



WORKING PAPERS

RESEARCH DEPARTMENT

WORKING PAPER NO. 07-34
THE REACTION OF CONSUMER SPENDING AND DEBT
TO TAX REBATES: EVIDENCE FROM
CONSUMER CREDIT DATA

Sumit Agarwal
Federal Reserve Bank of Chicago

Chunlin Liu
University of Nevada, Reno

Nicholas Souleles
The Wharton School
University of Pennsylvania and
Visiting Scholar, Federal Reserve Bank of Philadelphia

November 2007

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/econ/index.html

The Reaction of Consumer Spending and Debt to Tax Rebates: Evidence from Consumer Credit Data *

Sumit Agarwal
Federal Reserve Bank
of Chicago

Chunlin Liu
College of Business Administration
University of Nevada, Reno

Nicholas S. Souleles
Finance Department
The Wharton School
University of Pennsylvania
and NBER

November 2007

Abstract: We use a new panel data set of credit card accounts to analyze how consumers responded to the 2001 federal income tax rebates. We estimate the monthly response of credit card payments, spending, and debt, exploiting the unique, randomized timing of the rebate disbursement. We find that, on average, consumers initially saved some of the rebate by increasing their credit card payments and thereby paying down debt. But soon afterward their spending increased, counter to the canonical permanent-income model. Spending rose most for consumers who were initially most likely to be liquidity constrained, whereas debt declined most (so saving rose most) for unconstrained consumers. More generally, the results suggest that there can be important dynamics in consumers' response to "lumpy" increases in income like tax rebates, working in part through balance-sheet (liquidity) mechanisms.

Keywords: consumption, saving, life-cycle model, permanent-income hypothesis, liquidity constraints; fiscal policy, tax cuts, tax rebates, windfalls; credit cards, consumer credit, consumer balance sheets, household finance.

JEL classification: D91, E21, E51, E62, G2, H31

* We would like to thank the editor and two referees, Bob Hunt, Alan Krueger, Joanne Maselli, Larry Mielnicki, Jim Papadonis, Jonathan Parker, Anthony Santomero, Joel Slemrod, Melvin Stephens, and seminar participants at the European Central Bank, Princeton University, the ASSA meetings, Harvard University, the Federal Reserve Banks of New York and Chicago, the Federal Reserve Board of Governors, Singapore Management University, National Singapore University, the University of Wisconsin-Milwaukee, the University of Nevada-Reno, the 2003 NBER summer institute, the Federal Reserve Bank of Philadelphia, and the University of Pennsylvania, as well as the Wharton Financial Institutions Center and the Rodney L. White Center for Financial Research. We are grateful to Diana Andrade and Ron Kwolek for their excellent research assistance. The views expressed are those of the authors alone and do not necessarily represent those of the Federal Reserve Banks of Chicago and Philadelphia or the Federal Reserve System. This paper is available free or charge at www.philadelphiafed.org/econ/wps/. Corresponding author: Nicholas Souleles at souleles@wharton.upenn.edu.

I. Introduction

This paper uses a unique, new panel dataset of thousands of credit card accounts to analyze how consumers respond to “lumpy” increases in income like tax rebates. Specifically, to what extent did consumers use the 2001 Federal income tax rebates to increase spending or to pay down debt? About two-thirds of U.S. tax-filers received the rebates, typically \$600 for couples and \$300 for singles, for an average gain of about \$500 per recipient household.¹ This represents an historically significant tax cut, corresponding to about 5% of quarterly median family income. In aggregate about \$38 billion of rebates were disbursed, which corresponds to about 2% of quarterly personal consumption expenditures.² Our analysis will exploit a key feature of the rebate disbursement – its randomized timing. The rebate checks were disbursed over ten successive weeks from July through September 2001, depending on the second-to-last digit of the recipients’ social security numbers.³ Because this penultimate digit is randomly assigned, the timing of rebate receipt represents truly exogenous variation. Such randomization is quite rare in the history of fiscal policy and provides a unique natural experiment that cleanly identifies the causal effects of the rebates.

Although our estimation does not depend on any particular economic model, the results can be interpreted as a novel test of the canonical Life-Cycle/Permanent-Income [LCPI] model.

¹ The rebates were used to deliver the benefits of reducing the lowest Federal income tax bracket (which applied to the first \$12,000 of taxable income for joint returns, the first \$6,000 for individual returns, and the first \$10,000 for heads of households) from 15% to 10%. According to unpublished estimates from the Treasury, 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). The average gain at the household level was calculated by Johnson, Parker, and Souleles (2006) using the Consumer Expenditure Survey.

² These calculations draw on Shapiro and Slemrod (2003a). The rebates represented the dominant component (about 84%) of the tax cuts implemented in the first year of the Economic Growth and Tax Relief Reconciliation Act of 2001. 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 Act, see Auerbach (2002), Kiefer *et al.* (2002), and Shapiro and Slemrod (2003a and 2003b).

³ Taxpayers who had filed their year 2000 returns late could receive the rebate later than this, but typically about 92% of filers file on time [Slemrod *et al.* 1997]. Our analysis does not use any variation resulting from late returns.

In particular, in our high-frequency framework the rebates can be thought of as being pre-announced: Congress passed the Economic Growth and Tax Relief Reconciliation Act in May 2001, and expectations of some tax cut arose even earlier.⁴ Hence under the LCPI model, consumption should not significantly increase at the time of rebate receipt. Most previous tests of the model in micro data have had trouble identifying such predictable changes in income, which might have biased some of the tests against rejecting the model [Shea 1995]. Further, even studies that reject the model often find it difficult to identify the source of their rejection. The leading alternative model allows for liquidity constraints, but there is still no consensus about their actual importance. (For a useful review of the literature see Browning and Lusardi (1996).) Part of this disagreement is due to difficulties identifying which households in the data are in fact constrained. Most studies split the sample on the basis of net worth, but net worth conflates credit demand and credit supply. The fact that someone has low (even negative) net worth does not necessarily imply that he has reached his borrowing limit [Jappelli 1990]. One advantage of using credit card data is that they separately record credit limits and credit balances, which helps distinguish credit supply and demand [Gross and Souleles 2002].

Using innovative questions about the 2001 rebate that were added to the Michigan Survey of Consumers, Shapiro and Slemrod (2003a) found that about 46% of respondents who received (or expected to receive) a rebate reported that it would mostly lead them to pay down debt (and another 32% of respondents reported they would mostly save the rebate, in the sense of accumulating assets). This finding further justifies our focus on credit cards, whose debt carries higher interest rates than other forms of consumer debt. Rebate recipients who pay down debt should generally first pay down any credit card debt they hold. We use distributed lag models

⁴ Indeed, tax cuts were a central element of George W. Bush's platform in the 2000 election. Moreover, the Treasury sent taxpayers a letter in advance informing them of the size of their upcoming rebate and the particular week in which it would be disbursed.

that are interpretable as event studies to estimate the month-by-month response of credit card payments, spending, and debt to the tax rebates. This allows us to determine whether there are salient dynamics in consumers' response to the rebates, and, if so, helps identify the mechanisms behind the dynamics. For instance, if consumers initially use the rebates to pay down debt (or otherwise save) and thereby improve their balance sheets, what does this imply about their subsequent spending?

Because credit cards play an important role in consumer finances, they can be quite useful for studying consumer behavior. About 20% of aggregate personal consumption is being purchased using credit cards [Chimerine 1997]. Moreover, for most households, credit cards, in particular bankcards (*e.g.*, Visa, MasterCard, Discover, and Optima cards), represent the leading source of unsecured credit.⁵ About two-thirds of households have at least one bankcard, and of these households at least 56% are borrowing on their bankcards, that is, paying interest not just transacting [Survey of Consumer Finances (SCF) 1995].⁶ Conditional on borrowing, in the mid-to-late 1990s the typical bankcard account was borrowing about \$2000, with the account-holder having roughly another \$5000 of balances on other cards. These are large magnitudes relative to typical household balance sheets. They are also large in the aggregate: Total credit card balances currently amount to about \$900 billion [Federal Reserve Board 2007].

Previewing the results, we find that, on average, consumers initially saved some of the rebate, by increasing their credit card payments and thereby paying down debt and increasing their liquidity. But soon afterwards their spending increased, counter to the LCPI model and Ricardian Equivalence. For consumers whose most intensively used credit card account is in the

⁵ Moreover, Jappelli, Pischke, and Souleles (1998) found that households with bankcards are better able to smooth their consumption past income fluctuations than are households without bankcards.

⁶ As discussed by Gross and Souleles (2002), this figure probably understates the actual fraction of households borrowing on their bankcards, because SCF households appear to underreport their bankcard debt. This paragraph draws heavily on Gross and Souleles (2002). See also Yoo (1998).

sample, spending on that account rose by over \$200 cumulatively over the nine months after rebate receipt, which represents over 40% of the average household rebate. We also find other significant heterogeneity across different types of consumers. Notably, spending rose most for consumers who were, according to various criteria, initially most likely to be liquidity constrained. By contrast, debt declined most (so saving rose most) for unconstrained consumers. These results suggest that liquidity constraints are important. More generally, the results suggest that there can be important dynamics in consumers' response to "lumpy" increases in income like tax rebates, working in part through balance sheet (liquidity) mechanisms.

Sections II and III discuss the data and the econometric methodology. The main results appear in Section IV. Section V discusses related literature and Section VI concludes. The Appendices contain some additional results and data description.

II. Data

We use a unique, proprietary dataset from a large financial institution that issues credit cards nationally. The dataset contains a representative sample of about 75 thousand credit card accounts open as of June 2000, followed monthly for 24 months.⁷ The bulk of the data consists of the main billing information listed on each account's monthly statement, including total monthly payments, spending, balances, and debt, as well as the credit limit. Note that credit cards can be used for both transactions and borrowing purposes. "Debt" includes only interest-incurring balances that are rolled over, whereas "balances" also includes transactions balances that are paid off.

⁷ The sample excludes bankrupt/delinquent and dormant/closed accounts, which is consistent with our interest in accounts that might potentially respond to the rebates.

The dataset also contains some credit bureau data about the other credit cards held by each account-holder, in particular the number of other cards and their combined balances. (The credit bureaus do not separately record credit card debt, spending, or payments; they record only balances.) The credit card issuer obtained this data from the credit bureaus quarterly. Finally, there is some limited demographic data, *i.e.*, the age and marital status of the account-holders. (Account-holders are assumed to be married if there is a spouse also listed on the account.) An important advantage of the underlying data source is that it also included a variable indicating the penultimate digit of the account-holders' social security numbers. This variable was used to identify the time of rebate receipt.

This dataset has a number of additional advantages. Relative to traditional household datasets like the SCF, the sample is large with little measurement error. Also, because each account is observed over many months, it is possible to study high-frequency dynamics. On the other hand, using credit card data does entail a number of limitations. The main unit of analysis is a credit card account, not an individual (who can hold multiple accounts). We partially circumvent this limitation by using the available data about the account-holders from the credit bureaus. Also, we do not observe household assets or total spending (*i.e.*, spending via cash and checks).⁸ Table I provides summary statistics for the main variables used below. Appendix C provides further details about the data.

⁸ As discussed below, as a result of such limitations our main results are likely to understate the full effect of the rebates per account-holder and per household. Nonetheless, our results are broadly consistent with the consumption dynamics (described below, also over a nine-month horizon) found by Johnson, Parker, and Souleles (2006) using the Consumer Expenditure Survey, which records total household spending, but does not separately distinguish spending via credit cards, cash, or checks. Also, credit card data could capture a different fraction of the response of total household debt (and total debt payments) relative to the response of total spending. For example, if rebate recipients who pay down debt disproportionately pay down their high-interest credit card debt, but do relatively more of their increased spending via checks and cash, then the credit card data could capture more of the total response of debt than of spending. However, Johnson, Parker, and Souleles (2006) finds that the largest consumption response to the rebate came in apparel, which is relatively likely to be purchased using credit cards.

