

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

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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

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

“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 exploits the October 2013 U.S. Federal Government shutdown to examine how consumers respond to a short-lived and entirely reversed drop in income. The three most important findings are that first, despite the fact that it was brief and had no effect on lifetime earnings, consumers responded aggressively to the drop in income. Second, the granularity and integration of the data reveal the multitude of instruments that consumers use to adjust to shocks, most notably their delay of recurring expenses such as mortgage payments and credit card balances. The importance of these smoothing instruments highlights often overlooked consequences of interest rate policies for individual welfare. Last, even though the shock was brief and left lifetime earnings unchanged, it appears to have caused some individuals with less liquidity to accumulate still more expensive credit debt that lasted for many months after the shutdown ended.

Efforts to measure the response of individuals to changes in income have faced two 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; or they may measure many things well, but for too few individuals to perform disaggregated analysis.

To address the challenge of identifying income shocks and their persistence, this paper exploits the 2013 U.S. Federal Government shutdown, which produced a significant, unanticipated, temporary, and easily identified negative shock to the incomes of a large number of employees. To address the challenge of measuring a household's full range of responses to this shock, the paper uses a new dataset derived from the de-identified transactions and account data, aggregated and normalized at the individual level, of more than 1 million individuals living in the United States. The data are captured in the course of business by Mint Bills (formerly Check), a mobile banking application. While newly developed, this dataset has already proved useful for studying the high-frequency response of spending to anticipated income by levels of spending, income, and liquidity (Gelman et al. 2014).¹

The data allow us to identify Federal government employees subject to the shutdown by the description of the direct deposit of their paychecks to their bank accounts. Having identified those subject to the income shock, we can examine their responses in terms of spending and other variables before, during, and after the government shutdown. The paycheck of a typical affected employee was 40% below normal during the shutdown because the government was shut down for the last four days of the previous ten-day pay period. By the next pay period, however, government operations resumed and employees were reimbursed fully for the income lost during the shutdown. The Mint Bills aggregate transaction data clearly show this pattern for affected employees. This event combined with the distinctive features of the data, which link income and spending at a high frequency for each individual, provides an unusual opportunity to study the response to a relatively sizeable shock that affected the timing of income for

¹ Other studies that utilize similar types of account data include Baker (2014) and Kuchler (2014)

individuals across the income distribution, but had no net effect on their lifetime incomes. See the related literature section below for a discussion of the distinctions of our study.

The Mint Bills aggregate transaction records reveal a sharp drop in total spending during the week of missing paycheck income. Weekly spending drops by roughly half the decline in income and then recovers roughly equally over the two pay periods following the end of the shutdown. On its face, it is troubling that this brief shock to income, which had no effect on lifetime earnings, so importantly influenced the spending of affected employees during and after the shutdown. It suggests either that economic theories, which emphasize a taste for smoothing consumption, are wrong; that households are very inadequately buffered against even very temporary shocks; or that the financial markets that make smoothing possible are functioning poorly.

The comprehensive nature of the Mint Bills data is able to shed light on the way in which consumers depart from what simple theoretical benchmarks predict. First, we do identify a large drop in spending that coincides with the shutdown. Our econometric analysis reveals a marginal propensity to spend of about 0.57 as a response to the income shock. Most individuals reversed this drop in spending immediately after they receive the paychecks that reimburse them for their lost income.

Second, many affected employees used low-cost methods to shift the timing of payments for committed forms of expenditure. Mortgage payments, in particular, were shifted later; and many individuals delayed credit card balance payments. Hence, despite responding to the temporary shock by reducing expenditures, which seems to be a precautionary response that was quickly reversed for most consumers in our sample, a large part of their reaction was to push forward recurring payments that impose little to no penalty. This shows how consumers make

use of short-term instruments that are mostly overlooked in the literature on methods of smoothing consumption in response to income shocks. As such, it reveals an important welfare benefit of, especially, mortgages with low interest rates. Mortgages function as a (cheap) line of credit that can help smooth even large shocks to income at relatively low cost.

Last, for some more liquidity-constrained consumers, the temporary loss of income during the shutdown had lasting consequences. This is surprising because as the data show, the affected employees only had four days' worth of income delayed by a few weeks. These consumers, who were already carrying revolving debt on their credit cards, apparently lacked sufficient capacity to adjust using low-cost financial instruments. As a consequence, they turned to expensive credit to fund their spending during those two weeks, and appear to hold that excess debt at least nine months later. On average, individuals in the lower two thirds of the liquidity distribution, who relied disproportionately on credit card debt to smooth the shock, held about 2 more days of average expenditure in revolving debt six months following the shutdown. This accumulation of debt suggests that for especially constrained consumers, even small and income-neutral shocks can have long-lived consequences.

I. Related Literature

The literature concerned with individual responses to income shocks is large. Jappelli and Pistaferri (2010) offer a recent and insightful review. Relative to that large literature, the most distinctive feature of our paper is the integrated, administrative data that allow direct measurement of the income shock and reveal, in detail, the precise methods by which individuals smooth spending in response. Most studies of income shocks rely on the self-reports of survey respondents to provide critical information either about the shock or about the response of

spending and savings and debt. A newer literature turns to administrative records to measure spending and some debt. See, e.g., Agarwal et al. (2007), Broda and Parker (2014), and Agarwal and Qian (2014). Until recently, however, administrative data have provided only a slice of the economic activity of affected individuals. They have captured only some of the spending and debt. The data we use provide an accurate and high-resolution picture of both the spending and net saving response, at high frequency. They thus offer an unusually precise description of how individuals cope with a large, temporary drop in income. In a recent policy brief, Baker and Yannelis (2014) use data from a related mobile banking app to describe the response of affected government workers to the 2013 shutdown. That policy brief focuses on income and spending, but does not integrate those outcomes with other financial positions.

In the language of Japelli and Pistaferri (2010), our paper adopts a “quasi-experimental” approach to identify an income shock. That is, it uses an event observable to the econometrician, in this case the 2013 government shutdown, to isolate and measure income shocks. In this way, it is distinguished from an important literature, dating back at least to Hall and Mishkin (1982), which assumes a statistical model for individual income processes and uses that model to decompose income changes into transitory and persistent shocks. Relative to that literature, the quasi-experimental approach has the advantage of avoiding the need for specific assumptions about the income process. A disadvantage is that the shocks it can isolate typically involve somewhat uncommon circumstances and, thus, the responses to those shocks may not generalize to other, more common sources of income changes.

As a quasi-experimental study of a temporary shock, the paper is closest to the many studies of the responses to government stimulus efforts. See, for example, Shapiro and Slemrod (1995), Parker et al. (2013) Agarwal et al. (2007), Bertrand and Morse (2009), Broda and Parker

(2014), and Agarwal and Qian (2014). Different from all of these studies, the government shutdown represents a negative shock to income that has no effect on lifetime earnings. It thus provides an uncommon opportunity to study the response to a brief loss of income that is later recovered. Also important, the government shutdown caused a 40% drop in anticipated paycheck income for individuals across a wide range of the income distribution. Thus, unlike the stimulus payments, the shutdown represented a sizeable shock even to high-income households.

As a downward shock to income that importantly affects even high-income households, the government shutdown is related to unemployment events. See, for example, Dynarski and Sheffrin (1987), Gruber (1997), and Stephens (2001). Interpretation of the response to unemployment events is, however, more challenging because unemployment has wealth effects, its duration is importantly uncertain, it is often partially insured, and it involves a large increase in leisure time. As explained in more detail below, the government shutdown of 2013 limits many of these confounding effects, in part because furloughs were rare, and therefore offers evidence on the response to a shock in circumstances where economic theory offers an unusually clear prediction.

II. The 2013 U.S. Government Shutdown

A. Background

The U.S. government 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.² The shutdown was preceded by a series of legislative battles surrounding

² 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)

the Affordable Care Act (ACA), also known as Obamacare. Key events and their timing are described in Figure 1.

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.³ For most federal employees, therefore, the shutdown and its duration were likely difficult to anticipate. Therefore, while it was not a complete surprise, it was a shock to many that the shutdown did indeed occur.

B. Impact on Federal Employees

Our analysis focuses on the consequences of the shutdown for part 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 to work. They were not, however, paid during the shutdown for work that they did during the shutdown. The 800,000

³ Each survey was based on 1,000 U.S. adults. See YouGov/Huffington Post (2013a, b).

“non-essential” employees were simply furloughed without pay.⁴ Though employees in previous shutdowns had been paid retroactively (whether or not they were furloughed), 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. This vote suggested that total income would not change, but all federal employees would still face a temporary reduction in paycheck income if the shutdown continued beyond the next pay date. 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.

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 variety of skills and experiences in over 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

⁴ 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.

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 might 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

Mint Bills offers a financial aggregation and bill-paying computer and smartphone application that had approximately 1.5 million active users in the U.S. in 2013.⁵ 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 central 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 Mint Bills 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. Mint Bills does not collect demographic information directly and instead uses a third party business that gathers both public and private sources of demographic data, anonymizes them, and matches them back to the de-identified dataset. Table 1 from Gelman et al. (2014) compares the gender,

⁵ Mint Bills was formerly known as Check. Check was acquired by Intuit in May 2014 and was rebranded as Mint Bills in December 2014. 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.

age, education, and geographic distributions in the Mint Bills sample that matched with an email address to the distributions in the U.S. Census American Community Survey (ACS), representative of the U.S. population in 2012. The Mint Bills population is not perfectly representative of the US population, but it is heterogeneous, including large numbers of users of different ages, education levels, and geographic location.

We are able to identify paycheck types using the transaction description of checking account deposits. Among these paychecks, we can 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. In brief, 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 Mint Bills at roughly the same frequency as the general population.

