NBER WORKING PAPER SERIES

ASYMMETRIC CONSUMPTION RESPONSE OF HOUSEHOLDS TO POSITIVE AND NEGATIVE ANTICIPATED CASH FLOWS

Brian Baugh
Itzhak Ben-David
Hoonsuk Park
Jonathan A. Parker

Working Paper 25086
http://www.nber.org/papers/w25086

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2018

Early versions of this paper were entitled “A Test of Consumption Smoothing and Liquidity Constraints: Spending Responses to Paying Taxes and Receiving Refunds.” We thank the company for providing the data set. For helpful comments, we thank Manuel Adelino, Sumit Agarwal, Andrew Chen, Peter Ganong, Damon Jones, Olivia Kim, Emi Nakamura, Scott Nelson, Pascal Noel, Michaela Pagel, Michael Palumbo, Joel Slemrod, René Stulz, Annette Vissing Jorgensen, and the participants of the conferences and seminars at the Cleveland Federal Reserve Bank, Philadelphia Federal Reserve Bank Payments Conference, The Ohio State University, Annual Conference of the Western Finance Association, 2018, JPMorgan Chase Institute 2018 Academic Conference, and the 2015 AEA winter meetings. We are grateful for the financial support of the NBER Household Finance Grant. This work was supported in part by an allocation of computing time from the Ohio Supercomputer Center. Ben-David gratefully acknowledges the financial support of the Dice Center at the Fisher College of Business and the Neil Klatskin Chair in Finance and Real Estate. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Brian Baugh, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker. 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.

Asymmetric Consumption Response of Households to Positive and Negative Anticipated Cash Flows
Brian Baugh, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker
NBER Working Paper No. 25086
September 2018
JEL No. D12,E21,G4,H31

ABSTRACT

We use account-level data to document that households respond differently to expected transitory cash receipts than to cash payments. Consumers increase consumption spending when they receive tax refunds; however, they do not reduce their spending when they make expected tax payments. The central asymmetry in response and its pattern across liquidity and income levels is consistent with the behavior of rational consumers with liquidity constraints, but this canonical model cannot explain the lack of spending days before arrival of a refund or the lack of spending response to information about taxes around filing.

Brian Baugh
College of Business
University of Lincoln - Nebraska
baugh.41@osu.edu

Itzhak Ben-David
Department of Finance
Fisher College of Business
The Ohio State University
2100 Neil Avenue
Columbus, OH 43210
and NBER
bendavid@fisher.osu.edu

Hoonsuk Park
Nanyang Business School
Nanyang Technological University
hoonsukpark83@gmail.com

Jonathan A. Parker
MIT Sloan School of Management
100 Main Street, E62-642
Cambridge, MA 02142-1347
and NBER
JAParker@MIT.edu

1 Introduction

Under the assumption that people have diminishing marginal utility of consumption, they prefer to stabilize consumption over time and to insulate their standard of living from fluctuations in liquidity. In practice, studies show that households increase their consumption spending significantly in response to increases in income, even when the increase is predictable and transitory.[1] What is the role of financial frictions in this behavior? Is this excess sensitivity to cash flow due to liquidity constraints or is it instead due to the objectives and abilities of households such as in a range of behavioral models?

In this paper, we provide a test of household consumption behavior that distinguishes between these two explanations, or at least a class of models of each type. We compare the consumption responses of households to positive and negative predictable cash flows. If excess sensitivity arises due to financial frictions facing people who are otherwise smoothing consumption over time, there should be an asymmetry in the response of consumption to predictable changes in income or liquidity, as liquidity constraints can delay consumption spending until income arrives but not prohibit saving in advance of decreases in income. In contrast, if excess sensitivity arises instead due to myopia, inattention, mental accounting, or simple rules like consuming one's income, then consumption should decline with expected decreases in income in (roughly) the same way it rises with expected increases. In making such an inference it is critical that, as in our study, the increases and decreases in liquidity arise from the same source, since if they arose from different sources then differences in behavior could be due differences in mental framing (as in Thaler, 1990).

To compare the consumption responses of households to positive and negative cash flows which arise from the same source, we focus on tax refunds and tax payments associated with the US Federal individual income tax system.[2] We use a high-frequency administrative household-level dataset containing every transaction into or out of linked checking accounts, savings accounts, and credit cards for 2.7 million U.S. households from 2011 through 2015. We use a sub-sample of households that file their taxes with online tax-preparation services for which we observe the filing date, the tax refund/payment date, the refund/payment amounts, and prior tax information from which we can estimate expected taxes and the information acquired as returns are prepared.

Our results show a strong asymmetry: while household increase spending on consumption dramatically following the receipt of positive expected cash flows, they do

[1] See e.g., Bodkin, (1959), Parker (1999), Souleles (1999), Hsieh (2003), Johnson, Parker, and Souleles (2006), Agarwal and Qian (2014), DiMaggio, et al. (2017), and Kueng (forthcoming).
[2] In studying tax refunds, we follow Souleles (1999).

not reduce consumption following negative expected cash flows. Furthermore, we find little evidence for reaction of households to the news of forthcoming cash flows, whether positive or negative. Our results mostly, but not perfectly, can be explained by a rational model with liquidity constraints.

Our baseline analysis focusses on a small subset of household-years for which we observe the following spending and tax information. For spending, we require that we capture the entire spending profile of a household, and therefore remove any household with credit cards that are not linked to the aggregation service (dropping 93% of the sample). For tax information, we focus on a subset of households for which we observe both tax-filing events (as they pay a filing service fee) and consequently tax refund or payment (through electronic bank transfer, or credit/debit cards), for at least two consecutive years. We measure the information revealed during tax preparation and at filing as the residual from a regression of cash flow (refund or negative payment) on the previous year's tax information. Our simple predictive regression explains just over half the sample variation in refunds less payment.

These criteria (and a few relatively minor other filters) leave us with 11,138 households in our baseline sample. Our subsample is obviously not a random sample, because households first have to choose to be part of the aggregation service that makes the administrative data available, and because we require that households use a tax-preparation program, pay for it with a credit card or electronic payment, and finally pay any taxes due online rather than by check.

