

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

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

First Version: February 26, 2015  
This Revision: December 13, 2016

---

\*We thank Kyle Herkenhoff, Patrick Kline, Dimitriy Masterov, Melvin Stephens Jr., and Jeffrey Smith as well as participants at many seminars 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).

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

ABSTRACT

Using comprehensive account records, this paper examines how individuals adjusted spending and saving in response to a temporary drop in income due to the 2013 U.S. government shutdown. The shutdown cut paychecks by 40% for affected employees, which was recovered within 2 weeks. Because it affected only the timing of payments, the shutdown provides a distinct experiment allowing estimates of the response to a liquidity shock holding income constant. Spending dropped sharply implying a naïve estimate of the marginal propensity to spend of 0.58. This estimate overstates how consumption responded. While many individuals had low liquidity, they used multiple strategies to smooth consumption including delay of recurring payments such as mortgages and credit card balances.

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

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

Matthew D. Shapiro  
Department of Economics  
and Survey Research Center  
University of Michigan  
Ann Arbor, MI 48109  
and NBER  
shapiro@umich.edu

Dan Silverman  
Department of Economics  
Arizona State University  
Tempe AZ 85287  
and NBER  
Daniel.Silverman.1@asu.edu

Steven Tadelis  
Haas School of Business  
University of California, Berkeley  
Berkeley, CA 94720  
and NBER  
stadelis@berkeley.edu

“Unless we get our paychecks this coming Monday we don't have the money to cover our mortgage, car payment, and the rest of the bills that we need to pay.” – ABC 7 news

How consumers respond to changes in income is a central concern of economic analysis and is key for policy evaluation. This paper uses the October 2013 U.S. Federal Government shutdown and a newly developed dataset of financial account records to examine how consumers with different levels of liquidity, income, and spending respond to a short-lived and entirely reversed drop in income. For affected government employees, the shutdown caused a sharp decline in income that was recovered within two weeks. The new dataset, derived from the de-identified account records of more than 1 million individuals living in the United States, provides a granular and integrated view of how individuals in different economic circumstances adjusted spending, saving, and debt in response to the shock. Because the government shutdown affected the timing of payments to workers over a short horizon, but had no effect on the level of income, it provides a distinct experiment for directly estimating the effect of liquidity on spending. This evidence is hence distinct from that from other income shocks, such as tax rebates, which affect liquidity, but also affect income.

The most important findings are, first, that many workers routinely have very low levels of liquidity, especially in the days just before their regular paycheck arrives. Second, and consistent with low liquidity, spending by affected workers declined sharply in response to the drop in income caused by the shutdown—though the drop lasted at most two weeks and was then offset by an equal increase. Third, the granularity and integration of the data reveal the means used by affected workers to smooth consumption—if not spending—most notably their delay of recurring expenses such as mortgage payments and credit card balances. Last, though many

workers found very low-cost ways to weather the shock, some with low liquidity who were already relying on credit card debt accumulated still more credit card debt.

Prior studies that measure the response of individuals to changes in income or liquidity have faced multiple challenges. First, the optimal reaction to an income change depends both on whether the change is anticipated, and on its persistence; but standard data sources make it difficult to identify shocks to expected income and the longevity of these shocks. Second, analysis and policy prescriptions often require a comprehensive view of the heterogeneous responses to an income change. 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. Third, changes in income—whether one-time (such as tax rebates) or recurring (such as paychecks)—provide both resources and liquidity, so reactions to them combine the effects of liquidity and income.

We overcome the challenge of identifying liquidity shocks by using the 2013 U.S. Federal Government shutdown, which produced a significant, temporary, and easily identified negative shock to the liquidity of a large number of employees. We address the challenge of measuring a household’s full range of responses to this shock by exploiting a new dataset derived from the integrated transactions and balance data of more than 1 million individuals in the U.S.<sup>1</sup>

More specifically, the data allow us to distinguish Federal government employees subject to the shutdown. They are distinguished by the transaction description associated with direct

---

<sup>1</sup> The data are captured in the course of business by a mobile banking application. While newly developed, 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.

deposit of their paychecks to their bank accounts. 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 income lost during the shutdown. The transaction data clearly show this pattern for affected workers. Because of routine lags between end of pay period and crediting of paycheck, 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.