B. Design: Treatment and Controls

We use a difference-in-differences approach to analyze 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. The control group is mainly non-Federal employees, but also includes some Federal employees not subject to the shutdown.⁶ Table 2 shows summary statistics from the Mint Bills data for these groups of employees. As in the CBO study cited above, Federal

⁶ 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.

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 included aggregate shocks and seasonality in income and spending. Additionally, interactions of pay date, spending, and day of week are quite 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 way to control for these substantial, but subtle effects.

We normalize many variables of interest by average daily spending at the individual level. This normalization is a simple way to pool individuals with very different levels of income and spending. It also removes the differences in income levels between treatment and control seen in Table 2. Figure 2 shows the time series of weekly spending (normalized by average daily spending) averaged across the treatment and control groups. By showing a wide span of data before and after the shutdown, it confirms the adequacy of the control group. In subsequent figures we use a narrower window to highlight the effects of the shutdown. The figure shows that the employees not subject to the shutdown have nearly identical movement in spending except during the weeks surrounding the shutdown. In this way, the controls appear highly effective at capturing aggregate shocks, seasonality, payday interactions, etc. In particular, note that there is a 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. Much of the sensitivity of spending to receipt of paycheck arises from reasonable choices of individual to time recurring payments immediately after receipt of paychecks (see Gelman et al. 2014). 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.

IV. Main Results

To formalize the difference-in-difference apparent in Figure 2 and to provide for statistical inference, we estimate 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.⁷ 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 individual level and adjusted for conditional heteroskedasticity.

⁷Specifically, 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.

A. Paycheck Income and Spending

This section examines how measured paycheck income and spending were affected by the government shutdown. Recall that we normalize each of these variables of interest, measured at the individual level, by the individuals' average daily spending computed over the entire sample period. Because individual spending levels vary widely, this normalization allows us to compare deviations from average spending across users experiencing very different economic circumstances. The unit of analysis in our figures is therefore days of average spending. Figure 3 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) represent the week in which pay dropped due to the shutdown, and the third when pay was restored.

Panel A of Figure 3 shows, as expected, a drop in income equal to approximately 4 days of average daily spending during the first paycheck period after the shutdown.⁸ 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 Mint Bills data allow an accurate measure of those income changes.

⁸ 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).

Panel B of Figure 3 plots the results on total spending, showing the estimated β_k where y is normalized total spending. 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 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.

We now convert the patterns observed in Figure 3 into an estimate of the marginal propensity to spend (which we call the MPC as conventional). 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 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 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.⁹

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 0.57 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.

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. Following Gelman et al. (2014), recurring spending includes expenditures (including ATM withdrawals) of at least \$30 that recur, in the exact same amount (to the cent), at regular frequencies such as weekly or monthly. It identifies spending that, due to its regularity, is very likely to be a committed form of expenditure (cf. 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

⁹ 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.

priori categorization based on goods and services. This approach to categorization is made possible by the distinctive features of our data infrastructure.

Figure 4 presents estimates of the β_k from equation (1) where the outcome variable y takes on different expenditure categories. 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. Recall the results on total spending (Figure 3) 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 spending and rose only by 0.84 days once the lost income was recovered. Non-recurring spending exhibits the opposite tendency: it dropped relatively little (1.8 days) and rose more abruptly (2.0 days). Thus, recurring spending drops more and does not recover as strongly as non-recurring spending.

To better understand what is driving 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 4 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 (see Figure 5). 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 4 shows the response of account transfers to the income shock.¹⁰ 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.

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 4 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.

Indeed, as we see in Panel F of Figure 4, there was no average reaction of credit card *spending* to the shutdown. In this way, 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 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

¹⁰ 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.

smoothing spending because they had little cash on hand did not—either because they were constrained by credit limits or were not interested in additional borrowing. In the next section we will examine the consequences for revolving debt of these postponed credit card 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 not a drop in actual spending or 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 6 shows that categories of expenditure that are quite close to consumption, such as the fast food and coffee shops spending index shows a sharp drop during the week starting October 10 when employees had already been out of work for a long period of time. Given that a cup of coffee or fast food meal is highly 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-employees after the shutdown, perhaps to trade war

stories.¹¹ Hence, in a sense, a cup of coffee is not entering the utility function as an additively separate non-durable.

B. Debt

The previous section showed that users reduced their credit card balance payments during the shutdown. For users who are not using revolving debt, this action is costless as long as the balance is paid before its due date. For users who are using revolving debt, they will be charged interest until they pay the balance back. To better understand the consequences of this behavior, we therefore analyze revolving debt at the extensive as well as the intensive margin.

Figure 7 shows the estimated β_k from equation (1) where y_{it} is either an indicator function for whether a user borrowed on a credit card (extensive margin), or the weekly average revolving debt balance, normalized by average daily expenditure (intensive margin) in week t for user i .¹² The fraction of revolvers among the treatment group is similar to that of the controls before the shutdown, and does not seem to react to the shutdown. On average, that is, affected workers were not pushed by the loss of income into becoming revolving debt holders if they had not already been in the seven weeks prior to the shutdown. Consistent with the balance payment results, the average revolving balance increases during the shutdown and then returns to very near pre-shutdown levels after six or seven weeks. It thus appears that temporarily drawing on pre-existing and revolving credit lines was one way in which consumers smoothed their spending throughout the shutdown.

¹¹ 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.

¹² The sample period starts later because we only have balance data (both credit card debt and checking and savings account balances) starting in August 2013.

C. Liquid Savings

For users who have built up a liquid savings buffer, they may draw down on these reserves to help smooth income shocks. We define liquid savings as the balance on all checking and savings accounts. Figure 8 shows the estimated β_k from equation (1) where y_{it} is the weekly average liquid savings balance normalized by average liquid savings or by average daily spending in week t for user i . Because of the heterogeneity in balances, normalizing by average liquid balances leads to more precise results. Normalizing by spending is less precise but allows for a more meaningful interpretation of the results. 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. However, the drop is mitigated by the anticipatory rise as well as sharp drop in spending seen in the main results. The recovery of the lost income causes a large spike in the balances, which is mostly spent during the following two weeks. Figure 8 shows that liquid balances fell by around 1.8 days of average spending. Therefore, on average, users reduced spending by about 2 days and drew down about 1.8 days of savings to fund their consumption when faced with a roughly 4 days drop in income.

V. Heterogeneity

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 is relatively easy to smooth

through the shutdown. For those with little liquidity, even this brief drop in income may pose significant difficulties.

A. Measuring Liquidity

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 58 days of daily spending at 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 β_k s from equation (1), for various forms of spending, by 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. Revolving 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 instruments for those with lower levels of liquidity. Here we consider the consequences for revolving debt for the three liquidity groups.

Figure 10 explores the heterogeneity across liquidity groups in the response of their revolving debt. The fraction of revolvers among medium liquidity and high liquidity users remained fairly steady throughout the shutdown. However, we see an important difference in behavior in the size of the revolving balance. High liquidity users continue on a downward trend relative to controls. There is no evidence that the shutdown caused them to take on additional credit card debt. Medium and low liquidity users, however, took on about 1.5 additional days of average total expenditure in revolving debt during the shutdown. Low liquidity users seemed, on average, to bring these balances back down after six months; medium liquidity users took another three or four months to recover.

We also analyze the revolving balance of users who borrowed throughout the whole period. Once we condition on these users who were borrowing the whole period, there is a clearer and more substantial discrepancy between the long run levels of revolving debt of high liquidity users compared to the rest of the population. Using the shutdown as the base level, the bottom two thirds of the liquidity distribution seems to hold an additional 1.5 to 2 days of average daily expenditure in revolving debt. Hence, for a population that was already leaning on credit card debt, even this brief shock to income may have lasting consequences.

V. Conclusion

This paper uses a novel infrastructure linking transaction accounts that provide an integrated picture of spending, income, saving, and credit card debt on a daily basis for a large sample of individuals. The accuracy and resolution of these aggregate datasets provide a new and important

opportunity for studying economic events and for providing fresh insights into consumer behavior. We use this infrastructure to study the effects of the government shutdown in October 2013. For a pay period during the shutdown, government employees subject to the shutdown suffered a reduction of approximately 40 percent of their normal pay. By the time they received this reduced pay, it was quite clear that they would be reimbursed for the reduced pay at the end of the shutdown. Since this loss of income would be rapidly reversed, standard economic theories predict that the shutdown would have little to no effect on spending. Yet, spending fell by about half the amount of the drop in pay.

We find, however, that much of the drop in spending did not correspond to a drop in consumption. Instead, many individuals re-arranged the timing of payments to defer cash outlays in ways that likely had little effect on their consumption or well-being. The data infrastructure thus provides new insights into how individuals manage their affairs and use financial instruments to smooth through the drop in income. Many workers affected by the shutdown deferred their debt payments, notably mortgages and credit card balances. Others deferred transfers to asset accounts. Interestingly, there was essentially no incremental spending using credit cards. Those with unused capacity to borrow either had other means of smoothing or little desire to do so. Those with little liquid assets, however, often did not have sufficient capacity to borrow.

This paper provides direct evidence on the importance of deferring debt payments, especially mortgages, as an important 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 (2013), 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 paper also shows, however, that an important subset of the population is still too constrained to cope effectively with even a brief shock that leaves their lifetime income unchanged. For this subset of most constrained individuals, who deferred payments of credit card debt with revolving debt, the shock had lasting effects in that the deferred balance payments were not made up even months after the shutdown. Hence, those who are in precarious financial situations in advance of a shock, even if it is completely temporary, may suffer sustained consequences.

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 (2014) “Did the 2013 Government Shutdown Severely Damage the U.S. Economy?” SIEPR Policy Brief, September.
- 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.”
- Dynarski, Mark, and Steven M. Sheffrin (1987) “Consumption and Unemployment,” *The Quarterly Journal of Economics*, 102(2): 411-428.