Our main analysis measures spending activity around the tax filing date and around the refund/payment dates. We identify the spending response to cash flows by regressing consumption spending onto a distributed leads and lags of refunds and tax payments. A significant advantage of our data and study is that we observe and control for the arrival of information about refunds or payment, as well as measure its impact. We measure the effect of refunds or payment by including distributed leads and lags of an indicator of the day of filing, as well as their interactions with both the amount of negative news and the amount of positive news about cash flow learned during return preparation. These distributed lags of information measure the convolution of the impulse response of spending to information about taxes and the average pattern of its arrival relative to the date of filing. Because the time between the dates of filing and dates of refund/payment varies across households, we can identify households' response to each of the events separately.

When examining the response of households to tax refund using our high-frequency account-level data, our findings are consistent with those in prior research; that is, households increase spending when refunds arrive. People spend over 30% of their refunds on our consumption measure the month following receiving a refund, a number that increases slowly to nearly 40% over several months. We use daily variation in inflows and outflows to identify the effect, and find no increase in spending related to the timing of the refund arrival prior to the day of arrival.

More importantly, households do not cut spending when paying taxes due. The point estimate of the spending effect of making a payment is small, statistically indistinguishable from zero, and positive as opposed to negative. We also find no evidence of any effect around the day of payment. As we discuss later, these findings are inconsistent with most behavioral theories that operate symmetrically, and are instead consistent with consumption smoothing in the face of restrictions on, or high costs of, borrowing or accessing less-liquid wealth. Because refund status is not exogenously determined, we provide a variety of evidence that other differences in households do not drive our results.

Not only are our main findings consistent with behavior driven by financial constraints rather than preferences or objective functions, but we also find additional support for the model of optimizing consumption smoothing subject to liquidity constraints. Specifically, we show that the spending response to refund receipt is larger for households with lower account balances, as measured by lower interest inflows to the accounts in the three months before tax season. Further, the spending responses to refund receipt are also larger for households with lower previous income inflows into the account. We find no measurable consumption-spending responses to making a payment for households with either low account balances or low incomes.

Finally, we test another prediction of the rational theory, namely, that liquidity constraints make households worse off by delaying consumption; hence, households observed to file earlier should be more constrained and spend more from their refunds. We find exactly this behavior: households that file in February have higher propensities to spend from refunds than those that file in March, which in turn have higher propensities to spend than those that file in April. Looking at payments, households with little liquidity should file earlier in case they owe taxes, so that they have time to accumulate the liquidity before they have to make the payment at the tax deadline. If households failed to take this timing into account, we would expect to observe spending reductions for households filing near the deadline and making payments. But even for households filing late (or households paying late), we do not see spending declines around the time of payment.[3]

Although most of our findings are broadly consistent with the behavior of reasonably sophisticated, although potentially quite impatient, households, we also uncover two behaviors that do not appear to be consistent with rational consumption

[3] Both pieces of evidence are inconsistent with the prediction that a significant fraction of households suffer from self-control problems that lead them to both procrastinate and to file near the deadline and spend at high rates from liquidity.

smoothing subject to financial frictions. First, spending increases on (and after) the exact day of refund arrival, even for liquid households, those with credit cards, and high-income households. This behavior suggests either a large response to the small remaining amount of uncertainty over the refund amount and arrival date after filing, or that behavioral factors delay the increase in spending until the refund has actually arrived. Second, our data allow us to observe the date of tax filing; hence, we can measure the spending reaction to the news about the tax refund or payment. According to the baseline theory, households that receive the news that they must make larger-than-expected payments should reduce spending and increase saving in order to be able make the payment. We find no evidence of this behavior – point estimates consistently show no reductions in spending to bad news – but standard errors are large. Approximately 5% of our sample pays taxes, the average payment is one-half the size of the average refund, and our measures of the amount and timing of news are less precise than our measures of cash flow.[4]

In summary, we find that households spend more when they receive expected refunds, and they do not cut spending when they make expected payments. The spending responses are larger for households that are more likely to be constrained along a number of dimensions. Many, but not all, of these empirical patterns are consistent with a model in which households are motivated to smooth spending and are constrained by financial frictions.

This paper is most closely related to the literature on consumption smoothing (cited throughout the paper) and to recent attempts so distinguish among competing models in which consumption changes with expected changes in income. Research measuring the effect of liquidity constraints dates back to Zeldes (1989) and to more recent papers using better measures of liquidity constraints and household debt (e.g., Jappelli, Pishcke, and Souleles, 1998; Agarwal, Liu, and Souleles, 2007; Aaronson, Agarwal, and French, 2012). Some papers examine the response of spending only to predictable decreases in liquidity, but they tend to focus almost exclusively on permanent or highly persistent decreases. Ganong and Noel (forthcoming) and Jorring (2018) find substantial declines in spending when unemployment benefits expire and when mortgage payments rise, but Souleles (2000) and Aguiar and Hurst (2005) find

[4] This last factor – mismeasurement of news timing and amount – raises concerns about our measurement of the spending response to expected cash flows. However, as we discuss, our results are robust to a variety of more flexible and extensive controls for the arrival of information. Further, because refunds always arrive with delay after filing, the fact that we find the spending effect on the day of refund arrival suggests strongly that we are correctly measuring the effect of cash flow and not news for refunds. For payments, if bad news arrived at the same time as the cash flow, we would be biased toward finding larger spending reductions in response to making payments (because bad news is correlated with payment amount). In fact, we find smaller spending reductions.

that consumption is well smoothed when college expenses start and when people retire, respectively.

Our most immediate predecessors are Shea (1995) and Park (2017), who test for liquidity constraints as opposed to myopia, by investigating the asymmetric response of aggregate consumption to predictable changes in aggregate income. These papers find the opposite pattern to what we find at the household level: in aggregate, consumption responds more to predictable decreases in income than to predictable increases. Finally, we also are closely related to other papers that also use high-frequency data (Stephens 2003, 2006; Gelman et. al., 2014) or account-level data (Olafsson and Pagel, forthcoming; Baker, forthcoming) to study consumer spending behavior.