This event combined with the distinctive features of the data, which link income, spending, and liquid assets at a high frequency for each individual, provides an unusual opportunity to study the response to a relatively sizeable shock that affected just the timing of income for individuals across the income distribution, without any net effect on their lifetime incomes. See the related literature section below for a discussion of the distinctions of this study.

An important fact revealed by the balance records is that many affected employees 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, liquidity exhibits systematic changes over the pay-cycle. Just before payday, the median level of liquidity 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 liquidity distribution has, on average, a liquid asset balance of zero.

Given such low levels of liquidity, 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 income and then recovered roughly equally over the two pay periods following the end of the shutdown. Econometric analysis reveals a marginal propensity to spend of about 0.58 as a response to the income shock. Most individuals reversed this drop in spending immediately after they receive the paychecks that reimbursed them for their lost income.

On its face, it is troubling that so many affected workers maintained such low liquidity and exhibited such a sharp spending response to an unexpected but brief delay in income. It suggests either that benchmark theories founded on a taste for smoothing consumption are badly specified; that households are inadequately buffered against even very temporary shocks; or that the financial markets that make smoothing possible are functioning poorly.

Further examination of the data reveals, however, that even consumers with low liquidity can smooth consumption better than spending using low-cost methods to shift the timing of payments for committed forms of expenditure. More detailed analysis shows that affected workers delayed mortgage payments, in particular; and many individuals shifted credit card balance payments. At the same time, the data show no increase in spending on credit cards; average debt only increased due to delays in debt payments. Hence, while they responded to the

temporary shock by reducing spending, a large part of their reaction was to delay recurring payments that impose little to no penalty. This shows how consumers make use of short-term margins of adjustment that are mostly overlooked in the literature on methods of smoothing consumption. As such, it also reveals a potentially important welfare benefit of, especially, mortgages with low interest rates. Mortgages can function as a (cheap) line of credit that can help smooth even large, if brief, shocks to income at relatively low cost.

While the data show that many affected workers were able to use perhaps unconventional means to smooth consumption, if not spending, for some with low liquidity 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, when the negative shock to income was quickly offset by post-shutdown payments of lost income, this incremental debt was repaid quickly.

Deferral of debt payments—whether mortgage or credit card—is not, however a viable strategy for a negative income shock of longer duration. Especially among households with chronically low liquidity or high debt, using payment deferral to smooth spending could lead to financial distress in the face of shocks such as job loss that persistently reduce income.

The remainder of the paper proceeds as follows. Section I describes the paper's relationship to prior studies of individual responses to 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 liquidity 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 liquidity distribution and examines the consequences for credit card debt.

## I. Related Literature

The literature concerned with individual responses to income shocks is large. Jappelli and Pistaferri (2010) offer an insightful review. Relative to that large literature, a principal distinction of our paper is the integrated, administrative data that allow accurate observation of liquidity, of the income shock itself, and of several forms of response to the shock. These data thus provide measures of important constraints and outcomes that allow improved inference from the heterogeneous and multi-dimensional reactions to this change in income.

Prior studies of income shocks have mostly relied on self-reports of survey respondents to provide information either about the shock or about the response of spending and savings and debt.<sup>2</sup> Carroll, Crossley, and Sabelhaus (2013), Dillman and House (2013), Einav and Levin (2013), and others, have called for increased use of administrative records to augment survey research. So far, however, the administrative records available for research have typically represented just a slice of economic activity, either providing information about spending at just one retailer, or about the use of a few credit cards, or about just one form of spending.<sup>3</sup> Other approaches blend survey and administrative data. For example, Broda and Parker (2014) and Parker (2015) use consumer-based scanner data to study the response of spending to an income shock; and use surveys for measuring income.<sup>4</sup> This paper is different from these studies of

---

<sup>2</sup> See, for example, Souleles (1999), Browning and Crossley (2001), Shapiro and Slemrod (2003), and Johnson et al. (2006).

