

NBER WORKING PAPER SERIES

HOW INDIVIDUALS SMOOTH SPENDING:
EVIDENCE FROM THE 2013 GOVERNMENT SHUTDOWN USING ACCOUNT DATA

Michael Gelman
Shachar Kariv
Matthew D. Shapiro
Dan Silverman
Steven Tadelis

Working Paper 21025
<http://www.nber.org/papers/w21025>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2015, Revised July 2018

This research is part of a collaboration with Intuit Mint Bills (formerly Check). We are grateful to the executives and employees of Intuit who have made this research possible. We thank Patrick Kline, Dimitriy Masterov, Melvin Stephens Jr. and Jeffrey Smith and participants at seminars at the Bureau of Economic Analysis and the University of Southern California for helpful comments. This research is supported by a grant from the Alfred P. Sloan Foundation. Shapiro acknowledges additional support from the Michigan node of the NSF-Census Research Network (NSF SES 1131500). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Michael Gelman, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How Individuals Smooth Spending: Evidence from the 2013 Government Shutdown Using Account Data

Michael Gelman, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis
NBER Working Paper No. 21025

March 2015, Revised July 2018

JEL No. D12,D91,E21,H31

ABSTRACT

Using comprehensive account records, this paper examines how individuals respond to a temporary drop in income following the 2013 U.S. Federal Government shutdown. Affected employees saw their income decline by 40% on average, which was recovered within two weeks. Despite having no effect on lifetime earnings, spending dropped sharply, implying a naïve estimate of the marginal propensity to spend of 0.57. This estimate overstates how consumption responded. To smooth consumption, individuals adjusted by delaying recurring payments such as mortgages and credit card balances. Those with the least liquidity struggled most to smooth spending and were left holding more debt months after the shutdown.

Michael Gelman
Department of Economics
University of Michigan
Ann Arbor MI 48109-1220
mgelman@umich.edu

Shachar Kariv
Department of Economics
University of California, Berkeley
Berkeley, CA 94720
kariv@berkeley.edu

Matthew D. Shapiro
Department of Economics
University of Michigan
611 Tappan St
Ann Arbor, MI 48109-1220
and NBER
shapiro@umich.edu

Dan Silverman
Department of Economics
W.P. Carey School of Business
Arizona State University
P.O. Box 879801
Tempe, AZ 85287-9801
and NBER
dsilver3@asu.edu

Steven Tadelis
Haas School of Business
University of California, Berkeley
545 Student Services Building
Berkeley, CA 94720
and NBER
stadelis@haas.berkeley.edu

There is substantial evidence that households respond to income shocks by more than standard models would predict. A key explanation of the excess response of spending to income is that income shocks relax liquidity constraints. Interpreting these responses is challenging because shocks to income and shocks to liquid wealth usually arrive together, making it hard to distinguish the response to a change in liquidity from the response to a change in resources. Moreover, it is notoriously difficult to isolate from conventional data on income and assets a shock to expected income, its duration, and whether the shock reverts entirely and thus affects liquidity rather than lifetime income.

Measurement is also challenging. Estimating the response to a liquidity shock requires comprehensive data on households' income and financial position. Existing data typically capture only some dimensions with sufficient resolution. They may measure total spending with precision, but not savings or debt; or they measure spending and debt well, but do not measure income with similar accuracy.

In this paper, we address these challenges by identifying a pure shock to liquidity. The 2013 U.S. Federal Government shutdown produced a significant, temporary, and easily identified negative shock to the liquidity of a large number of affected government workers. We address the challenge of measuring a household's full range of responses to this shock by exploiting a new dataset derived from the de-identified and integrated transactions and balance data of more than 1 million households in the U.S.¹

For affected government workers, the shutdown caused a large, unexpected drop in pay that was fully recovered within two weeks. The integrated account data allow us to isolate

¹ The data are captured in the course of business by a mobile banking application. This dataset has already proved useful for studying the high-frequency response of spending to regular, anticipated income by levels of spending, income, and liquidity (Gelman et al. 2014). The related literature section below discusses other studies that use similar types of account data.

Federal government workers affected by the shutdown using the transaction description associated with the direct deposit of their paychecks to their bank accounts. By knowing who was subject to the liquidity shock, we can examine their responses in terms of spending and other variables before, during, and after the government shutdown. These responses are estimated by a difference-in-difference approach, where the outcomes of affected government workers are compared with those of a control group consisting of workers that have the same biweekly pay schedule as the Federal government, but who were not subject to the shutdown. The control group is mainly non-Federal workers, though also includes some Federal workers not subject to the shutdown.

The pay of a typical affected worker was 40% below normal during the shutdown because the government was closed from October 1 to October 16, 2013, thus including the last four days of the previous ten-day pay period. By the next pay period, however, government operations had resumed and workers were reimbursed fully for the pay lost during the shutdown. The transaction data show this pattern clearly for affected workers. Because of routine lags between end of pay period and receipt of paycheck funds, at the time affected workers received the reduced paycheck they had clear indications that the shutdown would be resolved with full compensation for lost pay. Hence, the shutdown produced a reduction in liquidity that was exactly offset within two weeks. It is a pure shock to liquidity because it carried no implications for lifetime resources, or indeed for resources beyond the duration of the shutdown.

An important fact revealed by the balance records is that many affected workers maintained low levels of liquid assets (checking and saving account balances), especially in the days just before their regular paychecks arrive. Prior to the shutdown, the median worker in the data held an average liquid assets balance sufficient to cover just eight days of average spending.

Moreover, liquid assets exhibit systematic changes over the pay-cycle. Just before payday, the median level of liquid assets is only five days of average spending. Indeed, a substantial fraction of this population barely lives paycheck-to-paycheck. On the day before their paycheck arrives, the bottom third of the liquid assets distribution has, on average, a combined checking and savings account balance of zero. Hence, this paper provides novel insight into how potentially financially fragile households cope with a large, temporary shock.

Given such low levels of liquid assets, it is perhaps unsurprising that the transaction records show a sharp drop in total spending by affected workers during the week of missing paycheck income. Weekly spending declined by roughly half the reduction in paycheck income and then recovered roughly equally over the two pay periods following the end of the shutdown. Econometric analysis reveals, at the margin, about \$0.58 less spending in response to each dollar of lost liquidity. Most individuals reversed this drop in spending immediately after they receive the paychecks that reimbursed them for their lost income.

It is perhaps puzzling that so many workers maintained such low liquid asset levels and exhibited such a sharp spending response to an unexpected but brief delay in income.² These facts could be taken to suggest either that benchmark theories founded on a taste for smoothing consumption are badly specified; or that households are inadequately buffered against even very

² These low levels of liquid assets and sensitivity of spending to a liquidity shock are broadly consistent with a survey-based literature about financial fragility. Caner and Wolff (2004), for example, analyze the Panel Study of Income Dynamics and find that the median household maintains less than \$5,000 (1999) combined in checking, saving, and investment assets. Lusardi, Schneider and Tufano (2011) asked more than 2,000 US households “How confident are you that you could come up with \$2,000 if an unexpected need arose within the next month?” Nearly 28% said they certainly would not be able to come up with that much and another 22% reported they probably would not be able to. In a 2013 Federal Reserve Board survey, just 53% of households said they could cover a \$400 emergency either with checking or savings account balances or with a credit card that would be paid off at the end of the month.

temporary shocks; or that the financial markets that make consumption smoothing possible are functioning poorly.

The key finding of this paper is, however, that even workers with surprisingly low liquid assets can smooth consumption using low-cost methods to shift the timing of payments for committed forms of expenditure. In particular, many affected workers delayed mortgage payments; and many shifted credit card balance payments. At the same time, affected workers show no increase in spending on credit cards. Average debt only increased due to delays in debt payments. Hence, a substantial part of their reaction was to delay recurring payments that impose little to no penalty. Therefore, households make use of short-term margins of adjustment that are mostly overlooked in the literature on methods of smoothing consumption.

While the data show that many affected workers were able to use unconventional means to smooth consumption, for some with low liquid assets these methods were either inadequate or unavailable. This group, who was carrying some credit card debt already, emerged from the shutdown with still more debt owing to failure to make payments rather than new borrowing. In the case of the government shutdown, however, when the missing paycheck income was offset by post-shutdown payments, this incremental debt was repaid quickly.

More generally, deferral of debt payments—whether mortgage or credit card—is not a viable strategy for managing a persistent negative income shock, or perhaps even a liquidity shock of longer duration. Especially among households with chronically low liquid assets or high debt, using payment deferral to smooth consumption could lead to financial distress in the face of shocks such as job loss or poor health that persistently reduce income.

The remainder of the paper proceeds as follows. Section I describes the paper's relationship to prior studies of household responses to liquidity and income shocks. Section II

provides key facts about the circumstances surrounding the shutdown. Section III describes the data and our research design. It establishes that many workers regularly have low liquid assets prior to receiving their paycheck. Section IV estimates the average response of spending and liquid assets to the shock. Section V considers heterogeneity in these responses across the liquid assets distribution and examines the consequences for credit card debt.

I. Related Literature

Several other papers have turned to account records, like those studied here, to augment survey research on liquid assets, and liquidity more generally. Baker (2018) uses account records from an online banking app and instruments for individuals' income changes with news about their employers. These are income shocks (from layoffs or plant closings, e.g.) that should have different implications for spending from the liquidity shock caused by the government shutdown. But, like the present paper, Baker (2018) finds evidence of that liquid assets (more than debt) play an important role in the spending response to a shock. Park (2016) uses transaction records from a financial aggregator to measure household liquid assets and estimate the relationship between liquid assets and mortgage delinquency, especially following the onset of unemployment. Olafsson and Pagel (2018) use data from a third financial aggregator to study the high-frequency dynamics of liquid assets and the spending response to both predictable and unexpected income changes. The present paper is different from these three analyses in its isolation of a liquidity shock (unexpected change) from a change in lifetime resources, and in its focus on the methods by which those with little liquid assets smooth consumption through a liquidity shock.

In its use of account records integrated across checking, saving, and credit card accounts, this paper joins a new and still growing small literature. Gelman, Kariv, Shapiro, Silverman, and Tadelis (2014) use a small subsample of the same data we use in this paper to study the spending response to the arrival of predictable (paycheck and Social Security benefit) income. That paper does not examine shocks either to income or liquidity. Instead, it asks the complementary question of how spending responds to payments—paychecks and Social Security payments—that are predictable and periodic. Zhang (2017) is similar to Gelman et al. (2014), but isolates the special circumstances of the “extra” paychecks that workers paid bi-weekly receive when three pay-periods occur in the same month. These, too, are not shocks to income or liquidity, but instead predictable and periodic payments. Kuchler (2014) studies integrated account records from an online financial management service that elicits from its customers plans for paying down credit card debt. Kuchler uses those plans, along with the spending responses to income changes, to evaluate a model of present-biased time preferences.

Baker and Yannelis (2017) is the closest study to ours. That paper also analyzes the response to the liquidity shock caused by the 2013 US government shutdown using data derived from account records linked to an online banking app. Baker and Yannelis (2017) emphasizes the heterogeneous spending response by category of spending, and also the differential response to the shutdown for workers who were likely furloughed in addition to missing a part of a paycheck. That paper thus gives special attention to time allocation and home production response to the shock. Baker and Yannelis (2017) measure savings before the shutdown by taking the difference between income and spending in the 9 months prior to the event and find that the spending response is muted for those who, by this measure, have done more saving. Our paper focuses on the precarious liquid asset position households find themselves in near the end

of the paycheck cycle and the different channels, including the drawdown of liquid assets, the use of credit card debt, and the delay of bill payments, through which households smoothed their consumption in response to the liquidity shock.

