Household Expenditure and the Income Tax Rebates of 2001

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

August 2005

Abstract

During 2001, most U.S. taxpayers were mailed a Federal tax rebate in a randomly assigned week between July and September. Using special questions added to the Consumer Expenditure Survey, we use 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. Households spent about 20-40 percent of their rebates on non-durable goods during the three-month period in which they received their rebates, and roughly two-thirds of their rebates cumulatively during the quarter of receipt and subsequent three-month period. The implied effects on aggregate consumption demand are substantial. Responses are larger for households with low liquid wealth or 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, D91

*Author affiliations: Johnson: Division of Consumer Prices and Price Indexes, 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 Pierre Bahizi, Miles Kimball, Jeffrey Kling, David Laibson, Matthew Shapiro; discussants Mel Stephens, Nicola Fuchs-Schuendeln, Julia Coronado, and Norbert Michel; and participants in seminars at the European Central Bank, the 2005 ASSA meetings, Boston College, Rochester, 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.

Policymakers often try to use tax policy to reduce the magnitude of economic fluctuations. They cut income taxes in recessions, assuming that the resulting increases in disposable income raise 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. Moreover, even for more permanent changes in taxes, consumption should respond as expectations of a tax change arise, not when consumers' disposable income actually changes.

This paper estimates the causal effect of the disbursement of the 2001 Federal income tax rebates on household expenditure, using unique data and features of the rebates. The Economic Growth and Tax Relief Reconciliation Act of 2001 sent tax rebates, typically $300 or $600 in value, to most 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 adequately study the 2001 tax rebates. In particular, the ongoing CE survey does not record the timing of taxes and transfers within the year, 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 CE survey. This module asked households about the timing and amount of each rebate check they received, and was included in the survey from shortly 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 survey.

[1]The 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 the staggered disbursement schedule is that it was difficult in practice to print and mail the rebate checks all at once. Accordingly, the schedule was keyed to the randomized social security digit largely for purposes of fairness.

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. This is in contrast to research relying on the time-series properties of the consumption Euler equation to test the null hypothesis of the Permanent Income Hypothesis (PIH). Strictly speaking such tests cannot estimate causal effects outside of the null hypothesis.

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

Summarizing our main findings, 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 cumulatively during the quarter of receipt and subsequent three-month period.

Although these findings do not depend on any particular theoretical model, we show that they constitute a rejection of the benchmark rational-expectations PIH, which predicts that any wealth effects from the rebates should be uncorrelated with the randomized timing of rebate receipt. We also discuss the implications of the PIH under alternative informational assumptions.

To shed further light on the reasons behind the estimated average response of spending, we contrast the responses across different types of households and different subcategories of nondurable goods. Households with low levels of liquid assets or low income spent significantly more of the rebate than typical, consistent with an important role for liquidity constraints. While not statistically significant, the point estimates also suggest somewhat larger responses among those with high levels of liquid assets or income (relative to households with intermediate levels of assets or 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 in the third quarter of 2001, our estimates imply that the rebates directly increased aggregate nondurable consumption expenditures by an economically significant amount: about 2.9 percent in the third quarter of 2001 and 2.1 percent in the fourth quarter. The full 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 study to the prior literature. Section II describes the relevant tax law changes and Section III our use of the CE survey data. Section IV discusses our empirical methodology. Section V presents the main results regarding the short-run response to the rebate, while Section VI examines the longer-run response and then considers the implications of the results for the PIH. Section VII examines differences in the response across different types of households and consumption goods. The final section discusses the aggregate impact of the rebates and concludes. Appendixes contain additional information about the data.

I. The Literature

To determine whether consumption responds to predictable (or transitory) changes in income, one must find clean measures of predictable (or transitory) income changes, and isolate their effect from other factors also impacting the consumption decision (e.g., concurrent changes in the stock market). However this is generally difficult. In particular, the previous literature relies on the potentially questionable assumption that the characteristics of a household used to determine the size or timing of its income change are uncorrelated with all other unobserved reasons for differential consumption growth rates (e.g., differences in preferences).[2] A key advantage of the present paper is that the random variation in the timing of the 2001 tax rebates is a priori known to satisfy this assumption, helping to avoid these recurrent difficulties in the literature.

[2]For 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 control variables such as education to Bodkin's original (pre- Euler equation) regression, Roger C. Bird and Bodkin (1965) find smaller spending responses.

Research using aggregate data to measure how much tax cuts increase consumption expenditures has 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 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 spending.[3]

Our paper builds more directly on the literature using household 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) examine windfalls like insurance dividends to WWII veterans and German restitution payments. More recently, a few studies examine more directly changes in fiscal policy, using larger, more representative samples (Jonathan A. Parker (1999), Nicholas S. Souleles (1999, 2002), and Chang-Tai Hsieh (2003)).[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 (2003a) find 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 implies an average marginal propensity to consume of about one third, consistent with the present paper's

[3]Franco Modigliani and Charles Steindel (1977), Alan S. Blinder (1981), and James M. Poterba (1988) study the (temporary) 1975 tax rebate. These papers find that consumption expenditure responded too much to the rebate relative to the PIH, though they come to somewhat different conclusions regarding the dynamics of the response. Blinder finds that about 55% of the rebate was cumulatively spent by the seventh quarter after its receipt, with much (about 16%) being spent within the quarter of receipt, and Modigliani and Steindel finds that almost 90% of the rebate was cumulatively spent by the fifth quarter after receipt, with much of the spending coming in the later quarters. Blinder and Angus Deaton (1985) finds smaller consumption responses analyzing the 1975 rebate and the 1968-70 tax surcharge together, but also finds that consumption is too sensitive to the pre-announced changes in taxes in the later phases of the Reagan tax cuts. The paper notes that these mixed results are "probably not precise enough to persuade anyone to abandon strongly held a priori views." Also using aggregate data, David W. Wilcox (1989, 1990) finds excess sensitivity of consumption expenditure to Social Security benefit increases and to Federal income tax refunds.