<sup>3</sup> See, for example, Gross and Souleles (2002), Agarwal, Liu, and Souleles (2007), and Aguiar and Hurst (2005).

<sup>4</sup> Einav, Leibtag, and Nevo (2010) discuss the challenges that even scanner technologies like Nielsen Consumer Panel (formerly Homescan), and Feenstra and Shapiro (2003) discuss the challenges of using store-based scanner data to measure expenditure and prices.



purely administrative data and from those of blended data sources: The integrated data we use provide an accurate, high-resolution, and high frequency picture of liquidity before the shock, and both the spending and net saving responses to the shock.

By using integrated account records, this paper is part of a new and still small literature that includes Baker (2014), Kuchler (2014), Gelman et al. (2014), and Baker and Yannelis (2015). Baker (2014) uses account records from an online banking app, links them to external data on employers, and instruments for individuals' income changes with news about their employers. Because they are persistent, theory suggests that some of these income shocks (from layoffs or plant closings, e.g.) should have different implications for spending from the one caused by the government shutdown. Nevertheless, like the present paper, Baker (2014) finds evidence of the importance of liquidity (more than debt) for the spending response to an income shock. The present paper is distinct in its study of the methods by which those with very little liquidity smooth consumption through a temporary income shock.

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 (2014) uses those plans, along with the spending responses to income changes, to evaluate a model of present-biased time preferences.

Gelman et al. (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 did not examine other outcomes besides spending. Finally, in a complementary study completed shortly after ours, Baker and Yannelis (2015) use data from the same banking app used in Baker (2014) to describe the response of affected government workers to the 2013 shutdown. Baker and Yannelis (2015) focuses on income and spending, but does not

integrate those outcomes with financial positions. Their analysis confirms that the spending and income response to the government shutdown is identifiable in these data sets. From these initial facts, their paper analyzes time allocation and home production. Our paper focuses on the precarious liquidity position individuals find themselves in near the end of the paycheck cycle and the different channels through which individual's smoothed their consumption.

The shutdown also 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 (2015), 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. While Federal government shutdowns have historical precedent, it was difficult to anticipate whether this shutdown would occur and how long it would last.<sup>5</sup> 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.

---

<sup>5</sup> 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)

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. 7% 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.<sup>6</sup> For most federal employees, therefore, the shutdown and its duration were likely difficult to anticipate at the outset. While it was not a complete surprise, it was a shock to many that the shutdown did indeed occur.

On the other hand, as we will discuss in the next subsection, the shutdown was essentially 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.

### *B. Impact on Federal Employees*

Our analysis focuses on the consequences of the shutdown for a group of the approximately 2.1 million federal government employees. The funding gap that caused the shutdown meant that most federal employees could not be paid until funding legislation was passed. The 1.3 million employees deemed necessary to protect life and property were required

---

<sup>6</sup> Each survey was based on 1,000 U.S. adults. See YouGov/Huffington Post (2013a, b).

to work. They were not, however, paid during the shutdown for work that they did during the shutdown. The 800,000 “non-essential” employees were simply furloughed without pay.<sup>7</sup> In previous shutdowns, employees were paid retroactively (whether or not they were furloughed). Of course, it was not entirely clear what would happen in 2013. 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 income risk to employees 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 employees, 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, employees did not receive payment for 5 days of the 14 day pay period. For most employees on a Monday to Friday work schedule, this would lead to 4 unpaid days out of 10 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 debits. Since the government shutdown ended before the next pay date, employees who only received a partial paycheck were fully reimbursed in their next paycheck.

Federal government employees are a distinctive subset of the workforce. According to a Congressional Budget Office report (CBO 2012), however, federal employees represent a wide

---

<sup>7</sup> Some federal employees were paid through funds not tied to the legislation in question and were not affected. The Pentagon recalled its approximately 350,000 employees on October 5, reducing the number of furloughed employees to 450,000.

variety of skills and experiences in more than 700 occupations. Compared to private sector employees, 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 employees. Breaking down the compensation difference by educational attainment shows that federal employees 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 employees in our data. The analysis must be interpreted, however, with the caution that Federal employees 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.<sup>8</sup> 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 data are de-identified and the analysis is performed on normalized and aggregated user-level data as described in the Appendix. The firm

---

<sup>8</sup>We 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.