More generally, the large literature that studies the spending response to changes in income has mostly relied on self-reports of survey respondents to provide information either about the income change or about the response of spending, savings and debt. See, for example, Souleles (1999), Browning and Crossley (2001), Shapiro and Slemrod (2003), Johnson et al. (2006), and Zhang (2014). This paper is different from these studies for two main reasons. First, the integrated account data we use provide an accurate, high-resolution, and high frequency picture of liquid assets before the shock, and both the spending and net saving responses to the shock. Second, as noted above, the shutdown is a distinctive shock. The shock is large, negative, proportional to income, and immediately reversed. These features stand in contrast to shocks arising from government stimulus payments, which are positive and often weakly related to income. In addition to providing a shock to liquidity, such shocks also provide incremental resources for consumption currently or in the future. See, for example, Shapiro and Slemrod (1995, 2003), Johnson et al. (2006), Parker et al. (2013), Agarwal et al. (2007), Bertrand and Morse (2009), Broda and Parker (2014), Parker (2017), and Agarwal and Qian (2014).

II. The 2013 U.S. Government Shutdown

A. Background

The U.S. government was shut down from October 1 to October 16, 2013 because Congress did not pass legislation to appropriate funds for fiscal year 2014. The shutdown was preceded by a series of legislative battles surrounding the Affordable Care Act (ACA), also known as Obamacare. Key events and their timing are described in Figure 1. While Federal government

shutdowns have historical precedent, it was difficult to anticipate whether this shutdown would occur and how long it would last.³

Opponents of the ACA in the House of Representatives sought to tie FY 2014 appropriations to defunding the ACA. They used the threat of a shutdown as a lever in their negotiations and thus generated considerable uncertainty about whether a shutdown would occur. Just days before the deadline to appropriate funding and avoid a shutdown, there was substantial uncertainty over what would happen. A YouGov/Huffington Post survey conducted on September 28-29, 2013 showed that 44% of U.S. adults thought Congress would reach a deal to avoid a shutdown while 26% thought they would not, and 30% were unsure. A similar survey taken after the shutdown began on October 2-3, 2013 showed substantial uncertainty over its expected duration. Seven percent thought the shutdown would last less than a week, 31% thought one or two weeks, 19% thought three or four weeks, and 10% thought the shutdown would last more than a month. 33% were unsure of how long it would last.⁴ For most federal workers, therefore, the shutdown and its duration were likely difficult to anticipate at the outset. While it was not a complete surprise, it was far from certain to occur.

On the other hand, as we will discuss in the next subsection, the shutdown was effectively resolved contemporaneously with the receipt of the paycheck affected by the shutdown. Hence, there was no reason based on permanent income—or indeed resources at the monthly horizon—to respond to the drop in income.

³ There have been 12 shutdowns since 1980 with an average length of 4 days. The longest previous shutdown lasted for 21 days in 1995-1996. See Mataconis (2011).

⁴ Each survey was based on 1,000 U.S. adults. See YouGov/Huffington Post (2013a, b).

B. Impact on Federal Workers

Our analysis focuses on the consequences of the shutdown for a group of the approximately 2.1 million federal government workers. The funding gap that caused the shutdown meant that most federal workers could not be paid until funding legislation was passed. The 1.3 million workers deemed necessary to protect life and property were required to work. They were not, however, paid during the shutdown for work that they did during the shutdown period. The 800,000 “non-essential” employees were furloughed without pay.⁵ In previous shutdowns, employees were paid retroactively whether furloughed, or not. At the onset of the 2013 shutdown, it was not entirely clear what would happen. On October 5, however, the House passed a bill to provide back pay to all federal employees after the resolution of the shutdown. While not definitive, this legislation was strong reassurance that the precedent of retroactive pay would be respected, as in fact it was when the shutdown concluded. After the October 5 Congressional action, most of the remaining risk to workers was due to the uncertain duration of the shutdown and to potential cost-cutting measures that could be part of a deal on the budget.

Unlike most private sector workers, Federal workers are routinely paid with a lag of about a week, so the October 5 House vote came before reduced paychecks were issued. For most government workers, the relevant pay periods are September 22 - October 5, 2013 and October 6 - October 19, 2013. Because the shutdown started in the latter part of the first relevant pay period, workers did not receive payment for 5 days of the 14-day pay period. For most workers on a Monday to Friday work schedule, this would lead to 4 unpaid days out of 10

⁵ Some federal workers were paid through funds not tied to the legislation in question and were not affected. The Pentagon recalled its approximately 350,000 workers on October 5, reducing the number of furloughed workers to 450,000.

working days, so they would receive 40% less than typical pay. The actual fraction varies with hours and days worked and because of taxes and other payments or deductions. Since the government shutdown ended before the next pay date, workers who received a partial paycheck were reimbursed fully in their next paycheck.

Federal government workers are a distinctive subset of the workforce. According to a Congressional Budget Office report (CBO 2012), however, federal workers represent a wide variety of skills and experiences in more than 700 occupations. Compared to private sector workers, they tend to be older, more educated, and more concentrated in professional occupations. Table 1 below reproduces Summary Table 1 in the CBO report. Overall, total compensation is slightly higher for federal workers. Breaking down the compensation difference by educational attainment shows that federal workers are compensated relatively more at low levels of education while the opposite holds for the higher end of the education distribution. In the next section, we make similar comparisons based on Federal versus non-Federal workers in our data. The analysis must be interpreted, however, with the caution that Federal workers may not have identical behavioral responses as the general population. We return to this issue in the discussion of the results.

III. Data and Design

A. Data

The source of the data analyzed here is a financial aggregation and bill-paying computer and smartphone application that had approximately 1.5 million active users in the U.S. in 2013.⁶

⁶We gratefully acknowledge the partnership with the financial services application that makes this work possible. All data are de-identified prior to being made available to the project researchers. Analysis is carried out on data aggregated and normalized at the individual level. Only aggregated results are reported.

Users can link almost any financial account to the app, including bank accounts, credit card accounts, utility bills, and more. Each day, the application logs into the web portals for these accounts and obtains key elements of the user's financial data including balances, transaction records and descriptions, the price of credit and the fraction of available credit used.

We draw on the entire de-identified population of active users and data derived from their records from late 2012 until October 2014. The analysis is performed on normalized and aggregated user-level data as described in the Appendix. The firm does not collect demographic information directly and instead uses a third party business that gathers both public and private sources of demographics, anonymizes them, and matches them back to the de-identified dataset. Appendix Table A (replicated from Table 1 of Gelman et al. 2014) compares the gender, age, education, and geographic distributions in a subset of the sample to the distributions in the U.S. Census American Community Survey (ACS) that is representative of the U.S. population in 2012. The app's population is not representative of the U.S. population, but it is heterogeneous, including large numbers of users of different ages, education levels, and geographic locations.

We identify paychecks using the transaction description of checking account deposits. Among these paychecks, we identify Federal workers by further details in the transaction description. The appendix describes details of the method for identifying paychecks in general and Federal paychecks in particular. It also discusses the extent to which we are capturing the expected number of Federal workers in the data. The number of federal workers and their distribution across agencies paying them are in line with what one would expect if these workers enroll in the app at roughly the same frequency as the general population.

B. Design: Treatment and Controls

Much of the following analysis uses a difference-in-differences approach to study how Federal workers reacted to the effects of the government shutdown. The treatment group consists of Federal workers whose paychecks we observe changing as a result of the shutdown. The control group consists of workers that have the same biweekly pay schedule as the Federal government who were not subject to the shutdown (see the Appendix for more details). The control group is mainly non-Federal workers, but also includes some Federal workers not subject to the shutdown.⁷ Table 2 shows summary statistics from the app's data for these groups of workers. As in the CBO study cited above, Federal workers in our sample have higher incomes. They also have higher spending, higher liquid asset balances, and higher credit card balances.

We use the control group of workers not subject to the shutdown to account for a number of factors that might affect income and spending during the shutdown: these include aggregate shocks and seasonality in income and spending. Additionally, interactions of pay date, spending, and day of week are important (see Gelman et al. 2014). Requiring the treatment and control to have the same pay dates and pay date schedule (biweekly on the Federal schedule) is a straightforward and important way to control for these substantial, but subtle effects.

There is substantial variability in economic circumstance across households both within and across treatment and controls. We normalize many variables by average daily spending, or where relevant by average account balances, at the household level. This normalization is a

⁷ Employees not subject to the shutdown include military, some civilian Defense Department, Post Office, and other employees paid by appropriations not involved in the shutdown. An alternative strategy would use just the Federal workers not affected by the shutdown as the control. We do not adopt this strategy because we believe, a priori, that it is less suitable because the workers exempted from the shutdown are in non-random agencies and occupations. This selectivity makes them potentially less suitable as a control. For completeness, however, the Appendix includes key results using the unaffected Federal workers as the control. The findings are quite similar though less precise because of smaller sample size.

simple, and given the limited covariates in the data, a practical way to pool households with very different levels of income and spending. In particular, it serves to equalize the differences in income levels between treatment and control seen in Table 2.

Showing a wide span of data before and after the shutdown, Figure 2 provides evidence of the adequacy of the control group and the effectiveness of using average daily spending as the normalization. Figure 2 shows that workers not subject to the shutdown have nearly identical movement in spending except during the weeks surrounding the shutdown. Thus the controls appear effective at capturing aggregate shocks, seasonality, payday interactions, etc. In particular, note the regular, biweekly pattern of fluctuations in spending. It arises largely from the timing of spending following receipt of the bi-weekly paychecks. There are also subtler beginning-of-month effects—also related to timing of spending. In subsequent figures we use a narrower window to highlight the effects of the shutdown.

Gelman et al. (2014) shows that much of the sensitivity of spending to receipt of paycheck, like that seen in Figure 2, arises from choices of households to time recurring payments—such as mortgage payments, rent, or other recurring bills—immediately after receipt of paychecks. Figure 2 makes clear that the control group does a good job of capturing this feature of the data and therefore eliminating ordinary paycheck effects from the analysis. The first vertical line in Figure 2 indicates the week in which workers affected by the shutdown were paid roughly 40% less than their average paycheck. There is a large gap between the treatment and control group during this week. Similarly, the second vertical line indicates the week of the first paycheck after the shutdown. The rebound in spending is discernable for two weeks. The figure thus demonstrates that the control group represents a valid counterfactual for spending that occurred in the absence of the government shutdown.

C. Liquid Assets Before the Shutdown

To understand how affected workers responded to the shutdown, it is useful to examine first how they and others like them managed their liquid assets prior to the shock. Analysis of liquid asset balances before the shutdown shows that, while some workers were well buffered, many were ill-prepared to use liquid assets to smooth even this brief and fully reverting pay shock.

We define liquid assets as the balance on all checking and savings accounts. The measure is based on daily snapshots of account balances. Hence, it measures the stock of liquid assets independent from the transactional data used to measure spending and income. Having such high-frequency data makes it possible to observe distinctive, new evidence on liquidity and how it interacts with shocks. Figure 3 shows median liquid assets over the pay-cycle, by terciles of the distribution of liquid assets. The measure is expressed as a ratio of checking and savings account balances to average daily total spending. The results are for the period prior to the shutdown and aggregate over both treatment and control groups.^{8,9}

While the optimal level of liquid assets is not clear, the figure shows the top third of the liquid asset distribution is well-positioned to handle the pay shock due to the shutdown. The median of this group could maintain more than a month of average spending with their checking and savings account balances, even in the days just before their paycheck arrives.

