure of the data and the fact that the rebates were disbursed over only a three-month period, as we focus in on timing alone, we reduce the sample size and amount of effective variation that identifies the key parameter, β_2 . This

¹⁵In preliminary work we also used $\Delta R_{i,t+1}$ as the key independent variable, to estimate the ‘instantaneous’ change in spending before and after rebate receipt. The distributed lag framework appropriately generalizes this idea in the presence of the lagged spending effects that we identified.

¹⁶To be clear, this test of the PIH applies under even weaker informational assumptions than usual. For our purposes the relevant aspects of the tax cut can be modeled in stylized form as follows. Consider five periods $t_0 < t_1 < t_2 < t_3 < t_4 < t_5$, where t_0 is earliest in time. Suppose that at time t_0 the government announces the following tax policy: It will cut tax rates starting at a specified future time t_1 . However, it will not reduce withholding rates until later, at t_5 ; meanwhile the reduction in tax liability before t_5 will instead be implemented by sending consumers a rebate check. Suppose the government randomly divides the population into three groups, and sends the checks to each group at t_2 , t_3 , and t_4 , respectively. (Our regressor $R_{i,t+1}$ captures this staggered disbursement.)

Notice that under the PIH, any wealth effect due to the tax cut should arise at t_0 , when the tax cut was announced. Even if, counter-factually, the tax cut were not announced in advance at t_0 , and instead was announced on implementation at t_1 , then the wealth effect should arise at t_1 . In neither case should the wealth effect be correlated with the actual rebate receipt at t_2 through t_4 , whose timing was randomized. Also note that in this framework it does not matter how temporary or permanent consumers expect the tax cuts to be. In particular, we do not have to take a stand on whether consumers expected the tax cuts to actually “sunset” after ten years, as specified under the original Tax Act.

substantially reduces the power of our estimator. We accordingly use Hausman tests to test whether the discarded variation, such as the magnitudes of the rebates, can be taken to be exogenous even though it is not randomized, and so can be validly utilized in order to maximize power and efficiency.

This identification strategy helps us avoid potential omitted variables bias and other confounding factors, at both the household and aggregate levels. By contrast, in most previous studies the income gain at issue (e.g., a windfall) was usually systematically related to various household characteristics, in ways that would often have been difficult to control for. For instance, suppose that high-income households, who are more likely to own stocks, receive larger windfalls (or larger predictable income gains); and that for other reasons the stock market happens to rise at the time of the windfall, leading high-income households to increase their consumption expenditures. In this case the estimated effect of the windfall on expenditure would be exaggerated by the stock market appreciation. During our sample period there were probably large changes in spending patterns induced by concurrent macroeconomic events, such as the recession, changes in monetary policy, the terrorist attacks of 9/11, etc. Nonetheless, all these events, even if their impact is correlated with household characteristics, are uncorrelated with the randomly assigned date at which households received their tax rebates.

We begin by estimating the short-run response of spending to the rebate, and then turn to the longer-run response. We later examine both the role of liquidity constraints, by interacting the rebate variables $R_{i,t+1}$ with indicators for illiquid households, and the response of different subcategories of spending, by changing the dependent variable in equation (2).

VI. The Short-Run Response of Expenditure

This section estimates the short-run change in consumption expenditures caused by rebate receipt, using just the contemporaneous rebate variables $Rebate_{t+1}$ and $I(Rebate_{t+1} > 0)$. These estimates of the contemporaneous effects of the rebate are nearly identical to the contemporaneous effects estimated in the following section after adding lagged rebate variables to equation (2). For ease of exposition, and in response to data limitations discussed below, we begin by focusing on the short-run effects alone.

In Table 2, the first three columns display the results of estimating equation (2) by OLS with the dollar change in consumption expenditures as the dependent variable and the contemporaneous amount of the rebate ($Rebate_{t+1}$) as the key independent variable, using all available rebate data. The resulting estimates of β_2 measure the average fraction of the rebate spent on each expenditure category, within the three-month reference-period in which the rebate was received. We find that, during the three-month period in which a rebate was received, relative to the previous three-month period, a household on average increased its expenditures on food by 11 percent of the rebate, its expenditures on non-durable goods strictly defined by 24 percent of the rebate, and its expenditures on non-durable goods (broadly defined) by 37 percent of the rebate. The latter two results are both economically and statistically significant, and counter to the PIH.

These results identify the effect of a rebate from variation in both the timing of rebate receipt and the dollar amount of the rebate. While the variation in the rebate amount is possibly uncorrelated with the residual in equation (2), it is not purely random. The amount of the rebate depends upon household characteristics, such as whether the household contains a married couple that filed jointly. Unlike most previous tests of consumption smoothing, which generally have no choice but to assume that the income change under investigation (the analogue to $R_{i,t+1}$) is exogenous, we can further explore this issue by progressively limiting the amount of variation that we utilize.

The remaining columns of Table 2 use only variation in whether a rebate was received at all in a given period, not the dollar amount of rebates received. The second triplet of columns displays the results of estimating equation (2) using the indicator variable $I(Rebate_{t+1} > 0)$. In this case, β_2 measures the average dollar increase in expenditures caused by receipt of a rebate. During the three-month period in which a rebate was received, relative to the previous three-month period, households on average increased their expenditures on food by \$51, their expenditures on non-durable goods strictly defined by \$96, and their expenditures on non-durable goods by \$179. Compared to an average rebate of about \$500, these results are quite consistent with those in the previous columns that include variation in the magnitude of the rebates received.