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 user 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 employees by further details in the transaction description. The appendix describes details 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 employees in the data. The number of federal employees and their distribution across agencies paying them are in line with what one would expect if these employees 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 employees reacted to the effects of the government shutdown. The treatment group consists of Federal employees whose paycheck income we observe changing as a result of the shutdown. The control group consists of employees 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 employees, but also includes some Federal employees

not subject to the shutdown.<sup>9</sup> Table 2 shows summary statistics from the app's data for these groups of employees. As in the CBO study cited above, Federal employees in our sample have higher incomes. They also have higher spending, higher liquid balances, and higher credit card balances.

We use the control group of employees 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, of course, substantial variability in economic circumstance across individuals both within and across treatment and controls. We normalize many variables by average daily spending, or where relevant by average account balances) at the individual level. This normalization is a simple, and given the limited covariates in the data, a practical way to pool individuals 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.

By showing a wide span of data before and after the shutdown, Figure 2 provides strong evidence of the adequacy of the control group and the effectiveness of using average daily spending as the normalization. Figure 2 shows that the employees not subject to the shutdown have nearly identical movement in spending except during the weeks surrounding the shutdown.

---

<sup>9</sup> Employees not subject to the shutdown include military, some civilian Defense Department, Post Office, and other employees paid by funds 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.

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 reasonable choices of individuals 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 employees 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. Liquidity Before the Shutdown*

To understand how affected employees 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 a brief income shock.

We define liquid assets as the balance on all checking and savings accounts. The measure of liquidity is based on daily snapshots of account balances. Hence, they are measures of 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 liquidity over the pay-cycle, by terciles of the distribution of liquidity. Liquidity 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.<sup>10,11</sup>

While the optimal level of liquidity is not clear, the figure shows the top third of the liquidity distribution is well-positioned to handle the income 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.

The lower two-thirds of the liquidity 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. Liquidity drops to only 5 days of average spending in the days just before their paycheck arrives. Thus, even in the middle of the liquidity distribution many would be hard pressed to use liquid 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,

---

<sup>10</sup> The distributions for treatment and control are similar. For example, in the control group the median liquidity ratio for the first, second, and third, terciles of the liquidity 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.

<sup>11</sup> Liquidity peaks two days after a payday. The balance data are based on funds available, so liquidity 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 liquidity 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, liquidity is, of course, net of inflows and outflows. Recurring payments made just after the receipt of paycheck will therefore lead daily balances to understate gross liquidity right after the receipt of the paycheck.

large fractions of consumers do not have the liquid assets to absorb a large, but brief, shock to income.

#### IV. Responses to the Shutdown

Having established that many (affected) workers had little liquidity prior to the shutdown, we now examine how their 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

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

where  $y$  represents the outcome variable (total spending, non-recurring spending, income, debt, savings, etc.),  $i$  indexes individuals ( $i \in \{1, \dots, N\}$ ), 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 individual  $i$  is in the treatment group and 0 otherwise, and  $X$  represents controls to absorb the predictable variation arising from bi-weekly pay week patterns.<sup>12</sup> The  $\beta_k$  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 individual level and adjusted for conditional heteroskedasticity.

##### A. Paycheck Income and Total Spending

We begin with an examination of how income, as measured in these data, was affected by the shutdown. External reports indicate that the paycheck income of affected employees should

---

<sup>12</sup>Specifically,  $X$  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.

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 individual level, by the individuals' 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  $\beta_k$  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 income equal to approximately 4 days of average daily total spending during the first paycheck period after the shutdown.<sup>13</sup> This drop quickly recovers during the first paycheck period after the shutdown ends, as all users are reimbursed for their lost income. Some users 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 income that was caused by the government shutdown, and that the account data allow an accurate measure of those income changes.