⁸ The distributions for treatment and control are similar. For example, in the control group the median liquid assets ratio for the first, second, and third, terciles of the distribution is, 2.9, 7.9, and 32.1, respectively. The analogous numbers for the treatment group are 3.3, 8.1, and 32.0.

⁹ Liquid asset balances peak two days after a payday. The balance data are based on funds available, so liquid assets should lag payday according to the banks funds-availability policy. There is at least one-day lag built into the data because the balances are scraped during the day, so will reflect a paycheck posted the previous day. Appendix Figure A4 shows that the two-day delay in the peak of liquid assets is due to funds availability, not delays in posting based on interactions of day of payday and delays in posting of transactions over the week-end. (As discussed in the Appendix, even within the government bi-weekly pay schedule, there is some heterogeneity in day of week of the payday.) Additionally, liquid asset balances are, of course, net of inflows and outflows. Recurring payments made just after the receipt of paycheck will therefore lead daily balances to understate gross asset balances right after the receipt of the paycheck.

The lower two-thirds of the liquid asset distribution has a substantially smaller cushion. Over the entire pay-cycle, the middle tercile has median liquid assets equal to 7.9 days of average spending. Liquid asset balances drop to only 5 days of average spending in the days just before their paycheck arrives. Thus, even in the middle of the liquid asset distribution many would be hard pressed to use their savings to smooth a temporary loss of 4 days' pay. The bottom third of this population is especially ill-prepared. Prior to the shutdown, the median of this group consistently arrives at payday with precisely zero liquid balances. (Balances can be negative owing to overdrafts.) These balance data thus reveal how, even among those with steady employment, large fractions of workers do not have the liquid assets to absorb a large, but brief, shock to their pay.

IV. Responses to the Shutdown

Having established that many workers had low liquid asset balances prior to the shutdown, we now examine how their paycheck income, and various form of spending responded to the shutdown. Our method is to estimate the difference-in-difference, between treatment and control, for various outcomes using the equation

$$y_{i,t} = \sum_{k=1}^T \delta_k \times Week_{i,k} + \sum_{k=1}^T \beta_k \times Week_{i,k} \times Shut_t + \Gamma'X_{i,t} + \varepsilon_{i,t} \quad (1)$$

where y represents the outcome variable (total spending, non-recurring spending, paycheck income, debt, savings, etc.), i indexes households ($i \in \{1, \dots, N\}$), and t indexes time ($t \in \{1, \dots, T\}$). $Week$ is a complete set of indicator variables for each week in the sample, $Shut$ is a binary variable equal to 1 if household i is in the treatment group and 0 otherwise, and X represents controls to absorb the predictable variation arising from bi-weekly pay week

patterns.¹⁰ The β_k coefficients capture the average weekly difference in the outcome variables of the treatment group relative to the control group. Standard errors in all regression analyses are clustered at the household level and adjusted for conditional heteroskedasticity.

A. Paycheck Income and Total Spending

We begin with an examination of how paycheck income, as measured in these data, was affected by the shutdown. External reports indicate that the paycheck income of affected employees should have dropped by 40% on average. The analysis of paycheck income here can thus be viewed, in part, as testing the ability of these data to accurately measure that drop. Once that ability is confirmed, we move to an evaluation of the spending responses.

Recall that we normalize each variable of interest, measured at the household level, by the household's average daily spending computed over the entire sample period. The unit of analysis in our figures is therefore days of average spending. Figure 4 plots the estimated β_k from equation (1) where y is normalized paycheck income. We plot three months before and after the government shutdown to highlight the effect of the event. The first vertical line (dashed-blue) represents the week that the shutdown began and the second vertical line (solid-red) represents the week in which pay dropped due to the shutdown, and the third when pay was restored.

Panel A of Figure 4 shows, as expected, a drop in paycheck income equal to approximately 4 days of average daily total spending during the first paycheck period after the

¹⁰Specifically, X contains dummies for paycheck week, treatment, and their interaction. This specification allows the response of treatment and control to ordinary paychecks to differ. These controls are only necessary in the estimates for paycheck income.

shutdown.¹¹ This drop quickly recovers during the first paycheck period after the shutdown ends, as all workers are reimbursed for their lost pay. Some workers received their reimbursement paychecks earlier than usual, so the recovery is spread across two weeks. The results confirm that the treatment group is indeed subject to the temporary loss and subsequent recovery of paycheck income that was caused by the government shutdown, and that the account data allow an accurate measure of those changes.

Panel B of Figure 4 plots the results on total spending, showing the estimated β_k where y is normalized total spending. On average, total spending drops by about 2 days of spending in the week the reduced paycheck was received. Hence, the drop in spending upon impact is about half the drop in pay. In the inter-paycheck week, spending is about normal. In the second week after the paycheck affected by the shutdown, spending rebounds with the recovery spread mainly over that week and the next one.

To aid interpretation we convert the patterns observed in Figure 4 into an estimate of the marginal propensity to spend from a shock to liquidity. Let τ be the week of the reduced paycheck during the shutdown. The variable $s_{i,\tau-k}$ denotes total spending for individual i in the k weeks surrounding that week. To estimate the marginal propensity to spend, we consider the relationship

$$s_{i,\tau-k} = \alpha_k + \beta_k(\text{paycheck}_{i,\tau} - \text{paycheck}_{i,\tau-2}) + \varepsilon_{i,\tau-k} \quad (2)$$

where $(\text{paycheck}_{i,\tau} - \text{paycheck}_{i,\tau-2})$ is the change in paycheck income. Both s and paycheck are normalized by household-level average daily spending as discussed above. We present estimates for the one and two week anticipation of the drop in pay ($k=1$ and $k=2$), the

¹¹ The biweekly paychecks dropped by 40 percent on average. For the sample, paycheck income is roughly 70 percent of total spending on average because there are other sources of income. So a drop of paycheck income corresponding to 4 days of average daily spending is about what one would expect (4 days \approx 0.4 x 0.7 x 14 days).

contemporaneous ($k = 0$), and one lagged ($k = -1$) marginal propensity to spend. We do not consider further lags because the effect of the lost pay is confounded by the effect of the reimbursed pay beginning at time $\tau + 2$.

There are multiple approaches to estimating equation (2). The explanatory variable is the change in paycheck. We are interested in isolating the effect on spending due to the exogenous drop in pay for workers affected by the shutdown. While in concept this treatment represents a 40 percent drop in pay for the affected employees and 0 for the controls, there are idiosyncratic movements in pay unrelated to the shutdown. First, not all workers affected by the shutdown had exactly a 40 percent drop in pay because of differences in work schedule or overtime during the pay period. Second, there are idiosyncratic movements in pay in the control group. Therefore, to estimate the effect of the shutdown using equation (2) we use an instrumental variables approach where the instrument is a dummy variable $Shut_i$. The IV estimate is numerically equivalent to the difference-in-difference estimator.¹²

Table 3 shows the estimates of the marginal propensity to spend. These estimates confirm that the total spending of government workers reacted strongly to their drop in pay and that this reaction was focused largely during the week that their reduced paycheck arrived. The estimate of the average marginal propensity to spend is 0.58 in this week, with much smaller coefficients in the two weeks just prior. Thus, at the margin, about half of the lost pay was reflected in reduced spending.

¹² Estimating equation (2) by least squares should produce a substantially attenuated estimate relative to the true effect of the shutdown if there is idiosyncratic movement in pay among the control group, some of which results in changes in spending. In addition, if the behavioral response to the shutdown differs across households in ways related to variation in the change in paycheck caused by the shutdown (e.g., because workers with overtime pay might have systematically different marginal propensities to spend), the difference between the OLS and IV estimates would also reflect treatment heterogeneity. This heterogeneity could lead the OLS estimate to be either larger or smaller than the IV estimate, depending on the correlation between the size of the shutdown-induced shock and the marginal propensity to spend. The OLS estimate of the marginal propensity to spend for the week the reduced paycheck arrived is 0.123, with a standard error of 0.004.

B. Spending and Payments by Type

Analyzing different categories of spending offers further insight into the response of these households to the liquidity drop. We separate spending into non-recurring and recurring components. Recurring spending is identified using patterns in both the amount and transaction description of each individual transaction.¹³ It identifies spending that, due to its regularity, is very likely to be a committed form of expenditure (see Grossman and Laroque (1990), Chetty and Szeidl (2007), and Postlewaite, et al. (2008)). Non-recurring spending is total spending minus recurring spending. These measures thus use the amount and timing of spending rather than an a priori categorization based on goods and services. This approach to categorization is made possible by the distinctive features of the data infrastructure.

Figure 5 presents estimates of the β_k from equation (1) where the outcome variable y takes on different spending, payment, or transfer categories. For each graph, the data are normalized by household-level averages for the series being plotted. In the top two panels we can compare the normalized response of recurring and non-recurring spending and see important heterogeneity in the spending response by category. The results on total spending (Figure 4) showed an asymmetry in the spending response before and after the liquidity shock; total spending dropped roughly by 2 days of average spending during the three weeks after the shutdown began and only rose by 1.6 days of average spending during the three weeks after the shutdown ended. The reaction of recurring spending drives much of that asymmetry; it dropped by 2.6 days of average recurring spending and rose only by 0.84 days once the lost paycheck

¹³ We identify recurring spending using two techniques. First, we define a payment as recurring if it takes the same amount at a regular periodicity. This definition captures payments such as rent or mortgages. Second, we also use transaction fields to identify payments that are made to the same payee at regular intervals, but not necessarily in the same amount. This definition captures payments such as phone or utility bills that are recurring, but in different amounts. See appendix for further details. Gelman et al. (2014) uses only the first technique to define recurring payments.

income was recovered. Non-recurring spending exhibits the opposite tendency: it dropped by 1.8 days of average non-recurring spending and rose by 2.0 days. Thus, recurring spending drops more and does not recover as strongly as non-recurring spending.

To better understand this pattern of recurring expenditure and its significance we focus on a particular, and especially important, type of recurring spending—mortgage payments.¹⁴ Panel C of Figure 5 shows that, while the mortgage spending data is noisier than the other categories, there is a significant drop during the shutdown and this decline fully recovers in the weeks when the workers’ missing paycheck income was repaid. In this way, we see that some households manage the shock by putting off mortgage payments until the shutdown ends.¹⁵ Indeed, many of those affected by the shutdown changed from paying their mortgage early in October to later in the month as shown in Figure 6. The irregular pattern of payment week of mortgage reflects interaction of the bi-weekly paycheck schedule with the calendar month. The key finding of this figure is that the deficit in payments of the treatment group in the second week of October is largely offset by the surplus of payments in the last two weeks of October.

Panel D of Figure 5 shows the response of account transfers to the pay shock.¹⁶ During the paycheck week affected by the shutdown, transfers fell and rebounded when the pay was reimbursed two weeks later. This finding implies a margin of adjustment, reducing transfers out

¹⁴ Among recurring expenditures, mortgage spending is identified by the transaction description field of each transaction. For this purpose, we searched for the terms “mortgage,” “mortg,” and “mtg.”