- 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.
- Gruber, Jonathan (1997) "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 87(1): 192-205.
- 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 (2013) "Foreclosure Delay and U.S. Unemployment," Working paper, University of Minnesota Department of Economics.
- Jappelli, Tullio and Luigi Pistaferri (2010) "The Consumption Response to Income Changes," *Annual Review of Economics*, 2: 479-506.
- 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 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.

Stephens, Melvin Jr. (2001) “The Long-Run Consumption Effects of Earnings Shocks,” *The Review of Economics and Statistics*, 83(1): 28-36.

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/>.

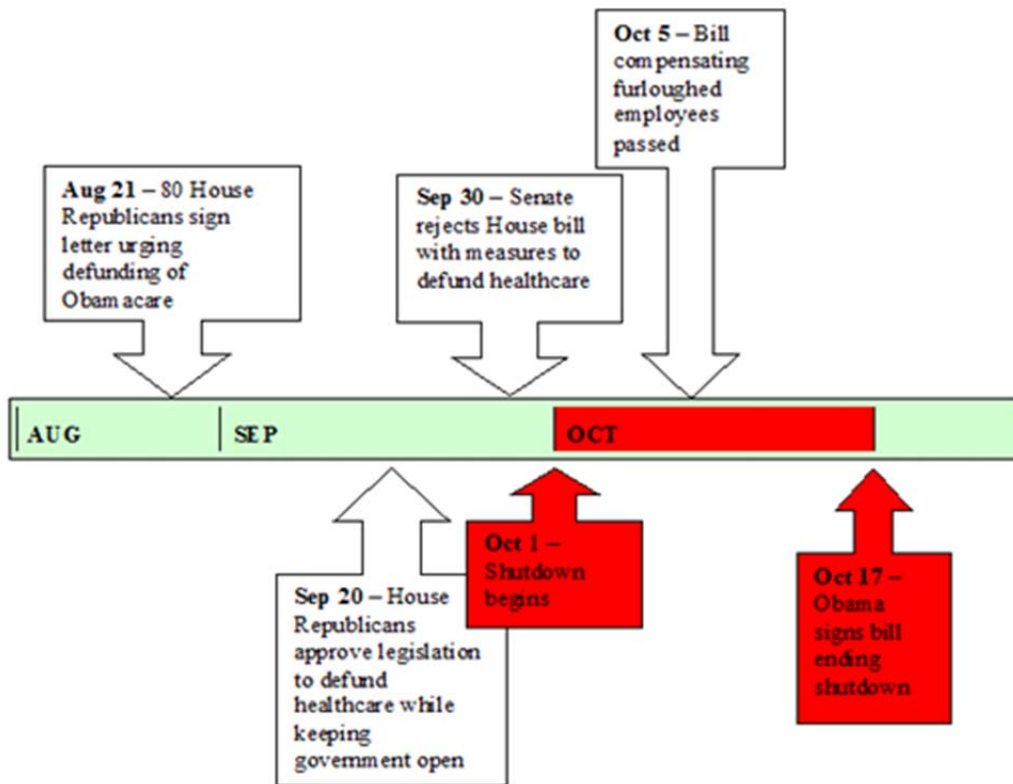


FIGURE 1. GOVERNMENT SHUTDOWN TIMELINE

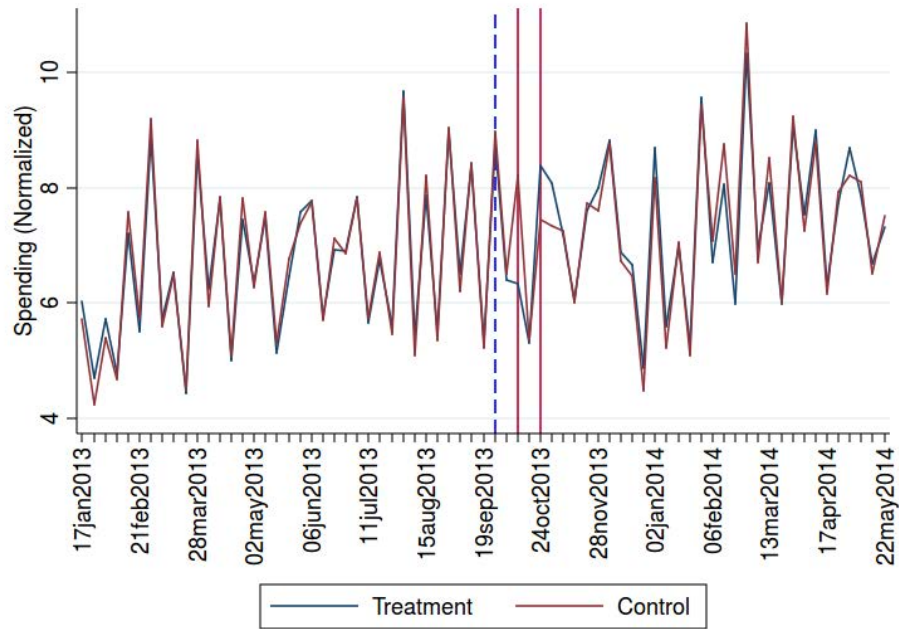
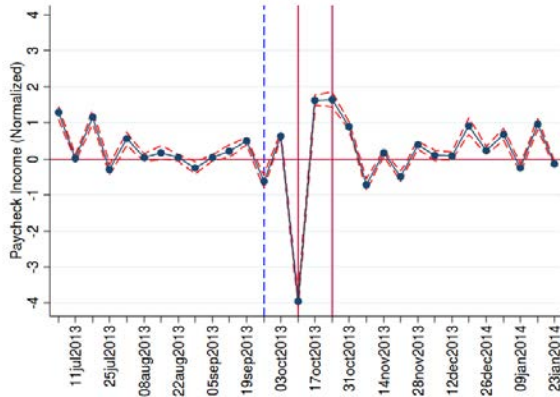


FIGURE 2. TIME SERIES OF SPENDING

Notes: The figure shows average weekly spending (expressed as a fraction of 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.

A. Paycheck Income



B. Total Spending

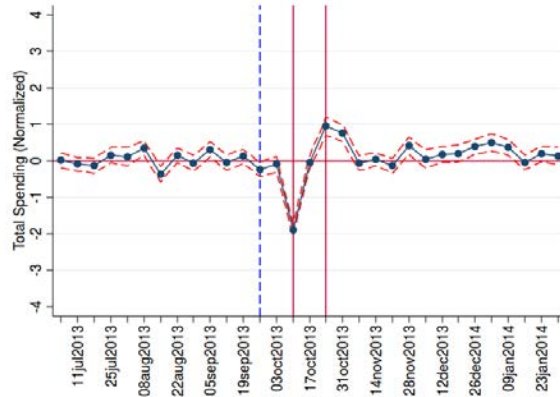


FIGURE 3. ESTIMATED RESPONSE OF NORMALIZED PAYCHECK INCOME AND NORMALIZED TOTAL SPENDING TO GOVERNMENT SHUTDOWN

Notes: Difference-in-difference estimates based on equation (1). The Paycheck income plot is estimated using additional controls which include paycheck week and treatment group interactions. $N = 3,762$ and $N = 94,722$ 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.