III. Methodology

We analyze the response of credit card account-holders to the tax rebates, beginning with the monthly account-level data and later turning to the credit bureau data. Specifically, we estimate distributed lag models of the following form:

$$Y_{i,t} = \boldsymbol{\alpha}'\mathbf{time}_t + \beta_0 R_{it} + \beta_1 R_{i,t-1} + \beta_2 R_{i,t-2} + \dots + \beta_9 R_{i,t-9} + \varepsilon_{i,t}, \quad (1)$$

where R_{it} is an indicator variable for whether the holder of account i received a rebate in month t . The dependent variable $Y_{i,t}$ variously represents either the spending ($S_{i,t}$) or payments ($P_{i,t}$) in account i in month t , or the amount of debt held by account i at the end of month t ($D_{i,t}$). Because rebate receipt is a temporary event and debt is a stock variable, to allow for potentially persistent effects of the rebate on debt, the specification for debt uses the change in debt as the dependent variable ($Y_{i,t} = \Delta D_{i,t} \equiv D_{i,t} - D_{i,t-1}$). Since payments and spending are flow variables linked to the change in debt via an accounting identity, payments and spending are accordingly analyzed in levels ($Y_{i,t} = S_{i,t}$ or $P_{i,t}$). The vector **time** represents a complete set of month indicator variables, *i.e.*, a separate indicator for each month in $t = 3/2001 - 5/2002$, which is the available sample period given the original dataset and the lags utilized in estimation. These indicator variables control for all aggregate effects, such as seasonality, the recession, changes in monetary policy, etc.

Following Gross and Souleles (2002), the results can be interpreted as an event study. The coefficient β_0 measures the immediate response of the dependent variable to rebate receipt, in dollar terms. The *marginal* coefficients $\beta_1, \beta_2, \dots, \beta_9$ measure the *additional* responses one month after rebate receipt, two months later, ..., and nine months later, respectively. (Allowing for nine lags is consistent with the available data period and turned out in the baseline analysis to be sufficient for the results to converge.) Therefore, for debt the *cumulative* coefficients

$b_s \equiv \sum_{t=0}^s \beta_t$ give the total change in debt after s months, $s = 0-9$. For the flow variables payments and spending, b_s gives the cumulative sum of changes in payments or spending over the first s months. For instance, if spending rises by $\beta_0 = \$10$ in the month of rebate receipt, and after one month spending is still $\beta_1 = \$5$ greater than it was before receipt, then the cumulative effect on spending after month one is $b_1 = \$15$.

To gauge the expansionary impact of the rebate, the response of spending is of central interest, especially the long-run cumulative response b_9 . The response of payments and debt are of independent interest, and can also help shed light on the response of spending. Under our specification the responses of all three variables are naturally related: The total effect of the rebate on debt after s months (b_s^{AD}) will approximately equal the difference of the cumulative effects of spending and of payments over the s months ($b_s^S - b_s^P$).

The key explanatory variable R depends only on the penultimate digit of the account-holders' social security numbers. Following the Treasury's disbursement timetable, we assume that rebates were disbursed in July 2001 to account-holders with a penultimate digit of 0 or 1, in August 2001 to those with a digit of 2 to 6, and in September 2001 to those with a digit of 7 to 9. Since the digit is randomly assigned, the resulting variation in the timing of rebate receipt is by construction exogenous. In an extension we also briefly consider the response to rebates of different sizes, using the data on marital status. However this variation is related to family structure and tax status (*e.g.*, since couples filing jointly generally received the largest rebates), and so cannot be guaranteed to be exogenous. By contrast R uses only the part of the potential variation that is guaranteed to be exogenous, and so allows for a clean test of whether there is a

causal effect of the rebate on credit card usage.⁹ Indeed, since the account-holder in the data is not necessarily the member of the household who actually filed the tax return (and so whose social security number determined the timing of rebate receipt), and some households (effectively those with minimal tax liability) did not even receive a rebate, the results will likely understate the full effect of the rebate.¹⁰ Hence, by using the limited variation in R , we are setting a high hurdle for finding significant effects of the rebate.

As an extension, we will also examine the response of the balances on the other credit cards held by account-holder i , using the credit bureau data. These data were collected only quarterly, however, which constrains their analysis. In interpreting these results we will accordingly focus on whether the response of other balances appears to reinforce, or to offset, the response of balances on the accounts in the main sample.

Unless indicated otherwise, equation (1) is estimated by ordinary least squares (OLS), with the standard errors adjusted for heteroscedasticity across accounts as well as serial correlation within accounts. We will also consider several alternative specifications. In particular, we will test whether the response to the rebate differs across different groups of account-holders, such as those who are potentially liquidity constrained. Indicator variables for these groups will be added to equation (1), both directly and interacted with R_{it} and all nine of its lags.

⁹ In previous studies of the response of consumption to changes in income, the income change at issue was usually systematically related to various household characteristics, and so the estimated effect of the income change might spuriously reflect these characteristics. For instance, suppose that high-income households, who are more likely to own stocks, receive larger windfalls (or other income gains), and also that for other reasons the stock market happens to rise at the time of the windfall, leading high-income households to increase their consumption. In this case the estimated (unconditional) effect of the windfall on consumption would be exaggerated by the stock market appreciation. Here, by contrast, since the timing of rebate receipt is independent of other personal characteristics, by comparing consumers who received rebates at different times, we avoid omitted variables bias and other confounding factors.

¹⁰ Nonetheless, the person who opened the bankcard account is probably most likely to be in charge of household finances and so disproportionately likely to be the tax-filer. Recall that about a third of households did not receive rebates. On the other hand, the population of consumers with credit cards is more likely to have received a rebate than the population without credit cards.

IV. Results

We begin by estimating the average response of payments, spending, and debt to the rebate, using the credit card accounts in the main sample. We subsequently analyze the heterogeneity in response across different types of account-holders. Because we find significant differences across account-holders, we discuss these results in detail. Finally, we also examine the response of the other credit cards held by the account-holders, using the credit bureau data.

Table 2 and Figure 1 show the baseline results from applying equation (1) to payments, spending, and the change in debt, in the main sample. The table reports the marginal coefficients β_s , $s = 0-9$, along with the final cumulative coefficients b_9 , which summarize the long-run effects of the rebate. The three graphs in Figure 1 show the entire paths of cumulative coefficients b_s , $s = 0-9$, along with their corresponding 95% confidence intervals. The results can be interpreted as an event study, with month 0 being the time of rebate receipt, $s = 0$ in event time.

Starting with the point estimates for payments, in the month of rebate receipt, (monthly) payments rise by $\beta_0 = \$12$ on average. One month after receipt, payments are still $\beta_1 = \$11$ larger than before receipt (so $b_1 = \$23$), but the subsequent marginal coefficients tend to decline in magnitude and significance (with only β_0 being significant at the 10% level or better). As is evident in the first graph, the path of cumulative payments plateaus after month four. Allowing for some time delay before consumers deposit their rebate check, make a larger payment to the issuer, and have that payment register on the monthly statement, the point estimates imply that, on average, consumers initially used some of the rebate to increase their credit card payments. The average long-run, cumulative increase in payments is $b_9 = \$49$ (bottom of Table 2).

As for spending, in the second graph in Figure 1, the path of cumulative spending coefficients is initially flatter than the path for payments. But after month two the path for spending starts to rise faster, overtaking the payments path after month six, before plateauing.¹¹ The average long-run, cumulative increase in spending is $b_9 = \$62$, which corresponds to about 12% of the average household rebate. Although not statistically significant, this baseline cumulative effect, like the corresponding b_9 for payments, is economically significant, considering that it reflects only the average response per credit card account, not per account-holder.

In general, the baseline results in Table 2 are imprecisely estimated. In part this is due to the limited random variation we are using. Since estimating high-frequency, monthly responses to this variation is demanding of the data, the individual marginal coefficients are difficult to estimate with precision. By summing across these coefficients, the paths of cumulative effects for spending and payments are smoother, but nonetheless their significance levels need not necessarily increase with the horizon (s). (For spending, the significance of the baseline cumulative coefficients does tend to increase with horizon, but only slowly.) This reflects the fact that the estimates of the underlying marginal coefficients tend to have positive covariance across horizons, which the standard errors for the cumulative coefficients take into account. Even so, we shall see shortly that there are significant differences in response across different account-holders, and that allowing for heterogeneity yields more significant marginal and cumulative

¹¹ The delay before spending starts to increase is further discussed below. It partly reflects the need for high-utilization account-holders to first make payments before they can make additional purchases (in addition to the short delay before payments and spending register on the credit card statement), plus potentially other mechanisms, such as habits. Gross and Souleles (2002) found qualitatively similar dynamics in average credit card spending in the months after exogenous increases in credit limits.

coefficients for certain types of account-holders.¹² Hence it is not surprising that the full-sample results, which assume similar responses across all account-holders, are less significant.

The baseline results for payments and spending are reflected in the results for debt in the third graph. Since payments rise before spending, debt initially declines, significantly so in month 0. The point estimates imply that debt subsequently increases, reflecting the lagged increases in spending. The estimated long-run change in debt b_9 is positive, but it is insignificant and small in magnitude (*e.g.*, relative to average debt levels of almost \$2000), both here and in most of the subsequent analysis, so this particular result should not be overemphasized.¹³ Another way to assess the significance of the dynamics of debt and of the other variables is to test the joint significance of the entire set of marginal coefficients $\{\beta_s \mid s = 0-9\}$, which is equivalent to testing the joint significance of the cumulative coefficients $\{b_s \mid s = 0-9\}$.¹⁴ As reported near the bottom of Table 2, for debt the estimated coefficients are jointly significant at the 7% level, with the initial coefficient being the most significant.¹⁵

¹² For instance, for many of the key sample-split groups analyzed later in Table 3 (*e.g.*, for young, low-limit, high-utilization, high-usage, and the “composite-constrained” account-holders), the significance levels of the cumulative coefficients for spending generally increase with the horizon more quickly than in Table 2, becoming statistically significant within the nine months after rebate receipt. Also, in Table A1, the cumulative significance levels for both spending and payments generally increase in horizon for all four groups of rebate recipients, becoming significant within a few months after receipt.

¹³ As evidenced by its relatively wide 95% confidence interval, which ranges from about -\$19 to \$74, b_9 for debt is insignificantly different from both zero and b_0 (and b_2 , the most negative cumulative coefficient). Thus, while one cannot reject the hypothesis that debt reverts (zero long-run change), one can also not reject the hypothesis that the initial decline in debt is permanent. Nonetheless, as discussed below, both hypotheses can be consistent with a significant increase in spending.

¹⁴ This is because the partial sums $\beta_0, \beta_0 + \beta_1, \dots, \beta_0 + \beta_1 + \dots + \beta_9$ are all zero iff the marginal coefficients $\beta_0, \beta_1, \dots, \beta_9$ are all zero.

¹⁵ To help further gauge the magnitudes of the results in Table 2 and subsequent tables, Table 1 reports the average levels and average monthly changes of payments, spending, and debt over the sample period. The dependent variables in equation (1) also vary significantly over time, as the month indicator variables (relative to the omitted first month of the sample) are jointly significant in each of the three regressions in Table 2 (for brevity, not reported). For payments, the estimated coefficients on the fourteen month indicators range from -\$19 to \$44, for an intra-sample swing of about \$63. For spending, all the month coefficients are positive, with the largest being about \$103, and for the change in debt the coefficients range from -\$35 to \$88, for a swing of about \$123. Note that the estimated b_9 's in Table 2 for spending and payments are sizeable fractions of these intra-sample swings. They are also sizeable relative to the annualized average monthly changes in Table 1.

To illustrate more broadly the potential importance of such balance sheet dynamics, consider someone who receives a \$600 rebate and, as a result, increases his total spending by the full \$600, using a credit card. Suppose he uses the rebate proceeds to increase his credit card payments by \$600 to fully pay for the extra spending. In this case there would be no longer-run effect on debt (ignoring any small differences due to interest), even though spending fully responds to the rebate. If the extra payments precede the extra spending, debt will first decline but then recover. If instead spending partially responds to the rebate, but by less than the full \$600 in extra payments, then there would be a persistent decline in debt. Hence a persistent decline in debt is also consistent with an expansionary effect of spending (so long as debt does not decline by the full amount of the extra payments). As a result, since we can directly estimate the response of spending, we do not need a precise estimate of the long-run change in debt in order to gauge per se the expansionary effect of the rebate.