To check that the functional form of our specification is not driving our findings, and to further help calibrate the size of the effect of the rebate, the third triplet of columns in Table

2 uses the change in log expenditures as the dependent variable. On average in the three-month period in which a rebate is received, relative to the previous three-month period, consumption expenditures increased by 2.7 percent, 1.8 percent, and 3.2 across the three categories of expenditure. Again, given the average amount spent on each of these expenditure categories, these estimates are consistent with the previous estimates.¹⁷

Finally, since it is interesting to estimate a value interpretable as a marginal propensity to spend upon the rebate's arrival, we estimate equation (2) by two-stage least squares (2SLS). We instrument for the rebate amount, *Rebate*, using the indicator variable, $I(\text{Rebate} > 0)$, along with the other independent variables. In this case, as in the first three columns, β_2 measures the fraction of the rebate that is spent within the three-month period of receipt. As shown in the last triplet of columns in Table 2, the estimated marginal propensities to spend (11 percent, 20 percent, and 37 percent) remain statistically significant and are very close to those estimated without treating *Rebate* as potentially non-exogenous. This suggests that the variation in the rebate amount that was used in the first three columns can be taken to be exogenous. Comparing these results formally (the first three columns to the last three columns), a Hausman test does not reject the hypothesis that the variation in the rebate amount is exogenous, for each definition of consumption expenditures.

Overall, the results across the various specifications in Table 2 are quite consistent, implying a statistically significant short-run effect of the rebate on spending. The estimated effect is also economically significant. As discussed in Section IX, these estimates imply a substantial increase in aggregate consumption expenditures.

These results identify the effect on spending by comparing the behavior of households that received rebates at different times to the behavior of households that did not receive rebates at those times. Recall that some households did not receive rebates at all, in any period, so our results to this point implicitly use some information that comes from comparing households that received rebates to those that never received rebates. We investigate the role of this variation using a number of approaches.

¹⁷Using average levels of expenditures and rebates, the percent changes in spending in the third triplet of columns implies dollar spending changes and propensities to consume of \$40 and 0.084 for food, \$56 and 0.12 for strictly non-durable goods, and \$131 and 0.27 for non-durable goods, again broadly consistent with the other columns.

First, we directly control for rebate receipt, by adding to equation (2) a separate intercept for households that received a rebate in *some* sample reference quarter, $I(\text{Total Rebates} > 0)$. This allows the expenditure growth of rebate recipients to differ from that of non-recipients. In this case the main regressor $I(\text{Rebate}_{t+1} > 0)$ captures only high-frequency variation in the timing of rebate receipt -- receipt in quarter $t+1$ in particular -- conditional on receipt in some quarter. In Table 3, Panel A shows that this approach has little effect on our baseline conclusions. The separate intercept $I(\text{Total Rebates} > 0)$ is never statistically significant. Hence the expenditure growth of rebate recipients is on average similar to that of non-recipients. Further, in each regression, the point estimates of the effect of the rebate ($I(\text{Rebate}_{t+1} > 0)$) are actually somewhat larger than before. That is, even controlling for whether a household ever received a rebate, spending significantly increases in the particular quarter of rebate receipt.

Our second approach is more severe. We exclude from our sample the households that did not receive a rebate (or, more precisely and conservatively, those that are not known to have received a rebate using the available data).¹⁸ We also exclude the relatively few households that received late rebates due to filing late tax returns in the previous year. Even though the timing of these rebates is unlikely to be endogenous, it was not randomized.¹⁹ The advantage of this approach is that it identifies the response of spending from only purely randomized variation in the timing of rebate receipt conditional on receipt. The cost of this approach is that it leads to a substantial loss of power due to the resulting decline in sample size and effective variation. Recall that the CE rebate module was in the field through December 2001. Hence β_2 is now identified from only two groups of rebate-recipients: those with CE interviews in August (covering about 3 percent of non-late rebates) and in November (29 percent); and those with interviews in September (19 percent) and in December (19 percent). We lose all information regarding the sizable number of rebate recipients interviewed in October (31 percent). Accordingly, we also drop these households

¹⁸ For example, consider households whose last CE interview is in September 2001. Even if they report no rebates in their reference period, which covers June-August 2001, we cannot tell whether they received a rebate after August. Thus we drop all interviews of such households.

¹⁹ We exclude observations for which *Rebate* includes rebates received in November or December; but not rebates received in October, since rebates mailed in September (the end of the randomized disbursement period) can arrive in October.

from the sample.²⁰ As a result of these exclusions, the sample size is only about one-third of its original size.

Panel B of Table 3 shows that, consistent with the reduction in power, statistical uncertainty rises substantially, such that a 95 percent confidence interval contains both no rebate response and much larger responses than our baseline estimates in Table 2. With the caveat that they are statistically insignificant, the point estimates are somewhat lower than before, but still show an economically significant impact of the rebate on spending.²¹ As formally confirmed by Hausman tests, these estimates are not statistically different from our baseline estimates: in no column can we reject the hypothesis that the coefficient in this restricted subsample is the same as the corresponding coefficient estimated using the baseline sample in Table 2. While these Hausman tests have limited power, as before they suggest that the greater variation in the baseline sample can be taken to be exogenous.²²

Overall, the results of these extensions provide little evidence against our baseline estimates and support our conclusion that the rebates had an economically significant short-run effect on spending. We now turn to estimating the longer-run effects, and subsequently study how the effects differ across households and subcategories of expenditure. Because these extensions are even more demanding of the data than the short-run effects estimated so far, we return to using all of the households available in the baseline sample.

²⁰ That is, given the time that the tax module was in the field, this approach effectively identifies the impact of the rebate from only the behavior of households that both have consecutive interviews covering the period of randomized rebate disbursement, and report a rebate only in the earlier interview or only in the later interview. We lose all information from October interviews because these households were surveyed only once about the rebate and so their indicator for rebate receipt, $I(\text{Rebate}_{t+1})$, is collinear with the October month dummy, and so has little effect on the estimated coefficient on the indicator.

²¹ We dropped the late rebates in order to be conservative and limit our variation to just the variation that was randomized. However, as noted above the lateness of a rebate is *a priori* unlikely to be endogenous since it depends on a household filing a late tax return in the *previous* year. Hence we also estimated our model excluding from the baseline sample of Table 2 only the households that did not receive rebates (without excluding the late rebates). While the results are statistically insignificant for the reasons just discussed, the resulting estimates of β_2 remain closer to those found in Table 2, e.g., $\beta_2 = 0.30$ for nondurable goods using 2SLS. If instead we exclude only the late rebates (without excluding the non-recipients), the corresponding β_2 remains both statistically and economically significant at 0.30.