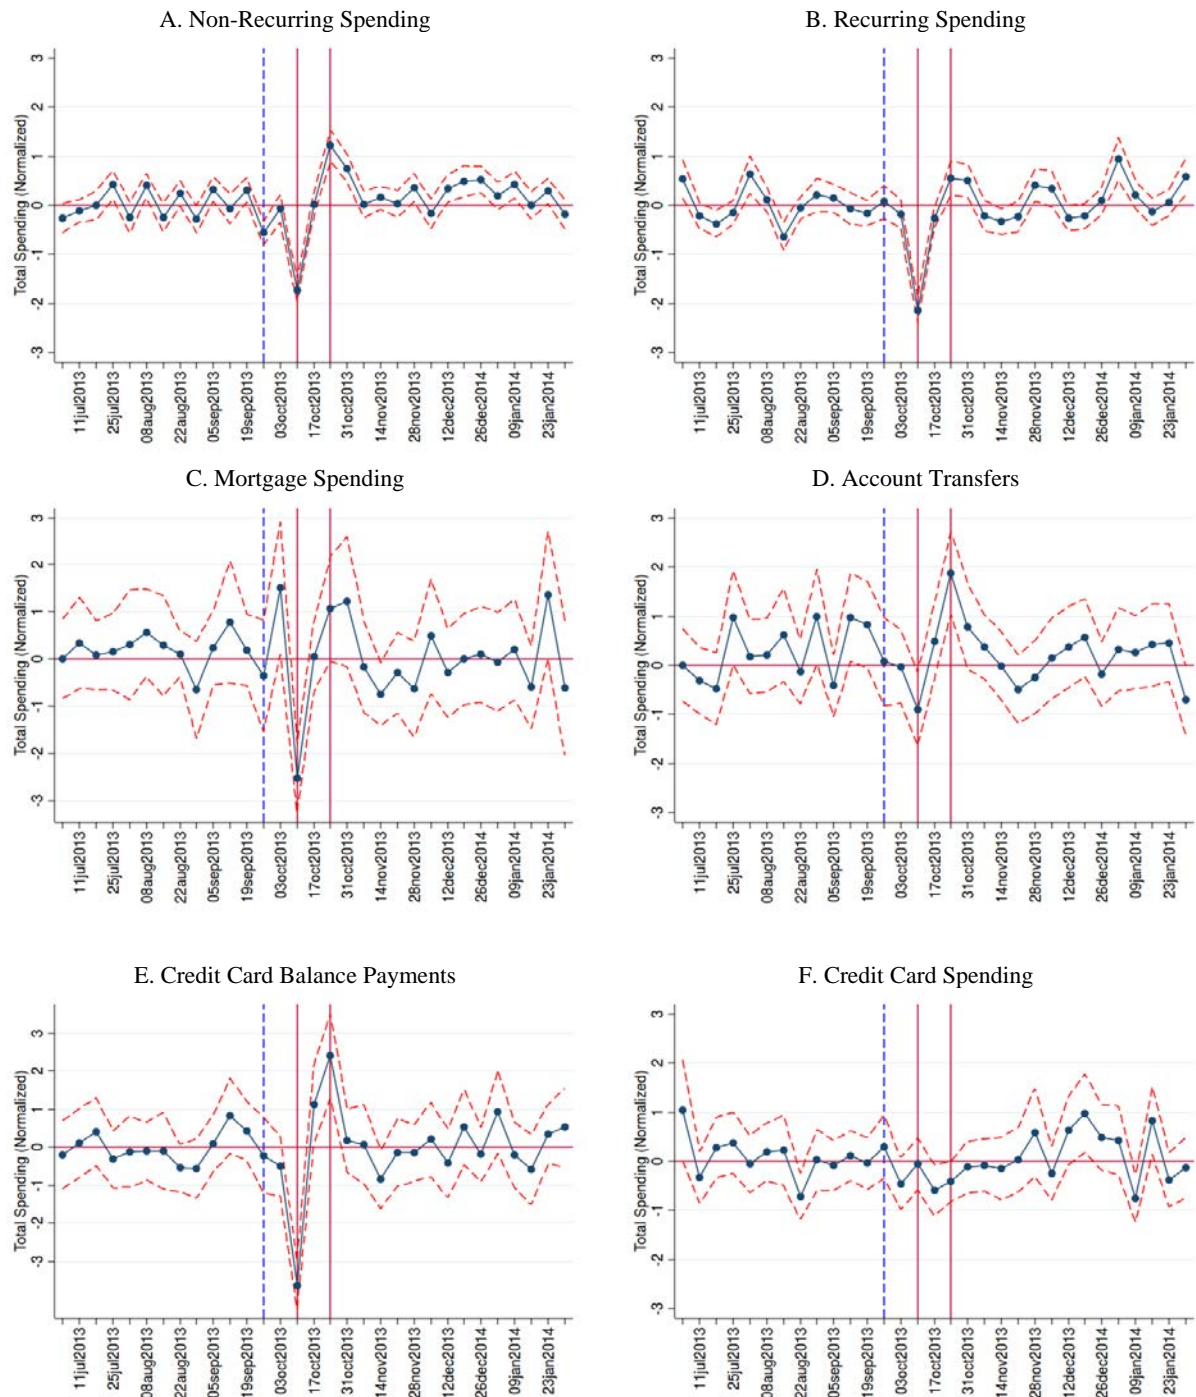


FIGURE 4. ESTIMATED RESPONSE OF SPENDING CATEGORIES TO GOVERNMENT SHUTDOWN

Notes: N = 3,762 and N = 94,722 for treatment and control group respectively.

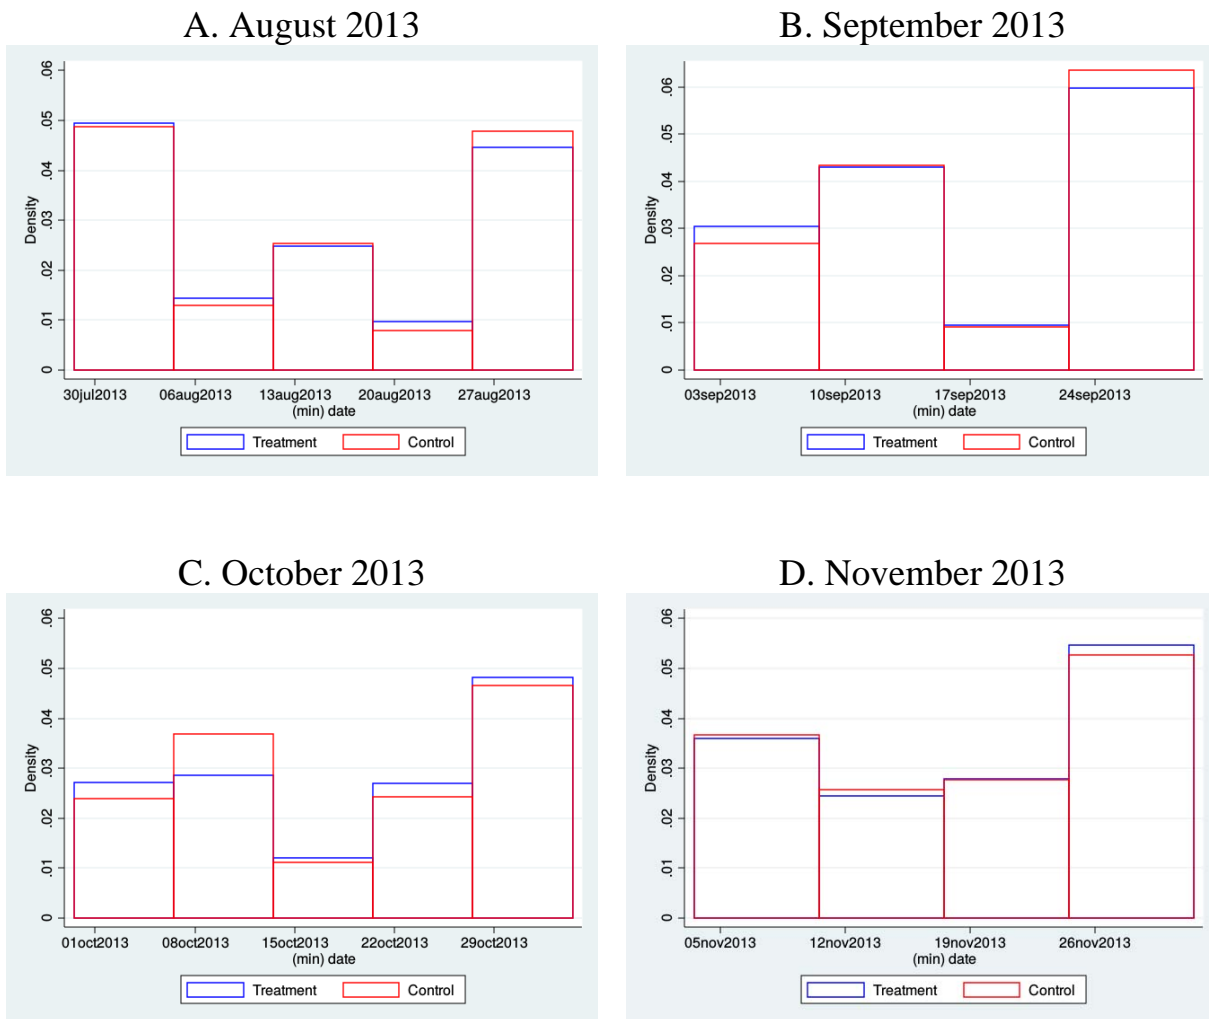


FIGURE 5. DISTRIBUTION OF WEEK MORTGAGE IS PAID

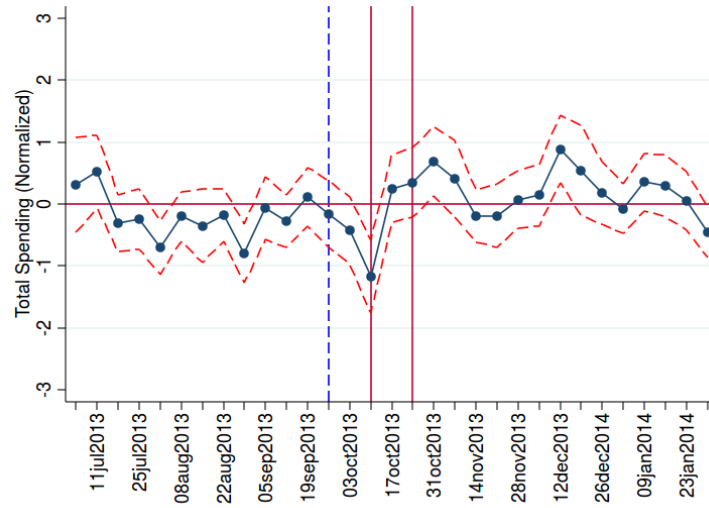


FIGURE 6. ESTIMATED RESPONSE OF COFFEE SHOP AND FAST FOOD SPENDING TO GOVERNMENT SHUTDOWN

Notes: N = 3,762 and N= 94,722 for treatment and control group respectively.

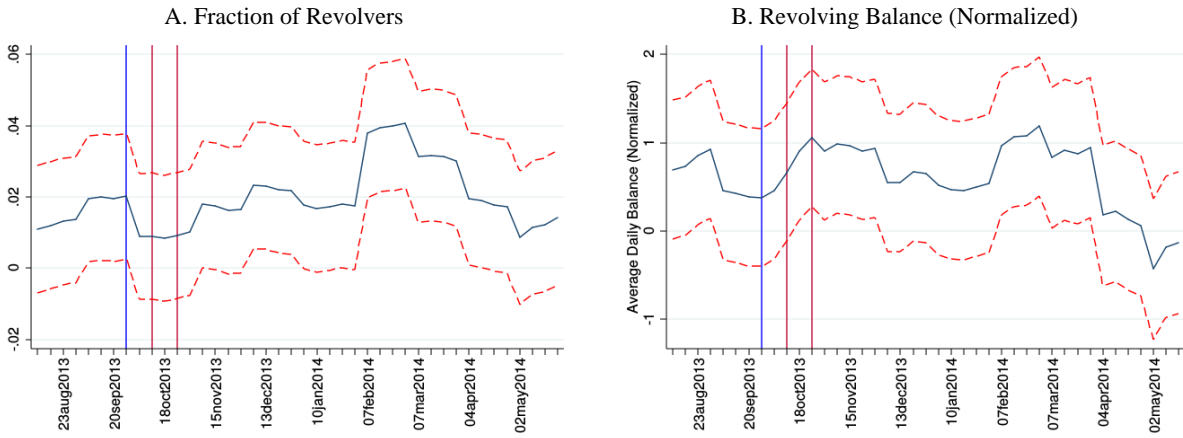


FIGURE 7. ESTIMATED RESPONSE OF REVOLVING DEBT TO GOVERNMENT SHUTDOWN

Notes: Outcome variables are winsorized at the upper and lower 1%. N= 72,584 and N = 2,327 for non-government and government employees respectively.