More generally, in the presence of significant spending dynamics, static specifications that allow for only an immediate spending response to the rebate would underestimate its full effect. In particular, without a flexible dynamic analysis, it would be difficult to identify a lagged response of spending to the extra liquidity arising after consumers initially save some of their rebate, whether by accumulating assets or paying down debt.

The distributed lag specification in Table 2 accommodates the average monthly dynamics of the dependent variables in a very general way. As an extension, we undertake some intra-monthly analysis, distinguishing account-holders according to how early in the month (whether in July, August, or September) their rebate was disbursed, with week 1 being the first week in the month, week 4 being the last. For example, if (marginal) spending takes place roughly evenly

over the month, we might be able to detect that the path of spending starts increasing slightly earlier on average for those receiving their rebates earlier in the month.

Since such intra-monthly analysis is even more demanding of the data than the baseline analysis in Table 2, we impose some additional structure for these results (and a few of the subsequent extensions). Note that, although the baseline estimation was totally non-parametric, the results turned out relatively well behaved. In particular, the baseline marginal coefficients for spending and payments are generally positive; and they generally decline in horizon for payments, and are hump-shaped for spending. Consequently, the cumulative effects for payments and spending are generally increasing (non-decreasing) in the horizon, before eventually plateauing. To increase precision in the intra-monthly specification, we impose *a priori* the constraint that the marginal responses of spending and payments be non-negative: $\beta_s \geq 0 \forall s$. In the simplest analysis, spending and payments should not in general decline in response to an increase in liquidity, so this should be a relatively minimal restriction. (By contrast, debt can decline, so we do not estimate this specification for debt.)¹⁶

The resulting point estimates, reported in Appendix Table A1, provide some rough indication that the increases in spending and payments begin slightly earlier for the early-in-the-month recipients (week 1) relative to the later-in-the-month recipients. However, these differences are not statistically significant, so these results are inconclusive. Further discussion of them is reserved for Appendix A. Not surprisingly, such intra-monthly comparisons cannot be

¹⁶ While the cumulative responses should not be negative, some marginal coefficients could potentially go negative in some situations. *E.g.*, suppose that in response to the rebate someone purchases a large durable good. As a result, in some subsequent months spending could potentially end up lower than it would have been in the absence of this purchase. However, in the baseline results in Table 2, few of the marginal coefficients for spending and payments are negative, and even then they are small in magnitude. Consistently, in Table A1 the non-negativity constraint binds in just a few cases (indicated by “n.a.” in the standard errors column). Also, Johnson, Parker, and Souleles (2006) finds significant responses to the rebates only in nondurable expenditures (including apparel), not in expenditures on larger durables.

made with much precision. Nonetheless, many of the results are significant in absolute terms. Notably, for both payments and spending, and for all four weeks of receipt, the long-run cumulative effects b_9 are statistically and economically significant, even though they do not significantly differ across the weeks. Since the variation underlying these results for spending, as well as the other significant results for spending below, is randomized, the results imply a causal link from the rebate to spending, counter to the LCPI model.

We now turn to a comparison of how consumers of different demographic and credit characteristics responded to the rebate. Because it is difficult to simultaneously estimate separate responses by week-of-receipt for each of the different groups of account-holders that we will examine, we return to our baseline monthly specification (without imposing the non-negativity constraint). Table 3 reports the long-run, cumulative coefficients b_9 for each group, as well as p -values for the significance of these coefficients and of the corresponding coefficients at an intermediate horizon of five months (b_5). Figure 2 graphs all of the cumulative coefficients, across all horizons $s = 0-9$, for payments, spending, and the change in debt in separate panels. For each regression, Table 3 also reports the joint significance of the marginal coefficients (equivalently, the cumulative coefficients) separately for each group in the regression, and also combined across all groups in the regression (under the label “combined test”).

Row A) of Table 3 starts with marital status. To equation (1) we added indicator variables for couples and singles and their interaction with the rebate indicator and its nine lags. The results for spending appear in the second set of columns in Table 3. As reported, spending by couples increases by somewhat more on average than spending by singles – by $b_9 = \$74$ versus $b_9 = \$61$ cumulatively over the nine months after rebate receipt, though neither result is statistically significant. Since the rebates that couples received were typically twice as large

(\$600 vs. \$300), their moderately greater spending in dollar terms represents a smaller share of the rebates they received. The intermediate dynamics are displayed in Figure 2, in the first two columns of graphs (labeled with the prefix “A”). Singles initially increase their payments, such that their debt significantly declines in the month of receipt, and further declines in the next two months (with the most negative cumulative debt coefficient being $b_2 = -\$47$ (23)). However, as their spending later increases, the point estimates for debt subsequently increase and become insignificant. These dynamics are reflected in the joint significance of the singles’ marginal coefficients for debt (p-value = 0.01 in Table 3).

Row B) of Table 3 instead considers age, contrasting young (<35 years old), middle-aged, and older (>60) account-holders. For spending, the long-run cumulative responses b_9 decline monotonically with age. The spending of the young account-holders increases on average by $b_9 = \$200$ cumulatively. This response is statistically and economically significant. It is also statistically significantly larger than the b_9 ’s of each of the two groups of older account-holders (despite the individual standard errors for each group). These results are suggestive of liquidity constraints, since the young are disproportionately likely to be constrained [Jappelli 1990]. If one re-estimates this specification adding the non-negativity constraint $\beta_s \geq 0 \forall s$, while the resulting b_9 remains significant and larger for the young, it is statistically significant for each of the older age groups as well. Moreover, even without the constraint, the reported p-value of 0.04 for the combined test of significance indicates that the marginal and cumulative coefficients for spending are jointly significant when all three age groups are considered together (*i.e.*, considering all the coefficients $\{\beta_{s,a}\}$ for all ages a and all s jointly). Thus, after allowing for heterogeneity in age, the coefficients for spending are jointly significant across the entire sample. Such results illustrate the importance of heterogeneity.

The coefficients for spending are not, however, jointly significant for the age groups separately, not even for the young whose b_9 is significant. While the young's cumulative spending coefficients start to increase in the first month after receipt, the increases (based on the underlying marginal coefficients and reflected in Figure 2) become consistently significant only after month five. (Nonetheless, the cumulative coefficient b_5 , although insignificant itself, is already significantly larger than the b_5 for the older account-holders, at the 6% level.) This lag in spending (both here and in later results below) can potentially be explained by a number of additional mechanisms, such as habits or other costs of adjusting consumption, precautionary motives (*e.g.*, Carroll 1992), time to search and buy, or heterogeneous inattention (*e.g.*, Reis 2006).¹⁷

By contrast, for older account-holders (and, to a lesser extent, the middle-aged), the point estimates suggest more of an initial decline in debt, though this decline is not statistically significant and the point estimates subsequently increase. Nonetheless, in month five after receipt, their change in debt is still negative and (although insignificant itself) significantly different from that of the young, at the 7% level.

Our data allow for even more direct tests for liquidity constraints. One advantage of credit card data is that they separately record credit limits and credit balances. On average, consumers whose balances start near their limits are expected to be more likely to be liquidity

¹⁷ As discussed below, Johnson, Parker, and Souleles (2006) also finds a persistent response of consumption over the nine months after rebate receipt. The response is somewhat backloaded for certain categories of expenditure, such as apparel, which as noted is the largest contributor to the overall response of consumption and is relatively likely to be purchased using credit cards.

One should also keep in mind the relatively wide confidence intervals around the estimated coefficients at different horizons (as evident in Figure 2). This partly reflects the imprecision associated with non-parametric estimation, especially in the smaller effective samples for the various subgroups of accounts in Table 3. If one imposes the non-negativity constraint $\beta_s \geq 0 \forall s$, the results tend to be even more significant than those reported in Table 3. *E.g.*, under the constraint, for the spending of the young, all of the marginal coefficients, other than β_0 and β_2 , are significant at the 10% level or better, and the corresponding cumulative coefficients are all significant after month two.

constrained [Gross and Souleles 2002]. Row C) starts by considering the credit limit alone, since it is more exogenous than balances. (To further minimize any endogeneity, we lag the credit limits, taking them from month $t-9$, the start of the distributed lag horizon in equation (1).) We divide the accounts into those with low credit limits ($\leq \$7,000$), which constitute about two-fifths of the sample; intermediate limits, another two-fifths; and high limits ($> \$10,500$), the remaining fifth. The low-limit accounts are most likely on average to actually be constrained by their limit.¹⁸

As expected, the low-limit accounts do exhibit the largest increase in spending. The long-run, cumulative increase is $b_9 = \$141$. This response is economically and statistically significant, and significantly larger than the b_9 of the intermediate-limit group. While the marginal spending coefficients for the low-limit group are jointly significant, the resulting increase in spending is again backloaded. Based on the point estimates, cumulative payments for this group initially increase faster than cumulative spending, so their debt initially declines. However, the decline in debt is insignificant and small in magnitude, and debt subsequently increases as spending later overtakes payments.¹⁹

By contrast, the high-limit accounts exhibit a substantial increase in payments, by almost \$200 cumulatively. This response is significant at the 6% level (and at the 5% level on imposing the non-negativity constraint). It is also significantly larger than the corresponding b_9 for the low-limit group (despite the individual standard errors for each group). On the other hand, the

¹⁸ For example, on average these accounts have the largest utilization rates. For the regressions that split the sample directly using the credit limit, we exclude the (relatively few) accounts with limits below \$1000, to accommodate the typical rebates of \$300-\$600. This exclusion leads to slightly sharper but similar results when splitting using the limit, but has very little effect if applied to the baseline results.

¹⁹ We shall see shortly that part of the explanation for these dynamics appears to be that account-holders who start with high utilization rates first need to make payments before they can increase spending. When looking directly at high-utilization accounts, the initial rise in spending takes place very soon after the initial rise in payments, so that the rise in spending is relatively more frontloaded (less backloaded) than for other groups. Since the other groupings of potentially constrained accounts, such as low-limit accounts, are not perfectly correlated with the high-utilization accounts, this relative frontloading gets attenuated for the other groupings.

high-limit accounts do not show much change in spending. Consequently their debt declines by \$145 by month nine, a substantial amount. While this decline is imprecisely estimated, it is significantly different from the corresponding b_9 for the low-limit group. Thus, the high-limit account-holders, who were expected most likely to be unconstrained, are in fact more likely to save their rebates by paying down debt. Their response is more consistent with the LCPI model. The contrast between this response and the substantial increase in spending by the low-limit account-holders suggests that the latter are indeed relatively likely to be liquidity constrained (with the noted qualification about the backloading of their spending).

Row D) groups accounts according to their initial utilization rate, defined as the ratio of balances to the credit limit. While low-limit account-holders are relatively likely to be constrained by their limit, using utilization directly should be even more likely to identify the account-holders whose limits are actually binding, and so should yield sharper results.²⁰ To minimize endogeneity, utilization is taken from month $t-9$.

Almost 10% of the accounts have utilization rates above 90%. For them liquidity constraints are most likely on average to be binding. In fact, as reported the cumulative increase in their spending is a substantial $b_9 = \$333$. This response is statistically significant, and also significantly larger than the b_9 's of each of the groups with lower utilization. For the high-utilization accounts, the marginal coefficients for spending are jointly very significant (p-value = .001), and many of the individual coefficients are large and significant, starting in the second month after receipt. Accordingly, starting in month two, all the cumulative coefficients are

²⁰ For example, Gross and Souleles (2002) found that, for credit card accounts with high utilization rates (above 90%), spending sharply rises after increases in liquidity due to increases in credit limits, significantly more than for lower-utilization accounts. These results, which split the sample using utilization directly, are sharper than analogous results using other indicators of liquidity constraints, such as age. Other credit bureau data suggest that account-holders who have high utilization rates on a given credit card also tend to have high utilization rates on their other cards, and so high total credit-card utilization rates at the account-holder level.

significant. As evident in Figure 2, for this group the response of spending is relatively sharp and less backloaded than the previous results for the young and low-limit groups.

Since the high-utilization account-holders started near their limits, how can they substantially increase their spending? They first make large initial payments, to create liquidity in terms of available credit. Indeed, for this group the marginal coefficients for payments are jointly significant, with the initial coefficients β_0 through β_2 being especially large and significant (at the 5% level for β_2).²¹ The cumulative \$225 increase in payments is substantial and significant at the 10% level. Overall, the results for the high-utilization group are indicative of binding liquidity constraints.