¹⁵ Deferring mortgage and other bill payments may be especially easy when the shock to liquidity is widespread and common knowledge. Indeed, there is some evidence that firms whose clients include many federal government employees advertised “payment options” or late fee waivers for workers affected by the shutdown. Examples include local power companies, credit unions, and car dealerships. For a summary see “Furloughed Workers Worry About Bills During Shutdown,” Washington DC NBC Channel 4, October 3, 2013. It is not clear whether these payment options were enacted in response to the shutdown, or merely marketed as a gesture of goodwill to government workers. Mortgages, for example, routinely have clauses for unemployment forbearance.

¹⁶ These are transactions explicitly labeled as “transfer,” etc. For linked accounts, they should net out (though it is possible that a transfer into and out of linked accounts could show up in different weeks). Hence, these transfers are (largely) to and from accounts (such as money market funds) that are not linked.

of linked accounts, during the affected week. One might have expected the opposite, i.e., an inflow of liquidity from unlinked asset accounts to make up for the shortfall in pay. That kind of buffering is not present on average in these data. The finding on lack of transfers from asset accounts to buffer the temporary pay shock is consistent with the model of Kaplan and Violante (2014) where individuals with ample assets may act as liquidity constrained over horizons where assets are not easily accessed as liquid resources. Distinct from Kaplan and Violante (2014), these data provide evidence that consumption is nonetheless smoothed over a short horizon via the timing of bill payments.

Similar behavior is seen in the management of credit card accounts. Another relatively low-cost way to manage cash holdings is to postpone credit card balance payments. Panel E of Figure 5 shows there was a sharp drop in credit card balance payments during the shutdown, which was reversed once the shutdown ended. For households that pay their bill early, this is an easy and cost-free way to finance their current spending. Even if households are using revolving debt, the cost of putting off payments may be small if they pay off the balance right away after the shutdown ends. We examine credit card balances in greater detail in the next section.

Indeed, as we see in Panel F of Figure 5, there was no average reaction of credit card *spending* to the shutdown. Thus, we find no evidence that affected workers sought to fund more of their expenditure with credit cards. Instead, they floated, temporarily, more of their prior expenditure by postponing payments on credit card balances. Affected households who had ample capacity to borrow in order to smooth spending, by charging extra amounts to credit cards, had other means of smoothing, e.g., liquid checking account balances or the postponement of mortgage payments. On the other hand, those who one might think would use credit cards for smoothing spending because they had little cash on hand did not—either because they were

constrained by credit limits or preferred to avoid additional borrowing.¹⁷ In the next section we will examine the consequences for credit card balances of these postponed balance payments, and later probe the heterogeneous responses of individuals by their level of liquid assets.

This analysis of different categories of spending reveals that workers affected by the shutdown reduce spending more heavily on recurring spending and payments compared to non-recurring spending. It is important to note that this behavior appears to represent, in many cases, a temporal shifting of payments and neither a drop in eventual spending over a longer horizon or a proportionate drop in contemporaneous consumption. These results thus provide evidence of the instruments that households use to smooth temporary shocks to liquidity that has not been documented before. The drop in non-recurring spending indicates, however, that this method of cash management is likely imperfect; spending categories that are more likely to reflect consumption are not entirely smoothed.

This view that consumption smoothing is imperfect finds further support in Figure 7. That figure shows that categories of expenditure that are quite close to consumption, such as a fast food and coffee shops spending index, exhibit a sharp drop during the week starting October 10. Given that a cup of coffee or fast food meal is non-durable, one would not expect these categories to rebound after the shutdown. Interestingly, however, there is significant rebound after the shutdown. Hence, in a sense, a cup of coffee seems not to enter the utility function as an additively separate non-durable.

C. Response of Liquid Assets

For households who have built up a liquid asset buffer, they may draw down on these reserves to help smooth consumption after the liquidity shock. Figure 8 shows the estimated β_k

¹⁷ Credit limits can do little to explain the lack of new credit card spending on average. Just 4% of households are using all of their available credit, and 86% are using less than 90% of their available credit.

from equation (1) where y_{it} is the weekly average liquid balance normalized by its individual-level average (Panel A) or normalized by individual-level average daily total spending (Panel B). Because of the heterogeneity in balances, normalizing by average liquid balances leads to more precise results. Normalizing by total spending is less precise but allows for a more meaningful interpretation because it is in the same units as Figure 4. Consistent with the spending analysis, relative savings for the treatment group rises in anticipation of the temporary drop in paycheck income. There is a steep drop in the average balance the week of the lower paycheck as a result of the shutdown. The drop in liquid assets is, however, substantially attenuated relative to the drop in income because of the drop in payments that is documented in the previous section. The recovery of the lost income causes a large spike in the balances, which is mostly run off during the following weeks. Figure 8B shows that liquid balances fell by around 2 days of average daily total spending. Therefore, on average, users reduced spending by about 2 days and drew down about 2 days of liquidity to fund their consumption when faced with a roughly 4 days drop in income. These need not add up because of transfers from non-linked accounts and because of changes in credit card payments, though they do add up roughly at the aggregate level. In the next section, we explore the heterogeneity in responses as a function of liquid asset positions where specific groups of households do use other margins of adjustment than liquid assets.

V. Decomposing the Response to the Liquidity Shock

To summarize the many dimensions of the response to the liquidity shock caused by the shutdown, we decompose the overall spending and saving reaction into its constituent parts. The budget constraint says that the change in spending in a period must be distributed into change in income and change in assets. Table 5 decomposes the components of this identity into its detailed components. Each term in the table represents a difference-in-difference estimate of the

impact of the shutdown on the component of the budget constraint normalized by average daily spending. The -3.90 for paycheck income represents the fact that paycheck income fell by 3.9 days of average daily spending in the treatment group relative to the control group in the week following the government shutdown. In response to the income drop in the week after the shutdown, individuals could either accumulate less liquid savings, reduce spending, or delay recurring spending such as credit card payments, mortgage payments, and outgoing transfers. Spending (net of credit card and mortgage payments) dropped by 1.73 days of average spending (44% of the total drop in paycheck income) while liquid balances fell by 1.83 days of average spending (47% of the total drop in paycheck income).¹⁸ The rest of the shortfall was absorbed by reducing mortgage spending (0.31 days), credit card balance payments (0.22 days) and outgoing transfers (0.09 days). Summing these components, the delay of recurring spending accounted for 16% of the drop in paycheck income.¹⁹ Not everyone affected by the liquidity shock used each of these margins. For example, not everyone has a mortgage, and not all affected households with a mortgage delayed its payment, so there are many zeros included in the average. As Figures 5 and 9 show, nonetheless, deferred mortgage payments do move substantially. The same reasoning applies to credit card balance payments and transfers.

The decomposition presented here demonstrates the strength of observing a complete and granular picture of the household balance sheet.²⁰

¹⁸ In the results in Table 5 to make the decomposition additive, each component is normalized by the households' average daily total spending. In the figures, normalization is by households' average of the particular component.

¹⁹ Because of aggregation (not everyone has the same average daily spending), the identity does not hold exactly. The average residual of 0.02 is close to zero.

²⁰ Other examples of the decomposition of an income or liquidity shock include Agarwal, Liu, Souleles (2007) and Ganong and Noel (2017). Agarwal, Liu, Souleles use credit card account data to study the response to the 2001 federal income tax rebate. They decompose the response into credit card spending and change in credit card balance, but do not observe liquid assets and have data from one account provider. Ganong and Noel use account data from Chase to study the spending and saving reaction to receiving unemployment insurance benefits. They observe liquid assets, but only observe accounts from Chase.

VI. Liquid Assets and the Heterogeneity in Response to the Liquidity Shock

The preceding results capture average effects of the shutdown. There are important reasons to think, however, that different households will react differently to this liquidity shock, depending on their financial circumstances. Although all may have a desire to smooth their spending in response to a temporary shock, some may not have the means to do so.

In this section we examine the heterogeneity in the response along the critical dimension of liquid assets. For those with substantial liquid balances relative to typical spending, it should be relatively easy to smooth through the shutdown. Section III showed, however, that many households in these data have little liquid assets, especially in the days just before their regular paycheck arrives. For those (barely) living paycheck-to-paycheck, even this brief drop in liquidity may pose significant difficulties.

We investigate the impact of the shutdown for those with varying levels of liquid assets by first further quantifying the buffer of liquid assets that different groups of workers had. Second, we return to each of the spending categories examined above and compare how different segments of the liquid assets distribution responded to the liquidity shock. Last, we study how the precise timing of the shock, relative to credit card due dates, influenced credit card balances coming out of the shutdown.

A. Liquid Assets and Spending

As before, we define the liquid assets ratio for each household as the average daily balance of checking and savings accounts to the household's average daily spending until the government shutdown started on October 1, 2013 and then divide households into three terciles. Table 4 shows characteristics of each tercile. Households in the highest tercile have on average 54 days of daily spending on hand while the lowest tercile only has about 3 days. This indicates

that a drop in liquidity equivalent to 4 days of spending should have significantly greater effects for the lowest tercile compared with the highest tercile.

Figure 9 plots the estimates of β_k s from equation (1), for various forms of spending, by terciles of liquid assets. The results are consistent with liquid assets playing a major role in the lack of smoothing. Households with little buffer of liquid savings are more likely to have problems making large and recurring payments such as rent, mortgage, and credit card balances. In terms of average daily expenditure, spending for these recurring payments drops the most for low liquid asset households. In contrast, the drop in non-recurring spending is similar across all liquid asset groups.²¹ Like those with more liquid assets, however, low liquid asset households refrained from using additional credit card spending to smooth the income drop.

B. Liquid Assets and Credit Debt

The preceding results indicate that the sharp declines in recurring spending (especially mortgages) and credit card balance payments induced by the shutdown were important strategies for those with lower levels of liquid assets. The granularity of the data shows, however, that fine differences in timing are consequential when liquidity levels are so low.

To examine how households manage credit card payments and balances, we carry out the analysis at the level of the individual credit card account, rather than aggregating across accounts as in the previous section. The account-level analysis allows us to examine the role of payment due dates in the response to the shutdown. These due dates may represent significant requirements for liquidity. That they are staggered and unlikely to be systematically related to

²¹ If, however, we restrict attention to the top 10% of the liquid asset distribution, the pattern of non-recurring spending is flatter, and there is no statistically significant decline during the shutdown.

the timing of the shutdown provides another means for identifying behavioral responses that exploits the high resolution of the data infrastructure.

In this analysis, however, attention is restricted to the accounts of “revolvers.” That is, we focus on accounts held by those who, at some point during the study period (including the period of the shutdown), incurred interest charges on at least one of their credit cards, indicating that they carried some revolving credit card debt. This represents 63% of the treatment group and 63% of controls; and 70% of these workers fall in the lower two-thirds of the liquid assets distribution. The complement of the revolver group is the “transactors.” Members of this group routinely pay their entire credit card balance, and have a distinct monthly pattern of balances that reflects their credit card spending over the billing cycle and regular payment of the balance at the end of the cycle. Only 44% of transactors fall in the lower two thirds of the liquid assets distribution. Including transactors would obscure the results for those who carry credit card debt.²²

Figure 10 shows the response of credit card balances, at the account level, to the loss of income due to the shutdown. The estimates again present the difference-in-difference between accounts held by revolvers in the treatment group and those held by revolvers in the control group. These estimates are specified in terms of days since the account’s August 2013 statement date instead of calendar time in order to show the effect of statement due dates. In Figure 10, Days 0 through 30 on the horizontal axis correspond to payment due dates in late August or in September 2013. (Payments are due typically 25 days after the statement date.) The different

²² We investigated those who shifted from being transactors to revolvers at the time of the shutdown. This group was relatively small (17%) and did not yield interesting results. By definition, transactors may become revolvers, but not vice-versa, and various circumstances cause some transitions each week. Estimates of the difference between treatment and control of transitions into revolver status show a very small relative increase among controls during the shutdown.

panels of Figure 10 show alternative cuts of the data that we will explain next. Focus, however, on the first 25 days since the August statement date, i.e., due dates that occur in advance of the shutdown. Regardless of cut on the data, the difference-in-difference between treatment and control is essentially zero.