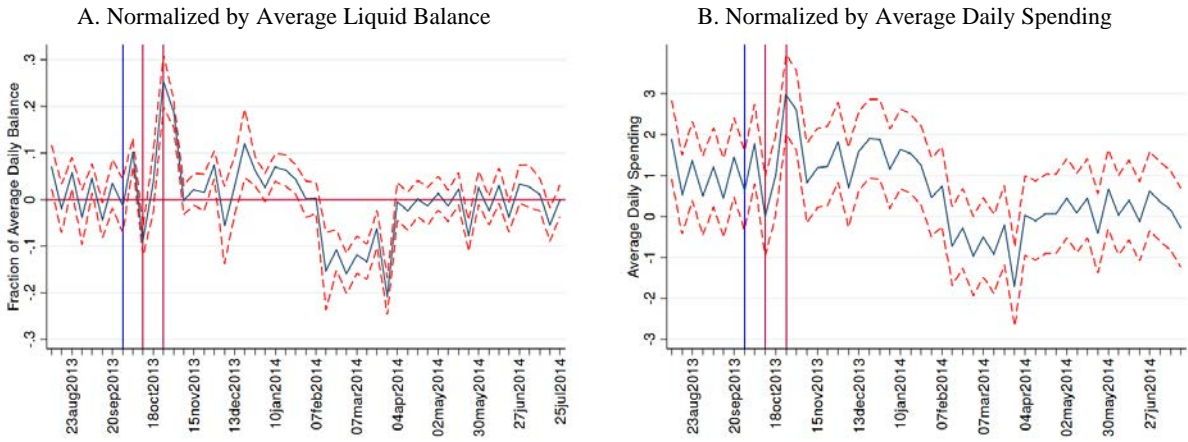


FIGURE 8. ESTIMATED RESPONSE OF LIQUID SAVINGS TO GOVERNMENT SHUTDOWN

Notes: Outcome variables are Winsorized at the upper and lower 1%.