[4]Parker (1999) shows that when high-income households hit the Social Security payroll-tax cap, their expenditures increase by about half of the resulting predictable increase in after-tax income. Souleles (1999) finds that when households receive their annual (pre-determined) Federal income tax refunds, their total expenditure rises by about 1/3 to 2/3 of the refunds within the quarter of receipt. These results, based on CE data, are broadly consistent with our baseline findings in Table 2. Hsieh (2003) finds that the expenditure of Alaskans responds more to their Federal income tax refunds than to the annual payments they receive from the Alaska Permanent (Oil) Fund. Other related studies finding expenditure to be excessively sensitive to income include John Shea (1995), Shapiro and Slemrod (1995), Souleles (2002), and Melvin Stephens Jr. (2003, 2004, 2005). By contrast, Souleles (2000) and Browning and M. Dolores Collado (2001) find less evidence of excess sensitivity in the context of tuition payments and seasonal income payments in Spain, respectively.

estimate of the short-run response of expenditure. However, they find no evidence that liquidity constraints play a role in this response and no evidence of a lagged effect on expenditure.[5]

A concurrent paper by Sumit Agarwal, Chunlin Liu, and Souleles (2004) also exploits the random timing of the rebate mailing, using credit-card data, to identify the dynamic response of credit-card payments, spending, and debt to the rebates. The paper finds that households initially used some of their rebates to increase credit-card payments and thereby pay down debt, but soon afterwards they increased their spending, such that within nine months their debt returned back near its pre-rebate levels. These dynamics of credit-card spending are consistent with the dynamics of consumption expenditure that we find in this paper.

II. The 2001 Tax Rebates

The Economic Growth and Tax Relief Reconciliation Act of 2001 enacted substantial reductions in Federal personal and estate tax rates, which were forecast 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 tax bracket from 15 percent to 10 percent, with the change applied retroactively to income earned from the start of 2001. The tax rebates represented an advance payment of this tax cut for 2001. 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.[6] Second, Congress passed

[5]Of 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 (2003b) use a novel follow-up survey in 2002 to try to determine whether there was a lagged response to the rebate. They find 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 find similar results for those who report they will save by accumulating assets.

[6] Households that filed their year 2000 tax return late may have been mailed their rebates after the ten-week period of randomized disbursement ended in September. Since 92 percent of 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,

the Tax Act in May, 2001, and of course expectations of some tax cut arose even earlier.[7] Given these features, we are able to treat the arrival of the rebates as pre-announced, as discussed in Section IV. Their pre-announcement matters for interpreting the results as a test of the rational-expectations PIH, but not for measuring per se the effect of the rebates on spending.[8]

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 (2003a, 2003b).

III. 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 up to four times, three months apart, to collect expenditure information. In each interview they 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 information about their demographic characteristics, income, and wealth.

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

---

and in any case is likely to be exogenous to the rebate. We present results below that exclude rebates received late in 2001.

[7] 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]."

[8] We focus on the behavior of expenditure after rebate receipt. A not-mutually-exclusive alternative approach would be to estimate the effect of the tax cuts on expected permanent income and measure the response of expenditure as soon as consumers expected the tax cuts, even before actual rebate receipt. However, such an approach would require many more assumptions to implement, as discussed in Section VIII. Also, consumers appear to have disagreed about the size and permanence of the tax cuts. (See the survey evidence in Shapiro and Slemrod (2003a).)

amount of each check received. These questions were asked at the end of the CE interview, after households completed their usual reporting of 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 relatively good. Only about 3% of the rebate amounts were flagged as invalidly missing (e.g., 'don't knows' or refusals), and another 4% of the months-of-receipt were flagged as invalidly missing.

We focus on three different aggregated 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 such expenditures on food, largely because of their 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, which is a broad measure of expenditures on nondurable goods and services, following previous research. However, this measure includes some semi-durable goods like apparel. Hence we also consider a subset of nondurable expenditures that excludes such goods, "strictly nondurable" expenditures, following Lusardi (1996). Appendix B provides more details about these expenditure aggregates, and Section VII provides a complete decomposition showing how the response of nondurable expenditure varies across its component subcategories of goods.

In preliminary analysis we also considered 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. Thus, in keeping with previous research, we focus on nondurable expenditures.[9]

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 (1) below covers expenditures in October 2000 to December 2000) and running through interviews in March 2002

[9]Generalizing across our baseline 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 estimates often (and implausibly) imply that less money is spent on total expenditures than on nondurable expenditures.