Panels A and B divide the sample of accounts into two groups based on the credit card statement date and, in particular, whether the statement date places them “at risk” for having to make a payment during the government shutdown. Panel A shows the accounts with statement dates on September 16-30, 2013. Panel B shows accounts that have statement dates on September 1-15. For those in the treatment group, the accounts with September 16-30 due dates (Panel A) are at risk. Based on our analysis of liquid assets over the paycheck cycle (Figure 3) it is likely that the mid-October paycheck that is diminished by the shutdown would have been a primary source of liquidity for making the payment on these accounts that come due during that pay period. Indeed, Panel A reveals this effect. Control and treatment accounts start to diverge about a week to 10 days into the October billing cycle (days 35-37). By the time the November statement arrives (days 58-60), a significant gap emerges; relative to controls, treatment account balances are now significantly above average. They return to average in a month, presumably as affected workers use retroactive pay to make balance payments. Panel B, those who made their payments before the shutdown, shows no such effect (the hump starting at day 30 is prior to the shutdown and is not statistically significant.)²³

The high-resolution analysis made possible with the data infrastructure reveals that, when liquid assets are so low, small differences in timing can matter. Workers whose usual credit card payment date fell before the shutdown adjusted on other margins; their balances did not rise. For

²³ We performed a similar analysis investigating the reaction of late fees and interest payments but did not find any quantitatively meaningful or statistically significant effects.

others, the shutdown hit just as they would have normally made their credit card payment; they deferred credit card payments and their balances were elevated for a billing cycle or two before returning to normal levels.

These findings for credit cards reinforce the findings for mortgage payments found in the previous section and Figure 6. For those who typically made payments on mortgages early in the month, that is, prior to the receiving the paycheck reduced by the shutdown, there is little effect of the shutdown on mortgage payments. For those who make payments in the second half of the month, they can and often did postpone the mortgage payments as a way to respond to the shock to liquidity.

VI. Conclusion

Living paycheck-to-paycheck, which is quite common among U.S. households, leaves these households vulnerable to liquidity and income shocks. The results of this paper reveal how workers use financial assets and debt, sometimes in unconventional ways, to reduce that vulnerability and adjust to shocks when they do occur. The findings indicate that to the extent a large but brief shock to liquidity is an important risk, a lack of liquid assets as a buffer is not necessarily a sign of myopia or unfounded optimism. Rather, the reactions to the 2013 government shutdown studied in this paper indicate that workers can defer debt payments and thus maintain consumption (at low cost) despite limited liquid assets. They may face higher costs to access less liquid assets. Such illiquidity may be optimal even if it leads to short- or medium-run liquidity constraints (see Kaplan and Violante 2014). This paper shows that the majority of households have such liquidity constraints as measured by low liquid assets, yet they

have mechanisms for coping with transitory shocks to income or liquidity so as to mitigate the consequences of such low liquid assets.

This paper provides direct evidence on the importance of deferring debt payments, especially mortgages, as an instrument for consumption smoothing. Mortgages function for many as a primary line of credit. By deferring a mortgage payment, they can continue to consume housing, while waiting for an income loss to be recovered. For changing the timing of mortgage payments within the month due, there is no cost. As discussed above, that is the pattern for the bulk of deferred mortgage payments. Moreover, the cost of paying one month late can also be low. Many mortgages allow a grace period after the official due date, in which not even late charges are incurred, or charge a fee that is 4-6 percent of the late payment. Being late by a month adds only modestly to the total mortgage when interest rates are low, and mortgage service companies cannot report a late payment to credit agencies until it is at least 30 days overdue. Even if there are penalties or costs, late payment of a mortgage is a source of credit that is available without the burden of applying for credit.

Thus, this paper's findings indicate that policies that encourage homeownership and low-interest mortgages may have under-appreciated welfare benefits to those mortgage holders. Our results suggest that expansion of mortgage availability not only finances housing, but has the added effect of making it easier to smooth through shocks to income, even absent a formal line of credit. As in Herkenhoff and Ohanian (2015), who show how skipping mortgage payments can function as a form of unemployment insurance, the results here reveal how the ability to defer mortgage payments can be an important source of consumption insurance in the face of large, temporary income fluctuations.

The timing of credit card balance payments provides another source of managing liquidity to buffer consumption against a temporary decline in income. For those with low levels of liquid assets, deferring or reducing credit card payments is a convenient and relatively low-cost way to address a temporary income shortfall. Among credit card borrowers who had payment due dates during the pay period with the reduced paycheck, we see significant deferral of payments. Their credit card balances rose, and stayed elevated for a billing cycle or two before returning to normal.

The distinctive findings of this paper derive high-frequency data on transactions and balances that provide new and distinctive evidence on consumer behavior. The precision and resolution of these data allow insights into behavior that are obscured by conventional data sources.

References

- Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles (2007) "The Reaction of Consumer Spending and Debt to Tax Rebates: Evidence from Consumer Credit Data," *Journal of Political Economy*, 115(6): 986-1019.
- Agarwal, Sumit, and Wenlan Qian (2014) "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore," *American Economic Review*, 104(12): 4205-30.
- Baker, Scott (2018) "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data," *Journal of Political Economy*, Forthcoming.
- Baker, Scott, and Constantine Yannelis (2017) "Income changes and consumption: Evidence from the 2013 federal government shutdown," *Review of Economic Dynamics* 23: 99-124.
- Bertrand, Marianne, and Adair Morse (2009) "Indebted Households and Tax Rebates," *American Economic Review*, 99(2): 418-23.
- Broda, Christian and Jonathan A. Parker (2014) "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption," *Journal of Monetary Economics*, 68(S): S20-S36.
- Browning, Martin, and Thomas F. Crossley (2001) "Unemployment insurance benefit levels and consumption changes," *Journal of Public Economics*, 80(1): 1-23.

- Caner, Asena, and Edward N. Wolff (2004) "Asset poverty in the United States, 1984–99: Evidence from the Panel Study of Income Dynamics," *Review of Income and Wealth*, 50(4): 493-518.
- Chetty, Raj and Adam Szeidl (2007) "Consumption Commitments and Risk Preferences," *Quarterly Journal of Economics*, 122(2): 831-877.
- Congressional Budget Office (2012) "Comparing the Compensation of Federal and Private-Sector Employees."
- Ganong, Peter and Pascal Noel (2017) "Consumer Spending During Unemployment: Positive and Normative Implications," Unpublished paper, University of Chicago.
- Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis (2014) "Harnessing Naturally Occurring Data to Measure the Response of Spending to Income," *Science*, 345(6193): 212-215.
- Grossman, Sanford J. and Guy Laroque (1990) "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 58(1): 25-51.
- Herkenhoff, Kyle F., and Lee E. Ohanian (2015) "The Impact of Foreclosure Delay on U.S. Employment," NBER Working Paper 21532.
- Johnson, David S., Jonathan A. Parker and Nicholas S. Souleles (2006) "Household Expenditure and the Income Tax Rebates of 2001" *American Economic Review* 96(5) 1589-1610.
- Kaplan, Greg, and Gianluca Violante. (2014) "A Model of the Consumption Response to Fiscal Stimulus Payments" *Econometrica* 82(4) 1199-1239.
- Kuchler, Theresa and Michaela Pagel (2018). "Sticking to Your Plan: The Role of Present Bias for Credit Card Paydown." Unpublished paper. New York University.

- Lusardi, Annamaria, Daniel Schneider, and Peter Tufano (2011) "Financially Fragile Households: Evidence and Implications," *Brookings Papers on Economic Activity*, 42(1): 83-134.
- Mataconis, Doug (2011) "A Brief History of Federal Government Shutdowns" *Outside the Beltway*, 08 April 2011. Web. 10 Feb 2015. <http://www.outsidethebeltway.com/a-brief-history-of-federal-government-shutdowns/>.
- Olafsson, Arna and Michaela Pagel (2018) "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software," *The Review of Financial Studies*, Forthcoming.
- Park, Hoonsuk (2016). "Financial Fragility and Mortgage Delinquency," Unpublished paper, Ohio State University.
- Parker, Jonathan (2017) "Why Don't Households Smooth Consumption? Evidence from a 25 Million Dollar Experiment," *American Economic Journal: Macroeconomics*, 9(4): 153-83.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland (2013) "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103(6): 2530-53.
- Postlewaite, Andrew, Larry Samuelson, and Dan Silverman (2008) "Consumption Commitments and Employment Contracts," *Review of Economic Studies*, 75(2): 559-578.
- Shapiro, Matthew D., and Joel Slemrod (1995) "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding." *American Economic Review*, 85(1): 274-283.
- Shapiro, Matthew D., and Joel Slemrod (2003) "Consumer Response to Tax Rebates," *American Economic Review* 93: 381-396

Souleles, Nicholas S. (1999) "The Response of Household Consumption to Income Tax Refunds." *American Economic Review*, 89(4): 947-958.

YouGov/Huffington Post (2013a) "Do you think President Obama and Republicans in Congress will reach a deal to avoid a government shutdown, or that they will not reach a deal and the government will shut down?" Question 4 of 6, September 28 – September 29, 2013. Retrieved from <https://today.yougov.com/news/2013/09/30/poll-results-shutdown/>.

YouGov/Huffington Post (2013b) "How long do you think the government shutdown will last before President Obama and Congress reach a new budget deal?" Question 5 of 5, October 2 – October 3, 2013. Retrieved from <https://today.yougov.com/news/2013/10/04/poll-results-federal-shutdown-hp/>.

Zhang, Yiwei (2017) , "Consumption Responses to Pay Frequency: Evidence From "Extra" Paychecks", in *NA - Advances in Consumer Research* Volume 45, eds. Ayelet Gneezy, Vladas Griskevicius, and Patti Williams. Duluth, MN: Association for Consumer Research, 170-174.