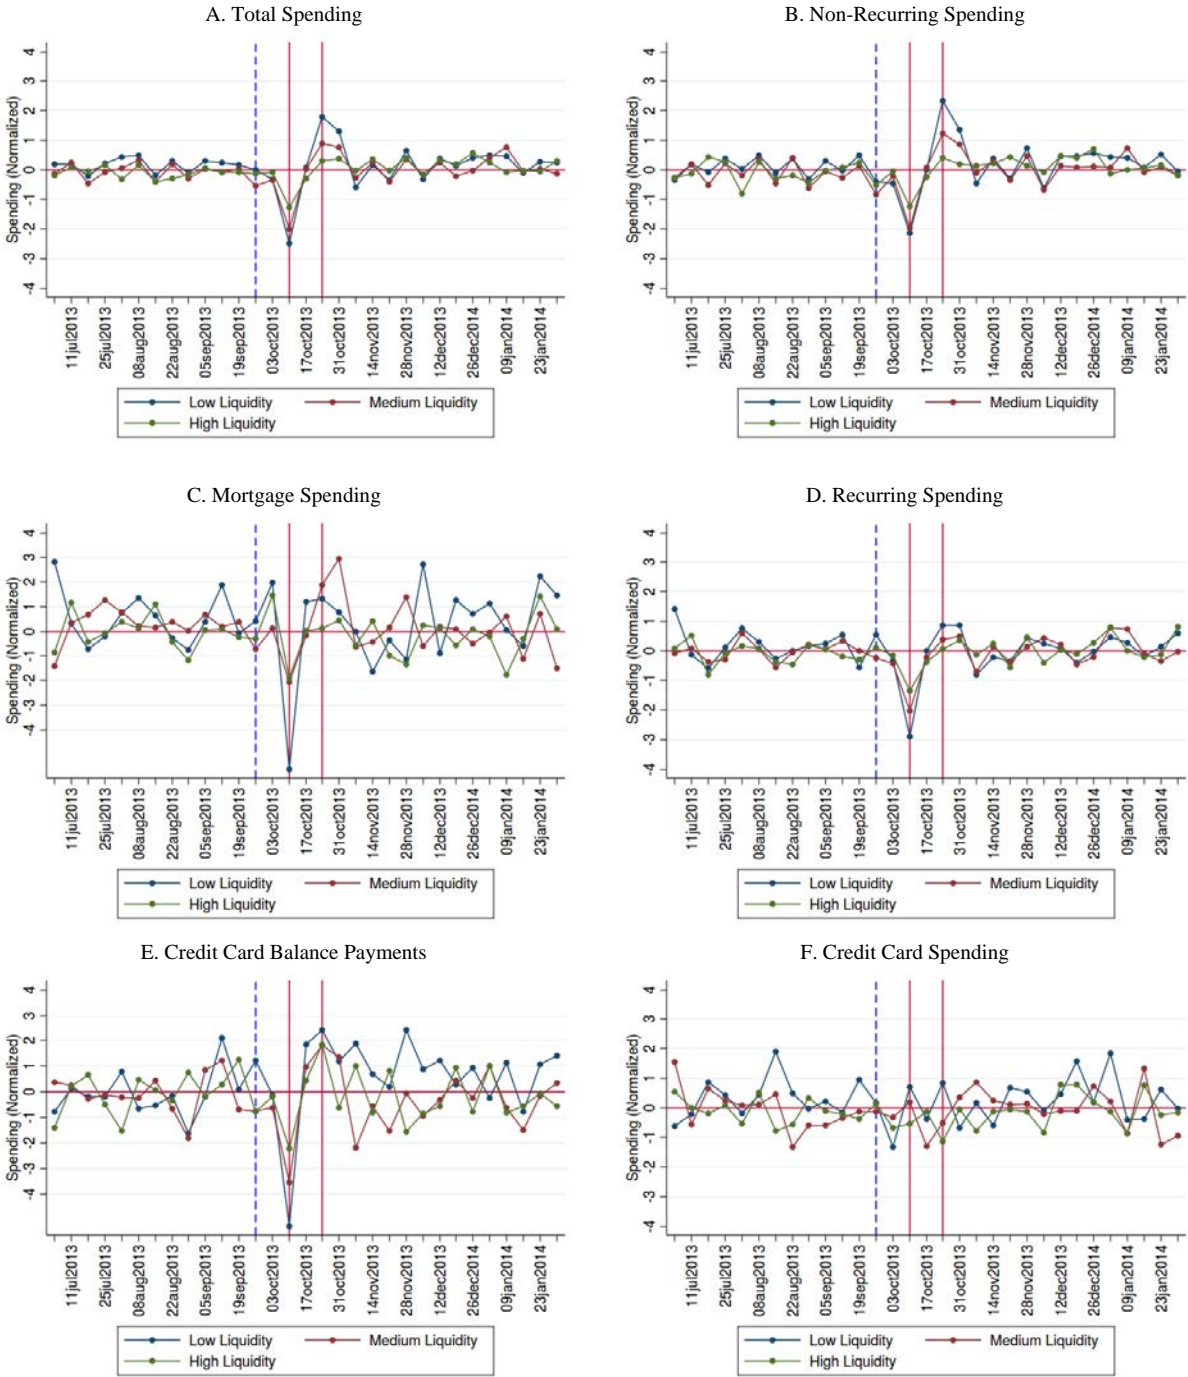


FIGURE 9. ESTIMATED RESPONSE OF SPENDING CATEGORIES TO GOVERNMENT SHUTDOWN BY LIQUIDITY TERCILE

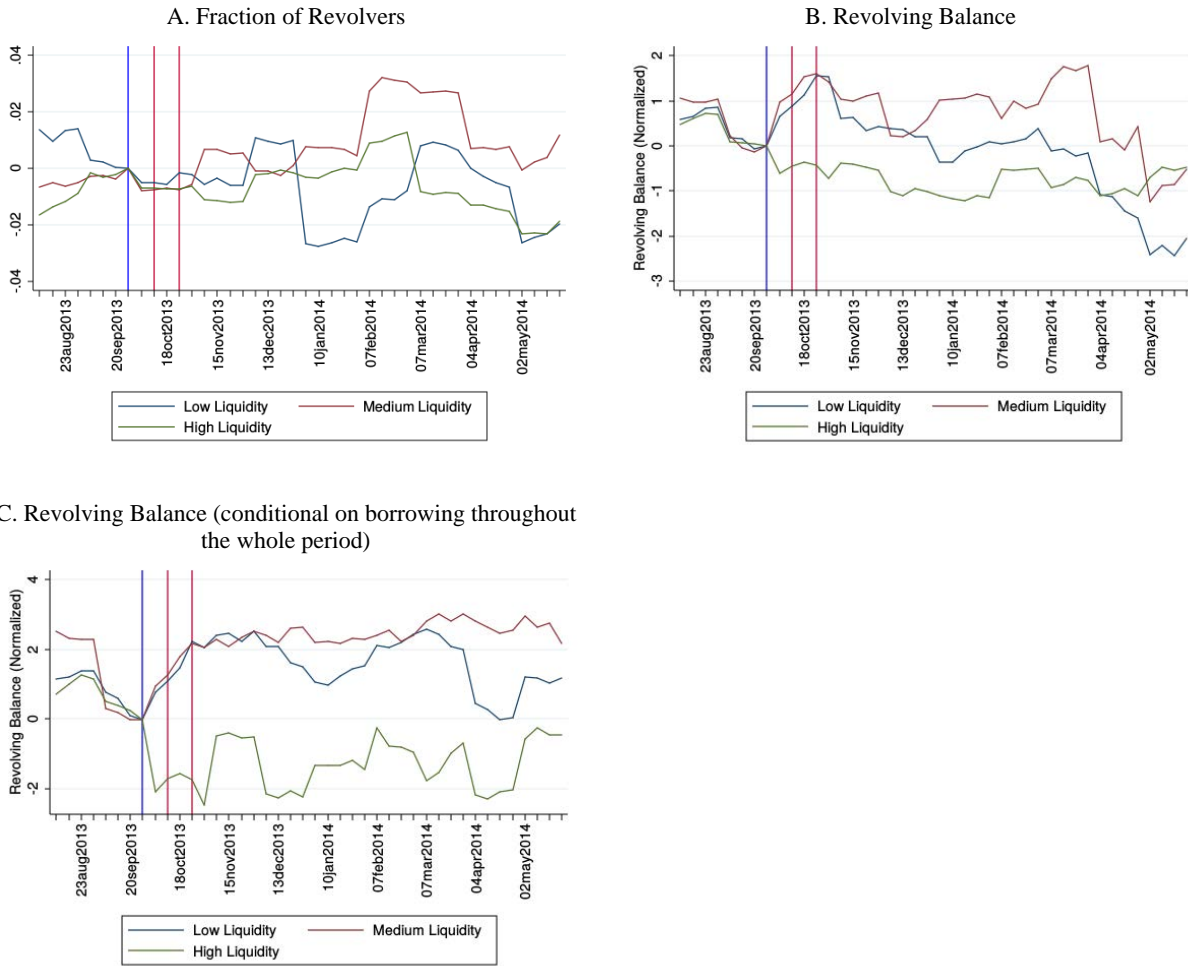


FIGURE 10. ESTIMATED RESPONSE OF REVOLVING DEBT TO GOVERNMENT SHUTDOWN

Notes: Outcome variables are winsorized at the upper and lower 1%.

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	6,710	\$1,726	\$1,418	\$1,837	28	\$3,673
Affected by the Shutdown	3,762	\$1,724	\$1,327	\$1,845	26	\$3,785
Not affected by the Shutdown	2,948	\$1,728	\$1,525	\$1,828	30	\$3,529
Non-Federal Employees	91,774	\$1,261	\$1,360	\$1,354	24	\$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
$\hat{\beta}_{MPC}$	0.0704 (0.0355)	0.0290 (0.0339)	0.5670 (0.0330)	0.0153 (0.0313)
Observations	98,477	98,477	98,477	98,477
SEE	7.112	6.778	6.603	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 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 heteroskedasticity, are in parentheses.

TABLE 4—LIQUIDITY RATIO

Liquidity Ratio Tercile	Treatment Group		Control Group	
	Mean (Days)	N	Mean (Days)	N
1	2.9	660	2.8	24,251
2	8.5	823	8.4	24,088
3	57.7	842	56.0	24,069

Notes: The sample is conditional on having accounts that are well linked. Variables are winsorized at 1%.

Appendix

A. Identifying Federal Employees

Most federal employees (88%) are employees of Cabinet Level Agencies such as the Department of Defense (DoD), the Department of Veterans Affairs (DVA), and the Department of Energy (DoE). The rest are employees of Independent Agencies such as the Environmental Protection Agency (EPA) or the Social Security Administration. We are able to identify employees who are paid via direct deposit by capturing the transaction description of their recurring paychecks.

Federal employees have the text “FED SAL” included in their transaction description. In total, we observe 12,214 users who we believe to be employed by the federal government during the time of the shutdown. Our data set identifies roughly 0.4% of the U.S. population over 18 (1,000,000 / 241,780,000). Since the Check population over samples younger working Americans, an upper bound would be an identification rate of 0.7% (1,000,000 / 144,303,000). Therefore, we would expect to observe between 8,400 to 14,700 federal employees. Our figure of 12,214 falls within this range. The transaction description also contains details about which federal organization the employee works for. Sometimes the description will list the department that the employee works for but there are cases where the description only lists the agency that processes the payroll of federal employees. There are four main agencies that process the payroll of federal employees. The largest agency is the Defense Finance and Accounting Service (DFAS) which provides payroll service for defense related departments such as the DoD and the DVA. However, they also service non-defense related organizations such as the EPA and the DoE. DFAS pays about 54% of federal employees. The second largest payroll service is the National Finance Center (NFC) which started off only servicing the Department of Agriculture but subsequently expanded to over 170 organizations including the Department of Commerce,

Department of Justice, Federal Bureau of Investigation, and the Congressional Budget Office. NFC pays about 31% of federal employees. We are unable to distinguish between departments paid by the NFC because they all use the keyword “AGRI.” The third largest service is the Interior Business Center (IBC) which started off supporting the Department of Interior (DoI) but expanded to service other agencies such as the Department of Transportation (DoT), and the National Science Foundation. The IBC services around 7% of federal employees. We are able to identify employees working for institutions serviced by the IBC such as the Department of Interior (DoI) and the Department of Transportation (DoT). Lastly, the General Service Administration (GSA) services many non-cabinet level agencies such as the Office of Personnel Management and the Railroad Retirement Board. The following table compares the fraction of employees employed in each agency in our data compared to the U.S. population for the largest agencies. The distribution of employees identified in the Check data roughly matches the U.S. population.

Table A1—Fraction of employees in each agency

	Check data	U.S. Population
DFAS (DoD, DVA, etc)	46%	54%
National Finance Center	31%	31%
Department of Interior	2%	3%
Department of	4%	3%
Department of State	1%	1%

Notes: The fractions for the Check data are calculated as the number of users under each agency divided by the total number of users identified as having the keyword “FED SAL” in a paycheck received in either September or October 2013. We are not able to further identify agencies under DFAS and NFC.

B. Identifying Effects of the Shutdown

There were many different situations that employees faced before and during the shutdown that potentially impacted their paycheck income. Some employees were not affected at all because their pay came from sources other than appropriations legislation. Military employees had their pay protected by the “Pay Our Military Act” signed into law on September 30, 2013 which appropriated funds even if a shutdown occurred. Others were furloughed and then recalled at various dates throughout the shutdown. Although we are not able to identify all of these possible situations, we are able to identify how paycheck amounts changed. Most federal government employees are paid on a bi-weekly basis. The pay periods and dates are usually set by the agency that handles payroll. For example, DFAS pays most employees on Thursday and Friday depending on the agency. For the NFC, the EFT date is Monday. For IBC, the EFT pay dates are on Tuesday. For GSA, the pay dates are Friday. In some situations, the paycheck may post one day early or late due to the characteristics of each financial institution. Although the actual pay dates may vary according to each agency, most employees share the same pay period dates. The relevant pay periods for this analysis are September 22 - October 5, 2013 and October 6 - October 19, 2013. Since the shutdown started during the latter part of the first relevant pay period, employees did not receive payment for 5 days out of the 14 day pay period. Their paycheck for that period is roughly 64% (9/14) of their regular paycheck because the shutdown was still in effect. Since the government shutdown ended before the next paycheck date, users who only received 64% of their paycheck were fully reimbursed by paychecks received after the shutdown ended. This event was a pure temporary shock in the sense that it had no impact on permanent income but simply reallocated income across time. The following figures show the pay dates for each payroll processing agency.