(when period *t+1* covers December 2001 to February of 2002). Where mentioned, we extend this baseline sample period by adding data from the 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 the 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 is consistent with the limited information about the rebates available 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.[10] 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.[11]

IV. Economic Theory and Empirical Methodology

To test the PIH, the recent literature starting with Robert E. Hall (1978) has typically relied on the time-series properties of the expectation errors in the consumption Euler equation. Motivated by the alternative hypothesis (which actually predates the PIH) that households to some extent consume income when it arrives, the tests typically focus on whether predictable changes in income are statistically significant when added to the Euler equation.

While this traditional approach to the Euler equation can test the null hypothesis of the PIH, it is not desirable for our purposes because it cannot estimate outside of the null hypothesis the causal impact of a predictable change in income on consumption growth. Moreover, there is insufficient time-series variation across our sample period to effectively exploit the usual time-

[10]The 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 typically received a $500 check.

[11]Despite 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).

series properties of the expectation error (Gary Chamberlain (1984), Souleles (2004)). By contrast, our approach does not rely on time-series asymptotics. We 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 (e.g., Stephen P. Zeldes (1989a), Lusardi (1996), Parker (1999), Souleles (1999)), our main estimating equation is

$$(1) \qquad C_{i,t+1} - C_{i,t} = \Sigma_s \beta_{0s} * month_{s,i} + \boldsymbol{\beta_1}' \boldsymbol{X}_{i,t} + \beta_2 R_{i,t+1} + u_{i,t+1},$$

where C is either consumption expenditures or their log; *month* is a complete set of indicator variables for every period in the sample, used to absorb the seasonal variation in consumption expenditures as well as all other concurrent aggregate factors; and X are control variables (here age and changes in family composition) included to absorb some of the preference-driven 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 $I(Rebate>0)$, to measure the longer-run effects of the rebates. We correct the standard errors to allow for arbitrary heteroskedasticity and within-household serial correlation.

The key coefficient β_2 measures the average causal effect of rebate receipt on expenditure. While this measurement does not depend on any particular model, the results can nonetheless be interpreted as a test of the benchmark rational-expectations PIH, which assumes that consumers are aware of all publicly available information. Under this model, the null hypothesis is that β_2 should equal zero, since consumption expenditures should be smoothed across rebate receipt, which was pre-announced. To be clear, whenever information about the tax cuts underlying the rebates became publicly available, whether preceding the actual passage of the Tax Act or not, any resulting wealth effects should have arisen at the same time(s) for all consumers, and so their average effects on expenditure would be picked up by the corresponding time dummies in equation (1). Even heterogeneity in these wealth effects would not be correlated with the timing of rebate receipt, which is randomized, so β_2 should still equal zero.[12] Note that

[12] This argument covers the situation where public information arrives after the start of period t in equation (1), which is possible for some households given the structure of the CE data. (The argument is unnecessary however for

as a test of this null hypothesis, it is irrelevant how temporary or permanent consumers expected the tax cuts to be. Any significant expenditure response constitutes a rejection of the null hypothesis. For instance, 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 Tax Act.

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 estimating equation (1) utilizing all available CE households and 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 tests of the PIH, we can go further than these tests and 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 significantly reduce the sample size and amount of effective variation that identifies β_2. This 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, for instance a windfall, is usually systematically related to various household characteristics, in ways that would generally 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 same time as 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

---

any heterogeneity in wealth effects that precedes period *t*, since it would not affect the consumption change in equation (1).) There is one minor exception. Households that received their rebates later did receive a slightly smaller real rebate in present value due to the delay, but over just a few weeks these differences are trivial.

undoubtedly were large changes in spending patterns induced by concurrent macroeconomic events, such as the recession, changes in monetary policy, the terrorist attacks of September 11th, 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.

As we discuss in the conclusion, one caveat is in order. While our empirical methodology provides a clean causal estimate of the effects of the 2001 tax rebates, without a complete structural model one cannot conclude that future tax rebates will necessarily have the exact same effect. Nonetheless, by enhancing our understanding of consumer behavior, the results provide useful guidance for analyzing future policies.

We now turn to our results. We first present estimates of the short-run response of spending to the rebate, and then turn to the longer-run response. We subsequently 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 (1).

V. 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 are nearly identical to the short-run effects estimated in the following section after adding lagged rebate variables to equation (1). For ease of exposition, we begin by focusing on the short-run effects separately.

In Table 2, the first three columns display the results of estimating equation (1) by ordinary least squares (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 the different expenditure aggregate in each column, 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 strictly non-durable goods 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 statistically and economically significant.

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 (1), 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 tests of the PIH, which generally had 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 uses the indicator variable $I(Rebate_{t+1}>0)$ in equation (1). 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 $52, their expenditures on strictly non-durable goods 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.

As a robustness 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 was 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 categories, these estimates are consistent with the previous estimates.

Finally, since it is interesting to estimate a value interpretable as a marginal propensity to spend upon the rebate's arrival, we estimate equation (1) by two-stage least squares (2SLS). We instrument for the rebate amount, *Rebate*, using the indicator variable, $I(Rebate>0)$, 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 38 percent) remain statistically significant and are very close in magnitude to those estimated in the first three columns that do not treat *Rebate* as potentially non-exogenous. This again suggests that the variation in the rebate amount that was used in the first columns can be taken to be exogenous, a conclusion formally confirmed by Hausman tests.[13]

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 effects are also economically significant, implying a substantial increase in aggregate consumption expenditures, as discussed in Section VIII.[14]

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 these results implicitly use some information that comes from comparing households that received rebates to those that never received rebates. Table 3 investigates the role of this variation using a number of approaches, for brevity focusing on strictly nondurable goods and nondurable goods.