Panel B of Figure 4 plots the results on total spending, showing the estimated  $\beta_k$  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

---

<sup>13</sup> 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 ( $4 \text{ days} \approx 0.4 \times 0.7 \times 14 \text{ days}$ ).

the drop in income. That implies a propensity to spend of about one-half—much higher than most theories would predict for a drop of income that was widely expected to be made up in the relatively near future. 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 ease interpretation we convert the patterns observed in Figure 4 into an estimate of the marginal propensity to spend (which we call the MPC as conventional). Let  $\tau$  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 MPC, we consider the relationship

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

where  $(\text{paycheck}_{i,\tau} - \text{paycheck}_{i,\tau-2})$  is the change in paycheck income. Both  $s$  and  $\text{paycheck}$  are normalized by individual-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 contemporaneous MPC ( $k=0$ ), and one lagged MPC ( $k=-1$ ). 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 employees affected by the shutdown. While in concept this treatment represents a 40 percent drop in income for the affected employees and 0 for the controls, there are idiosyncratic movements in income unrelated to the shutdown. First, not all employees 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.<sup>14</sup>

Table 3 shows the estimates of the MPC. These estimates confirm that the total spending of government employees reacted strongly to their drop in income and that this reaction was focused largely during the week that their reduced paycheck arrived. The estimate of the average MPC 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 income was reflected in reduced spending.

### *B. Spending and Payments by Type*

Analyzing different categories of spending offers further insight into the response of these users to the income 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.<sup>15</sup> 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

---

<sup>14</sup> 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 income among the control group, some of which results in changes in spending. In addition, if the behavioral response to the shutdown differs across individuals in ways related to variation in the change in paycheck caused by the shutdown (e.g., because employees with overtime pay might have systematically different MPCs), 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 MPC. The OLS estimate of the MPC for the week the reduced paycheck arrived is 0.123, with a standard error of 0.004.

<sup>15</sup> 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.

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  $\beta_k$  from equation (1) where the outcome variable  $y$  takes on different spending, payment, or transfer categories. For each graph, the data are normalized by individual-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 income 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 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. 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 employees' missing income was repaid. In this way, we see that some users manage the shock by putting off mortgage payments until the shutdown ends. 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 income shock.<sup>16</sup> 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 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 income shock supports 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.

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 users who pay their bill early, this is an easy and cost-free way to finance their current spending. Even if users 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 employees sought to fund more of their expenditure with credit cards but instead floated, temporarily, more of their prior

---

<sup>16</sup> 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.

expenditure by postponing payments on credit card balances. Affected individuals 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. 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 liquidity.

This analysis of different categories of spending reveals that users 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 individuals use to smooth temporary shocks to income that has not been documented before. The drop in non-recurring spending shows, however, that this method of cash management is not perfect; it does not entirely smooth spending categories that better reflect consumption.

Spending could have fallen in part because employees stayed home and engaged in home production instead of frequenting establishments that they encounter during their work-day. Recall, however, that many employees affected by the shutdown were not, in fact, furloughed. They worked but did not get paid for that work on the regular schedule. In addition, Figure 7 shows that categories of expenditure that are quite close to consumption, such as a fast food and coffee shops spending index, show a sharp drop during the week starting October 10 when



employees were out of work. Given that a cup of coffee or fast food meal is non-durable, one would not expect these categories to rebound after the shutdown. Yet, there is significant rebound after the shutdown. We interpret this spending as resulting from going for coffee, etc., with co-workers after the shutdown, perhaps to trade war stories.<sup>17</sup> Hence, in a sense, a cup of coffee is not entering the utility function as an additively separate non-durable.

### *C. Response of Liquid Assets*

For users who have built up a liquid asset buffer, they may draw down on these reserves to help smooth income shocks. Figure 8 shows the estimated  $\beta_k$  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 liquidity 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 day 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

---

<sup>17</sup> Interestingly, the rebound is highest for the most liquid individuals (figure not reported) who are also higher income. This finding supports our notion that the rebound in coffee shop and fast food arises from post-shutdown socializing.

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 individuals do use other margins of adjustment than liquid assets.