Turning to the other groups in row D), the accounts with lower utilization rates, of 50%-90%, 1%-50%, and zero, constitute about 20%, 60%, and 10% of the sample, respectively. For the accounts with 50%-90% utilization, the estimated b_0 for spending is still substantial in magnitude, albeit insignificant and smaller than for the high-utilization group. This is consistent with the idea that the intermediate-utilization accounts face a material, albeit smaller, probability of being constrained.²² Their debt does not decline in the months after receipt (apart from a tiny coefficient $b_0 = -\$4$ for month zero). By contrast, for the accounts with positive but small utilization (1%-50%), debt significantly declines in the month of receipt. Although the later point

²¹ Focusing on the paths for payments and spending in Figure 2, the initial rise in payments resembles the initial rise in spending shifted one month earlier. Thus debt declines in months zero and one, with $b_1 = -71(41)$ being significant at the 10% level. But subsequently the point estimates for debt increase and become insignificant. Account-holders who start constrained by their credit card limit could of course spend their rebate without using their credit card. But if for various reasons (*e.g.*, convenience, safety, perks, etc.) they prefer to use their card for their marginal spending (*e.g.*, on apparel), then they need to first increase their payments.

²² Gross and Souleles (2002) also found an intermediate response to liquidity by accounts with utilization rates of 50%-90%. They note that this response is consistent with models with precautionary motives, in which liquidity constraints matter even if they do not currently bind, so long as there is a possibility that they bind in the future. Our utilization cutoffs are based on those in Gross and Souleles (2002), but of course the particular cutoffs are somewhat arbitrary. In general, the tighter the cutoff the more likely the resulting group includes consumers with a high probability of being constrained. If we tighten the definition of the second utilization group to 60%-90% utilization, the resulting b_0 becomes larger (though remains smaller than the b_0 for those with utilization above 90%) and significant at the 6% level. About 20% of the sample accounts have utilization above 60%.

estimates are insignificant, they continue to decline through month four (with $b_4 = -\$42$ (40)), before subsequently increasing. Regarding the accounts that start with zero utilization (and so no balances), one might expect them to be relatively less likely to respond at all to the rebate. Indeed, their results suggest little response. Their cumulative coefficients are all relatively small in magnitude and insignificant, with somewhat larger confidence intervals because of the smaller sample size. These results might be thought of as characterizing, in a sense, the amount of underlying noise in the data.

The average household has over two credit cards [1995 SCF], so by estimating at the account-level we have so far probably understated the full effect of the rebate per account-holder. As a starting point we could look at account-holders who have only one credit card in total – the card in our sample – but they are probably not representative. However one can generalize the notion of having only one card, since consumers with multiple cards can choose to concentrate their usage on a subset of their cards. As one way of measuring the relative intensity of usage of the card in our sample, we define the “usage ratio” as balances on the sample account relative to all other credit card balances held by the account-holder, based on the most recent credit bureau data. Consumers with a large usage ratio are relatively intensive users of the accounts in the sample. If the intensively used accounts tend to be the marginal accounts in terms of spending, then these accounts would be expected to respond the most to the rebate, and vice-versa.²³

Row E) distinguishes accounts with usage ratios of above 3, 0 to 3, and 0 (approximately 30%, 60%, and 10% of the sample, respectively), taking the usage ratio from month $t-9$ to minimize endogeneity. Starting with the high usage-ratio accounts, their spending increases by

²³ To illustrate the opposite case, suppose that after consumers put significant balances on one account they use their other accounts with smaller balances for their marginal spending. Then the latter accounts might respond the most to the rebate. This case is not, however, supported by the results below.

over \$200 cumulatively. This response is statistically and economically significant, corresponding to over 40% of the average household rebate. The response is also significantly larger than that of the accounts with a usage ratio of zero.²⁴ For the high-usage accounts, the underlying marginal coefficients for spending are also jointly very significant (p-value = .001). The resulting path of cumulative spending coefficients is somewhat backloaded, but nonetheless the coefficients are significant by month five (as reported) and subsequently.

By contrast, for the middle-usage accounts, the point estimates show little change in spending and a moderate increase in payments, and so a moderate decline in debt. This decline is significant in the first few months, and while it subsequently becomes less significant, it is relatively persistent, with $b_9 = -41$.²⁵ Regarding the accounts with a usage ratio of zero, they generally show little response to the rebate, like the accounts with zero utilization.²⁶

Finally, building on the previous results, row F) utilizes a composite sample-split that uses multiple characteristics to distinguish account-holders that are relatively likely to be constrained versus unconstrained. The composite split identifies potentially constrained account-holders (labeled as “constrained” in the table) as those who are young (≤ 35 years old), or have a small credit limit ($\leq \$7000$) that is relatively likely to be binding due to high utilization (above 60%). This group is expected to increase spending the most. The potentially “unconstrained” group includes account-holders who have large limits ($> \$10,500$), or who are older (> 55 years

²⁴ The results based on the number of credit cards held are consistent: Consumers who also have other credit cards (outside the sample) spend less on their account in the sample than do consumers who do not have other cards.

²⁵ For the middle-usage accounts, the b_9 for payments becomes significant on imposing the non-negativity constraint. The decline in their debt becomes more pronounced if one restricts the group to accounts with a usage ratio between 0 and 0.4, which is approximately the median ratio. Accordingly, this modified cutoff is used in the composite split in row F) below.

²⁶ One small difference between these groups is that the lowest-utilization group includes the (relatively few) accounts with negative utilization (*e.g.*, due to overpayment or returns), whereas the usage ratio drops observations with negative balances. The point estimates for the group with a usage ratio of zero show a moderate initial decline in debt. However, this is driven by an (implausible) estimated decline in spending, but this decline is small in magnitude and insignificant after the month of receipt.

old) with low utilization (below 40%) and low but positive usage (usage ratio between 0 and 0.4). This group is expected to increase payments and reduce debt the most. The composite split puts accounts with a usage ratio of zero into a separate group (labeled “no balance”), which is expected to show relatively little response to the rebate. These three groups constitute about 25%, 25%, and 10% of the sample, respectively. The remaining, harder-to-classify, 40% of accounts are grouped together as “other”.²⁷

Beginning with the potentially constrained account-holders, as reported spending significantly increases, by a cumulative $b_9 = \$212$. This response is statistically and economically significant, and significantly larger than that of each of the three other groups. The response is again somewhat backloaded. According to the point estimates, payments initially increase faster than spending, leading to a small initial decline in debt (though the cumulative coefficients for payments are all insignificant, and for debt only b_0 is significant at the 10% level or better). Debt subsequently increases, however, as spending later overtakes payments. By contrast, for the potentially unconstrained account-holders, payments significantly increase, by a substantial cumulative amount $b_9 = \$281$. This response is statistically significant and significantly larger than the corresponding b_9 for the constrained group. Many of the initial marginal payment coefficients for the unconstrained group are significant, so all of their cumulative coefficients after the month of receipt are significant (and even the immediate response b_0 is significant at the 10% level). On the other hand, their spending increases by a

²⁷ Some of the cutoffs used in defining these groups were slightly relaxed relative to the cutoffs in the previous “univariate” splits based on a single characteristic, in order to keep the composite groups from getting too small. Previous notes commented on some of these modifications. *E.g.*, accounts with utilization between 60%-90% show a significant increase in spending (at the 6% level), albeit smaller than for accounts with utilization above 90%. We also considered some even more expansive cutoffs. In general, we could expand the size of the composite constrained group and still find a significant, albeit smaller, increase in spending. Presumably this is because such expansions tend to include in the group more accounts with a smaller probability of being constrained. By contrast, we were unable to expand the unconstrained group by much and still find a significant decline in debt. As before, the characteristics used to form the groups are taken from month $t-9$ to minimize endogeneity.

much smaller and insignificant amount. Hence their debt persistently decreases, by a substantial amount $b_9 = -\$225$, which is statistically significant.

Overall, the results for these two groups are qualitatively similar to, but sharper than, the results for the corresponding groups in row C) using just the limit. Hence, subject to the same qualifications, the results again suggest that liquidity constraints are indeed more likely to be binding for the account-holders that were identified as constrained. Conversely, the account-holders identified as unconstrained behave more consistently with the LCPI model. More generally, these results, along with the other results in Table 3, show that there is significant heterogeneity in the response of different consumers to the rebates.

Briefly turning to the remaining groups in row F), the results for the “no balance” group are like those in row E) for the group with a usage ratio of zero. For the “other” group, the cumulative coefficients are all insignificant.²⁸

As an extension, we also directly examined the response of the account-holders’ other credit cards, using the credit bureau data. In sum, the response of balances on the other cards is qualitatively similar to that of balances on the accounts in the main sample, so these results generally reinforce (or at least, do not offset) our previous results. However, the estimates are insignificant and so the results are inconclusive. Identification of the response of other balances is complicated by the fact that the credit bureau data are available only quarterly. Accordingly, further discussion of these results is reserved for Appendix B, with the results appearing in Table A2.

²⁸ The point estimates for the “other” group show a (implausible) late decline in cumulative payments, but these estimates are not statistically significant. The results for the “no balance” group are not exactly identical to those for the group with a usage ratio of zero in row E), because of differences in the other groups used in the two regressions (*e.g.*, due to different missing values).

V. Related Literature

A few previous papers have studied consumers' response to tax rebates and refunds. Modigliani and Steindel (1977), Blinder (1981), and Poterba (1988) found that consumption responded too much to the 1975 tax rebate, relative to the prediction of the LCPI model, though they came to somewhat different quantitative conclusions regarding the timing and overall magnitude of the response. All three studies used aggregate time-series data, but there are a number of advantages to using micro-level data as well. First, it is difficult to analyze infrequent events like tax cuts using time-series data.²⁹ Second, with micro data one can use cross-sectional variation to investigate consumer heterogeneity, including issues such as liquidity constraints. Among recent studies using micro data, one of the most closely related to this paper is Souleles (1999), which found that consumption responds significantly to the Federal income tax refunds that most taxpayers receive each spring. That paper also found evidence of liquidity constraints.³⁰

Two recent papers analyzed the response to the 2001 tax rebates in particular. First, as noted in the introduction, Shapiro and Slemrod (2003a) found that the majority of their survey respondents reported that they would mostly save their rebate, most commonly by paying down debt. Only 22% of the respondents reported that they would mostly spend their rebate, a finding the authors calculate to imply an average marginal propensity to consume of about one third. These results are consistent with our finding that, on average, consumers initially used some of the rebate to increase their credit card payments and pay down debt. The Michigan survey results

²⁹ Blinder and Deaton (1985) found smaller consumption responses when considering the 1975 rebate together with the 1968-70 tax surcharge. However, consumption was found to be too sensitive to the pre-announced changes in taxes in the later phases of the Reagan tax cuts. The authors note that their mixed results are "probably not precise enough to persuade anyone to abandon strongly held a priori views".

³⁰ Other related studies of the response of consumption to income include Bodkin (1959), Kreinin (1961), Wilcox (1989, 1990), Parker (1999), Souleles (2000, 2002, 2004), Browning and Collado (2001), Hsieh (2003), and Stephens (2003, 2005, 2006), among others.

provide no evidence, however, of a lagged response of spending or of liquidity constraints.³¹ Second, a concurrent paper by Johnson, Parker, and Souleles (2006) finds, using the Consumer Expenditure Survey, that consumers spent on average about a third of their rebates during the three-month period in which they were received, counter to the LCPI model. This finding implies that consumers initially saved most of the rebate, though their data do not allow the authors to distinguish whether the saving took place by paying down debt or by accumulating assets. Moreover, they also find significant evidence of a substantial lagged consumption response over the next two quarters, with the long-run cumulative response being roughly two-thirds of the rebate on average. Illiquid households exhibited the strongest response. Despite the differences between consumption expenditure and credit card spending, these results are broadly consistent with the dynamics of credit card usage we estimated above.³²

VI. Conclusion

This paper used a unique, new panel dataset of credit card accounts to analyze how consumers responded to the 2001 Federal income tax rebates. We used distributed lag models to estimate the month-by-month response of credit card payments, spending, and debt to the rebates, exploiting the randomized timing of the rebates' disbursement to cleanly identify their causal effects. By limiting ourselves to the subset of potential variation that is by construction exogenous, we set a high hurdle for finding significant effects of the rebate.