²² Additional support regarding the validity of these results comes from Agarwal, Liu and Souleles (2004), who find statistically significant responses of credit-card spending to rebate receipt. As noted above, they identify rebate receipt using just an indicator for the corresponding digit of the card-holders' SSNs, variation which is purely exogenous. That is, they do not use variation in the magnitude of rebates received, nor in any difference between rebate recipients and non-recipients.

VII. The Longer-Run Response of Expenditures

Relative to the PIH, the estimated fraction of the rebate spent in the first quarter of receipt is large. But since this fraction is substantially less than one, it remains an open question as to what happens to spending in subsequent quarters.

In Table 4, Panel A shows the results of estimating our main specifications when we include the first lag of the rebate variable, R_t , as an additional regressor. First, note that the presence of the lagged variable does not alter our previous conclusions about the contemporaneous impact of the rebate. The coefficients on R_{t+1} are quite similar to those in Table 2. Second, the receipt of a rebate causes a *change* in spending one quarter later (i.e., from the three-month period of rebate receipt to the next three-month period) that is negative and smaller in absolute magnitude than the contemporaneous change. The *net* effect of the rebate on the *level* of spending in the later quarter is given by the *sum* of the negative lagged coefficients (on R_t) and the positive contemporaneous coefficients (on R_{t+1}). While the lagged coefficients are typically not statistically significant themselves, for nondurable goods the net effect is often significantly positive. This implies that, after increasing in the three-month period of rebate receipt, spending remains high (statistically significantly greater than before receipt) in the subsequent three-month period. For example, the second column shows that expenditures on nondurable goods rise by 39% of the rebate in the quarter of receipt. The expenditure change in the next quarter is -8%, so that expenditures in the second three-month period are still higher on net than before the arrival of the rebate by $39\% - 8\% \approx 30\%$ (due to rounding) of the rebate. This net 30% result is significant at the 95 percent level. Accordingly, the *cumulative* change in expenditures on nondurable goods over both three-month periods is estimated to be $39\% + 30\% = 69\%$ of the rebate, and is statistically significant (bottom row of Panel A). Similar calculations for the final column using 2SLS suggest that nondurable expenditures in the second three-month period are higher on net by 27% of the rebate (statistically significant at the 7% level), with a significant cumulative change over both periods of 66% of the rebate.

To estimate whether the rebate increases consumption expenditures for a longer period, we also add a second lag of the rebate variable (R_{t-1}) to our regression. To do so we extend the sample period of our data by three months by adding interviews from April through June 2002 from the 2002 CE data.

Panel B of Table 4 shows that the additional data and regressor do not change our previous conclusions about the impact of the rebate on expenditures in the contemporaneous three-month period or the subsequent three-month period (using R_{t+1} and R_t). More interestingly, the coefficients on the second lag of the rebate variable are all negative, implying that expenditures continue to decline following their initial increase upon rebate receipt. However, these coefficients are all imprecisely estimated. The net level of expenditures in the second three-month period following rebate receipt is no longer statistically significantly different from the level before receipt. For example, for nondurable goods in the second column, expenditures in the second three-month period are higher than before rebate receipt by only 16% net of the rebate ($\approx .39\% - .10\% - .12\%$, with rounding), and this figure is not nearly statistically significant. (The corresponding net effect in the final column is only 6%, and is also statistically insignificant.) Further, while the cumulative share of the rebate spent during all three periods (bottom row of Panel B) is still large and sometimes statistically different from zero (only in the second column), it is not significantly different from the share that we estimated was spent during the first two periods (Panel A), and the statistical uncertainty of the three-period estimate is much larger. For example, in the second column of Panel B, the 95 percent confidence interval for the cumulative response of nondurable goods over all three three-month periods extends from 5 percent of the rebate to 162 percent of the rebate.

In sum, the pattern of coefficients suggests a large increase in expenditure at the time of rebate receipt, then a decaying but still substantial effect in the subsequent quarter or two. Households spent about two thirds of their rebates on nondurable consumption goods in the quarter of receipt and subsequent three months. Since the net response in the second three-month period after rebate receipt is much smaller and imprecisely estimated, the balance of the paper focuses on the contemporaneous rebate variable and its first lag, using our baseline data sample.

VIII. Differences in Responses Across Households and Goods

This section analyzes heterogeneity in the response to the rebate, across different types of households and different subcategories of consumption goods. While it is independently interesting to know who bought what with the rebates, the results also provide some evidence

as to why household expenditure responded to the rebate. For brevity, we report only results from the 2SLS specification, instrumenting the rebate and its lag (and any interaction terms) with the corresponding indicator variables for rebate receipt (and their interactions, along with the other independent variables).

The presence of liquidity constraints is a leading explanation for why household spending might increase in response to a previously expected increase in income.²³ To investigate this explanation, we test whether liquid or illiquid households were more likely to increase their spending upon arrival of a rebate. Households with low liquid wealth may be unable or unwilling to increase their spending prior to the rebate arrival. On the other hand, households with high liquid wealth may find the costs of not smoothing consumption across the arrival of the rebate to be small.²⁴

Expanding equation (2), we interact the intercept, rebate and lagged rebate variables with indicator variables (*Low* and *High*) based on various household characteristics (all from households' first CE interview). We use three different variables to identify households that are potentially liquidity constrained: age, income (family income before taxes), and liquid assets (the sum of balances in checking and saving accounts). While liquid assets is the most directly relevant of the three variables for measuring liquidity constraints, it is the least well measured and the most often missing in the CE data, so we start with the other two variables. For each variable, we split households into three groups, *Low*, *Middle*, and *High*, with the cutoffs between groups chosen to include about a third of the rebate recipients in each group.

We begin by testing whether the propensity to spend the rebate differs by age. Because young households typically have low liquid wealth and high income growth, and because old households may be living pension check to pension check, both young and old households might be disproportionately likely to be liquidity constrained. In Table 5, in the first pair of columns, *Low* refers to young households (younger than 40) and *High* refers to older households (older than 55), and the coefficients on these variables represent differences relative to the households in the (baseline) middle age group. While the point estimates

²³Precautionary motives can generate observationally similar results. See Tullio Jappelli (1990), Stephen P. Zeldes (1989, 1989a), Christopher D. Carroll (1992), Wilcox (1989), and Stephens (2003).