First, we directly control for rebate receipt, by adding to equation (1) an indicator for households that received a rebate in *any* reference quarter, *I(Total Rebates >0)*. This allows the expenditure growth of rebate recipients to differ on average from that of non-recipients. In this case, the main regressor $I(Rebate_{t+1}>0)$ 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 Panel A, the indicator *I(Total Rebates >0)* is never statistically significant. Hence, apart from the effect of the rebate, the expenditure growth of rebate recipients is on average similar to that of non-recipients. Moreover, the estimated coefficients for the effect of the rebate (on $Rebate_{t+1}$ and $I(Rebate_{t+1}>0)$) are somewhat larger in size than before adding the additional control and have the same pattern of significance. Thus the baseline results in Table 2 are not driven by average differences between rebate recipients and non-recipients. That is, controlling for whether a household ever received a rebate, spending significantly increases in the particular quarter of rebate receipt.

[13] The 2SLS standard errors are also close to the OLS standard errors. This reflects the fact that *I(Rebate>0)* is a very good instrument for *Rebate*, explaining over four-fifths of its total variation.

[14] These implications are robust across a number of additional sensitivity checks. For instance, including education expenditures in nondurable goods, controlling for additional household characteristics like marital status that are correlated with the rebate amounts, or weighting the sample, all lead to similar results. Using median regressions leads to smaller but still significant results.

Our second approach is more stringent. We exclude from our sample all 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).[15] 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.[16] 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 (27 percent); and those with interviews in September (19 percent) and in December (20 percent). We lose all information regarding the sizable number of rebate recipients interviewed in October (31 percent). Accordingly, we also drop the latter households from the sample.[17] 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 the 95 percent confidence intervals contain both no rebate response and much larger responses than our baseline estimates in Table 2. The point estimates are somewhat lower than before, but still show an economically significant impact of the rebate on spending.[18] As confirmed by Hausman tests, these estimates are not statistically different from our baseline estimates: in no column can we reject the hypothesis that the

[15] 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.

[16] 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.

[17] 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, $I(Rebate_{t+1})$, is collinear with the October month dummy.

[18] We drop the late rebates to be conservative and limit our variation to just the variation that is 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. Estimating our baseline model excluding only the households that did not receive rebates (without excluding the late rebates) results in estimates that are even closer in magnitude to those in Table 2, e.g., $\beta_2 = 0.30$ (0.30) for nondurable goods using 2SLS. Instead excluding only the late rebates (without excluding the non-recipients), the corresponding β_2 remains statistically significant at 0.30 (0.12).

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.[19]

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 goods. Because these extensions are even more demanding of the data than the short-run effects estimated so far, we return to using all the households available in the baseline sample.

VI. The Longer-Run Response of Expenditure

In Table 4, Panel A shows the results of estimating our main specifications when the first lag of the rebate variable, R_t, is included as an additional regressor in equation (1). 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 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 (relative to the level in the quarter before receipt) 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 (but less high) in the subsequent three-month period (i.e., statistically significantly greater than before receipt). 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 rebate receipt by 39%-8% ≈ 30% (due to rounding) of the rebate. This 30% figure is significant at the 95

[19] Supporting results come from Agarwal, Liu and Souleles (2004), who find statistically significant responses of credit-card spending to rebate receipt, with a dynamic pattern consistent with that in Section VI below. They identify the timing of rebate receipt using just an indicator for the relevant random digit of the card-holders' SSNs, variation which is clearly exogenous. See also Gross and Souleles (2002).

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 (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 have little effect on the estimated coefficients on the contemporaneous and once-lagged rebate variables (R_{t+1} and R_t). The coefficients on the second lag of the rebate variable are all negative, implying that expenditures continue to decline in the second three-month period following rebate receipt (relative to the first three-month period following receipt). However, these coefficients on R_{t-1} are all imprecisely estimated. The net level of expenditures in the second three-month period following 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 (≈ .39% -.10% -.12%, with rounding), and this figure is not statistically significant. The corresponding net effect in the final column is only 6%, and is also insignificant. Further, while the cumulative share of the rebate spent during all three periods (bottom row of Panel B) is large and sometimes statistically different from zero (only in the second column), it is not significantly different from the cumulative share that 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 about 5 percent of the rebate to over 160 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

cumulatively 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 rest of this paper focuses on the contemporaneous rebate variable and its first lag, using our baseline data sample.[20]

These results are inconsistent with the benchmark rational-expectations version of the PIH. Under this model there should be no response of expenditure to the receipt of the (publicly pre-announced) rebates because the arrival of a rebate check brings no new information about lifetime wealth. However, as noted by one of our referees, there are departures from the rational-expectations informational assumption that could potentially better reconcile the PIH with the main results so far.

For example, consider a household that is completely unaware of the public information regarding the 2001 Tax Act and rebates, and so is surprised by the arrival of its rebate check (say at the start of September). Under this assumption, the β_2 implied by the PIH could be greater than zero, but its actual magnitude would depend on additional assumptions, in particular the household's beliefs about the persistence of the underlying tax cut. If, for instance, the household thinks that the unexpected rebate check is a one-time windfall, then β_2 should still be quite small, approximately equal to the annuitization factor '*1/N*', which is typically taken to be around 0.05. This is much smaller than our estimates above.