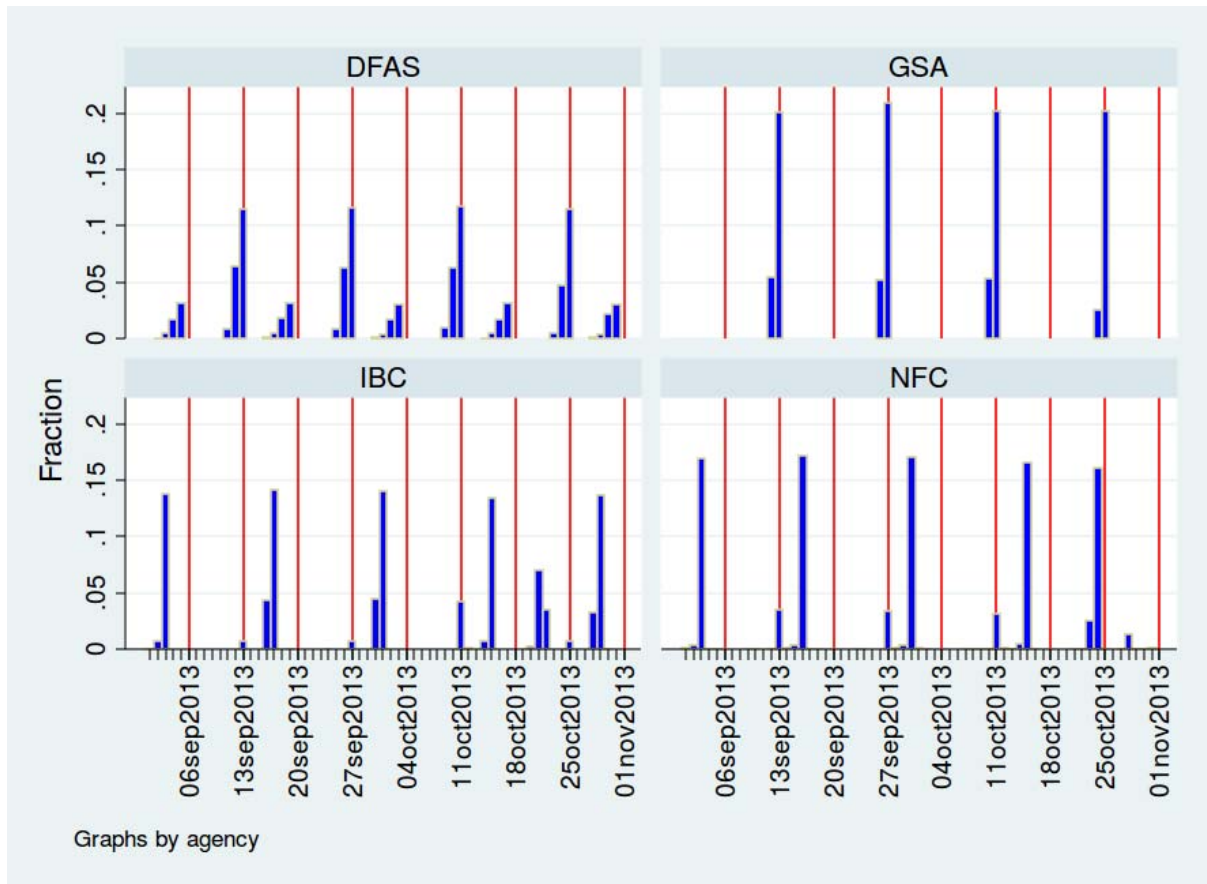


Figure A1. Paycheck histogram by agency

Notes: N = 5,778, 158, 838, and 4,227 for DFAS, GSA, IBC, and NFC respectively. Red lines represent Fridays.

Back pay for accrued hours was handled differently by each agency. DFAS rolled the back pay into the first paychecks for each employee after the shutdown ended and did not appear to change the timing. NFC and GSA also incorporated back pay into the first paycheck after the shutdown but paid people early the Thursday before their usual Monday pay date. IBC processed back pay during the week after the shutdown ended in a separate check. This paycheck was received during a week in which paychecks are not usually received. The figures below plot the time series of pay dates for each payroll agency.

Temporary loss of paycheck income — We further analyze the paycheck data to look for signs of the temporary loss of paycheck income as a result of the shutdown. We take all federal employees and look at the fraction between their current and last paycheck. For a typical employee affected by the shutdown, the paycheck will be around 60% of the previous paycheck. This fraction will vary by hours worked and withholding rates. I define a binary variable that equals 1 if the fraction is between 0.5 and 0.7. This variable is meant to be a rough check for employees affected by the shutdown. The following figure shows the fraction of users who experienced a fraction of paycheck change between 0.5 and 0.7 by pay dates.

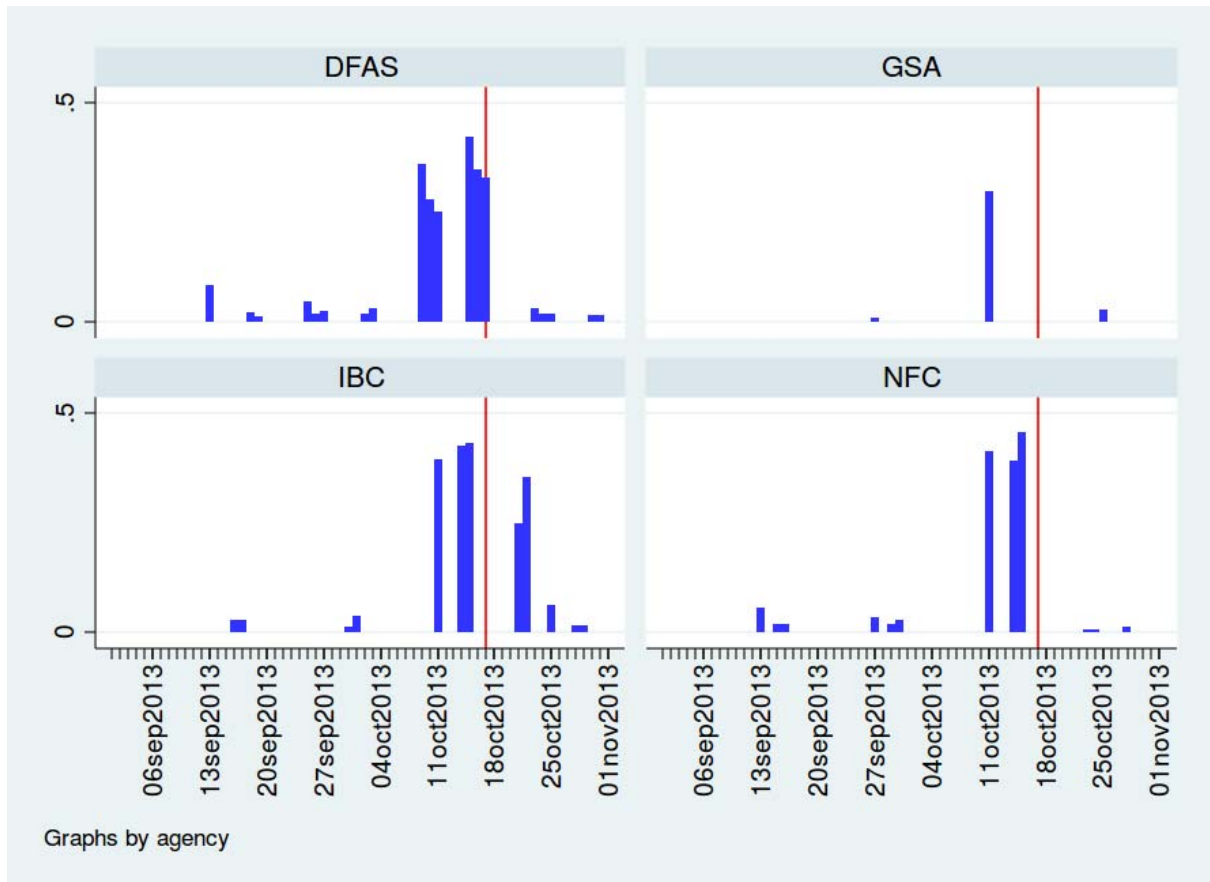


Figure A2. Fraction of users with paycheck RATIO OF 0.5-0.7 by agency

Notes: Only for dates on which there are 30 or more paychecks. The red line represents the end of the shutdown.

There is a large jump in the fraction of users that experience a drop in paycheck around 0.6 during the first paycheck received after the shutdown began. DFAS has a wider range of pay dates so the smaller paychecks are received from October 9 to October 17. NFC and GSA have more uniform pay dates that range from October 11 to October 15. For IBC, the first range of dates from October 11 to October 15 represents the smaller paycheck. The second range of dates from October 21 to October 23 represents the back pay received after the shutdown ended. The following figure shows a histogram for the fraction of each paycheck change for the first paycheck received after the shutdown began. In all agencies, there is a bimodal distribution with mass around 0.6 and 1. The mass at 0.6 represent those affected by the shutdown while the mass at 1 represents those who were unaffected. The results are consistent with the fact that employees paid by DFAS and GSA are more likely to have been paid out of funds that were not affected by the shutdown. In particular many employees paid by DFAS were protected by the “Pay Our Military Act” signed on September 30 right before the shutdown began.

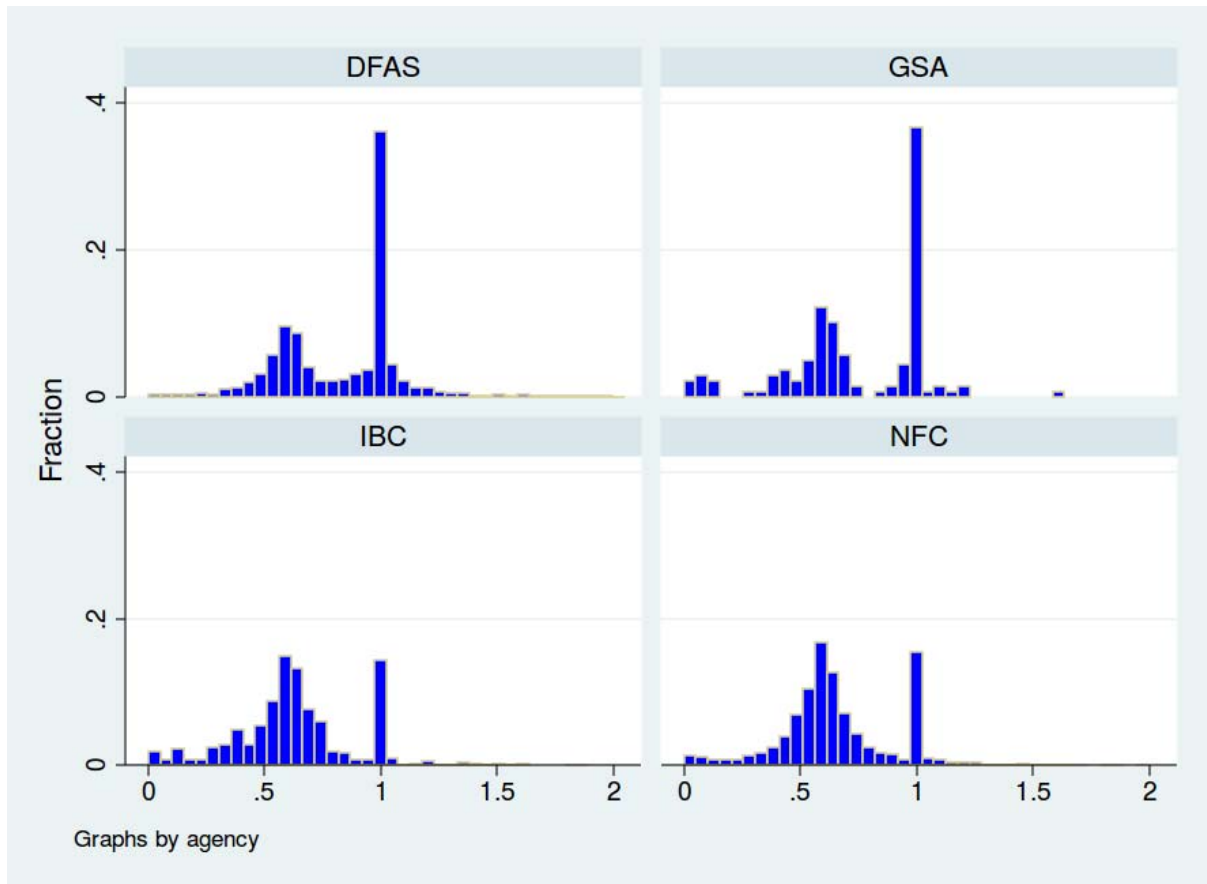


Figure A3. Histogram of fraction OF PAYCHECK by agency

Notes: The period analyzed is October 9 through October 17.