2 U.S. Individual Income Tax Returns

A key feature of our empirical setting is that we can separate the timing of the arrival of news about future after-tax income from the timing of the arrival of the change in income. We identify news about future cash flow and the timing and amount of cash flow using the structure of the US individual income tax system. The US individual income tax covers all sources of household income in each calendar year. For most labor income, employers withhold income taxes from household pay during the calendar year, typically following IRS guidelines based on pay and family structure. The employer remits these funds to the Internal Revenue Service (IRS) during the year.[5] By the end of January of the following year, people receive information on the previous year's income and tax withholding from their employers, banks, and investment firms, and fill out and submit (file) tax returns – some variant of the IRS 1040 and additional schedules – to calculate their total tax owed.[6] We use the fact that many households use online tax preparation companies such as TurboTax to help them fill out and file their tax forms. If taxes owed exceed total taxes withheld during the previous year, the taxpayer is responsible for paying the difference by the mid-April tax deadline.[7] If taxes owed are less than withholding, the IRS remits the difference to

[5] No corresponding system exists for most capital income, so that as interest and dividends are earned and as capital gains are realized, taxpayers accrue liabilities without withholding, which leads many higher-wealth taxpayers to make additional estimated tax payments during the year.
[6] People with low incomes or no taxes due do not have to file. Married individuals can file taxes jointly.
[7] Late filing or late payment leads to penalties and interest costs. Taxpayers can file for extensions, which can delay the legal requirement to file until October 15, but people are responsible for interest charges from April 15.

the taxpayer as a tax refund, typically a few weeks after the taxpayer files his or her return.[8]

Most households receive refunds, and we expect this pattern for three reasons. First, simple inertia would lead to a refund status for most households because default withholding rates and estimated-tax worksheets are structured so that most households following these guidelines receive a refund. Inertia seems prevalent (Jones, 2012). Second, households seeking to optimize their withholding have an incentive to choose lower withholding and pay taxes later, but also a countervailing incentive to avoid significant underpayment and the associated penalties and interest. Finally, the earned income tax credit leads many low-income households to have a negative tax for the calendar year and so to receive a refund.[9]

We treat these tax payments as reductions in after-tax income, and tax refunds as increases in after-tax income. We construct our measure of news based on the fact that households uncover information about their refund or taxes due as filling out their tax forms before filing. Thus, information about future cash flow arrives during the period before filing, and the cash flow happens at, or more typically during, the period after filing. The timing of the arrival of information is based on a household choice, at least within the main tax-filing season. The timing of the arrival of any tax refund is partly based on the endogenous filing date and partly due to the largely random delay between filing and disbursement by the IRS. Finally, payment of taxes is determined by the household, subject to the binding April 15 deadline.[10]

3 Theory: Taxes Refunds, Payments, and Consumption

In models in which agents are optimizing, forward-looking, and not credit constrained, the agents increase spending as they prepare to file their tax returns if they learn their refund is larger than expected, or they decrease spending if they learn their refund is lower than expected (and the reverse for taxes due). Then, when they receive their refund or make their payment, consumption spending does not increase or decrease. Given that tax refunds and payments are small relative to lifetime income, the adjustments to spending while preparing taxes should be small, on the order of a few percent of the new information in spending each year.

In models with liquidity constraints, these responses may be asymmetrically constrained. Agents can always increase savings to prepare to make a tax payment.

[8] In 2012, the IRS indicated that 90% of refunds were processed within 21 days of filing: http://www.irs.gov/uac/2012-Tax-Season-Refund-Frequently-Asked-Questions. See also Slemrod et al. (1997).
[9] Although households cannot choose negative withholding or estimated tax payments, households with children who qualify for the EITC can file a W-5 with their employer and receive up to 60% of the EITC credit early.
[10] We observe very few households that make payment before formally filing.

But households with limited liquidity cannot always borrow to increase spending in anticipation of a tax refund.[11] Thus, a model with liquidity constraints predicts different responses to news about refunds and payments. People should decrease spending in response to bad news about tax payments due, but some people cannot adjust spending in response to news about refunds. The response of spending to the inflow of a refund or the outflow of a payment is similarly predicted to be asymmetric. People should not decrease spending contemporaneously with making a payment, because they have prepared for it, but some people will increase their spending contemporaneously with receiving a refund.

The canonical theory with liquidity constraints also makes predictions about the timing of filing and its correlation with the reaction to tax-related information and cash flows. Because the timing of filing and payment are endogenous, optimizing households that are short on liquidity and expect a refund should file earlier than households that are not as short on liquidity. Thus, the optimizing theory predicts that households filing early and receiving refunds should have larger spending responses than households receiving refunds later. Furthermore, households that are short on liquidity and thus are concerned about making significant tax payments should file earlier and pay at the deadline. Households with plenty of liquidity should file and then pay whenever convenient.

Other theories of consumer behavior have quite different implications, and in particular, predict symmetric spending responses. If households are hand-to-mouth households (as in Campbell and Mankiw, 1989) or behave as target savers in the Reis (2006) model of inattention, they consume their income (or some constant fraction thereof). In this case, spending should increase with refund receipt but also fall with tax payment. Further, if households are target spenders, consumption spending should not respond to news, refund, or tax payment. Although in Reis (2006) these rules are time dependent, households might instead follow state-dependent or more sophisticated rules in which their propensity to spend on arrival is related to the size of the utility loss caused by spending behavior that deviates from that of the fully attentive model.

4 Data and variable construction

4.1 Data Source

The data we use were provided by an online account aggregator. This aggregation service allows households to view their various financial information in one place, allowing one to view spending by category, monitor investments, and so on. The

[11] Or if they can borrow, they may choose not to make the effort or pay the fixed cost to obtain credit, or choose not to pay the higher interest rate of unsecured borrowing, or choose not to take on the costs associated with turning less liquid assets into consumption.

service also provides alerts for upcoming bills and for approaching credit limits. Households join the service for free and provide their usernames and passwords to various financial institutions so that the service can extract relevant bank and credit card information.

The data we use consist of daily transactions for 2.7 million households from July 2010 to May 2015, and include both banking (i.e., checking, savings, and debit card) and credit card transactions. We observe the date, amount, and description of each transaction. Thus, our dataset contains transaction-level data similar to those typically found on monthly bank or credit card statements. Because each household is assigned a unique identifier, we are able to follow each household through time.