Alternatively, suppose instead that the household infers from the arrival of the check that it had received a permanent tax cut, i.e. that the lowest tax bracket had been reduced and that this reduction would be permanent (and not otherwise offset). In this case, β_2 should be about 2/3 (.05) + 1/2 (3/4) ≈ 0.28. [21] That is, under this alternative informational scenario, the PIH would predict responses similar in magnitude to those in Table 2.[22] This alternative scenario however

[20] We also tried some higher-frequency, monthly analysis, but found that it was too demanding of the data. While maintaining the CE expenditure reference period of three months (since the timing of only relatively few individual purchases is reported monthly and not reallocated), one can still exploit monthly variation in rebate receipt by comparing the responses of households who receive their rebates in the first vs. second vs. third months of the expenditure reference periods. These responses are imprecisely estimated however, and do not significantly differ across the three groups of households, reflecting the greater imprecision associated with estimating more parameters at higher frequencies, and perhaps various survey issues such as respondent telescoping within the reference period.

[21] Recall that the new 10% tax bracket applied retroactively from the start of 2001. Since the rebate check arrives in September, one can think of 2/3 of the check as a one-time windfall, which should increase expenditure permanently by about 1/N ≈ .05 percent; and the remaining 1/3 of the check as the first installment of the permanent reduction in tax liability, which should be spent evenly over the remaining four months of the year.

[22] Under this scenario the PIH also predicts that the household should maintain the calculated higher level of expenditure permanently into the future. Most of the estimated coefficients on the lagged rebates in Table 4 are

does not predict any difference in spending across liquid and illiquid consumers, evidence to which we now turn.

VII. 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 learn who bought what with the rebates, this analysis also provides evidence about *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. 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 (Ricardo J. Caballero (1995), Parker (1999), Christopher A. Sims (2003), and Ricardo Reis (2004)).

Expanding equation (1), 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*, *High*, and the baseline intermediate group, with the cutoffs between groups chosen to include about a third of the rebate recipients in each group.

---

insignificant, which is consistent with this prediction, though a few are significantly negative and the point estimates suggest that the response of spending decays over time. The survey evidence in Shapiro and Slemrod (2003a) suggests that many consumers in 2001 were pessimistic about the size and permanence of the future tax cuts.

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, they are disproportionately likely to be liquidity constrained (e.g., Tullio Jappelli (1990) and Jappelli et. al. (1998)).[23] 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 interaction terms with these variables represent differences relative to the households in the baseline, middle-age group. While the point estimates suggest that both young and old households spent somewhat more of the rebate than the typical (baseline middle-aged) household, in both the quarter of receipt and cumulatively including the subsequent three month period (as reported at the bottom of the table), 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 (baseline 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 63 percentage points more of their rebate on nondurable goods than typical, about 76 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.

Turning to the cumulative response of spending over both three-month periods, the cumulative spending of low-income households is economically and statistically significant, but relatively imprecisely estimated: the corresponding confidence interval ranges from about 0.5 to over 2. Despite the large standard errors, the cumulative response of low-income households is also statistically significantly larger than that of the baseline group. This difference however is driven by the coefficients for the quarter of receipt; while the net level of spending by low-income households in the three-month period after the quarter of receipt is higher than that for the baseline group, that difference is not statistically significant.

The last pair of columns in Table 5 tests for differences by liquid asset holdings. The conclusions are similar to those based upon differences by income, despite somewhat larger

[23] Also, there is evidence that some older households increase their spending on receiving their (predictable) pension checks (Wilcox (1989) and Stephens (2003)). Outside the rational-expectations PIH, older households might also spend relatively more because they have shorter time horizons.

standard errors resulting from the smaller sample sizes due to missing asset values. In particular, households with few liquid assets spent a significantly greater share of their rebates than the typical household, in both the quarter of receipt and cumulatively.[24]

In sum, we find that households with low income or low liquid wealth consumed more of their rebates than typical, which is consistent with their facing binding liquidity constraints. These households are consuming most of their rebates soon after receipt and not saving much of them for 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]

What did households buy with their rebates? Table 6 re-estimates our main dynamic regression (including *Rebate* and one lag) with different dependent variables in each column measuring spending across the different subcategories within the broad measure of nondurable expenditures. The columns also report the relative importance of each subcategory as a share of nondurable expenditures. The results provide a complete decomposition of the main results in Table 4. Note that few of the resulting estimates are on their own 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 aggregates 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 evidence that expenditures on food, both at home and away from home, respond to rebate receipt, although these results are statistically insignificant. In particular, expenditures on food away from home initially rise by more than their share in nondurable expenditures and our previous estimates would suggest. Turning to the remainder of nondurable goods, relative to their shares in nondurable expenditures, the point estimates suggest larger responses, both contemporaneous and cumulative, in personal care (and

[24]In unreported analysis we also considered other demographic characteristics. We did not find statistically significant differences in the response across education groups or marital status.

[25]Precautionary motives can generate observationally similar results as liquidity constraints. In particular, buffer stock models can generate large propensities to consume in response to transitory income gains (e.g. Zeldes (1989b) and Christopher D. Carroll (1992).) Adding hyperbolic discounting of the sort studied by David Laibson et. al. (2001) can generate even larger propensities to consume liquid wealth for reasonable parameterizations.

miscellaneous items), apparel (and apparel services), health expenditures, and reading materials. The cumulative responses of apparel and health expenditures are also statistically significant.