## **V. Liquidity and the Heterogeneity in Response to the Income Shock**

The preceding results capture average effects of the shutdown. There are important reasons to think, however, that different employees will react differently to this income 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 liquidity. 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 workers in these data have little liquidity, especially in the days just before their regular paycheck arrives. For those (barely) living paycheck-to-paycheck, even this brief drop in income may pose significant difficulties.

We investigate the impact of the shutdown for those with varying levels of liquidity 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 liquidity distribution responded to the income 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. Liquidity and Spending*

As before, we define the liquidity ratio for each user as the average daily balance of checking and savings accounts to the user's average daily spending until the government shutdown started on October 1, 2013 and then divide users into three terciles. Table 4 shows characteristics of each tercile. Users 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 income 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  $\beta_k$ s from equation (1), for various forms of spending, by terciles of liquidity. The results are consistent with liquidity playing a major role in the lack of smoothing. Users 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 liquidity users. In contrast, the drop in non-recurring spending is similar across all liquidity groups. Like those with more liquidity, however, low liquidity users refrained from using additional credit card spending to smooth the income drop.

### *B. Liquidity 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 particularly important strategies for those with lower levels of liquidity. The granularity of the data shows, however, that fine differences in timing are consequential when liquidity levels are so low.

To examine how individuals 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 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.” We focus, that is, 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 liquidity 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 liquidity distribution. Including transactors would obscure the results for those who carry credit card debt.<sup>18</sup>

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

---

<sup>18</sup> We investigated those who shifted from being transactors to revolvers at the time of the shutdown. This group was so small (17%) that it did not yield interesting results. Given that transactors tend to have high liquidity (15.9 median ratio vs 7.7 for revolvers), the lack of such transitions is not surprising.

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 liquidity 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.)

The high-resolution analysis made possible with the data infrastructure reveals that, when liquidity is 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 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 lets workers consume at higher levels, but would seem to leave them quite vulnerable to income shocks. The results of this paper reveal how workers use financial assets and markets, 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 income is a primary 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, yet they have mechanisms for coping with shocks to income or liquidity so as to mitigate the consequences of such illiquidity.

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 many mortgage service companies will not 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 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 expansion of mortgage availability not only finances housing, but has the added effect of making it easier to smooth through shocks to income. 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 shocks.

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 (2014) “Debt and the Consumption Response to Household Income Shocks,” manuscript, Kellogg School of Management.
- Baker, Scott, and Constantine Yannelis (2015) “Income Changes and Consumption: Evidence from the 2013 Federal Government Shutdown.” SSRN Working paper, March 2015.
- 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-S3.
- 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.”
- 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.
- Hall, Robert E., and Frederick S. Mishkin (1982) “The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households.” *Econometrica*, 50(2): 461-481.
- Herkenhoff, Kyle F., and Lee E. Ohanian (2015) “The Impact of Foreclosure Delay on U.S. Employment,” NBER Working Paper 21532.
- Jappelli, Tullio and Luigi Pistaferri (2010) “The Consumption Response to Income Changes,” *Annual Review of Economics*, 2: 479-506.
- 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 (2014) “Sticking to Your Plan: Hyperbolic Discounting and Credit Card Debt Paydown” manuscript, Stern School of Business.
- 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/>.
- Parker, Jonathan (2015). “Why Don’t Households Smooth Consumption? Evidence from a 25 Million Dollar Experiment” unpublished paper, MIT.
- 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/>.