The sample is selected but appears to be broadly representative of the population with some exceptions. Table 1 illustrates how our final sample is located geographically relative to the US Census. As shown, households in our sample are well dispersed geographically, though we have high concentrations of households in California, New York, and Texas. Figure 1 illustrates the income distribution of our final sample relative to the US Census. As shown, households in our sample vary widely in income, and the distribution maps the broader US distribution quite well. However, we appear to have a number of very low-income accounts, where we might be failing to identify income correctly.

4.2 Panel Construction

To construct our panel, we first identify federal tax refunds and payments by querying the transaction descriptions. Such transactions are easily identified via queries for "us treasury des tax" and "irs treas tax," among other terms.[12] In an effort to remove outlying tax activity such as that occurring through business owners, we exclude any household-year containing more than one tax refund or payments. We further remove any household that has ever incurred a tax payment or tax refund of over $20,000.

Likewise, we identify tax-preparation transactions by querying for payments to TurboTax, H&R Block, TaxAct, or TaxSlayer using electronic payment, debit cards, or linked credit cards. We do not observe a payment for households that elected to deduct the preparation charges directly from their refunds. We also exclude households that

[12] Conventional tax refunds and payments are easily identified in the data using the keywords "us treasury des tax", "irs treas des tax," "irs treas tax," and "irs usataxpymt." Our main analysis uses only these. However, when we predict refund, we also use refunds paid directly to households by tax preparers. Many tax-preparation software companies, such as TurboTax, allow customers to pay their tax-preparation fees directly from their refund rather than beforehand at the time of filing. In this event, the government first deposits the funds to TurboTax, which extracts the customer's normal tax-preparation fee plus an additional service charge and deposits the remaining balance to the customer. Such transactions are identified in the data by querying for "sbtpg," "tax products p," "block bank des hrbb," "block bank hrbb," and "republic trs."

have tax-preparation transactions on multiple days (as would be the case for a family filing separately on different days). The transaction date of the tax-preparation software is designated as the filing date of the household.

In general, we require that the filing date precede the tax-refund or tax-payment date. Due to small differences between financial institutions in how quickly transactions post to different accounts, we allow the tax-payment date to precede the tax-filing date by no more than two days. Such a scenario would occur in which an individual pays a tax-preparation fee with a credit card that posts immediately while paying their taxes with a direct debit from their checking account, which takes two days to post. On average, refunds are received and payments are made 10.3 days after filing.

To limit our sample to more typical refunds, we require that both the filing date and refund or payment date occur before June 1. Additionally, we require that we observe tax refunds or payment in the year prior to the year of data for each household.

In some specifications, we alleviate the filing requirement and only require the observation of tax payment or tax refund. Doing so increases our sample size by an order of magnitude.

We arrange our data by household years running from November 1 to October 31. Each year consists of 365 days, with the exception of 2015, when our sample ends, which consists of 237 days.

To measure spending accurately, our baseline sample is restricted to households for which we do not observe any payments made to credit cards that are not linked to the account aggregator. We drop households that we observe making payments to credit card accounts that are not linked, because we cannot categorize the payments made on these cards or determine the timing of spending on these cards. In untabulated results, the inclusion of such individuals dramatically reduces the observed consumption response to tax refunds because the vast majority of consumption is unobserved due to the unlinked nature of the account. By contrast, for accounts without unlinked cards, we observe and categorize all spending on a daily basis, whether these transactions occur via debit or credit card. The requirement to have no unlinked cards is by far our most limiting filter. We lose 93% of our sample with this single filter, because the vast majority of accounts have not linked all of their credit cards to the account aggregator.

To ensure we have active account users rather than dormant account holders, we impose a simple activity filter for households. We require that households have non-zero transactions in any category for at least 25% of days in a given year, which is equivalent into an average requirement of at least seven days with non-zero transactions per month.

After applying the above filters, our baseline sample contains 15,456 household-years from 11,138 unique households, which leads to a dataset with 5 million household-day observations for our regression analysis. When we alleviate the requirement to observe the filing date, our sample size grows to 154,507 household-years from 80,747 households, leading to a dataset of 51 million household day-day observations for our regression analysis.

4.3 Variable Construction

In our main analysis, we are interested in the consumption and savings responses to tax refunds and tax payments. We construct our main consumption variable as the sum of spending in the following categories: gas, restaurant, retail, groceries, cash, entertainment, healthcare, travel, utilities, miscellaneous bills (e.g., gym memberships), and insurance. We likewise construct a savings (and debt payment) measure as the sum of outflows on the following categories: mortgages, auto loans, net investing (flows to investing accounts – flows from investing accounts), net credit card payments (credit card payments minus net credit card expenditures), and other loan repayments (e.g., student loans). We also construct a measure of miscellaneous payments that captures payments we cannot assuredly categorize into either consumption or savings. This variable is equal to the sum of the following categories: checks and net uncategorized transactions (uncategorized inflows – uncategorized outflows). Checks are inherently difficult to categorize, because they may be a payment to an investment brokerage or a payment to a travel agency. Given that we classify checks as miscellaneous payments rather than consumption, we are likely understating true consumption.

Because we are interested in the extent to which spending reacts prior to the arrival of cash, we are particularly interested in spending that occurs in credit cards. We also measure consumption spending that occurs on credit cards. This consumption measure is not mutually exclusive to the other variables. To ensure this credit card spending is not mechanically related to taxes, we exclude from it any filing fees or tax payments.

We also measure income based on direct deposits of income. We measure income as the sum of all income receipts in the months of November, December, and January so that our measurement of income predates tax filing and refund within the year. Figure 1 shows the distribution of annualized income. The mean monthly household income in our baseline sample is $4,648, and the median is $3,997, corresponding to average and median annual household incomes of $55,776 and $47,724, respectively. The US Census Bureau estimates the average and median annual household incomes as $75,195 and $53,585, respectively, for 2013.[13] However, the income is observed after withholdings, which include federal taxes, state taxes, social security taxes, Medicare taxes, 401k contributions, healthcare premiums, and health savings account contributions. When accounting for the different items that reduce the observe income in our sample, it aligns fairly well with that of the general population.