²⁴See the arguments of Ricardo J. Caballero (1995), Parker (1999), Christopher A. Sims (2001), and Ricardo Reiss (2004).

suggest that both young and old households spent somewhat more of the rebate than the typical (middle-aged) household, most of these differences are not statistically significant.

The second pair of columns in Table 5 tests for differences in spending across income groups. Low income households spent a much larger fraction of their rebate during the three-month period of receipt than the typical (middle-income) household. For nondurable goods, these differences are both statistically and economically significant. In the three months in which the rebate arrived, low income households spent about 62 percentage points more of their rebate on nondurable goods than typical, about 75 percent of their rebate in absolute terms. Further, based on the point estimates, high income households also seem to have spent a somewhat greater fraction of the rebate on receipt, although this difference is not statistically significant. For both high and low income households, the net lagged effects are not statistically significantly different than typical.

The last pair of columns tests for differences by liquid assets. The conclusions are the same as those for income, despite the smaller sample size due to missing asset values. In particular, households with few liquid assets spent a significantly greater share of their rebates than the typical household. Based on the point estimates high liquid wealth households also spent somewhat more than typical, but again this difference is not significant.

In sum, we find evidence that households with low income and low liquid wealth consumed more of their rebates than typical, which is consistent with the existence of liquidity constraints. These households are consuming most of their rebates soon after receipt, not saving much of them to smooth expenditure in future periods. This could be either because they expect to have higher income in the near future (e.g., due to an economic recovery) or because they have a high propensity to consume one-time or highly liquid funds.^{25, 26}

What did households buy with their rebates? Table 6 displays the results of estimating our main dynamic regression (including *Rebate* and one lag) with different dependent variables measuring expenditure across the different subcategories within nondurable goods.

²⁵Buffer stock models can generate large propensities to consume in response to transitory income gains. Adding hyperbolic discounting of the sort studied by David Laibson, George-Marios Angeletos, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg (2001) can generate even larger short-run responses.

²⁶In unreported analysis we also considered other demographic characteristics. For instance, we did not find statistically significant differences in the response across education groups or marital status.

Few of the resulting estimates are statistically significant. For these narrow subcategories of goods there is much more variability in the dependent variable that is unrelated to the rebate regressor. Our previous results, by summing the subcategories into broader categories of nondurable goods, averaged out much of this unrelated variability (such as for example whether a trip to the supermarket happened to fall just inside or outside the expenditure reference-period).

Based on the point estimates there is some (statistically insignificant) evidence that expenditures on food, both at home and away from home, respond to rebate receipt. Expenditures on food away from home initially rise by more than their share in nondurable expenditures and our previous estimates would suggest. Within the remainder of nondurable goods, the point estimates suggest larger responses (relative to their shares in nondurable expenditure) in personal care (and miscellaneous items), apparel (and apparel services), health expenditures, and reading materials, although again we note that the statistical significance of these results is low.

IX. Conclusion

This paper finds significant evidence that households spent much of their 2001 income tax rebates. Specifically, households spent about 20-40 percent of their rebates on non-durable consumption goods during the three-month period in which the rebates arrived, depending on the specification. We also find additional, smaller but still substantial, lagged effects on spending. Roughly two-thirds of the rebates were spent during the quarter of receipt and subsequent three-month period. The expenditure responses are largest for households with relatively low liquid wealth and low income, which is consistent with liquidity constraints.

What do these results imply in terms of the economic stimulus provided by the rebate checks? In aggregate, the rebates totaled 38 billion dollars, or about 2.2 percent of aggregate personal consumption expenditures (PCE), and 7.5 percent of nondurable PCE, in the third quarter of 2001. Applying our estimated propensities to spend from Table 4 (Panel A), this implies that the receipt of the tax rebates directly raised total PCE by about 0.8 percent in the third quarter of 2001 and 0.6 percent in the fourth quarter, and raised nondurable PCE by 2.9 percent and 2.0 percent in the third and fourth quarters. Since these calculations do not

include any potential effect on durable goods or any multiplier effects, the full impact of the rebates on the economy is possibly even larger. On the other hand, these calculations assume that the rebates did not increase prices, but only real consumption expenditure.

While we measure the causal impact of the rebates using cross-sectional variation, without using movements in aggregate consumption expenditures, the behavior of both aggregate consumption expenditure and aggregate saving data is broadly consistent with our findings. Figure 1 shows the growth rate of real total and nondurable PCE in the quarters surrounding the rebate disbursement. In the first half of 2001, the economy was in a recession, and both the latter half of 2000 and the first half of 2001 had low PCE growth. After the rebates were mailed out, PCE growth rose substantially and the recession ended in November of 2001. This is consistent with our results. Further, the aggregate data also suggest that the rebates were not all spent immediately. The personal saving rate rose from 1.9 percent and 1.2 percent in the first two quarters of 2001 to 3.4 percent in the third quarter when the rebates were mailed out, a pattern and magnitude consistent with households initially saving about two-thirds of the rebates (see Shapiro and Slemrod (2003; 2002, forthcoming)). The household saving rate then fell to 0.5 percent in the fourth quarter of 2001 and rose to 2.7 percent in the first two quarters of 2002. The behavior of the saving rate in the third and fourth quarters of 2001 is consistent with our finding of a substantial lagged effect of the rebate on spending.²⁷

While we focused on consumers' response to the receipt of the rebates, we cannot directly estimate whether there was an earlier response in anticipation of the rebate. The passage of the Tax Act itself cannot be separated from aggregate effects captured by our time dummies, such as seasonality. Moreover, there is no single point in time at which a tax cut went from being entirely unexpected to being entirely expected; rather, expectations of some tax cut grew over a long period, starting at least as early as the 2000 election. Nonetheless, our results suggest that the anticipatory response is likely to be small, since we already find large responses at the time the rebate checks arrived.

²⁷Shapiro and Slemrod note that some of the decline in saving in 2001:Q4 is due to increased spending on durable goods, particularly cars. But since durable goods are a small share of total consumption expenditures, the rise in the saving rate was still caused mostly by changes in other, nondurable components of consumption expenditures.