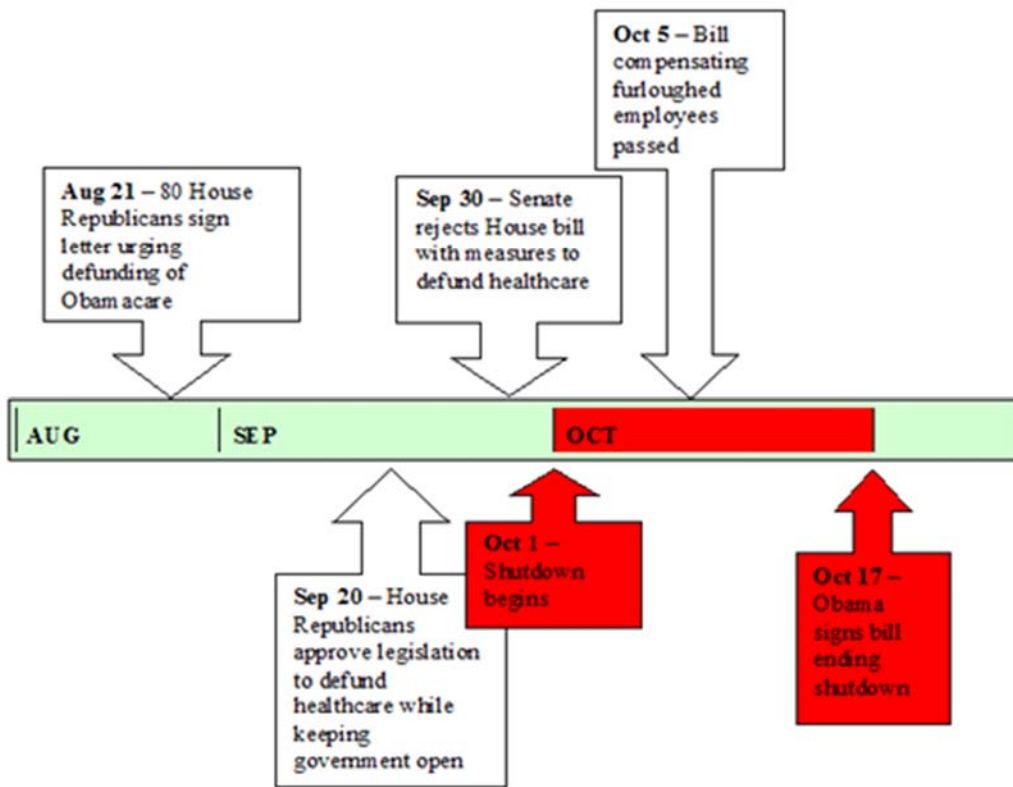


FIGURE 1. GOVERNMENT SHUTDOWN TIMELINE




FIGURE 2. TIME SERIES OF SPENDING

*Notes:* The figure shows average weekly spending (normalized by individuals' daily spending over the entire sample) for government employees subject to the shutdown (treatment) and other employees 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 where government most affected employees received retroactive pay.





FIGURE 3. PRE-SHUTDOWN MEDIAN LIQUIDITY OVER THE PAYCHECK CYCLE

*Notes:* Liquidity ratio is defined as checking and savings account balances normalized by average daily spending. The figure shows median liquidity 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 individual-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 individual-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




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

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




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

Notes: The spending and payment category in each panel is normalized by the individual-level daily average for that category. The treatment group includes 3,804 individuals and the control group includes 94,669 individuals. Liquidity is expressed as a ratio of checking and savings account balances to average daily spending. Average liquidity is 3, 8, and 54 days for the low, medium, and high liquidity 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 individual 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. In panel A, the control group observations represent 22,515 individuals, 45,712 accounts, and 4,084,450 days and the treatment group observations represent 1,040 individuals, 2,300 accounts, and 205,746 days. In panel B, the control group observations represent 22,914 individuals, 45,334 accounts, and 4,030,846 days and the treatment group observations represent 999 individuals, 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 EMPLOYEES RELATIVE TO THAT OF PRIVATE-SECTOR EMPLOYEES, 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 employees and for private-sector employees 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 employees than for similar private-sector employees. 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—EMPLOYEE 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 Employees	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 Employees	91,692	\$1,261	\$1,360	\$1,362	23	\$2,461

*Notes:* Sample is employees 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 users have data for the entire period.

TABLE 3— MPC ESTIMATES

Lag( $k$ )	2	1	0	-1
$\hat{\beta}_{MPC}$	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 ( $\tau$ ) relative to two weeks earlier. The left-hand side variable is weekly spending. Both variables are normalized by the individual-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 an individual being affected by the shutdown as the instrument. Standard errors, corrected for conditional heteroskedacity, are in parentheses.



TABLE 4—LIQUIDITY RATIO

Liquidity 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%.