Because we observe transactions rather than account balances, we extract interest flows to proxy for account balances and household liquidity. To avoid a mechanical relationship between interest transactions and refund or payments, we limit our search of interest transactions to the first three months of our year (November, December, and January). To be included in our financial-slack calculations, households need to have either interest received or paid or both.

Even though we do not directly observe account balances in our data, we are able to observe changes in account balances over time by simply integrating net flows to the household's accounts. We define our net-flow variable as the signed sum of inflows and outflows. When we integrate our net-flow measure, we begin with a value of $0 at the beginning of each year to illustrate changes through the year.

Table 2 shows summary statistics for filing, refunds, and payments. Figure 1a shows that returns are filed throughout tax season, but with a slight bimodal distribution, presumably consisting in February of people impatient for the funds, such as EITC recipients, and in April, people who postpone until the deadline. Figure 1b shows that returns with taxes due are filed significantly closer to the deadline. Figures 1c and 1d show the delay in days between filing and refund receipt and filing and tax payment, respectively. This delay for refunds is a function of IRS processing, determined in part by regional processing center delays at different times and by the complexity of the given return. This delay for payment is largely a function of whether households simply pay when filing or choose instead to pay right before the deadline, although many payments fit neither scenario.

The refund amounts represent a substantial amount of income for households. Approximately half of the households receive refunds greater than or equal to half a month's salary, and a quarter of households receive refunds greater than one month's salary. Figure 1e shows the distribution of refunds – payments, which has a mean of $2,170 and standard deviation of 2,448 (Table 2). The distribution is skewed, with 94% of returns leading to a refund in our sample.

4.4 Summary Statistics

Table 2 shows the summary stats for households in our sample. In our baseline sample (Panel A), the average monthly income is $4.648 per month, whereas the

[13] 2013 Current Population Survey from the US Census (HINC01).

median income is $3,977 per month. The average household files at the end of February, though the standard deviation of filing date is 30 days. The average household makes a payment or receives a payment on March 10, though the standard deviation of the refund or payment date is 29 days. The average distance between filing and refund is 10.3 days, though the standard deviation is 9.7 days. In this sample, 94% of households receive a refund. Conditional on receiving a refund, the average refund is $2,395. Conditional on making a payment, the average payment is $1,197. Sixteen percent of households have linked all of their credit cards, whereas 84% of households have no linked credit cards. As a reminder, we remove from our sample any household with unlinked credit cards. On average, households receive $1.89 in net interest per month across all accounts, though the median amount of net interest is $0.07. We observe an average consumption of $64.35 per day.

Panel B presents the unrestricted sample, which does not require the observation of a filing date. Compared to our baseline sample, which requires the filing date, the unrestricted sample has a slightly lower monthly income of $4,186 as compared to $4,648 for our baseline. This difference could be partially explained by lower-income households qualifying for free tax-preparation services, which would not show up on a credit card statement. Also consistent with this explanation, 97% of the unrestricted sample receives a refund as opposed to 94% for the baseline sample. Further, the size of the refund is about $500 larger for the unrestricted sample. Sixteen percent of households in both the unrestricted and restricted samples have linked credit cards.

Interesting differences are notable between households with no linked credit cards (Panel C) and those with linked credit cards (Panel D). Households with no linked credit cards have an average monthly income of $4,513, whereas those with linked credit cards have an average monthly income of $5,379. Similarly, households with no linked credit cards earn an average of $1.28 per month in net interest, whereas those with linked credit cards earn an average of $4.69 in net interest. We observe similar levels of consumption across both groups, with an average daily consumption of $65.22 for those without linked credit cards and $60.01 for those with credit cards. Households with credit cards file for taxes and receive refunds four days after those without credit cards. Ninety-four percent of households without credit cards receive refunds, whereas 91% of those with credit cards receive refunds. The average refund size for those without credit cards is $2,250, whereas the average refund size for those with credit cards is $1,762.

Interesting differences are also notable between those who receive refunds (Panel E) and those who make payments (Panel F). On average, those who receive refunds file on February 25, whereas those who make payments file on March 30. On average, those receiving refunds do so on March 8, whereas those who make payments do so on April 5. Sixteen percent of households receiving refunds have linked credit cards, whereas 25% of households making payments have credit cards. Households receiving refunds have an average monthly income of $4,525, whereas those making payments have an average monthly income of $6,654. Similarly, households receiving refunds earn an average of $1.63 per month in net interest, whereas households making payments earn an average of $4.77 in net interest. We observe an average of $63.98 in daily consumption for those receiving refunds and an average of $70.16 in daily consumption for those making payments.

In untabulated results, an interesting trend arises when comparing individuals across net-interest-received terciles. Ninety-five percent of households in the lowest tercile of net interest receive refunds, whereas only 87% of households in the highest tercile of net interest receive refunds. The average refund size for the lowest tercile of net interest received is $2,338, whereas that of the highest tercile is $1,920. Twenty-seven percent of households in the top tercile of net interest received have credit cards compared to only 13% of households in the bottom tercile of net interest received.

5 Estimation method

5.1 Information Acquired During Tax Preparation and Filing

We measure the news about future tax refund or payment as the difference between the actual tax refund/payment and the expected refund/payment. To compute the expected refund/payment, we predict refund payment using information on the previous year's income and take the residual from this equation as a measure of the information revealed by tax preparation. Specifically, we project refund payment in year t onto the previous year's refund (zero if taxes due), the previous year's taxes due (zero if refund), and an indicator variable for refund in the previous year using linear regression. This regression has a fit goodness (R^2) of 50%.[14] Our measure of information about tax information uncovered during filing, or news, is the residual in this regression, which we denote by E_{y-1}[Refund − Payment]. The average absolute value of the news is $1,092 relative to an average refund payment of $2,170 with a standard deviation of $2,448. This estimate is unlikely to perfectly match the true news that each household received during tax preparation, and we discuss how this difference affects the interpretation of our main results in the section below.