We conclude with a caveat. While the 2001 tax rebates stimulated consumer spending, without knowing the full structural model underlying these results, we cannot conclude that future tax rebates will have quantitatively the same effect. In 2001 the rebates were part of countercyclical stabilization policy. The spending response to other tax rebates may differ across time and circumstances. For instance, the response might be smaller outside of a recession or given a different situation for household balance sheets and liquidity.

Appendix A: CE Tax Rebate Survey Instrument

INTRO: Earlier this year a Federal law was passed cutting income tax rates and expanding certain credits and deductions. This year many households will receive a tax rebate check in the mail.

1. Since the 1st of (month, 3 months ago) have you (or any members of your CU [consumer unit]) received a tax rebate?
 1. YES—go to 2
 2. NO—end of interview

2. For each check received:

	check1	check2	check3	check4	check5
a. In what month did you receive the rebate?	___	___	___	___	___
b. What was the amount of the rebate?	\$_____	\$_____	\$_____	\$_____	\$_____
	-	-	-	-	-

3. For Interview number 2 and 5 and New Consumer Units: Did you already report the amount of this rebate in Section 22 , question 13, which asks about tax refunds?

Appendix B: CE Sample and Construction of Rebate Variables

We first construct a rebate variable from the raw data from question 2 of the CE tax rebate module (Appendix A). We then use the flags and other information to set the sample so that observations for which we are unsure about the validity of the rebate variable are not used. The variable *Rebate* is the sum of all rebates reported during the three-month expenditure reference period. If any of these magnitudes is missing, *Rebate* is set to missing.

Second, to maximize sample size, we use some rebate information from later interviews to fill in missing data in earlier interviews. Specifically, for interviews with no raw tax data, and for which the subsequent interview reports a rebate as having been received during the first interview's reference period, we treat the later interview's information as valid. (In particular, this completes the data for some of the households that were interviewed in early August before the tax rebate module was in the field.) Third, we use some rebate information from earlier interviews to create rebate measures for the reference period of the subsequent interview. For example, occasionally the first interview with tax data records a rebate received within the interview month itself (i.e., after the corresponding reference period), and the following interview reports no rebate for that same month. We treat this as a valid rebate response for the second reference period, since it is more likely to have been received then. Finally, the first interview sometimes reports no rebate, but the second interview records a rebate received during the first interview's reference quarter. In this case we assume that the household made a recall error in the second (more distant) interview and that the timing of the rebate reported in that interview is off. We therefore treat the rebate as if it occurred in the second interview reference-period if there is no other rebate already recorded for that period.

Rebate is set to zero for all observations covering reference periods ending June 2001 or earlier and starting October 2001 or later (unless a late rebate was reported) – periods during which the rebate questions were not on the CE survey.

We drop a rebate observation when: a) the lead-in question 1 states that a rebate was not received but there is a rebate reported in question 2; b) the lead-in question states that a rebate was received but there is no rebate reported for any month; c) there is a valid positive rebate amount but the associated month-of-receipt is either missing or flagged as invalid; d) a rebate is reported as received in a certain month but the rebate amount is missing, invalid or zero.

We use the following definitions of variables. Age is the average age of the head and spouse when the household is a married couple, otherwise it is just the age of the head. The number of children is calculated as the number of members of the household younger than 18.

Following Lusardi (1996), expenditures on nondurable goods strictly defined include expenditures on food (away from home, at home and alcoholic beverages), on utilities (fuels and public services) and household operations, on public transportation and gas and motor oil, on personal care, on tobacco, and on miscellaneous goods. Nondurable goods (broadly defined) adds expenditures on apparel goods and services, on health care expenditures (excluding payments by employers or insurers), and on reading materials.

Turning to the sample, we omit observations missing any of the key data that we use in our regressions. Our sample omits the bottom one percent of nondurable consumption expenditures in levels (after adjusting for family size and allowing for a time trend), since this data implies implausibly small consumption expenditures. Finally, we drop household

observations that report living in student housing, that report age less than 21 or greater than 85, that report age changing by more than one or a negative amount between quarters, or that report changes in the number of children or adults greater than three in absolute magnitude. When we split the sample based on income, we drop households flagged as incompletely reporting income.

References

- Agarwal, Sumit; Liu, Chunlin and Souleles, Nicholas S., 2004, "The Response of Consumer Spending and Debt to Tax Rebates – Evidence from Consumer Credit Data," working paper, University of Pennsylvania, May.
- Auerbach, Alan J., 2002, "The Bush Tax Cut and National Saving," National Tax Journal, LV, September, pp. 387-407.
- Bird, Roger C., and Bodkin, Ronald G., 1965, "The National Service Life-Insurance Dividend of 1950 and Consumption: A Further Test of the "Strict" Permanent-Income Hypothesis," Journal of Political Economy, 73, pp. 499-515.
- Blinder, Alan S., 1981, "Temporary Income Taxes and Consumer Spending," Journal of Political Economy, 89, pp. 26-53.
- Blinder, Alan S., and Deaton, Angus 1985, "The Time Series Consumption Function Revisited," Brookings Papers on Economic Activity, 2, pp. 465-511.
- Bodkin, Ronald G., 1959, "Windfall Income and Consumption," American Economic Review, September, 49(4), pp. 602-614.
- Browning, Martin, and Collado, M. Delores, 2003, "The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates," American Economic Review, 91, pp. 681-92.
- Browning, Martin and Lusardi, Annamaria, 1996, "Household Saving: Micro Theories and Macro Facts," Journal of Economic Literature, 34(4) , pp 1797-1855.
- Caballero, Ricardo J., 1995, "Near Rationality, Heterogeneity, and Aggregate Consumption," Journal of Money, Credit, and Banking, February, 27(1), pp. 29-48.
- Carroll, Christopher D., 1992, "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," Brookings Papers on Economic Activity, 2, pp. 61-156.
- Chamberlain, Gary, 1984, "Panel Data," in Griliches, Z., and Intriligator, M., Handbook of Econometrics, Amsterdam: North-Holland.
- Deaton, Angus, 1992, Understanding Consumption, Oxford: Clarendon Press.
- Gross, David, and Souleles, Nicholas S., 2002, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," Quarterly Journal of Economics, February, pp. 149-185.
- Hall, Robert E., 1978, "Stochastic Implications of the Life Cycle-Permanent Income