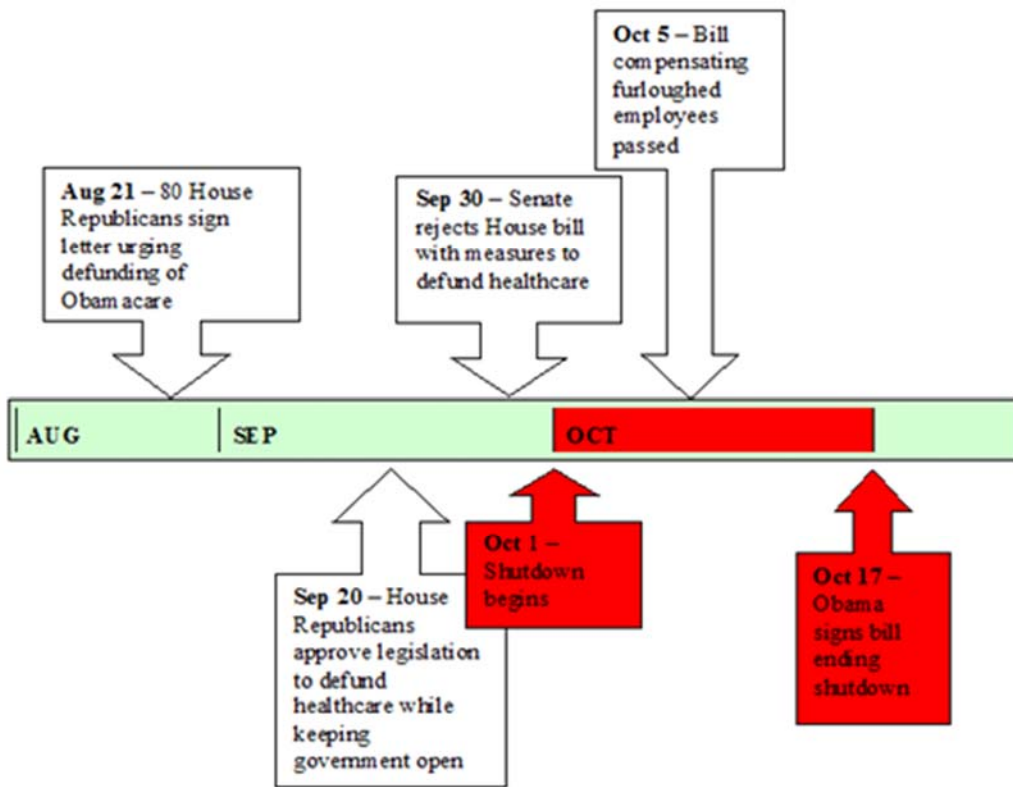


FIGURE 1. GOVERNMENT SHUTDOWN TIMELINE

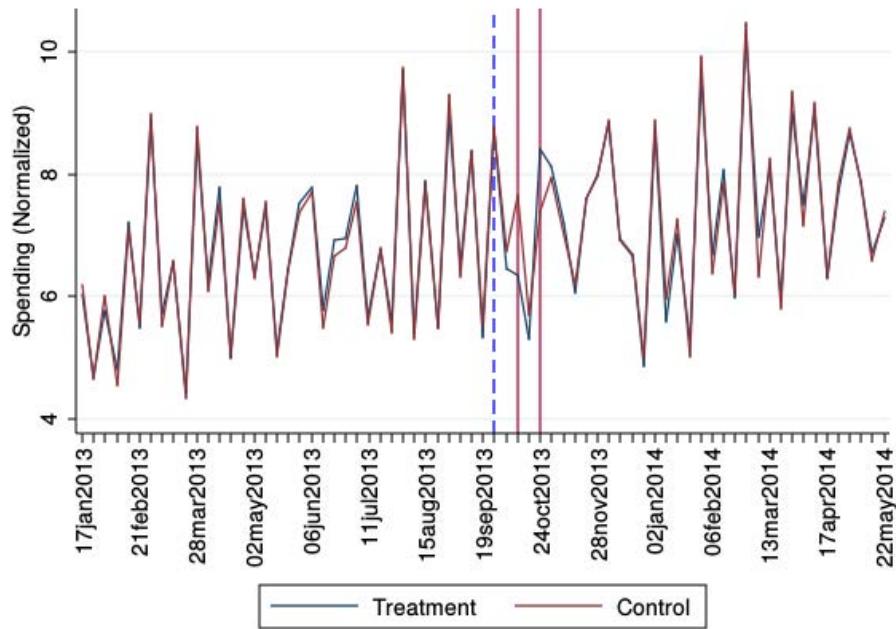


FIGURE 2. TIME SERIES OF SPENDING

Notes: The figure shows average weekly spending (normalized by households' daily spending over the entire sample) for government workers subject to the shutdown (treatment) and other workers on the same biweekly pay schedule (control). The first vertical line is the week in which paychecks were reduced owing to the shutdown. The second vertical line indicates the week when most affected workers received retroactive pay.

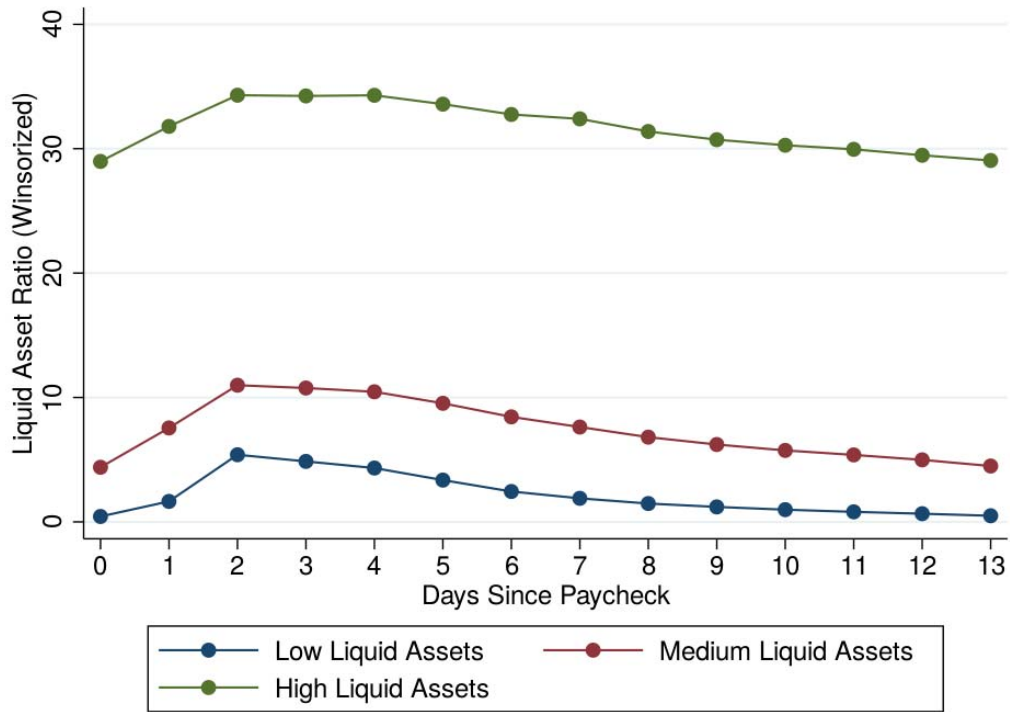
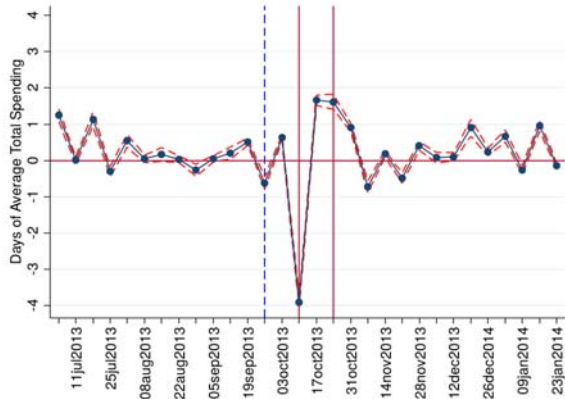


FIGURE 3. PRE-SHUTDOWN MEDIAN LIQUID ASSETS OVER THE PAYCHECK CYCLE

Notes: Liquid assets ratio is defined as checking and savings account balances normalized by average daily spending. The figure shows median liquid assets ratio with terciles by days since receipt of paycheck.

A. Paycheck Income



B. Total Spending

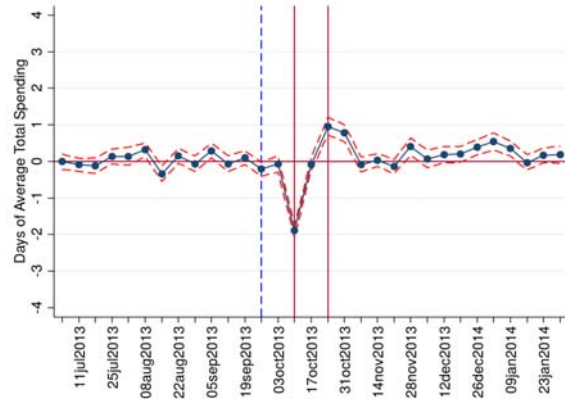


FIGURE 4. ESTIMATED RESPONSE OF NORMALIZED PAYCHECK INCOME AND NORMALIZED TOTAL SPENDING TO GOVERNMENT SHUTDOWN

Notes: Difference-in-difference estimates based on equation (1). Both paycheck income and total spending are normalized by household-level average daily total spending. The paycheck income plot is estimated using additional controls which include paycheck week and treatment group interactions. $N = 3,804$ and $N = 94,680$ for treatment and control group respectively. The estimation period is January 17, 2013 to May 22, 2014. The figures, however, display only the period from July 4, 2013 to January 30, 2014.