³¹ McNees (1973) analyzed similar surveys of refund-recipients in 1972. 45% of his sample said they spent their refund, 24% saved it, and 28% used it to pay off debt and bills. Shapiro and Slemrod (1995) analyzed a similar survey after the change in withholding rates in 1992. Shapiro and Slemrod (2003b) used a novel follow-up survey in 2002 to try to determine whether there was a lagged response to the 2001 rebate. Of the survey 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".

³² While our results do not use aggregate time-series variation, they are also consistent with the aggregate data on consumption expenditure and saving discussed by Johnson, Parker, and Souleles (2006). Aggregate spending rose substantially in the three quarters during and after which the rebates were disbursed, whereas the saving rate rose substantially in the quarter of disbursement, but then dropped in the next quarter.

We found that, on average, consumers initially saved some of the rebate, by increasing their credit card payments and thereby paying down debt and increasing their liquidity. But soon afterwards their spending increased, counter to the canonical LCPI model and Ricardian Equivalence. For consumers whose most intensively used credit card account is in the sample, spending on that account rose by over \$200 cumulatively over the nine months after rebate receipt, which represents over 40% of the average household rebate. Because these results relied exclusively on exogenous, randomized variation, they represent compelling evidence of a causal link from the rebate to spending.

We also found other significant heterogeneity in the response to the rebate across different types of consumers. Notably, spending rose most for consumers who were, according to various criteria, initially most likely to be liquidity constrained, by up to over \$300 depending on the criterion and its tightness. By contrast, debt declined most (so saving rose most) for unconstrained consumers. These results suggest that liquidity constraints are important. More generally, the results suggest that there can be important dynamics in consumers' response to "lumpy" increases in income like tax rebates, working in part through balance sheet (liquidity) mechanisms.

VII. Appendices

A. Intra-monthly Analysis

Table A1 reports the results of the intra-monthly analysis described in Section IV. As noted, the point estimates in the table provide some rough indication that the increases in spending and payments begin slightly earlier for the early-in-the-month recipients relative to the later-in-the-month recipients: The initial marginal coefficients β_0 are somewhat larger for the recipients in week 1 relative to the other recipients, though these coefficients are not fully monotonic across the weeks of receipt. Also, these differences in month zero are not statistically significant, and jointly across all horizons the paths of spending and of payments do not significantly differ across the weeks.³³ Nonetheless, many of the results are significant in absolute terms. For spending, for all four weeks, many of the marginal effects at intermediate horizons are significant. Thus the cumulative effects become significant within a few months (for brevity, not reported). For both payments and spending, and for all four weeks, the long-run cumulative effects b_9 are statistically and economically significant (bottom of Table A1), even though they do not significantly differ across the weeks.³⁴

The intermediate dynamics are qualitatively similar to those in Table 2. For most of the weeks of receipt, the point estimates suggest that payments begin to increase before spending

³³ To test for such differences, we created three indicator variables for accounts whose rebates were disbursed in the first three weeks of the month (whether in July, August, or September), and added them and their interactions with R and all of its lags to equation (1), omitting the fourth week. (Note, however, that Table A1 reports the implied total effects for each group, not the differences relative to week 4.) For each week 1 to 3 and each of spending and payments, the interaction terms are jointly insignificant. The interaction terms are also jointly insignificant when considered jointly across all three weeks together, for each of spending and payments.

³⁴ While the initial responses of payments and spending can differ across early- versus late-in-the-month recipients, the long-run responses should not differ. As just noted, they do not. For comparison, distinguishing the week of receipt without imposing the non-negativity constraint yields substantially less precise results than in Table A1: $b_9 = 122$ (93), 15 (97), 63 (89), and 1 (91) for payments for weeks 1 to 4 respectively; $b_9 = 129$ (84), 9 (89), 115 (80), and 73 (82) for spending across weeks 1 to 4; and $b_9 = 31$ (93), 11 (99), 67 (91), and 69 (95) for debt across weeks 1 to 4. We also tried constraining the marginal coefficients in Table A1 to lie on a low-order polynomial, but in general this did not help increase precision much.

(less so in week 3), but soon afterwards spending increases faster, such that the point estimates for b_9 for spending are generally similar in size or larger than the b_9 estimates for payments. However, given the statistical uncertainty around the estimated b_9 's for spending and payments, one cannot make from these results strong inferences about their implications for the long-run change in debt, since the latter depends on the difference between the former.

B. Balances on Other Credit Cards

Table A2 analyzes the response of the account-holders' other credit cards, using the credit bureau data. Because the credit bureau data on other balances are available only quarterly, we cannot identify the average response of these balances over time separately from month indicator variables.³⁵ Nonetheless we can still examine whether, in any given month of data, the other balances are larger for account-holders who received their rebates earlier (*i.e.*, in July and August 2001) relative to those who received their rebates later (in September 2001).

To interpret the results in the table, note that other balances are available in four months of the sample (June, September, and December 2001, and March 2002), and the specification includes the corresponding month indicators. Omitting R_t , the independent variable R_{t-1} [or R_{t-2}] then measures how much larger or smaller are the balances in September 2001 of those who received their rebates in August [July] 2001, relative to the balances in September 2001 of those who received their rebates in September 2001. The resulting point estimates for R_{t-1} and R_{t-2} are positive. This suggests that other balances increased between the month of receipt and the next two months, though the estimates are not statistically significant. Similarly, the point estimates for R_{t-4} and R_{t-5} , and for R_{t-7} and R_{t-8} , are also positive, suggesting that the balances of the earlier recipients remain higher through December 2001 and March 2002. However, these estimates are

³⁵ *E.g.*, since the rebates were disbursed in July, August, and September 2001, in the credit bureau data from September 2001, the combination of R_t , R_{t-1} , and R_{t-2} would be collinear with a month indicator for September 2001.

again insignificant, and jointly all the differences together $\{R_{t-s} \mid s = 1, 2, 4, 5, 7, 8\}$ are also insignificant. Hence these conclusions must be qualified.³⁶ For comparison, Table A2 applies the same specification to the balances of the accounts in the main sample, using the data only for the same four months for which other balances are available from the credit bureaus. The results are qualitatively similar, with all the regressors R_{t-s} being positive though insignificant. Hence, with the noted qualification, the response of other credit cards appears if anything to reinforce (or at least, to not offset) the responses estimated above for the accounts in the main sample.

C. The Data

The main unit of analysis in the data is an individual credit card account. The central account billing-statement data (*i.e.*, total payments and spending, and debt) are available by cycle-month. Debt includes only interest-incurring balances (*i.e.*, balances rolled over into the next month), not transactions balances (*i.e.*, balances paid off).

The credit bureaus store their information by individual borrower. The credit bureau data on *other* balances record total month-end balances across all other credit cards held by the same account-holder. This includes transactions balances – the credit bureaus do not separately record spending and debt. (For additional discussion of credit bureau data, see *e.g.* Musto and Souleles (2006).) The issuer obtained the credit bureau data every three months. For consistency, the usage ratio is defined as month-end balances on the account in the sample divided by (other balances + \$1), using the most recent credit bureau data on other balances. If both the numerator and other balances in the denominator of this ratio are zero, or if balances are negative, the ratio

³⁶ Eventually the balances of all three groups of rebate recipients should converge. However, as noted the estimated differences between the groups are not significant even in the short-run. Also, these results do not pin down the overall level of balances, and so cannot rule out a decline in balances in the month of receipt ($s = 0$). Note, however, that balances include both debt and transactions balances, so their behavior can differ from that of debt.

is set to missing. In Table 3, if a variable used to split the sample is missing, the corresponding observation is dropped from the corresponding regression.

The dataset contains a representative sample of accounts open as of June 2000, with the following exclusions: a) accounts that are bankrupt or two or more months delinquent, or otherwise frozen; and b) accounts that are dormant/closed, including *e.g.* credit cards that were issued but never activated. Following the data provider's standard practice given the available data fields, dormant/closed accounts were identified as those without any retail activity in the previous three quarters.

References

- Auerbach, A.J., 2002, "The Bush Tax Cut and National Saving," National Tax Journal, LV, pp. 387-407.
- Blinder, A.S., 1981, "Temporary Income Taxes and Consumer Spending," Journal of Political Economy, 89, pp. 26-53.
- Blinder, A.S., and Deaton, A., 1985, "The Time Series Consumption Function Revisited," Brookings Papers on Economic Activity, 2, pp. 465-511.
- Bodkin, R., 1959, "Windfall Income and Consumption," American Economic Review, 49, pp. 602-14.
- Browning, M., and Collado, M.D., 2001, "The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates," American Economic Review, 91, pp. 681-92.
- Browning, M., and Lusardi, A., 1996, "Household Saving: Micro Theories and Micro Facts," Journal of Economic Literature, 34, pp. 1797-1855.
- Carroll, C., 1992, "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," Brookings Papers on Economic Activity, 2, pp. 61-156.
- Chimerine, L., 1997, "Americans in Debt: The Reality," MasterCard International.
- Federal Reserve Board of Governors, 2007, "G.19: Consumer Credit."
- Gross, D., and Souleles, N.S., 2002, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," Quarterly Journal of Economics, 117, pp. 149-85.
- Hsieh, C., 2003, "Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund," American Economic Review, 93, pp. 397-405.
- Jappelli, T., 1990, "Who Is Credit Constrained in the U.S. Economy?" Quarterly Journal of Economics, 105, pp. 219-34.
- Jappelli, T., Pischke, J.-S., and Souleles, N.S., 1998, "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," The Review of Economics and Statistics, 80, pp. 251-62.
- Johnson, D.S., Parker, J.A., and Souleles, N.S., 2006, "Household Expenditure and the Income Tax Rebates of 2001," American Economic Review, 96, pp. 1589-1610.
- Kiefer, D., *et al.*, 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, M.E., 1961, “Windfall Income and Consumption: Additional Evidence,” American Economic Review, 51, pp. 388-90.
- McNees, S.K., 1973, “Tax Refunds and Consumer Spending,” New England Economic Review, January/February, pp. 3-11.
- Modigliani, F., and Steindel, C., 1977, “Is a Tax Rebate an Effective Tool for Stabilization Policy?” Brookings Papers on Economic Activity, 1, pp. 175-209.
- Musto, D.K., and Souleles, N.S., 2006, “A Portfolio View of Consumer Credit,” Journal of Monetary Economics, 53, pp. 59-84.
- Parker, J.A., 1999, “The Reaction of Household Consumption to Predictable Changes in Social Security Taxes,” American Economic Review, 89, pp. 959-73.
- Poterba, J.M., 1988, “Are Consumers Forward Looking? Evidence from Fiscal Experiments,” American Economic Review, 78, pp. 413-18.
- Reis, R., 2006, “Inattentive Consumers,” Journal of Monetary Economics, 53, pp. 1761-1800.
- Shapiro, M.D., and Slemrod, J.B., 1995, “Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding,” American Economic Review, 85, pp. 274-83.
- Shapiro, M.D., and Slemrod, J.B., 2003a, “Consumer Response to Tax Rebates,” American Economic Review, 93, pp. 381-96.
- Shapiro, M.D., and Slemrod, J.B., 2003b, “Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys,” in Tax Policy and the Economy, ed. Poterba, J.M., Cambridge: MIT Press.
- Shea, J., 1995, “Union Contracts and the Life-Cycle/Permanent-Income Hypothesis,” American Economic Review, 85, pp. 186-200.
- Slemrod, J. B., Christian, C., London, R. and Parker, J. A., 1997, “April 15 Syndrome,” Economic Inquiry, 35, pp. 695-709.
- Souleles, N.S., 1999, “The Response of Household Consumption to Income Tax Refunds,” American Economic Review, 89, pp. 947-58.
- Souleles, N.S., 2000, “College Tuition and Household Savings and Consumption,” Journal of Public Economics, 77, pp. 185-207.