While comparisons of different subsets of nondurable expenditure must be interpreted cautiously because of potential non-separabilities across goods, it is noteworthy that the largest cumulative response comes in apparel. Although apparel represents only about 8 percent of nondurable expenditures on average, it accounts for a much larger fraction of the total estimated cumulative response of nondurable expenditure to the rebate, about one third of the total. This could reflect a relatively large intertemporal elasticity of substitution (IES) for apparel. That is, expenditures on apparel are relatively less costly to postpone or accelerate, so its large response is consistent with theories of near rationality, as argued in Parker (1996). However some of the rest of the overall response to the rebate comes in goods like food and health care that presumably have a smaller IES.[26]

VIII. Discussion and Conclusion

This paper finds significant evidence that households spent much of their 2001 income tax rebates shortly after they arrived. Specifically, households spent about 20-40 percent of their rebates on non-durable consumption goods during the three-month period in which the rebates were received, depending on the specification, and roughly two-thirds of their rebates cumulatively during the quarter of receipt and subsequent three-month period. These results are inconsistent with the benchmark rational-expectations PIH. The expenditure responses are relatively large for households with relatively low liquid wealth or low income, which is consistent with their facing binding liquidity constraints.

Our findings imply that the rebates provided a substantial stimulus to the national economy, helping to end the recession of 2001. 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

[26] While nondurable goods tend to have a smaller IES than durable goods, all goods of course allow some scope for intertemporal substitution. For example, within health care, the largest response comes in spending on medical services, which includes relatively higher-IES expenditures like eye exams. In contrast to clothing, we found no significant evidence of response in durables such as autos or large household equipment like furniture and televisions, which again might reflect the relatively small size of the average refund per household and the greater volatility of expenditure on such durables.

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 2.1 percent in the third and fourth quarters.[27] Since these calculations do not include any potential 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.

Although we measured the causal impact of the rebates using only cross-sectional information, without using variation in aggregate consumption expenditure, the behavior of both aggregate consumption expenditure and aggregate saving 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 entirely 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 (Shapiro and Slemrod (2003a, 2003b)). The household saving rate then fell to 0.5 percent in the fourth quarter of 2001 before rising 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.[28]

Since our empirical approach focuses on consumers' response to the receipt of their rebates, it cannot directly estimate the magnitude of any earlier response that may have occurred in anticipation of the rebates. The passage of the Tax Act itself cannot be separated from other 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

[27] The NIPA and CE(BLS) measures of expenditure do not entirely coincide. Since our measure of nondurable expenditure in the CE includes many services, the calculation of the percentage effect on NIPA nondurables is likely an upper bound. On the other hand, our estimates do not include any spending in total PCE that is not in our measure of nondurable expenditure, so the calculation of the percentage effect on NIPA total PCE is likely a lower bound.

[28] Shapiro and Slemrod (2003b) notes 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.

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 after the rebate checks arrived.

We conclude by returning to 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 necessarily have quantitatively the same effect. The response of spending to tax rebates may differ across time and circumstances. In 2001 the rebates were part of countercyclical stabilization policy, but the response might be smaller outside of a recession or given a different situation for household balance sheets and liquidity. Nonetheless, our results provide a starting point for analyzing the impact of future tax rebates on expenditure.

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 in question 2; c) there is a valid positive rebate amount but the associated month-of-receipt is either missing or flagged as invalid (in which case we drop all rebate observations of the household, for all months); 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 strictly nondurable goods include expenditures on food (away from home, at home and alcoholic beverages), utilities (and fuels and public services), household operations, public transportation and gas and motor oil, personal care, tobacco, and miscellaneous goods. Nondurable goods (broadly defined) adds expenditures on apparel goods and services, health care expenditures (excluding payments by employers or insurers), and reading materials, following Lusardi (1996) but excluding education.

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 (often negative) 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. When we split based on liquid assets, we drop households if the asset information used in computing initial assets (as the difference between final assets and the change in assets) is topcoded.

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. Dolores, 2001, "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.

Jappelli, Tullio, Pischke, Jörn-Steffen, and Souleles, Nicholas S., 1998, "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," The Review of Economics and Statistics, 80, pp. 251-262.

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.

Kreinin, 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, 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, 1977, "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.

Reis, 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., 2003a, "Consumer Response to Tax Rebates," American Economic Review, 85, pp. 274-283.

Shapiro, Matthew D., and Slemrod, Joel B., 2003b, "Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys," Tax Policy and the Economy, ed. James Poterba., Cambridge: MIT Press.

Sims, Christopher A., 2003, "Implications of Rational Inattention," Journal of Monetary Economics, April, 50(3), pp. 665-690.

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 tha Month: Do Social Security Recipients Smooth Consumption Between Checks?" American Economic Review, 93, pp. 406-422.

Stephens, Melvin, Jr., 2004, "Paycheck Receipt and the Timing of Consumption," working paper, Carnegie-Mellon University, March.

Stephens, Melvin, Jr., 2005, "The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans," working paper, Carnegie-Mellon University, May.

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., 1989a, "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy, 97, pp. 305-346.

Zeldes, Stephen P., 1989b, "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," Quarterly Journal of Economics, 104(2), pp. 275-298.