This empirical model is identified from cross-sectional variation and has a short time dimension and so effectively endows agents with knowledge of the increase in

[14] Adding the previous year's income and its interaction with the previous year's indicator variable leads to only a trivial increase in fit. Adding two years' prior income as well leads to a slightly greater increase in fit (about 1%) but a large decline in sample size.

average refund over the few years we study. This period, however has a reasonably stable tax law. According to the IRS, average refunds declined reasonably steadily by $82 per year from their peak in 2010.[15] To the extent that households did not anticipate these declines, as our empirical model assumes, our measure of news could be slightly upward biased on average.

5.2 Estimation of Impulse Responses to Cash Flow and Information

We estimate the impulse response of household consumption spending (and other account flows, e.g., savings, income, and interest) to the arrival of a refund or the making of a payment. We model the spending response as linear in amount but different for refunds than for payments (linear with a kink at zero). We also control for the arrival of information by estimating in the same regression the impulse response of household consumption to the news uncovered prior to and at filing, allowing the spending response to be affine in the amount of positive news and affine in the amount of negative news.

To be precise, define the following variables:

- *Refund* = refund amount on day received; otherwise, 0
- *Payment* = payment amount on day paid; otherwise, 0
- *News* = *Refund* − *Payment* − E_{y-1}[*Refund* − *Payment*] on filing day, else 0
- *PosNews* = Max[*News*, 0]
- *NegNews* = Max[−*News*,0])
- *File* = 1 on day of filing; otherwise,0

Letting k index days, our main estimating equation is

$$Y_{h,t} = \sum_{k=-29}^{K} \beta_k^+ \, PosNews_{h,t+k} + \sum_{k=-29}^{K} \beta_k^- \, NegNews_{h,t+k} + \sum_{k=-29}^{K} \phi_k \, File_{h,t+k} + \sum_{k=-29}^{K} \gamma_k^+ Refund_{h,t+k} + \sum_{k=-29}^{K} \gamma_k^- Payment_{h,t+k} + \alpha_h + \tau_t + u_{h,t} \,, \quad (1)$$

where $Y_{h,t}$ is an inflow or outflow measure for household h on day t and α_h is a household-specific intercept and τ_t a day-specific intercept. K is set to the maximum identifiable lag. The β_k, γ_k, and ϕ_k coefficients measure the impulse responses—the prior, contemporaneous, and lagged response of the dependent variable across weeks—to news about refund or payment and the date of filing (where event time is day of filing), and to getting the refund or making the payment, respectively (where event time is the day of refund or payment).

[15] IRS Statistics of Income, Tax Stats, http://www.irs.gov/uac/SOI-Tax-Stats-Amount-of-Refunds-Issued,-Including-Interest,-by-State-and-Fiscal-Year-IRS-Data-Book-Table-8.

We smooth the daily impulse responses by imposing that the daily coefficients are constant within weeks from $k = -29$ to -15 days, and for $k > 14$ days. Standard errors allow for arbitrary heteroscedasticity, within-day correlation, and within-household correlation in $u_{h,t}$.[16] We report cumulative spending effects, and standard errors for cumulated daily total are calculated for the endpoint of each discrete interval (correctly from the variance-covariance matrix of the coefficients).

6 The Consumption Response to Tax Refunds and Tax Payments

This section first establishes our main result, namely, that households increase spending when tax refunds arrive but do not decrease spending when making tax payments. Subsections 6.1 and 6.2 provide evidence that this result is unlikely due to the endogeneity of refund or tax-payment status (section 6.2) or due to bias stemming from mismeasurement of the news about tax status (section 6.3).

6.1 Main Result

Figure 3 shows estimates of cumulated coefficients, $\sum_{k=-29}^{T} \gamma_k^+$ and $\sum_{k=-29}^{T} \gamma_k^-$, for different horizons T from the estimation of equation (1) on our main measure of spending on consumption. The figure thus shows the cumulative increase in spending as a percent of refund and as a percent of payment, since 29 days before the refund arrived or payment was made.

First, in Figure 3.a, we observe that, on average, people increase spending starting the day on which their refunds arrive. People spend about 30% of their refunds on our main measure of consumption over the month following receiving a refund; over the four months following receipt, they spend about 38%.

We find no evidence of increases in spending prior to the day of arrival, at least related to the timing of refund arrival. We also find no response to information uncovered during filing, a result we return to in section 9.

Our second main finding is that households do not cut spending when making a payment. Figure 3.b shows the change in spending around the time when households make a tax payment is small, statistically indistinguishable from zero, and positive as opposed to negative. That is, if anything, people tend to increase spending slightly a while after making a payment. We also find no evidence of any decline in spending around the day of payment as we might have expected given the strong reaction to cash flow in response to refunds.

[16] These variables are consistent as N and T go to infinity at the same rate.

Both as a robustness check and because it allows us to use a much larger sample, we also present results that control for the news that arrives but with the timing related to the cash flow rather than filing. That is, we re-estimate our main regression but replace *News* with:

$$Refund - Payment - \mathrm{E}_{y\text{-}1}[Refund - Payment] \text{ on refund or payment day else } 0$$

An additional benefit of this alternative is that it increases the sample size by an order of magnitude since we do not require that we observe the date of filing.

That said, in this alternative specification, the news variables have control-function interpretations, so that if the true model has a positive spending response to cash flow and non-negative spending response to news, this specification biases downward the estimated spending response to cash flow. On the other hand, omitting the possibility of a spending response to news would bias the coefficient of interest upwards. But, based on estimation on simulated data, the downward bias of our alternative appears to be small, on the order of 1% of the refund amount. We proceed noting that a small downward bias likely occurs in the spending response to refund, which is a cost of distinguishing the response of spending to cash flow from the response to news when we do not separately measure the timing of the arrival of news.

Figures 3.c and 3.d show that we find almost the same asymmetric spending response but with greater statistical precision. Households do not increase spending before receiving their rebate, and then they spend more than 30% within a month that we can readily classify as consumption spending, and over 40% within three months. We also continue to see no evidence of any decline in spending when people make payments.