- Hypothesis: Theory and Evidence,” Journal of Political Economy, 86(6), December, pp. 971-987.
- Hsieh, Chang-Tai, 2003, “Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund,” American Economic Review, 99, pp. 397-405.
- Jappelli, Tullio, 1990, “Who is Credit Constrained in the U.S. Economy?” Quarterly Journal of Economics, February, 105, pp. 219-234.
- Kiefer, Donald; Carroll, Robert; Holtzblatt, Janet; Lerman, Allen; McCubbin, Janet; Richardson, David and Tempalski, Jerry, 2002, “The Economic Growth and Tax Relief Reconciliation Act of 2001: Overview and Assessment of Effects on Taxpayers,” National Tax Journal, LV, pp. 89-117.
- Krein, Mordechai E., 1961, “Windfall Income and Consumption: Additional Evidence,” American Economic Review, June, 51(3), pp. 388-390.
- Laibson, David; Angeletos, George-Marios; Repetto, Andrea; Tobacman, Jeremy and Weinberg, Stephen, 2001, “The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation,” Journal of Economic Perspectives, 2001, 15(3), pp. 47-68.
- Lusardi, Annamaria, 1996, “Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets,” Journal of Business Economics and Statistics, January, 14(1), pp. 81-90.
- Modigliani, Franco and Steindel, Charles, 1997, “Is a Tax Rebate an Effective Tool for Stabilization Policy?” Brookings Papers on Economic Activity, 1, pp. 175-209.
- Parker, Jonathan A., 1999, “The Reaction of Household Consumption to Predictable Changes in Social Security Taxes,” American Economic Review, September, 89(4), pp. 959-973.
- Poterba, James M., 1988, “Are Consumers Forward Looking? Evidence from Fiscal Experiments,” American Economic Review (Papers and Proceedings), May, 78(2), pp. 413-418.
- Reiss, Ricardo, 2004, “Inattentive Consumers,” working paper, Harvard University, February.
- Shea, John, 1995, “Union Contracts and the Life-Cycle/Permanent-Income Hypothesis,” American Economic Review, 85, pp. 186-200.

- Shapiro, Matthew D., and Slemrod, Joel, 1995, "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding," American Economic Review, March, 85, pp. 274-283.
- Shapiro, Matthew D., and Slemrod, Joel B., 2002 (forthcoming), "Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys," Tax Policy and the Economy, ed James Poterba. Cambridge: MIT Press.
- Shapiro, Matthew D. and Slemrod, Joel B., 2003, "Consumer Response to Tax Rebates," American Economic Review, 85, pp. 274-283.
- Sims, Christopher A., 2001, "Implications of Rational Inattention," working paper, Princeton University, June.
- Slemrod, Joel B.; Christian, Charles; London, Rebecca and Parker, Jonathan A., 1997, "April 15 Syndrome," Economic Inquiry, October, 35(4), pp. 695-709.
- Souleles, Nicholas S., 1999, "The Response of Household Consumption to Income Tax Refunds," American Economic Review, September, 89(4), pp. 947-958.
- Souleles, Nicholas S., 2000, "College Tuition and Household Savings and Consumption," Journal of Public Economics, 77(2), pp. 185-207.
- Souleles, Nicholas S., 2002, "Consumer Response to the Reagan Tax Cuts," Journal of Public Economics, 85, pp. 99-120.
- Souleles, Nicholas S., 2004, "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys," Journal of Money, Credit, and Banking, 36(1), February, pp. 39-72.
- Stephens, Melvin, Jr., 2003, "3rd of the Month: Do Social Security Recipients Smooth Consumption Between Checks?" American Economic Review, 93, pp. 406-422.
- Wilcox, David W., 1989, "Social Security Benefits, Consumption Expenditures, and the Life Cycle Hypothesis," Journal of Political Economy, 97, pp. 288-304.
- Wilcox, David W., 1990, "Income Tax Refunds and the Timing of Consumption Expenditure," working paper, Federal Reserve Board of Governors, April.
- Zeldes, Stephen P., 1989, "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy, 97, pp. 305-346.
- Zeldes, Stephen P., 1989a, "Optimal Consumption with Stochastic Income: Deviations from

Certainty Equivalence,” Quarterly Journal of Economics, 104(2), pp. 275-298.

Table 1: Summary statistics

<u>Panel A: Sample statistics (N=13,066 observations)</u>				
Variable	Mean	Standard Deviation	Minimum	Maximum
Expenditures on:				
Food	1,482	1,115	0	24,799
Strictly nondurables	3,168	3,984	175	284,216
Nondurables	4,149	4,481	385	300,155
Change in Expenditures on:				
Food	0	936	-16,392	22,360
Strictly nondurables	30	1,684	-20,363	49,860
Nondurables	62	2,052	-31,884	52,978
Change in:				
Number of Adults	0.0	0.3	-3.0	3.0
Number of Children	0.0	0.2	-3.0	3.0
Age	50.2	16.6	21	85
<i>Rebate</i>	86.8	199.0	0	1,500
<i>Rebate</i> <i>Rebate</i> > 0	480.0	173.8	2	1,500
<i>I(Rebate > 0)</i>	0.181	0.385	0	1
Income (N=9,443)	47,020	36,806	-500	320,800
Liquid Assets (N=6,060)	7,877	16,661	-17,000	136,200

<u>Panel B: Distribution of Positive Rebate Values</u>		
Rebate value	Number of Observations	Percent of Positive Rebates
$0 < \text{Rebate} < 300$	171	7.2
$\text{Rebate} = 300$	638	27.0
$300 < \text{Rebate} < 600$	233	9.9
$\text{Rebate} = 600$	1,275	53.9
$\text{Rebate} > 600$	47	2.0

<u>Panel C: Means of Rebate Variables by Interview Period</u>			
Three month period	<i>Rebate</i>	<i>I(Rebate)</i>	<i>Rebate</i> <i>Rebate</i> > 0
May - July, 2001	30.6	0.07	444.7
June - Aug, 2001	152.5	0.33	467.7
July - Sept, 2001	279.6	0.57	489.5
Aug - Oct, 2001	254.7	0.52	487.8
Sept - Nov, 2001	167.1	0.36	470.3

Note: based on sample for baseline regression using nondurable goods (Table 2).