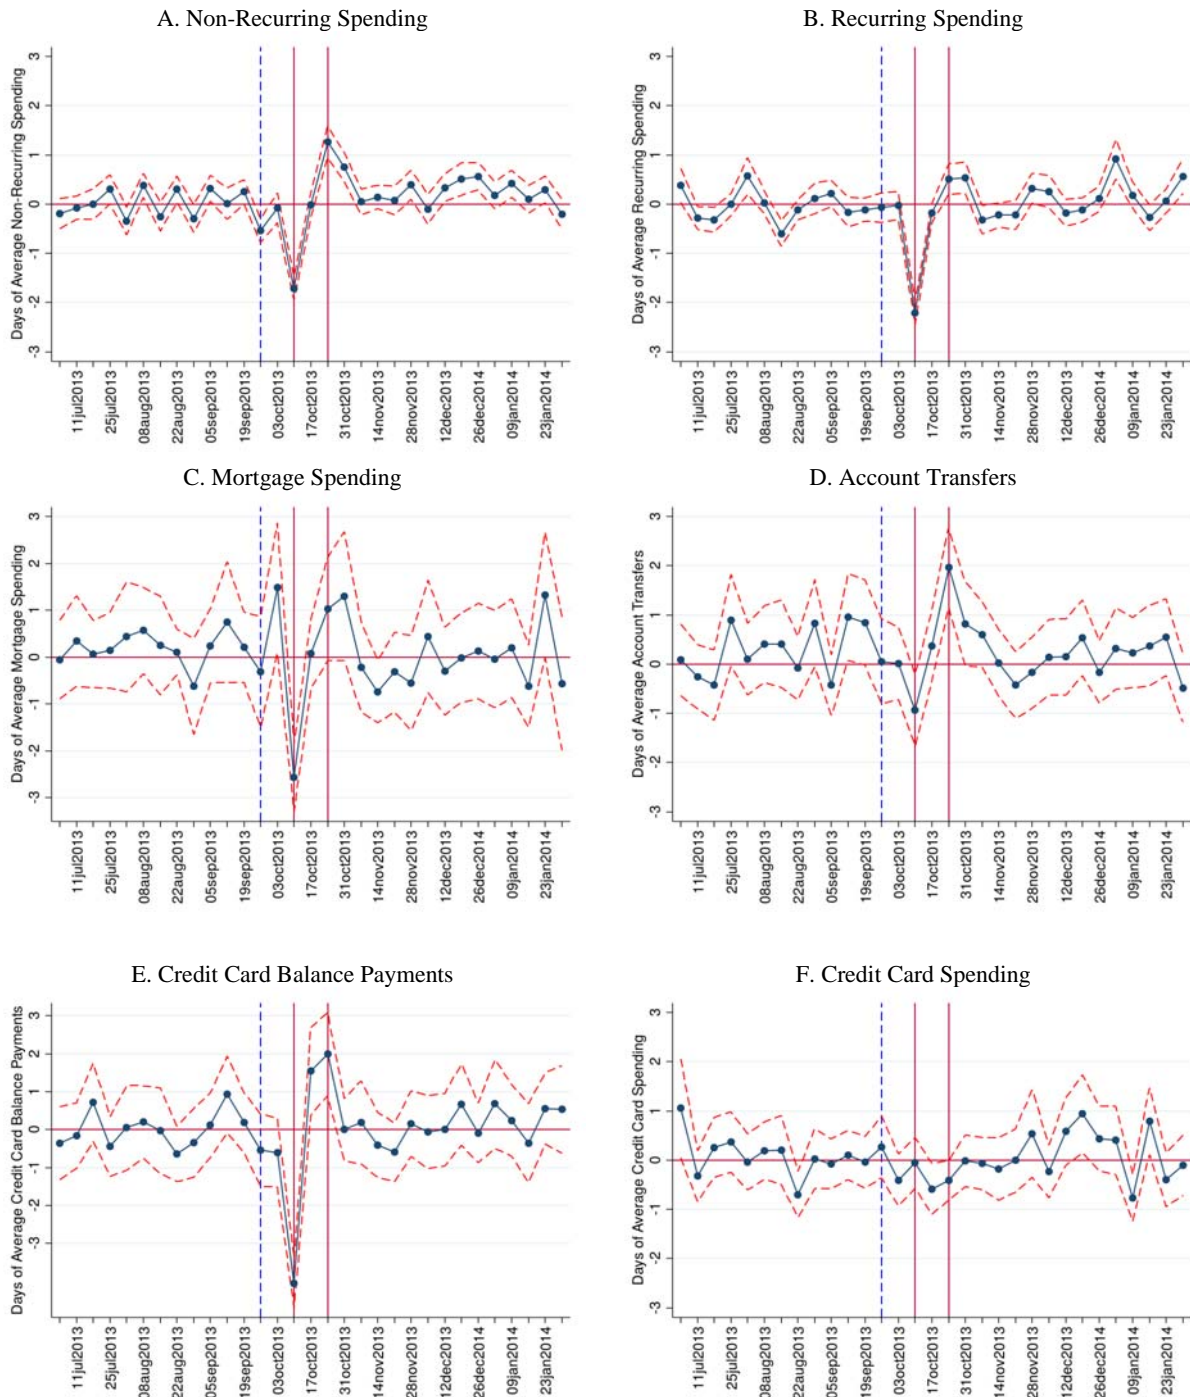
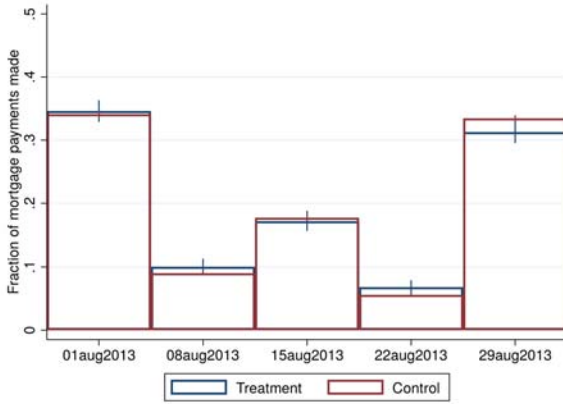


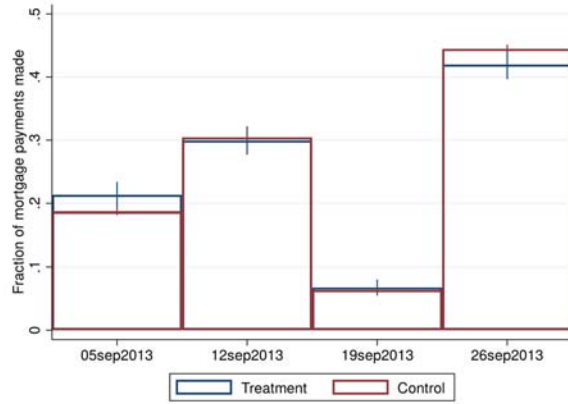
FIGURE 5. ESTIMATED RESPONSE OF SPENDING CATEGORIES TO GOVERNMENT SHUTDOWN

Notes: The spending, payment, or transfer category in each panel is normalized by the household-level daily average for that category. N = 3,804 and N = 94,680 for treatment and control group respectively.

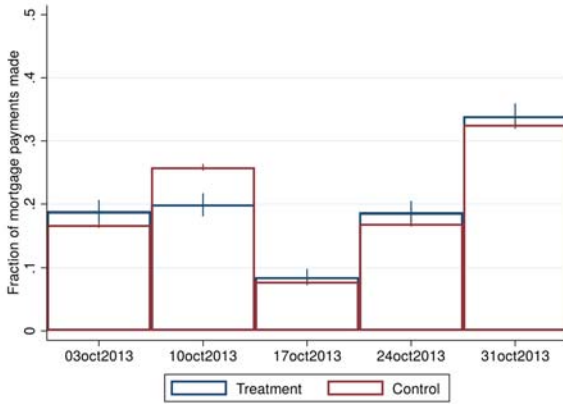
A. August 2013



B. September 2013



C. October 2013



D. November 2013

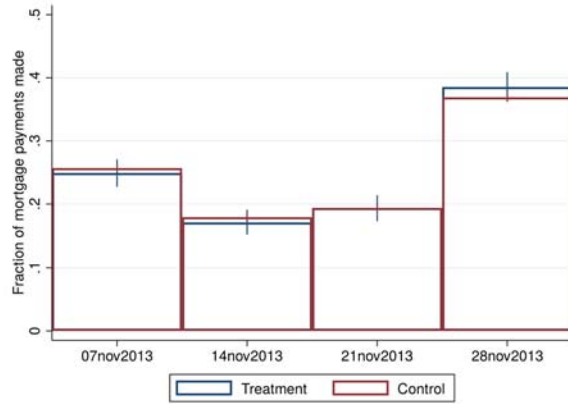


FIGURE 6. DISTRIBUTION OF WEEK MORTGAGE IS PAID

Notes: Blue and red vertical lines in the center of each bar depict the 95% confidence intervals for the weekly fraction for treatment and control groups, respectively.

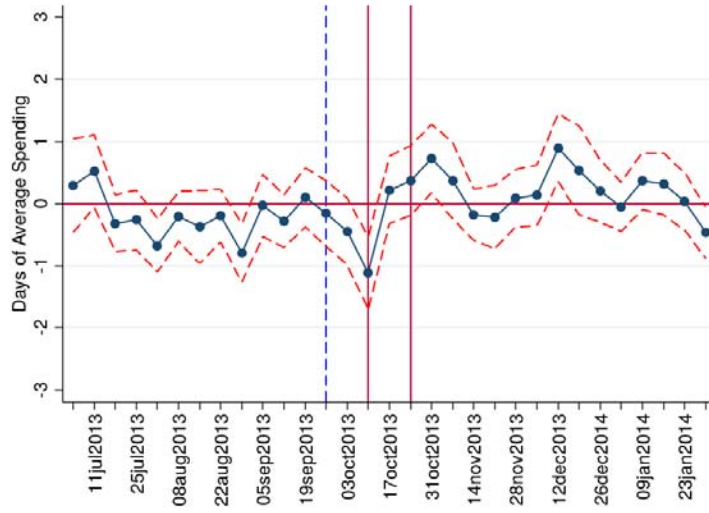


FIGURE 7. ESTIMATED RESPONSE OF COFFEE SHOP AND FAST FOOD SPENDING TO GOVERNMENT SHUTDOWN

Notes: Normalized by household-level average daily coffee shop and fast food spending. N = 3,804 and N= 94,680 for treatment and control group respectively.

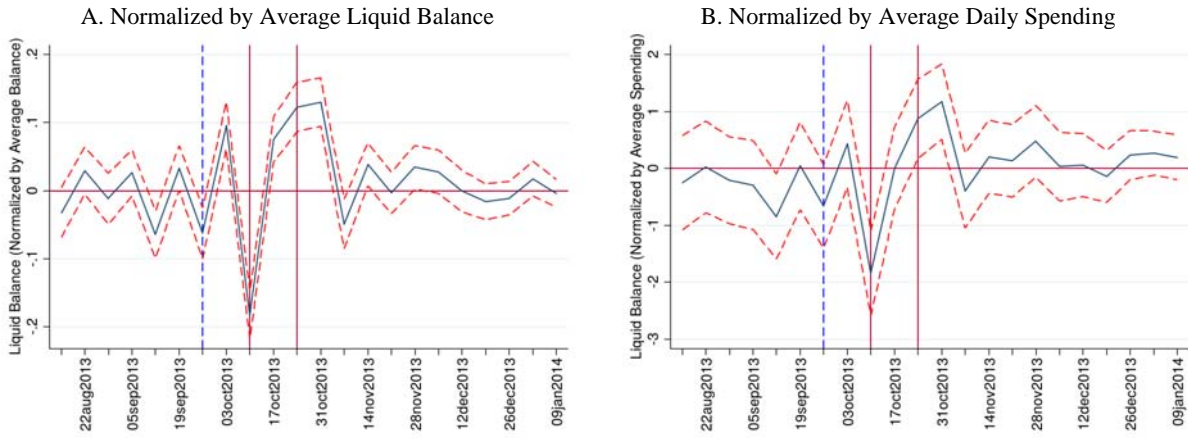


FIGURE 8. ESTIMATED RESPONSE OF LIQUID ASSETS TO GOVERNMENT SHUTDOWN

Notes: Panel A shows end-of-week liquid assets (checking plus saving balances) normalized by household-level average liquidity. Panel B shows end-of-week liquidity normalized by household-level average daily total spending (same normalization as Figure 4). The treatment group includes 3,804 households and the control group includes 94,669 households. Outcome variables are winsorized at the upper and lower 1%.