These findings are not spuriously driven by different typical seasonal patterns of spending around tax season or different refund and payment amounts across households. The day fixed effects (τ) capture the average spending on a particular calendar day, so that the typical fluctuations on weekends, holidays, spring months, and during tax season do not bias our results.[17] The household fixed effects (α) capture the household-specific level of outflows, both due to differences in standard of living across households and differences in the share of spending that we measure in our account-level data. These effects ensure we do not misestimate spending responses because higher spending households tend to have larger refunds.

Before providing more evidence on theories by looking at different subsamples, such as those with smaller and larger refunds, we deal with three important issues of measurement. First, the tax status of a household—its refund or payment amount—is

[17] And, as we show subsequently, we find similar results estimating equation (1) on subgroups of the sample in which the time effects are averaged over fewer households. And we find the asymmetry estimating with a log dependent variable and indicators of refund, payment, positive news, and negative news.

not randomly assigned. Could our differential responses reflect differences in spending propensities across different households rather than asymmetric responses? Second, the dates at which people file and pay taxes are both choices. Could our asymmetry be due to the endogenous timing rather than differences in spending propensities? Finally, could our results be driven by mismeasurement in the amount and timing of information about refund or taxes due, so that spending responses partly reflect reactions to this news?

6.2 Are Estimates Due to the Endogeneity of Tax Status and Payment or Refund Amount?

We are interpreting the difference in estimated propensities to spend between refunds and payments as being due to the sign of the cash flow. However, an alternative interpretation of our results so far is that they are instead due to differences between households that receive refunds and those that make payments. We would find an asymmetry if households that make tax payments smooth consumption through expected changes in cash flow, whereas those that receive refunds have high propensities to spend from cash flow.

From a theoretical perspective, a difference in this direction seems unlikely, at least as might be driven by differences in discount rates or liquidity. All other things equal, impatient or illiquid households should withhold less and thus be more likely to make payments and have high spending responses. More patient or liquid households should be less concerned about over-withholding and thus be more likely to have refunds and lower spending responses. Thus, differences across households in impatience or liquidity would lead households with lower spending reactions to be more likely to get refunds.

Turning to evidence on this point, we focus on a subsample of households that are similar: those that expect to either make payments or to receive small refunds. We rank households by their expected tax refund less payment, and then run our baseline regression on only the bottom 20% of households. These households on average have a refund less payment of \$490, with a standard deviation of \$1,673. Three-quarters of these households receive refunds.

Figure 4 shows that these households still display a large asymmetry in spending response, although statistical uncertainty increases substantially. Panels a and b display the results for our baseline sample; panels c and d display the specification without filing date and thus with the larger sample. Consumption spending increases rapidly only after the arrival of a refund to about nearly 30% (panel a) or 45% (panel c) of the refund in the first month. Thereafter, we rapidly lose statistical precision in the smaller sample and find no evidence of any continued increase in spending. In the larger sample, spending continues to increase, reaching nearly 65% of the refund.

Turning to payments, Figure 4 still shows no evidence that consumption spending declines in response to making a payment. In the larger sample, some statistically insignificant evidence points to households slowly lowering consumption over time following a payment.[18]

As a second approach, we analyzed the sample of households that receive a refund in at least one year and that make a payment in at least one year. In this sample, every household is used to identify both the response to refunds and the response to payments. This sample represents only 6% of observations, and standard errors are too large to distinguish behaviors (see Figures 5.a and 5.b). Thus, we expand our sample to include households that have unlinked credit cards, which increases the sample size by more than an order of magnitude. In this sample, we find the same asymmetry. Figures 5.c and 5.d show that we find the same asymmetric response in consumption spending in this larger sample. Because we are omitting consumption spending on unlinked cards, we estimate a much lower average consumption response to refunds (and still a slight *increase* in spending in response to making a payment).

A third set of evidence comes from the fact that if we are correct that this asymmetric spending behavior is driven by liquidity constraints rather than by sample selection, then we should see larger asymmetries for households that are more constrained (or more likely to be constrained). And this larger asymmetry should be driven by larger spending responses for households that are more likely to be liquidity constrained rather than by different responses to payments. In section 8, we investigate heterogeneity in behavior in the population by differences in the strength or likelihood of liquidity constraints and confirm these predictions. These findings provide further evidence that our results are not driven by endogenous sorting across tax status.

In sum, the differential spending response to refunds and payments does not appear to be due to persistent differences between households that receive payments and those that make payments. These findings constitute evidence against our results being driven by the endogeneity of tax status.

6.3 Are Results Due to the Endogeneity of Filing or Payment Date?

Although a priori unlikely, could the asymmetry in response be driven by the difference in when people file rather than by liquidity constraints? For example, people who file later have little time to save to maintain consumption when making their tax

[18] Not to make too much of borderline statistical significance, but this decline in not present in the sample of Figure 3, which is puzzling from a theoretical perspective. The households that make payments and are dropped from this smaller sample are those that expected large refunds. As such, we would expect the whole sample to contain households that were less prepared to make payments and so are more likely to have lowered consumption more.

payment. As with our previous concern, some evidence comes from the heterogeneity in responses across people.

As Figure 6 shows, we find the same asymmetry in spending response when we look separately at households filing in different months. (Figure A.1 shows the same pattern for the baseline sample.) Households filing in April still do not cut consumption when making tax payments and yet do increase spending when receiving refunds. It is also worth noting that these differences in spending behavior that we find across filing months are consistent with liquidity constraints and not with selection, as we argue in Section 8.

Closely related, could the asymmetry in response be driven by the fact that people can always postpone payment until the deadline in April?[19] This possibility seems unlikely for the following reason. Our finding for refunds is that spending increases persistently and starting the day of arrival, for many different subgroups of our sample, even those expecting large refunds ex ante, those filing in different months, and so on. If behavior were symmetric, we would expect to find some decreases in spending on the day of payment in some of these populations.

More specifically, consider households that file in April. For these households, which cannot vary the timing of payment much, we find the same asymmetry (Figures 6.e and 6.f). Households that file in April are not randomly selected. These households may be more liquid, which may be why we see little decline in spending when they make a payment. But if this were true, we would also expect little increase in spending when they receive a refund. In fact, we find a substantial increase in spending the day the refund arrives, although smaller than for households that file earlier, which is consistent with the date of filing being partly driven by differences in liquidity.