Table 2: The contemporaneous response of expenditures to the tax rebate

Dependent Variable:	ΔC Dollar change in			ΔC Dollar change in			$\Delta \ln C$ Percent change in			ΔC Dollar change in		
	Food	Non-durable goods (strict)	Non-durable goods	Food	Non-durable goods (strict)	Non-durable goods	Food	Non-durable goods (strict)	Non-durable goods	Food	Non-durable goods (strict)	Non-durable goods
Estimation method:	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS
<i>Rebate</i>	0.109 (0.056)	0.239 (0.115)	0.373 (0.135)							0.108 (0.058)	0.202 (0.112)	0.375 (0.136)
<i>I(Rebate>0)</i>				51.5 (27.6)	96.2 (53.6)	178.8 (65.0)	2.72 (1.36)	1.76 (1.05)	3.16 (1.02)			
<i>Age</i>	0.570 (0.320)	0.449 (0.550)	1.165 (0.673)	0.552 (0.318)	0.391 (0.548)	1.106 (0.670)	0.035 (0.020)	0.005 (0.016)	0.023 (0.015)	0.569 (0.320)	0.424 (0.549)	1.166 (0.671)
<i>Change in adults</i>	130.3 (57.8)	285.8 (90.0)	415.8 (102.8)	131.1 (57.8)	287.7 (90.2)	418.6 (102.9)	6.16 (2.08)	6.22 (1.58)	7.55 (1.50)	130.3 (57.7)	286.2 (90.0)	415.7 (102.7)
<i>Change in children</i>	73.7 (45.3)	98.3 (82.4)	178.4 (98.3)	74.0 (45.3)	98.7 (82.5)	179.2 (98.3)	3.99 (2.36)	3.73 (1.66)	4.59 (1.66)	73.7 (45.3)	98.3 (82.5)	178.4 (98.3)
<i>N</i>	13,066	13,066	13,066	13,066	13,066	13,066	13,007	13,066	13,066	13,066	13,066	13,066

Notes: All regressions include a full set of month dummies. Reported standard errors are adjusted for arbitrary within-household correlations and heteroskedasticity. The third triplet of three columns is multiplied by 100 so as to report a percent change. The last three columns report results from two-stage least squares regressions where $I(Rebate>0)$ with the other regressors are used as instruments for *Rebate*.

Table 3: The contemporaneous response of expenditures: extensions

Dependent Variable:	ΔC Dollar change in		$\Delta \ln C$ Percent change in		ΔC Dollar change in	
	Non-durable goods (strict)	Non-durable goods	Non-durable goods (strict)	Non-durable goods	Non-durable goods (strict)	Non-durable goods
<u>Panel A: All households (N=13,066), Controlling for Rebate Receipt</u>						
Estimation method:	OLS	OLS	OLS	OLS	2SLS	2SLS
<i>Rebate</i>	0.281 (0.136)	0.425 (0.162)			0.244 (0.141)	0.444 (0.173)
<i>I(Rebate>0)</i>			1.99 (1.37)	3.67 (1.33)		
<i>I(Total Rebates>0)</i>	-30.6 (30.0)	-37.3 (36.4)	-0.28 (0.69)	-0.63 (0.68)	-24.8 (31.8)	-40.4 (38.7)
<u>Panel B: Only households receiving rebates (N=4,662)</u>						
Estimation method:	OLS	OLS	OLS	OLS	2SLS	2SLS
<i>Rebate</i>	0.153 (0.183)	0.247 (0.214)			0.079 (0.225)	0.189 (0.264)
<i>I(Rebate>0)</i>			1.36 (2.18)	1.94 (2.11)		

Notes: All regressions also include the change in the number of adults, the change in the number of children, the age of the household, and a full set of month dummies. Reported standard errors are adjusted for arbitrary within-household correlations and heteroskedasticity. The final pair of columns report results from two-stage least squares regressions where $I(Rebate>0)$ with the other regressors are used as instruments for *Rebate*. $I(Total Rebates>0)$ is an indicator for households that received a rebate in some sample quarter, whereas $I(Rebate>0)$ indicates receipt in the contemporaneous quarter ($t+1$) in particular.

Table 4: The dynamic response of expenditures to the tax rebate

Rebate Variable: Dependent Variable:	<i>Rebate</i>		<i>I(Rebate)</i>		<i>Rebate</i>	
	ΔC_{t+1} Dollar change in		$\Delta \ln C_{t+1}$ Percent change in		ΔC_{t+1} Dollar change in	
	Non-durable goods (strict)	Non-durable goods	Non-durable goods (strict)	Non-durable goods	Non-durable goods (strict)	Non-durable goods
Panel A: lagged rebate and baseline sample (N=12,730)						
Estimation method:	OLS	OLS	OLS	OLS	2SLS	2SLS
<i>Rebate</i> _{t+1} or <i>I(Rebate</i> _{t+1}) > 0)	0.248 (0.114)	0.386 (0.135)	1.86 (1.05)	3.29 (1.01)	0.208 (0.111)	0.386 (0.135)
<i>Rebate</i> _t or <i>I(Rebate</i> _t) > 0) (first lag)	-0.156 (0.099)	-0.082 (0.115)	-1.89 (1.06)	-1.44 (1.02)	-0.190 (0.101)	-0.113 (0.118)
Implied cumulative fraction of rebate spent over both three-month periods	0.340 (0.219)	0.691 (0.260)	NA	NA	0.226 (0.212)	0.659 (0.262)
Panel B: two lags of rebate and extended sample (N=15,022)						
Estimation method:	OLS	OLS	OLS	OLS	2SLS	2SLS
<i>Rebate</i> _{t+1} or <i>I(Rebate</i> _{t+1}) > 0)	0.247 (0.114)	0.386 (0.135)	1.85 (1.04)	3.29 (1.01)	0.208 (0.111)	0.386 (0.135)
<i>Rebate</i> _t or <i>I(Rebate</i> _t) > 0) (first lag)	-0.172 (0.097)	-0.099 (0.113)	-2.17 (1.05)	-1.72 (1.01)	-0.212 (0.099)	-0.139 (0.115)
<i>Rebate</i> _{t-1} or <i>I(Rebate</i> _{t-1}) > 0) (second lag)	-0.034 (0.121)	-0.123 (0.141)	-0.32 (1.23)	-1.67 (1.21)	-0.055 (0.122)	-0.191 (0.142)
Implied cumulative fraction of rebate spent over all three three-month periods	0.363 (0.323)	0.837 (0.391)	NA	NA	0.145 (0.316)	0.690 (0.397)