Table 1: Summary statistics

Panel A: Sample statistics (N=13,066 observations)

Variable	Mean	Standard Deviation
Expenditures on:		
Food	1,482	1,115
Strictly nondurables	3,168	3,984
Nondurables	4,149	4,481
Change in Expenditures on:		
Food	0	936
Strictly nondurables	30	1,684
Nondurables	62	2,052
Change in:		
Number of Adults	0.0	0.3
Number of Children	0.0	0.2
Age	50.2	16.6
Rebate	86.8	199.0
Rebate | *Rebate* >0 (N=2,364)	480.0	173.8
I(Rebate>0)	0.181	0.385
Income (N=9,443)	47,020	36,806
Liquid Assets (N=6,060)	7,877	16,661

Panel B: Distribution of positive rebate values (N=2,364)

Rebate value	Number of Observations	Percent of Postive Rebates
0<*Rebate* <300	171	7.2
Rebate =300	638	27.0
300<*Rebate* <600	233	9.9
Rebate =600	1,275	53.9
Rebate >600	47	2.0

Panel C: Means of rebate variables by interview period (N=2,364)

Three month period	*Rebate*	*I(Rebate)*	*Rebate* | *Rebate* >0	Number of positive rebates
May - July, 2001	30.6	0.07	444.7	58
June - Aug, 2001	152.5	0.33	467.7	442
July - Sept, 2001	279.6	0.57	489.5	742
Aug - Oct, 2001	254.7	0.52	487.8	649
Sept - Nov, 2001	167.1	0.36	470.3	473

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			ΔlnC Percent change in			ΔC Dollar change in		
	Food	Strictly Non-durable goods	Non-durable goods	Food	Strictly Non-durable goods	Non-durable goods	Food	Strictly Non-durable goods	Non-durable goods	Food	Strictly Non-durable goods	Non-durable goods
Estimation method:	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS
Rebate	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(Rebate>0)				51.5 (27.6)	96.2 (53.6)	178.8 (65.0)	2.72 (1.36)	1.76 (1.05)	3.16 (1.02)			
Age	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)
Change in adults	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)
Change in children	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)
RMSE	934	1680	2047	934	1680	2047	0.50	0.37	0.36	934	1680	2047
R^2 (percent)	0.6	0.6	0.6	0.6	0.6	0.6	0.4	0.8	0.8	0.6	0.6	0.6

Notes: All regressions include a full set of month dummies, following equation (1). Reported standard errors are adjusted for arbitrary within-household correlations and heteroskedasticity. The coefficients in the third triplet of three columns are multiplied by 100 so as to report a percent change. The last three columns report results from 2SLS regressions where *I(Rebate>0)* with the other regressors are used as instruments for *Rebate*. All regressions have N=13,066 except percent change in food expenditures which has N=13,007.

Table 3: The contemporaneous response of expenditures: extensions

Dependent Variable:	ΔC Dollar change in Strictly Non-durable goods	 Non-durable goods	ΔlnC Percent change in Strictly Non-durable goods	 Non-durable goods	ΔC Dollar change in Strictly Non-durable goods	 Non-durable goods
	Panel A: All households (N=13,066), controlling for rebate receipt					
Estimation method:	OLS	OLS	OLS	OLS	2SLS	2SLS
Rebate	0.294 (0.136)	0.438 (0.161)			0.262 (0.141)	0.462 (0.173)
I(Rebate>0)			2.07 (1.37)	3.73 (1.33)		
I(Total Rebates>0)	-39.9 (30.0)	-46.8 (36.3)	-0.37 (0.70)	-0.70 (0.68)	-34.8 (31.8)	-50.6 (38.6)
	Panel B: Only households receiving rebates (N=4,739)					
Estimation method:	OLS	OLS	OLS	OLS	2SLS	2SLS
Rebate	0.152 (0.183)	0.247 (0.213)			0.079 (0.225)	0.190 (0.264)
I(Rebate>0)			1.35 (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 coefficients in the second pair of coumns are multiplied by 100 so as to report a percent change. The final pair of columns report results from 2SLS 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 reference quarter, whereas *I(Rebate>0)* indicates receipt in the contemporaneous quarter (*t+1*) in particular. The regression R^2's range from 0.6 percent to 0.9 percent and the RMSE are similar or slightly smaller than those reported in Table 2.

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

Dependent Variable:	ΔC_{t+1} Dollar change in		$\Delta \ln C_{t+1}$ Percent change in		ΔC_{t+1} Dollar change in	
	Strictly Non-durable goods	Non-durable goods	Strictly Non-durable goods	Non-durable goods	Strictly Non-durable goods	Non-durable goods
	Panel A: Lagged rebate and baseline sample (N=12,730)					
Estimation method:	OLS	OLS	OLS	OLS	2SLS	2SLS
Rebate $_{t+1}$ *or I(Rebate* $_{t+1}$ *>0)*	0.248	0.386	1.86	3.29	0.208	0.386
	(0.114)	(0.135)	(1.05)	(1.01)	(0.111)	(0.135)
Rebate $_{t}$ *or I(Rebate* $_{t}$ *>0)*	-0.156	-0.082	-1.89	-1.44	-0.190	-0.113
	(0.099)	(0.115)	(1.06)	(1.02)	(0.101)	(0.118)
Implied cumulative fraction of rebate spent over both three-month periods	0.340	0.691	NA	NA	0.227	0.659
	(0.218)	(0.260)			(0.212)	(0.262)
	Panel B: Two lags of rebate and extended sample (N=15,022)					
Estimation method:	OLS	OLS	OLS	OLS	2SLS	2SLS
Rebate $_{t+1}$ *or I(Rebate* $_{t+1}$ *>0)*	0.247	0.386	1.85	3.29	0.208	0.386
	(0.114)	(0.135)	(1.04)	(1.01)	(0.111)	(0.135)
Rebate $_{t}$ *or I(Rebate* $_{t}$ *>0)*	-0.172	-0.099	-2.17	-1.72	-0.212	-0.139
	(0.097)	(0.113)	(1.05)	(1.01)	(0.099)	(0.115)
Rebate $_{t-1}$ *or I(Rebate* $_{t-1}$ *>0)*	-0.034	-0.123	-0.32	-1.67	-0.055	-0.191
	(0.121)	(0.141)	(1.23)	(1.21)	(0.122)	(0.142)
Implied cumulative fraction of rebate spent over all three three-month periods	0.362	0.838	NA	NA	0.145	0.690
	(0.322)	(0.392)			(0.315)	(0.396)