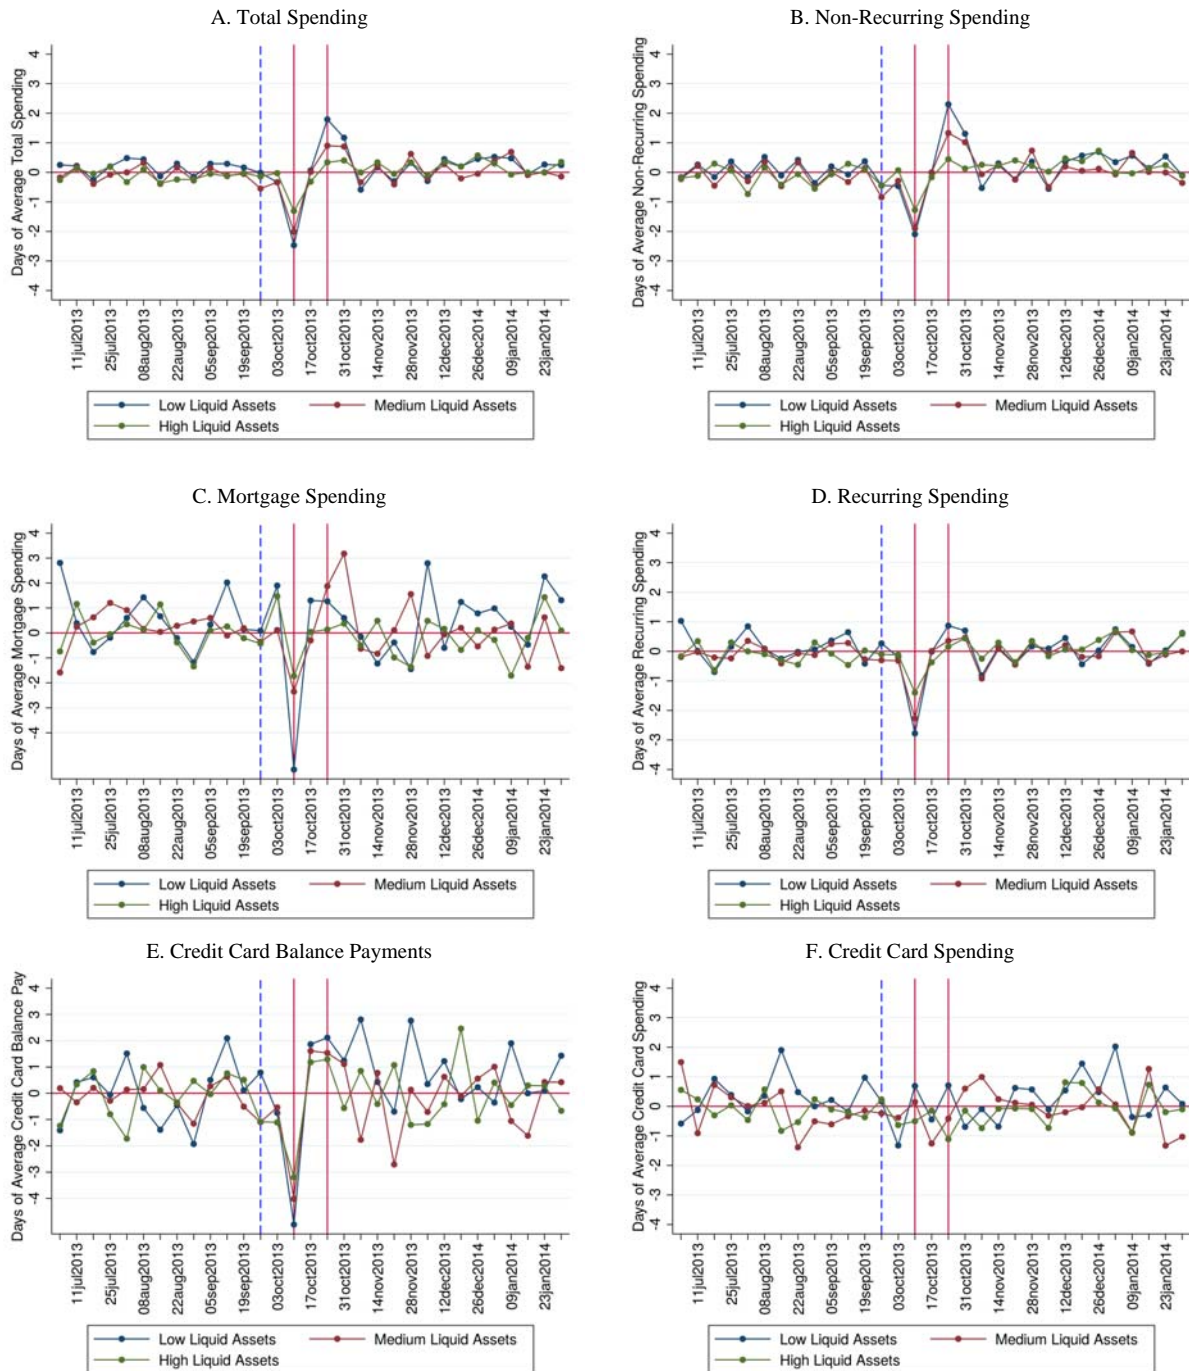


FIGURE 9. ESTIMATED RESPONSE OF SPENDING CATEGORIES TO GOVERNMENT SHUTDOWN BY LIQUID ASSETS TERCILE

Notes: The spending and payment category in each panel is normalized by the household-level daily average for that category. The treatment group includes 3,804 households and the control group includes 94,669 households. Liquid assets are expressed as a ratio of checking and savings account balances to average daily spending. Average liquid assets are 3, 8, and 54 days for the low, medium, and high groups respectively.

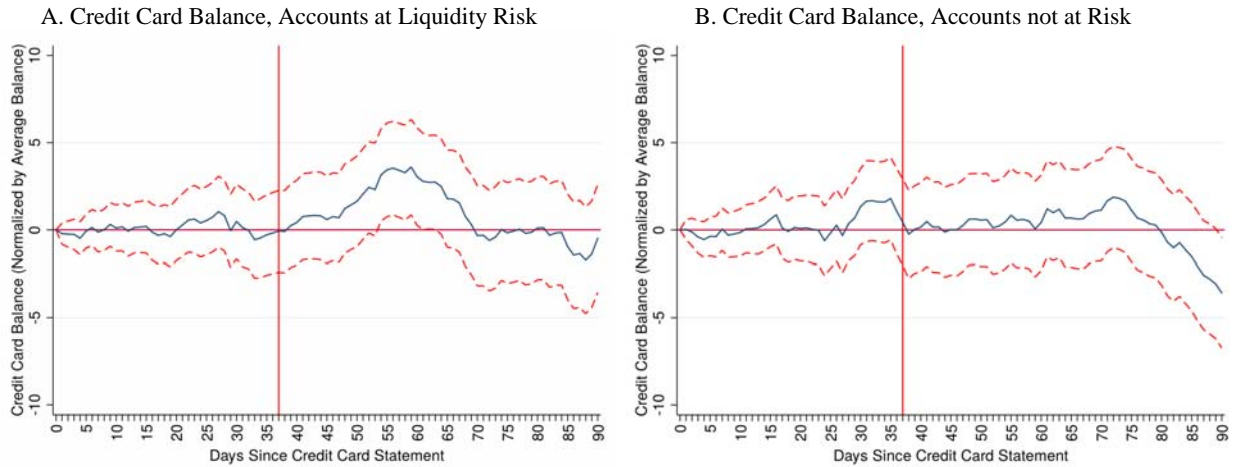


FIGURE 10. ESTIMATED RESPONSE OF CREDIT CARD DEBT TO GOVERNMENT SHUTDOWN

Notes: The sample excludes accounts which never carried revolving credit card debt. Analysis is at the account, not household level. The figure shows daily account balance normalized by the account-level average balance. Standard errors are clustered at the account level. The horizontal axis is the days since the August 2013 credit card statement. Panel A includes accounts with payment due dates during the pay period affected by the shutdown. Panel B includes accounts with due dates before that pay period. The red line represents 10 days into the October pay cycle. This is the point at which we would expect to see the short run effects of the shutdown for those that had their balances come due during the lower paycheck. In panel A, the control group observations represent 22,515 households, 45,712 accounts, and 4,084,450 days and the treatment group observations includes 1,040 households, 2,300 accounts, and 205,746 days. In panel B, the control group observations represent 22,914 households, 45,334 accounts, and 4,030,846 days and the treatment group observations includes 999 households, 2,203 accounts, and 194,972 days. The outcome variables are winsorized at the upper and lower 2%. Data are winsorized at the 2% level rather than the 1% level in other results to control for the greater outliers in the daily balance data.

TABLE 1—AVERAGE HOURLY COMPENSATION OF FEDERAL WORKERS RELATIVE TO THAT OF PRIVATE-SECTOR WORKERS, BY LEVEL OF EDUCATIONAL ATTAINMENT

	Difference in 2010 Dollars per Hour			Percentage Difference		
	Wages	Benefits	Total Compensation	Wages	Benefits	Total Compensation
High School Diploma or Less	\$4	\$7	\$10	21%	72%	36%
Bachelor's Degree	-	\$7	\$8	-	46%	15%
Professional Degree or Doctorate	-\$15	-	-\$16	-23%	-	-18%

Notes: CBO compared average hourly compensation (wages, benefits, and total compensation, converted to 2010 dollars) for federal civilian workers and for private-sector workers with similar observable characteristics that affect compensation—including occupation, years of experience, and size of employer—by the highest level of education that employees achieved. Positive numbers indicate that, on average, wages, benefits, or total compensation for a given education category was higher in the 2005–2010 period for federal workers than for similar private-sector workers. Negative numbers indicate the opposite. Source: Congressional Budget Office based on data from the March Current Population Survey, the Central Personnel Data File, and the National Compensation Survey.

TABLE 2—WORKER CHARACTERISTICS

	N	Average Weekly Income	Standard Deviation of Weekly Income	Average Weekly Spending	Average Normaliz ed Liquid Balance (days)	Average Credit Card Debt
All Federal Workers	6,792	\$1,728	\$1,415	\$1,855	27	\$3,673
Affected by the Shutdown	3,804	\$1,727	\$1,326	\$1,861	26	\$3,785
Not affected by the Shutdown	2,988	\$1,729	\$1,521	\$1,849	29	\$3,529
Non-Federal Workers	91,692	\$1,261	\$1,360	\$1,362	23	\$2,461

Notes: Sample is workers with biweekly paychecks on the same schedule as the government. See text for further details. Normalized Liquid Balance = Average Daily Liquid Balance / Average Daily Spending. The sample is conditional on having accounts that are well linked. Variables are winsorized at the 0.1% upper end. All values are calculated using data from December 2012 to September 2013. Not all households have data for the entire period.

TABLE 3— MARGINAL PROPENSITY TO SPEND ESTIMATES

Lag(k)	2	1	0	-1
$\hat{\beta}_k$	0.0653 (0.0308)	0.0213 (0.0335)	0.5765 (0.0271)	0.0248 (0.0248)
Observations	98,476	98,476	98,476	98,476
SEE	7.107	6.780	6.602	6.263

Notes: Estimates of equation (2). The right-hand side variable is the change in paycheck in the week starting October 10, 2013 (τ) relative to two weeks earlier. The left-hand side variable is weekly spending. Both variables are normalized by the household-level average daily spending calculated over the entire sample. Separate regressions are estimated for lags and leads of the LHS variable.

Estimation is by instrumental variables with a dummy for a worker being affected by the shutdown as the instrument. Standard errors, corrected for conditional heteroskedacity, are in parentheses.

TABLE 4—LIQUID ASSETS RATIO

Liquid Assets Ratio Tercile	Treatment Group		Control Group	
	Mean (Days)	N	Mean (Days)	N
1	2.9	851	2.8	25,105
2	8.4	1131	8.4	24,824
3	54.2	1201	54.7	24,754

Notes: The sample is conditional on having accounts that are well linked. Variables are winsorized at 1%. In the first tercile, the maximum ratio of liquid assets to average daily expenditure is 4.95. In the second tercile, the maximum ratio is 13.99. The highest ratio in the third tercile is 264.

TABLE 5 — DECOMPOSITION OF GOVERNMENT SHUTDOWN SHOCK

Component	Change
Paycheck Income	-3.90 (0.06)
Non-paycheck income	-0.27 (0.12)
Incoming transfers	-0.00 (0.01)
Spending (net credit card and mortgage spending)	-1.73 (0.11)
Credit card spending	-0.01 (0.04)
Mortgage payments	-0.31 (0.04)
Credit card balance payments	-0.22 (0.05)
Outgoing transfers	-0.09 (0.12)
Change in liquid balances	-1.83 (0.38)
Residual	0.02 (0.45)
Mean daily spending	\$214
Median daily spending	\$153
N	98,462

Notes: Each entry in the table is a difference-in-difference estimate based on equation (1) of the average response of the component during the pay period where pay was reduced by the shutdown. All components are normalized by household-level average daily total spending. Standard errors corrected for conditional heteroskedacity and clustered at the household level are in parentheses.