- Souleles, N.S., 2002, "Consumer Response to the Reagan Tax Cuts," Journal of Public Economics, 85, pp. 99-120.
- Souleles, N.S., 2004, "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys," Journal of Money, Credit, and Banking, 36, pp. 39-72.
- Stephens, M., 2003, "'3rd of the Month': Do Social Security Recipients Smooth Consumption Between Checks?" American Economic Review, 93, pp. 406-22.
- Stephens, M., 2005, "The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans," Carnegie-Mellon University, Working Paper.
- Stephens, M., 2006, "Paycheque Receipt and the Timing of Consumption," The Economic Journal, 116, pp. 680-701.
- Wilcox, D.W., 1989, "Social Security Benefits, Consumption Expenditure, and the Life Cycle Hypothesis," Journal of Political Economy, 97, pp. 288-304.
- Wilcox, D.W., 1990, "Income Tax Refunds and the Timing of Consumption Expenditure," Federal Reserve Board of Governors working paper, April.
- Yoo, P., 1998, "Still Charging: The Growth of Credit Card Debt between 1992 and 1995," Federal Reserve Bank of St. Louis Review, January/February, pp. 19-27.

Table 1: Sample Statistics

variable	mean (\$)	s.d. (\$)
payments	349	939
spending	327	895
debt	1788	2866
balances	2144	2958
other balances	7871	13030
credit limit	8584	3353
Δ payments	-5.3	1234.7
Δ spending	-1.5	1098.7
Δ debt	1.2	1045.3
# observations	739945	

Notes: The data come from the monthly billing statement of credit card accounts, except for “other balances” on the other credit cards held by the account-holders, which is obtained quarterly from the credit bureaus. All values are averaged over the sample period (3/2001 - 5/2002) used in the baseline results in Table 2, and are quoted in current dollars. The differences in the final three rows are average monthly changes over the sample period.

Table 2: Consumer Response to Rebates

	Payments		Spending		Δ Debt	
	coef.	s.e.	coef.	s.e.	coef.	s.e.
β_0	11.6	6.2 *	-2.8	5.6	-14.3	7.2 **
β_1	11.4	8.1	6.2	7.2	-6.0	9.0
β_2	8.2	9.2	2.7	8.5	-7.2	10.1
β_3	3.8	10.1	7.6	9.3	5.0	10.9
β_4	8.9	10.9	8.3	10.0	0.3	11.7
β_5	-1.9	11.4	8.3	10.4	11.3	12.1
β_6	2.9	11.7	13.8	10.4	15.9	12.3
β_7	-2.8	11.4	9.4	9.9	16.1	12.0
β_8	-0.6	10.7	6.4	9.0	10.1	11.2
β_9	7.2	11.0	1.9	9.2	-3.7	11.3
test: $\{\beta_s\}$ joint sig.	0.14		0.15		0.07	
Implied Long-Run Cumulative Effects						
b_9	48.7	54.5	61.7	68.9	27.5	23.3

Notes: Sample size $N = 739945$. This table reports the marginal effects β_s , $s = 0-9$, of receiving a rebate, corresponding to the indicator variable R_t and its nine lags, R_{t-1} to R_{t-9} , in equation (1). b_9 gives the corresponding implied long-run cumulative effect. All values are in current dollars (3/2001 - 5/2002). The reported p-values are from tests of the joint significance of the marginal effects β_0 to β_9 . (These are equivalent to tests of the joint significance of the cumulative effects b_0 to b_9 .) Each regression also includes a full set of month indicator variables. The standard errors are adjusted for heteroscedasticity across accounts as well as serial correlation within accounts.

** = significantly different from zero at the 5% level, * = at the 10% level

Table 3: Heterogeneity in the Response to Rebates

row	Payments					Spending					Δ Debt				
	b_9	s.e.	test: $b_9=0$	test: $b_5=0$	test: $\{\beta_s\}$ sig	b_9	s.e.	test: $b_9=0$	test: $b_5=0$	test: $\{\beta_s\}$ sig	b_9	s.e.	test: $b_9=0$	test: $b_5=0$	test: $\{\beta_s\}$ sig
A) Marital Status (N=739945)															
singles	73.6	78.1	0.35	0.25	0.28	60.7	70.0	0.39	0.69	0.18	9.6	78.9	0.90	0.49	0.01
couples	16.2	85.4	0.85	0.60	0.21	73.6	76.3	0.33	0.21	0.71	57.9	85.4	0.50	0.60	0.60
combined test					0.11					0.20					0.02
B) Age (N=720818)															
age <35	48.0	94.7	0.61	0.87	0.95	199.6	84.7	0.02	0.17	0.17	131.2	94.5	0.17	0.48	0.44
age 35-60	69.5	81.6	0.39	0.27	0.50	56.8	73.1	0.44	0.36	0.63	0.9	82.3	0.99	0.83	0.61
age >60	28.3	85.8	0.74	0.35	0.54	23.3	77.8	0.76	0.89	0.19	29.7	87.5	0.73	0.42	0.03
combined test					0.78					0.04					0.13
C) Credit Limit (N=731616)															
limit \leq \$7000	41.2	79.4	0.60	0.56	0.27	141.0	71.4	0.05	0.42	0.01	109.0	79.7	0.17	0.94	0.14
limit \$7-10.5k	-17.5	83.5	0.83	0.65	0.45	30.4	75.2	0.69	0.69	0.85	75.2	84.0	0.37	0.97	0.33
limit > \$10500	193.2	101.8	0.06	0.08	0.07	39.7	93.2	0.67	0.34	0.39	-145.3	104.6	0.16	0.45	0.13
combined test					0.03					0.001					0.002
D) Utilization (N=739923)															
util >90	224.5	131.0	0.09	0.12	0.04	332.8	106.2	0.002	0.02	0.001	99.3	145.8	0.50	0.83	0.12
util 50-90	103.5	102.4	0.31	0.70	0.60	124.3	87.3	0.15	0.19	0.41	67.2	110.0	0.54	0.38	0.73
util 1-50	25.1	78.5	0.75	0.41	0.35	19.7	70.6	0.78	0.92	0.88	4.9	78.0	0.95	0.48	0.34
util \leq 0	4.0	92.5	0.97	0.57	0.92	48.7	89.7	0.59	0.52	0.70	43.9	96.4	0.65	0.93	0.45
combined test					0.23					0.06					0.34

row	Payments					Spending					Δ Debt				
	b_9	s.e.	test: $b_9=0$	test: $b_5=0$	test: $\{\beta_s\}$ sig	b_9	s.e.	test: $b_9=0$	test: $b_5=0$	test: $\{\beta_s\}$ sig	b_9	s.e.	test: $b_9=0$	test: $b_5=0$	test: $\{\beta_s\}$ sig
E) Usage Ratio (N=707248)															
ratio >3	75.4	88.1	0.39	0.22	0.28	235.8	78.9	0.003	0.05	0.001	127.0	84.8	0.13	0.85	0.08
ratio 0-3	54.7	82.3	0.51	0.34	0.17	-20.4	73.2	0.78	0.88	0.43	-40.5	83.5	0.63	0.40	0.04
ratio 0	1.7	103.1	0.99	0.81	0.93	-34.7	98.8	0.73	0.67	0.46	-9.9	107.8	0.93	0.64	0.14
combined test					0.34										0.002
F) Composite Split (N=707162)															
constrained	103.5	89.8	0.25	0.57	0.54	211.9	79.8	0.01	0.14	0.13	112.2	90.9	0.22	0.56	0.31
unconstrained	280.9	96.7	0.004	0.01	0.01	36.5	87.7	0.68	0.68	0.42	-225.5	98.6	0.02	0.04	0.03
no balance	49.9	103.2	0.63	0.64	0.96	12.8	98.8	0.90	0.86	0.36	-5.6	107.8	0.96	0.67	0.16
other	-96.3	82.9	0.25	0.92	0.31	-5.7	74.4	0.94	0.97	0.55	104.3	82.9	0.21	0.98	0.09
combined test					0.001					0.01					0.001

Notes: This table shows the long-run cumulative response to the rebate, b_9 , across different decompositions of the sample account-holders. (The corresponding paths of cumulative coefficients b_s , $s = 0-9$, are graphed in Figure 2, grouped under the same row labels “A” to “F”.) Each row A) to F) corresponds to a different regression. Each regression adds to equation (1) an indicator variable for each group of account-holders in the regression, and its interaction with rebate receipt R_t and its nine lags. In each line of a given regression, p-values are reported for tests of the joint significance of the marginal coefficients $\{\beta_s\}$ (equivalently, the cumulative coefficients $\{b_s\}$) across all horizons, for the group in that line separately. The “combined test” gives the p-values for the joint significance of these coefficients jointly across all groups in the regression. The utilization rate is balances on the account normalized by the credit limit. The usage ratio measures balances on the account relative to balances on all other credit cards. In row F), the “constrained” group includes account-holders who are young, or have a small credit limit resulting in high utilization; the “unconstrained” group includes account-holders who have large limits, or who are older with low utilization and low but positive usage; “no balance” includes accounts with a usage ratio of zero; and “other” represents the remaining accounts. Each regression also includes a full set of month indicators. The standard errors are adjusted for heteroscedasticity across accounts as well as serial correlation within accounts.

Notes: $N = 739945$. This table reports the marginal effects β_s depending on the week within the month the rebate was received. Week 1 represents account-holders who received their rebates in the first week of a month (whether July, August, or September 2001), week 4 represents rebates received in the last week. The specification adds to equation (1) an indicator variable for each week and its interaction with rebate receipt R_t and its nine lags. The specification also includes a full set of month indicators. The marginal coefficients are constrained to be non-negative, $\beta_s \geq 0 \forall s$. (The cases in which the constraint binds are identified by “n.a.” in the standard errors column.) The reported p-values are from tests of the joint significance of the marginal coefficients $\{\beta_s\}$ jointly across all four weeks combined. The standard errors are adjusted for heteroscedasticity across accounts as well as serial correlation within accounts.

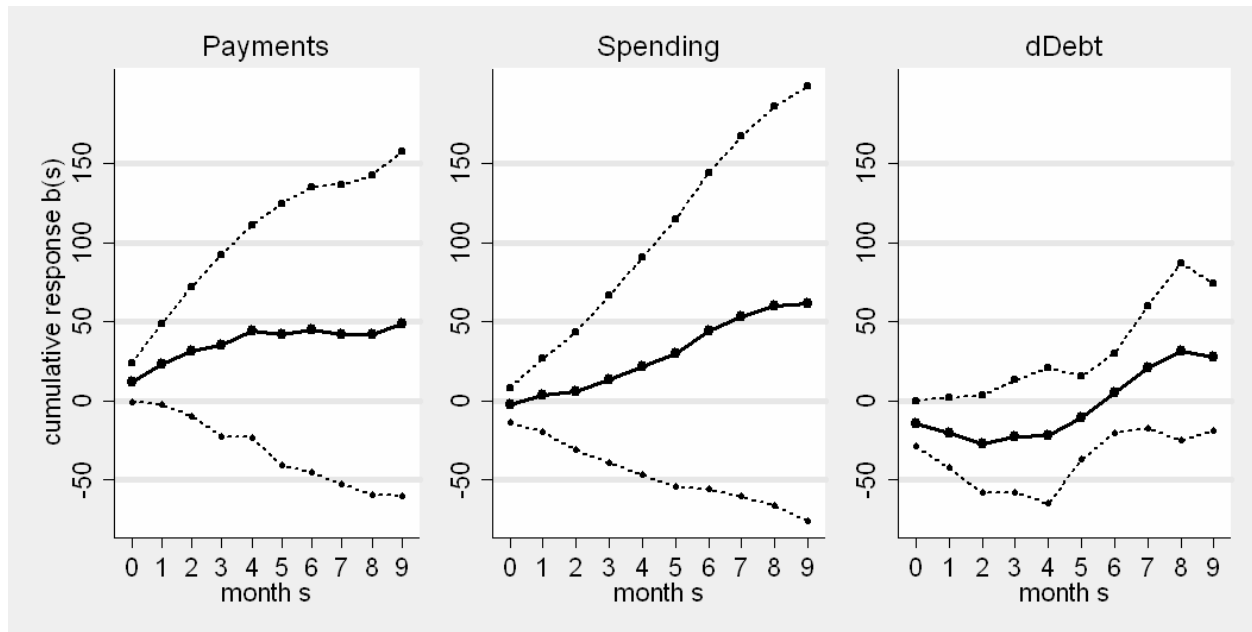
** = significantly different from zero at the 5% level, * = at the 10% level

Table A2: The Response Across Other Credit Card Accounts

	Other Balances		Balances	
	coef.	s.e.	coef.	s.e.
R_{it-1}	9.1	128.7	13.2	30.0
R_{it-2}	141.7	163.3	24.7	37.6
R_{it-4}	148.3	138.3	12.3	31.3
R_{it-5}	120.1	173.6	39.2	39.4
R_{it-7}	167.0	152.0	41.5	32.8
R_{it-8}	150.2	192.8	50.0	41.0
test: $\{R_{t-s}\}$ joint sig.	0.57		0.81	

Notes: N = 204747. Other balances are month-end balances on all other, non-sample credit cards held by the sample account-holders, using the quarterly credit bureau data. These data are available in June, September, and December 2001, and March 2002. (The specification includes the corresponding month indicator variables.) For each month of data, the coefficients R_{t-s} show whether other balances are larger or smaller for account-holders who received their rebates earlier (in July and August 2001), relative to those who received their rebates later (in September 2001). See the text for further discussion. The reported p-values are from tests of the joint significance of the lagged rebate indicators $\{R_{t-s} \mid s = 1, 2, 4, 5, 7, 8\}$. Balances represent month-end balances on the accounts in the main sample. For comparability, this variable is used only in the same months for which other balances are available from the credit bureau. The standard errors are adjusted for heteroscedasticity across accounts as well as serial correlation within accounts.