Recovery of temporarily lost income — We perform a similar analysis to analyze paycheck trends after the shutdown ended. Here we look for the fraction of users who receive a paycheck two or more times the previous paycheck. This is a rough check for the back pay of income owed but not paid during the shutdown. As expected, there is a jump in large paychecks during the first paycheck week after the shutdown ended.

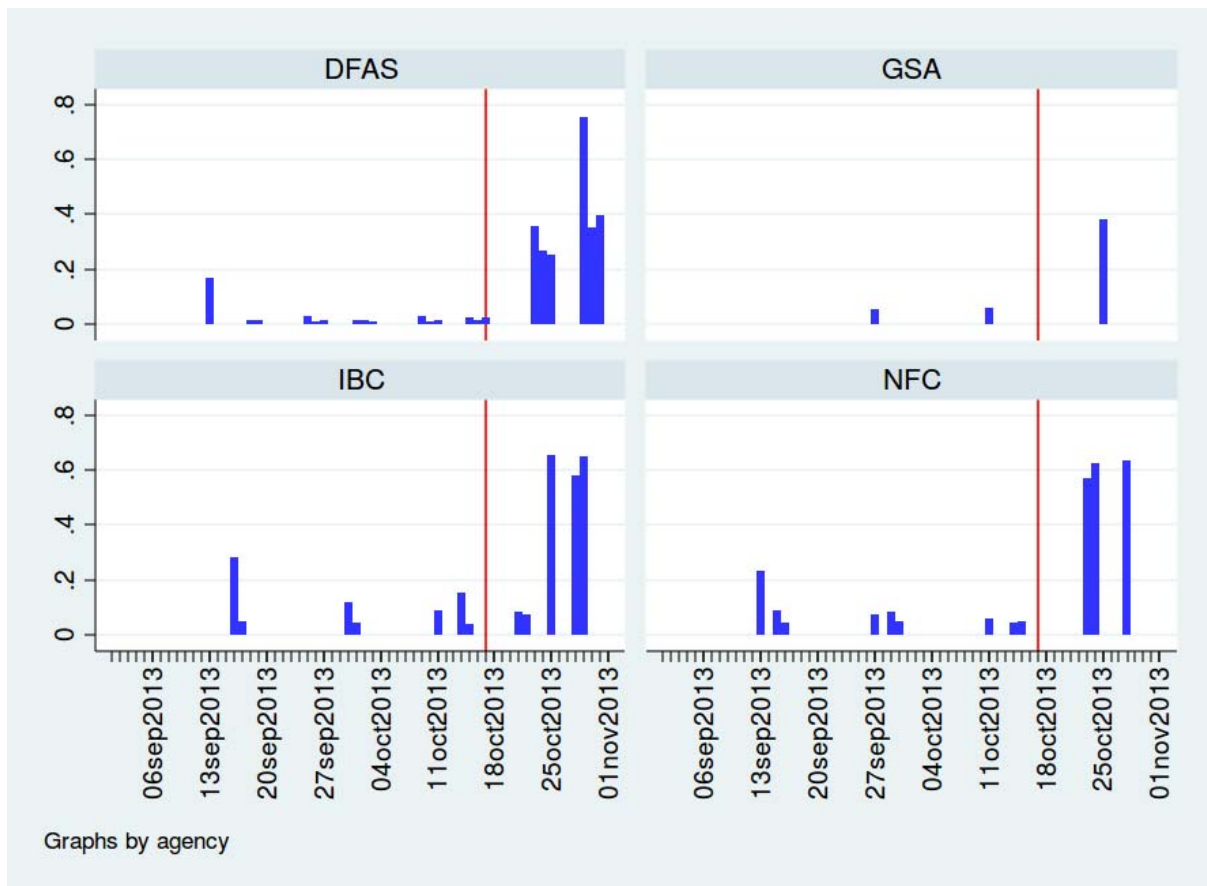


Figure A4 Fraction of users with paycheck RATIO > 2 by agency

Notes: Only for dates on which there are 30 or more paychecks. The red line represents the end of the shutdown.

The following figure shows a histogram of the fraction of the paycheck income of the first paycheck received after the shutdown relative to the paycheck received two pay periods ago. For most users, the paycheck received two pay periods ago should be the regular pre-shutdown paycheck. Since employees were compensated for the temporary loss of income suffered during the shutdown, the fraction should be centered on 1.4 (1.4/1) for employees impacted by the shutdown. The exception to the rule is for users paid through IBC. For those organizations, the reimbursement pay was received in a separate check the week after the shutdown ended.

Therefore, the ratio between the paycheck received during their regular pay period and the paycheck two pay periods ago should be 1.6. This represents the ratio of the return to the standard paycheck (1) divided by the smaller paycheck received as a result of the shutdown (0.6). There is a similar pattern to the previous section where fewer employees who were paid by DFAS and GSA were affected by the shutdown.

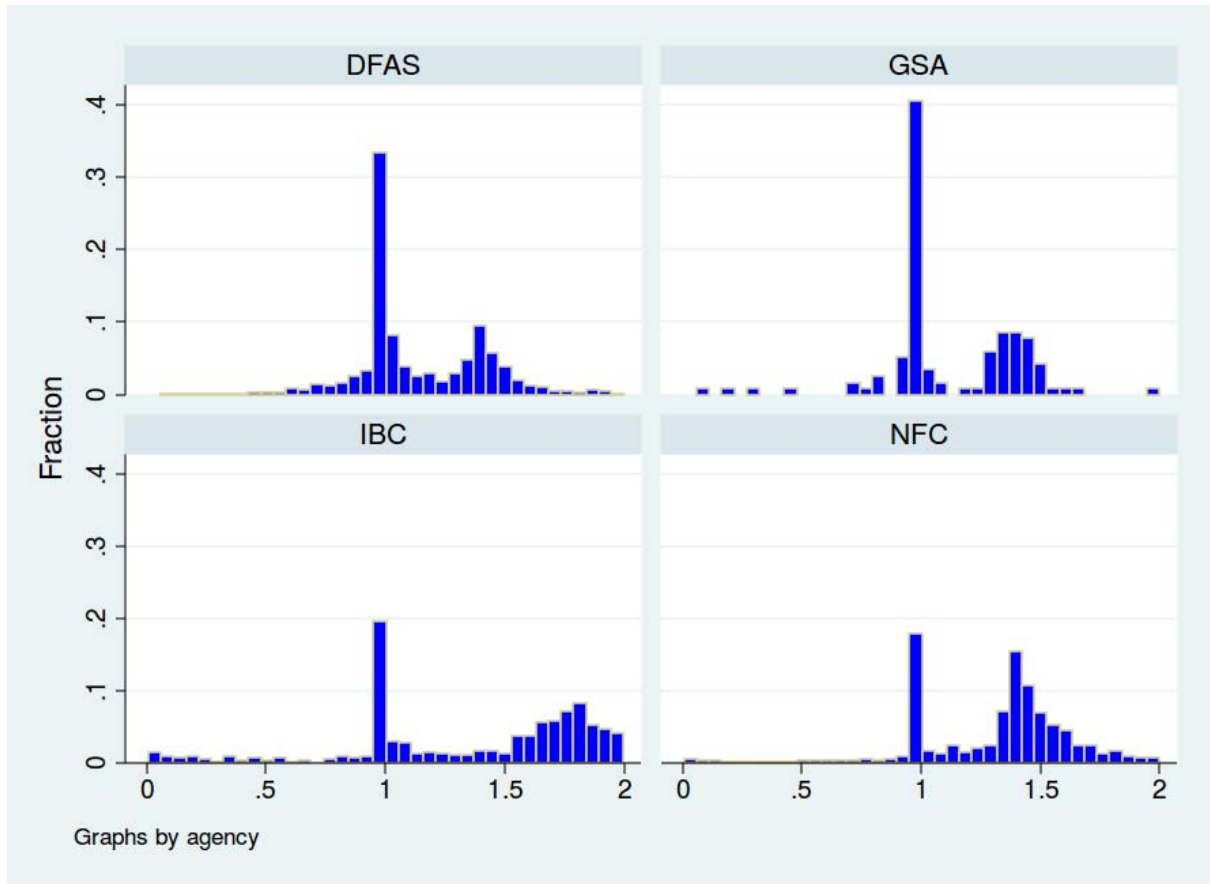


Figure A5. Histogram of fraction of paycheck by agency

Notes: The period analyzed is October 23 through October 31.