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. Standard errors are adjusted for arbitrary within-household correlations and heteroskedasticity. The coefficients in the second pair of coumns are multiplied by 100 so as to report a percent change. The final pair of columns report results from 2SLS regressions where *I(Rebate>0)* and its lags, along with the other regressors, are used as instruments for *Rebate* and its lags. The regression R^2's range from 0.7 percent to 0.9 percent and the RMSE are similar or slightly smaller than those reported in Table 2.

Table 5: The propensity to spend across different households

Dependent variable: ΔC_{t+1}

Dollar change in:	Strictly Non-dur. goods	Non-durable goods	Strictly Non-dur. goods	Non-durable goods	Strictly Non-dur. goods	Non-durable goods
	Interaction: Age Low: age ≤ 39 High: age ≥ 56		Interaction: Income Low: ≤ 34,298 High: > 69,000		Interaction: Liquid Assets Low: ≤ 1,000 High: > 8,000	
Rebate $_{t+1}$	0.249	0.363	0.050	0.129	-0.284	-0.243
	(0.177)	(0.209)	(0.163)	(0.184)	(0.177)	(0.217)
Rebate $_{t+1}$ **Low*	-0.063	0.033	0.319	0.627	0.569	0.876
(Low group difference)	(0.210)	(0.238)	(0.224)	(0.266)	(0.239)	(0.284)
Rebate $_{t+1}$ **High*	-0.095	0.034	0.275	0.256	0.312	0.404
(High group difference)	(0.264)	(0.304)	(0.251)	(0.291)	(0.299)	(0.364)
Rebate $_{t}$	-0.266	-0.250	-0.080	-0.064	0.201	0.283
	(0.142)	(0.167)	(0.148)	(0.172)	(0.226)	(0.261)
Rebate $_{t}$ **Low*	0.271	0.425	-0.053	-0.067	-0.290	-0.292
(Low group difference)	(0.190)	(0.223)	(0.198)	(0.248)	(0.253)	(0.302)
Rebate $_{t}$ **High*	-0.042	0.010	-0.310	-0.246	-0.659	-0.670
(High group difference)	(0.228)	(0.270)	(0.235)	(0.275)	(0.298)	(0.358)
N	12,730	12,730	9,233	9,233	5,951	5,951
	Implied cumulative fraction spent over both three month periods for each group					
Baseline Group	0.232	0.476	0.020	0.194	-0.367	-0.203
	(0.359)	(0.431)	(0.363)	(0.410)	(0.405)	(0.501)
Low group	0.377	0.967	0.604	1.380	0.481	1.256
	(0.323)	(0.370)	(0.347)	(0.428)	(0.364)	(0.425)
High group	-0.001	0.554	0.259	0.461	-0.403	-0.065
	(0.395)	(0.476)	(0.421)	(0.507)	(0.569)	(0.704)

Notes: All regressions also include separate intercepts for the High and Low groups, 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 2SLS 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. The R^2 's range from 0.9 percent to 1.5 percent, with the highest fits for the splits using liquid assets; the RMSE are somewhat smaller than those reported in Table 2, smallest for the splits using liquid assets.

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 strictly nondurable goods				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	0.27	0.08	0.02	0.24	0.04	0.10	0.02	0.08	0.14	0.01
Rebate $_{t+1}$	0.054	0.045	0.004	0.036	0.067	0.002	0.000	0.074	0.098	0.005
	(0.038)	(0.038)	(0.011)	(0.027)	(0.058)	(0.044)	(0.007)	(0.044)	(0.040)	(0.004)
Rebate $_{t}$	0.005	-0.039	0.003	-0.005	-0.070	-0.079	-0.004	0.085	-0.009	0.000
	(0.038)	(0.046)	(0.011)	(0.025)	(0.052)	(0.040)	(0.008)	(0.033)	(0.040)	(0.005)
Implied cumulative fraction spent over both 3-month periods	0.114	0.051	0.010	0.067	0.064	-0.074	-0.005	0.234	0.187	0.011
	(0.085)	(0.079)	(0.022)	(0.056)	(0.107)	(0.083)	(0.015)	(0.090)	(0.082)	(0.008)
RMSE	624	635	173	697	833	668	136	715	698	73
R^2 (percent)	0.3	0.6	0.2	0.6	0.2	0.8	0.2	1.5	0.2	0.5

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, the 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 2SLS 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