Figure 1: Consumer Response to Rebates



Notes: These figures graph the cumulative effects b_s , $s = 0-9$, implied by the baseline results in Table 2, along with their corresponding 95% confidence intervals (dashed lines), in current \$.

Figure 2: Heterogeneity in the Response to Rebates

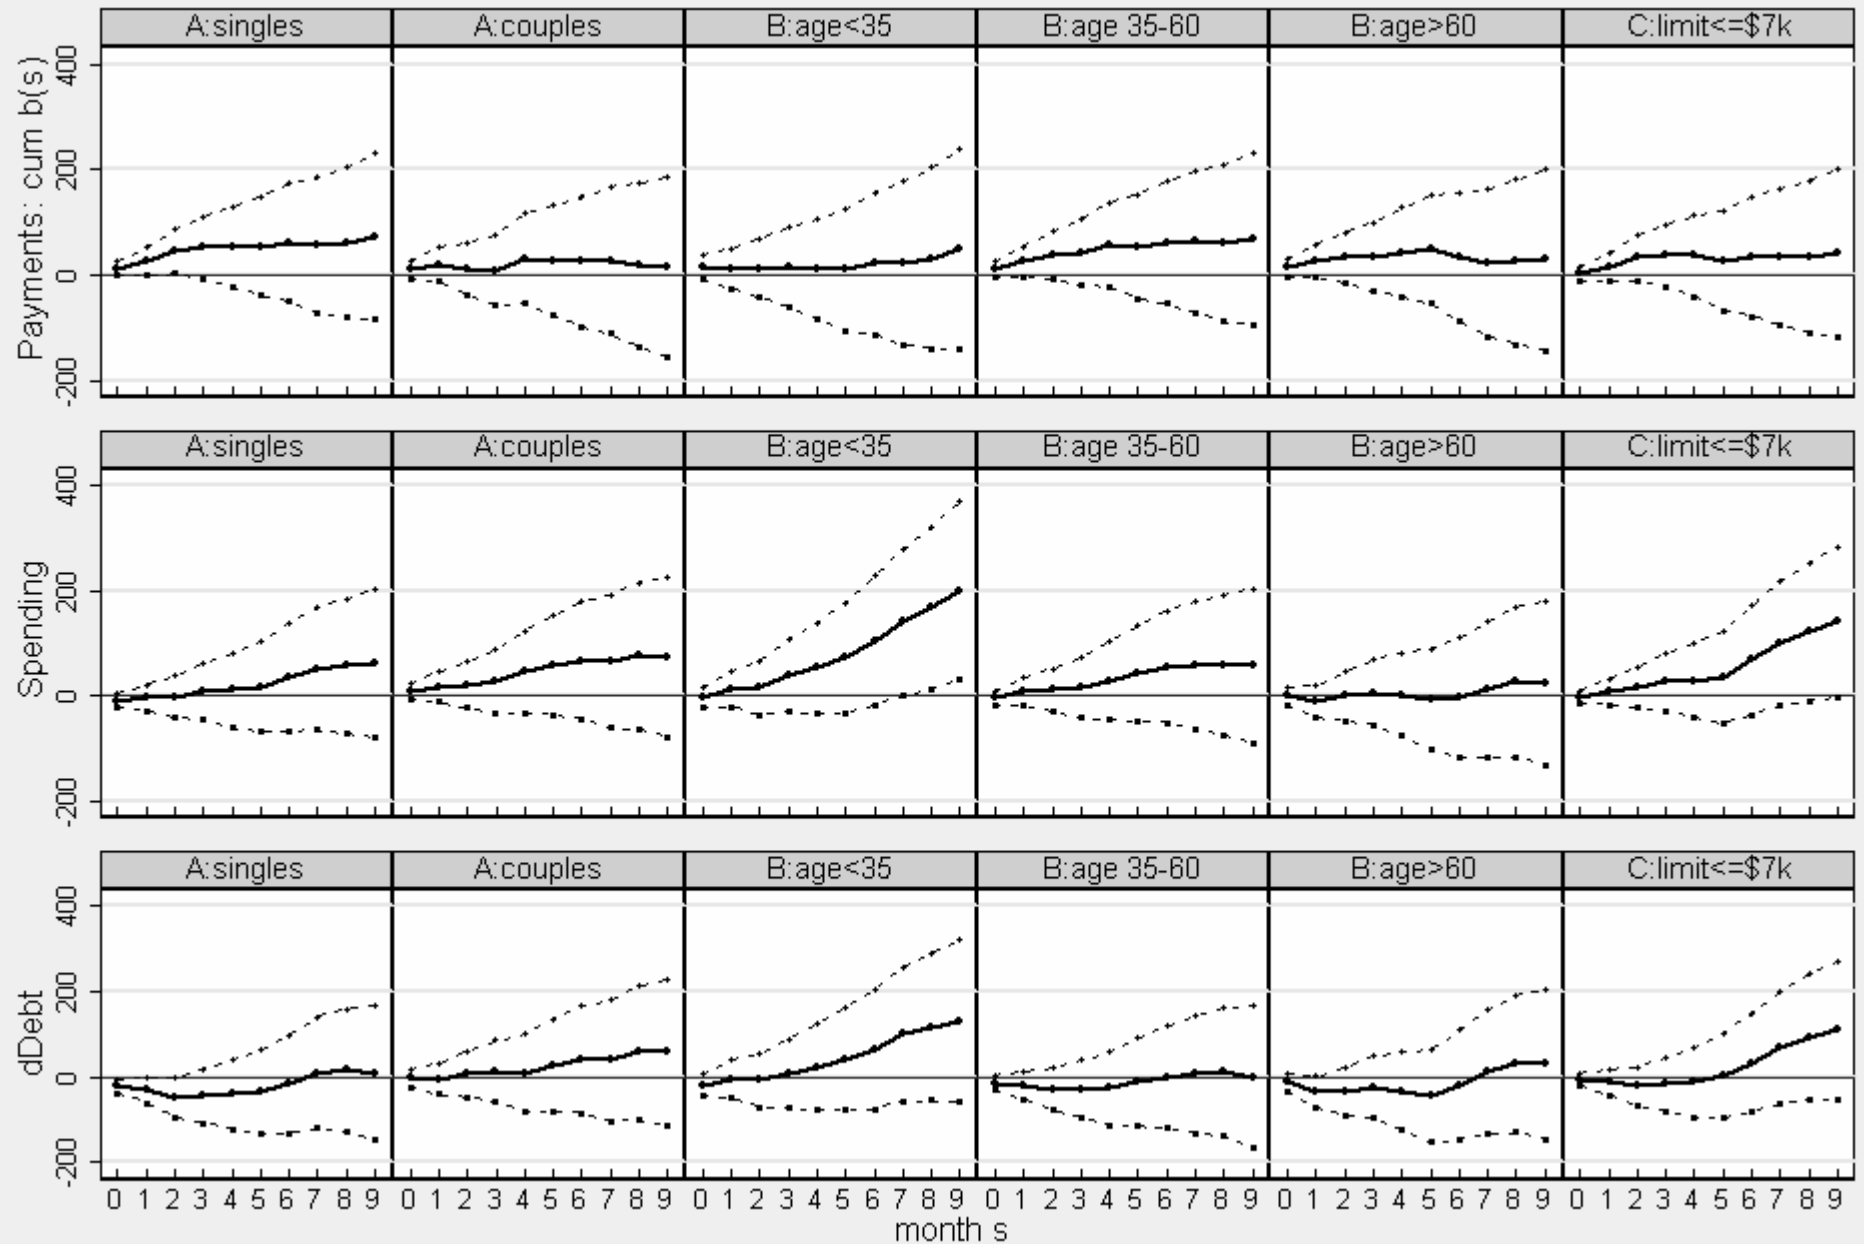


Figure 2 (ctd): Heterogeneity in the Response to Rebates

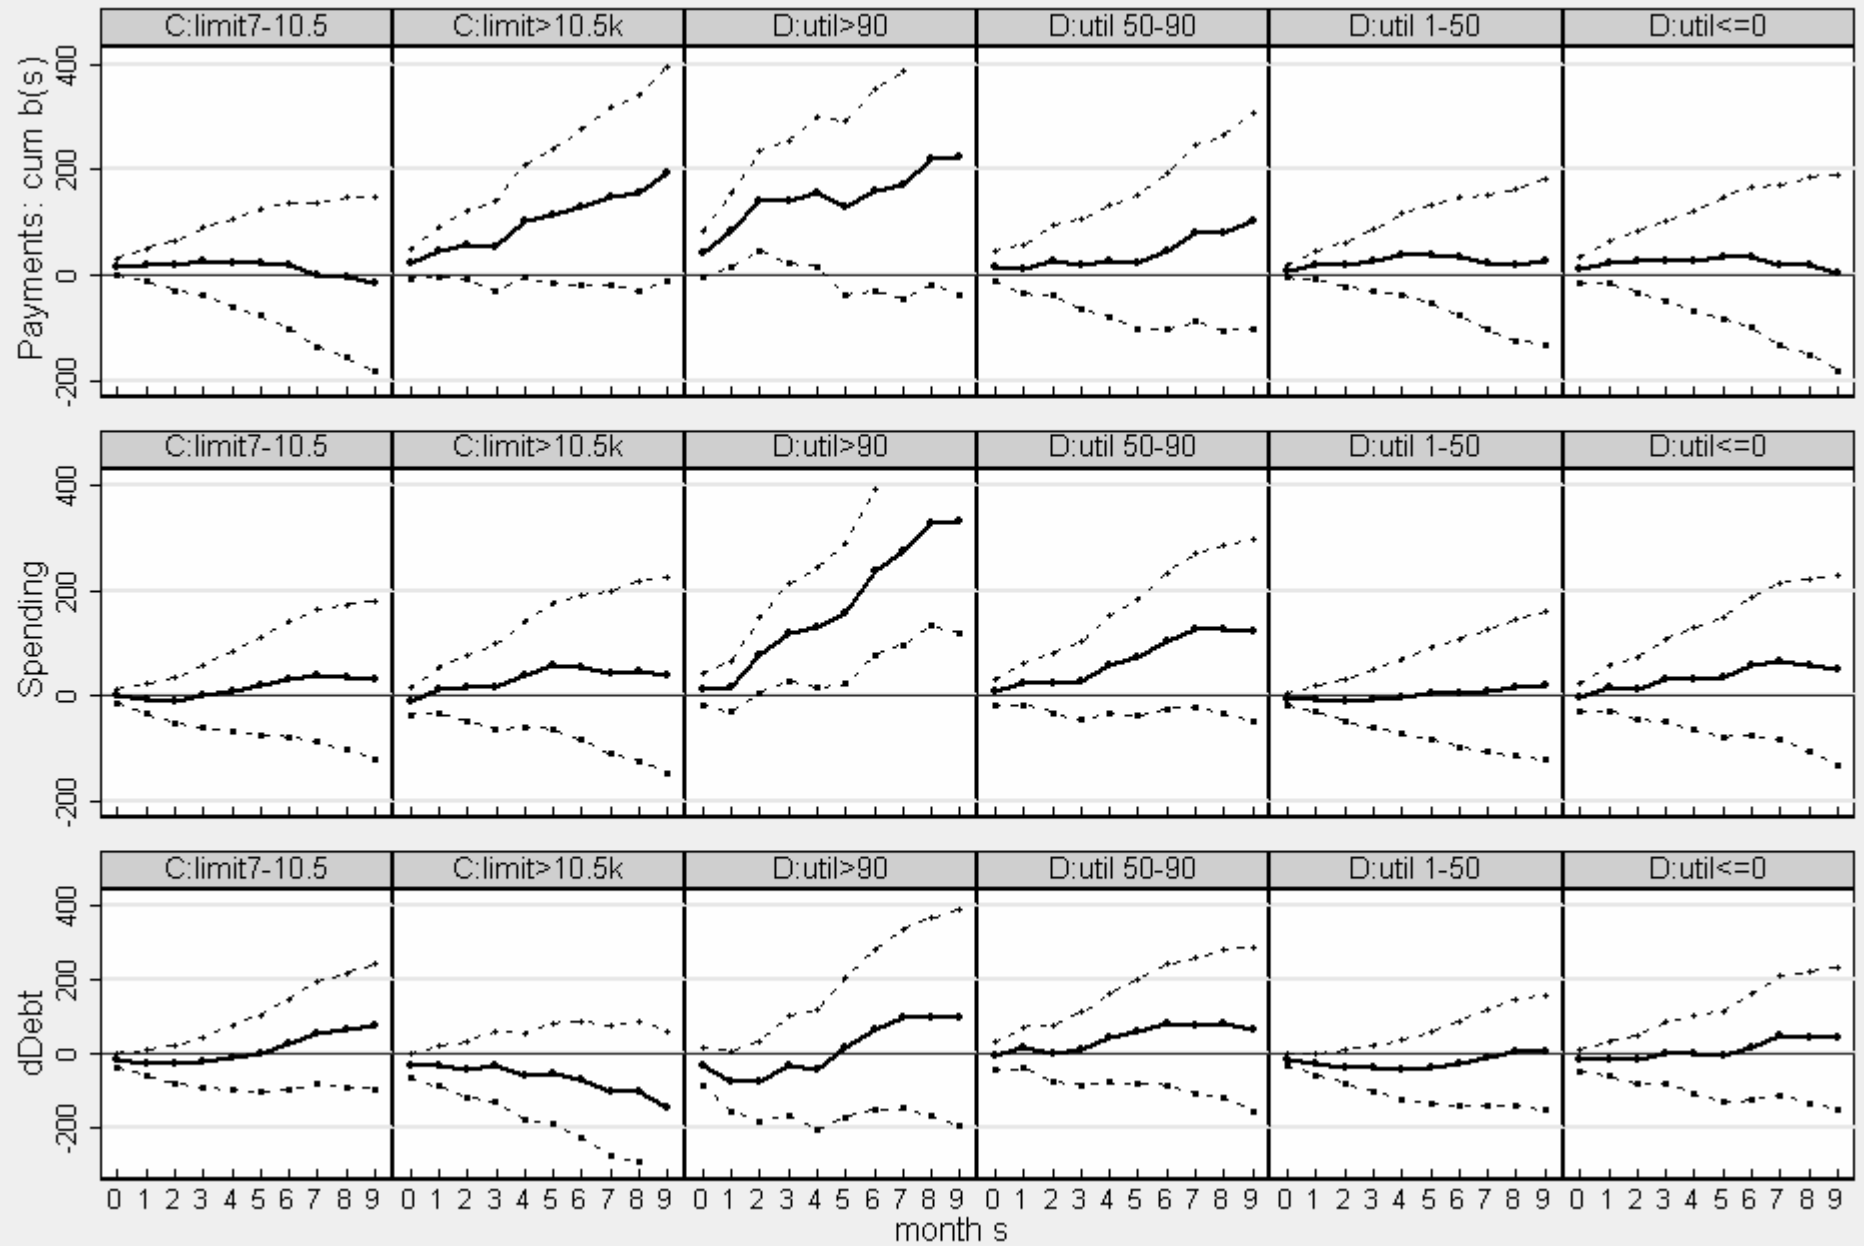
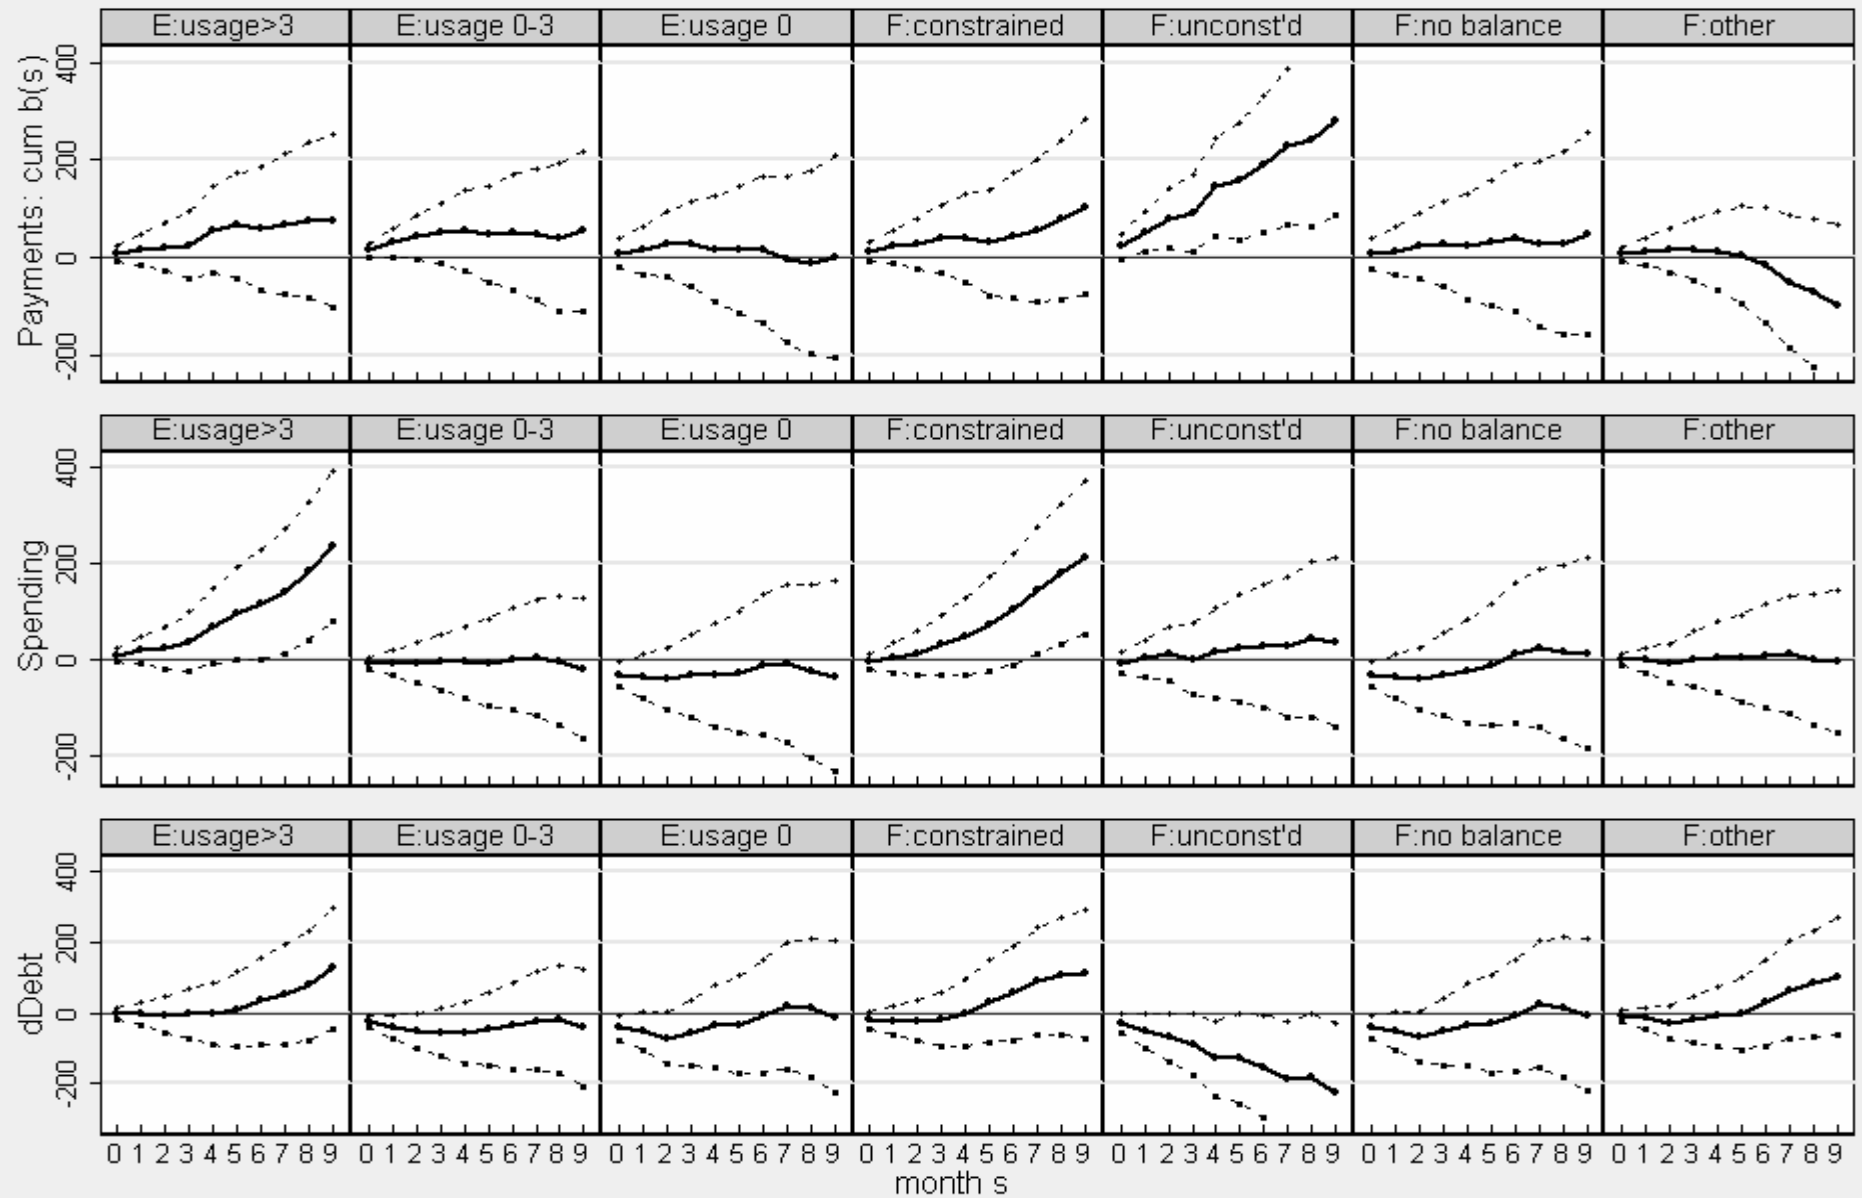


Figure 2 (ctd): Heterogeneity in the Response to Rebates



Notes: These graphs (for payments, spending, and the change in debt in separate panels) show the cumulative coefficients b_s , $s = 0-9$, corresponding to the various groups of account-holders in Table 3, along with 95% confidence intervals (dashed lines), in current \$. See Table 3 for group definitions and other details. For readability, a few of the final, largest (in absolute value) confidence-interval points are omitted; Table 3 reports all the final standard errors (for b_9).