Notes: All regressions also include the change in the number of adults, the change in the number of children, age of the household, and a full set of month dummies. Standard errors are adjusted for arbitrary within-household correlations and heteroskedasticity. The final pair of columns report results from two-stage least squares regressions where *I(Rebate*_t) > 0) and its lags, along with the other regressors, are used as instruments for *Rebate* and its lags.

Table 5: The propensity to spend across different households

Dependent variable: ΔC_{t+1}						
<u>Dollar change in:</u>	<u>Non-durable goods (strict)</u>	<u>Non-durable goods</u>	<u>Non-durable goods (strict)</u>	<u>Non-durable goods</u>	<u>Non-durable goods (strict)</u>	<u>Non-durable goods</u>
	<u>Interaction: Age</u>		<u>Interaction: Income</u>		<u>Interaction: Liquid Assets</u>	
	Low: age ≤ 39 High: age > 55		Low: $\leq 34,300$ High: $> 69,000$		Low: $\leq 1,000$ High: $> 8,000$	
<u>Fraction of rebate spent in first three-month period</u>						
<i>Rebate_{t+1}</i> (Middle group)	0.222 (0.177)	0.326 (0.211)	0.050 (0.163)	0.130 (0.185)	-0.284 (0.177)	-0.243 (0.217)
<i>Rebate_{t+1} *Low</i> (Low group difference)	-0.035 (0.211)	0.071 (0.239)	0.317 (0.224)	0.624 (0.266)	0.569 (0.239)	0.876 (0.284)
<i>Rebate_{t+1} *High</i> (High group difference)	-0.038 (0.263)	0.109 (0.302)	0.274 (0.251)	0.255 (0.291)	0.312 (0.299)	0.404 (0.364)
<u>Change in fraction spent in second three-month period</u>						
<i>Rebate_t</i> (Middle group)	-0.296 (0.132)	-0.291 (0.159)	-0.080 (0.148)	-0.067 (0.172)	0.201 (0.226)	0.283 (0.261)
<i>Rebate_t *Low</i> (Low group difference)	0.301 (0.183)	0.465 (0.218)	-0.052 (0.198)	-0.059 (0.248)	-0.290 (0.253)	-0.292 (0.302)
<i>Rebate_t *High</i> (High group difference)	0.018 (0.231)	0.092 (0.273)	-0.309 (0.235)	-0.243 (0.275)	-0.659 (0.298)	-0.670 (0.358)
<i>Low intercept</i>	-110.6 (50.5)	-113.93 (59.60)	-20.8 (35.2)	-26.92 (41.62)	-56.3 (37.5)	-80.5 (47.0)
<i>High intercept</i>	131.7 (79.4)	151.28 (93.60)	-45.3 (52.3)	-31.51 (65.22)	15.3 (56.3)	-8.5 (73.8)
<i>N</i>	12,730	12,730	9,233	9,233	5,951	5,951

Notes: All regressions also include the change in the number of adults, the change in the number of children, the age of the household, and a full set of month dummies. All results are from two-stage least squares regressions where $I(Rebate > 0)$ and its lag and interactions, along with the other regressors, are used as instruments for *Rebate* and its lag and interactions. Reported standard errors are adjusted for arbitrary within-household correlations and heteroskedasticity. All sample splits are chosen to include about 1/3 of rebate recipients in each grouping.

Table 6: The propensity to spend on different categories of goods

Dependent variable: ΔC_{t+1} Dollar change in:	Panel A: Food			Panel B: Additional nondurable goods (strict)				Panel C: Additional nondurable goods		
	Food at home	Food away from home	Alcoholic beverages	Utilities, Household operations	Personal care and misc.	Gas, motor fuel, public transportation	Tobacco products	Apparel	Health	Reading
Average share of Nondurable Goods	0.28	0.09	0.02	0.23	0.04	0.10	0.02	0.08	0.13	0.01
$Rebate_{t+1}$	0.054 (0.038)	0.045 (0.038)	0.004 (0.011)	0.036 (0.027)	0.067 (0.058)	0.002 (0.044)	0.000 (0.007)	0.074 (0.044)	0.098 (0.040)	0.005 (0.004)
$Rebate_t$	0.005 (0.038)	-0.039 (0.046)	0.003 (0.011)	-0.005 (0.025)	-0.070 (0.052)	-0.079 (0.040)	-0.004 (0.008)	0.085 (0.033)	-0.009 (0.040)	0.000 (0.005)
Implied cumulative fraction spent over both 3-month periods	0.113 (0.084)	0.051 (0.078)	0.010 (0.022)	0.067 (0.056)	0.063 (0.105)	-0.075 (0.083)	-0.005 (0.016)	0.234 (0.090)	0.188 (0.082)	0.010 (0.008)

Notes: N=12,730 for all regressions. All regressions also include the change in the number of adults, the change in the number of children, age of the household, and a full set of month dummies. Reported standard errors are adjusted for arbitrary within-household correlations and heteroskedasticity. All results are from two-stage least squares regressions where $I(Rebate)$ and its lag, along with the other regressors, are used as instruments for $Rebate$ and its lag.

Figure 1: Growth rates for personal consumption expenditures