Alternatively, we might expect to find no asymmetry for households expecting large refunds or expecting small refunds or expecting to have to make payments. But in each case, we find an asymmetry (see Figure 11 in section 8 for the response to refunds and Figure A.3 for the response to payments).

6.4 Are Results Due to Mismeasurement of the Arrival of Information?

The consumption-spending response to refunds and payments are estimated from regressions that include controls for a distributed lag of filing date and its interactions with the dollar amounts of positive and negative news about tax cash flow. But, presumably, we measure both the timing and amount of news with error. Could the different responses we uncover actually be due to news about refunds rather than the cash flow caused by arrival? In particular, roughly half of payments are made on

[19] Of households that file in February and owe taxes, 28% pay in April, and the average time between filing and payment is 20 days. Of those that file in March, 45% pay in April and the mean time between filing and payment is 11 days.

the day of filing or the day after, whereas no refunds arrive on the day of or after filing, and most arrive after two or more weeks delay.

Focusing first on timing, our investigation of other dependent variables provides evidence that our use of daily variation is sufficient to correctly measure the response to cash flows. We estimated equation (1) with two different dependent variables: tax refunds less payments and the tax-filing fee. Appendix Figure A.2 panels a to e show that our estimation methodology measures the effect of the news and cash flows on the tax-induced cash flows with near perfect accuracy. We correctly estimate that the effect of a refund (payment) is a one-time permanent cumulative increase (decrease) in tax-related cash flow equal to 100% of the refund (payment) on the day of the refund (payment) (panels d and e). And we find no measured changes related to the timing of filing (panels a to c with very small scales). We do find slightly less accurate measurement of the effect of the filing fee, but these effects are fractions of a percent, which is very small relative to the spending responses displayed in Figure 3.[20] As an aside, we note that we found that impulse responses smoothed to be the same across weeks of event time did not cleanly separate the effects of news from cash flow in these two regressions. Thus we find significant benefits of high-frequency data for identification.

Despite this evidence, could mismeasurement of the timing of the arrival of news bias upward our estimated spending responses to a refund relative to a payment? If people increase spending in response to good news, and goods news is primarily associated with refunds, we could exaggerate the spending response to refund arrival. But this scenario is unlikely. Our sample always contains a temporal delay between filing and refund, so we have much more power to separately identify the response to news and the response to cash flow for refunds. We also precisely identify the increase in spending on the day of refund arrival.

By contrast, payments are on average associated with bad news. If our statistical procedure attributed the spending decline in response to bad news to the spending decline in response to making a payment, we would be biased toward finding larger spending reductions in response to making payments, and not the small, insignificant, and often positive spending changes we find. Thus, this type of bias cannot account for the non-response of spending to making a payment.

[20] The small effects occur presumably because the cross-sectional heterogeneity in filing fee is only somewhat related to the timing and/or amount of news and cash flow. Panel (f) shows that the filing fee is estimated to rise by $45 on the day of filing, or almost exactly the average filing fee, but the effect is not a permanent cumulative impact – instead the effect is estimated to decay over time. Also, a small amount of the payment is estimated to be an effect of the arrival of news (less than one quarter of a percent of the news in dollars), and a small persistent amount is estimated to be due to the cash flow of making a payment (only -20 basis points of the payment).

To further rule out these concerns, we examined two specifications that include the additional controls related to the timing of filing. Specifically, we add to equation (1) the refund amount and payment amount interacted with the distributed lag of filing. The estimated responses to refunds and payments are almost the same in this specification as in our baseline specification.

We conclude that our main finding is unlikely due to mismeasurement bias: the consumption responses to cash flows are asymmetric. People increase expenditures on consumption substantially after refunds arrive, but do not reduce expenditures when and after they have to make a payment.

7 The Effect of Refunds and Payments on Other Account Inflows and Outflows

Miscellaneous net payments are uncategorized outflows minus uncategorized inflows. This category includes checks, transfers, and expenditures that are not readily classified into other categories as income, consumption, interest, or tax payments. Figures 7.a and 7.b show that miscellaneous expenditures respond with a similar temporal pattern and magnitude as consumption expenditures, rising on impact, increasing to about 20% of the refund after a month and to more than 30% of the refund by the end of six months. In contrast to consumption, miscellaneous expenditures fall before a payment is made, by about 10% of the payment, and then continue to decline by another 10% after payment, though remaining statistically insignificant. This is clearly consistent with inflows coming into the account ahead of making a payment, possibly representing transfers from less liquid accounts or payments for occasional work that we cannot confidently categorize as income or dissaving.

Note that these responses only include the change in behavior related to the timing of the payment (we discuss the changes in account flows related to the information about the refund of payment in section 9). Thus, nearly two-thirds of refunds end up being spent on consumption and uncategorized outflows, and roughly 10% to 20% of payments are met by reductions in miscellaneous expenditures prior to and after making payment.

Figures 7.c and 7.d show the response of saving and debt payment, made up primarily of payments to (linked) credit card accounts, loan payments, and transfers to (identified) investment accounts. Refunds lead to small increases in net savings. About 5% of a refund goes to saving or debt payment after a week, an amount that remains steady over the four months. Figure 7.d shows that we have very little ability to measure the response to making a payment, but what evidence we have suggests that households borrow or save less following making a payment. We find some evidence of reduced payment of credit cards, loans, or transfers to saving beforehand. Furthermore, we find no detectable changes in household income to receiving a refund or making a payment (Figures 7.e and 7.f). These last two results suggest some households smooth spending better than others and suggest an important role of liquidity, issues to which we now turn.

8 Theories of Behavior and Heterogeneous Responses

Our main finding of asymmetric response is consistent with consumption smoothing in the face of restrictions on, or high costs of, borrowing or accessing less-liquid wealth. This section first investigates further the roles of liquidity constraints and then considers the role of near-rationality.

If financial constraints are driving the asymmetric response to tax refunds and payments, we expect the spending responses of households to be larger for households that have lower account balances or are borrowing at high interest rates on their credit cards. We do not directly observe account balances or credit card limits, so we construct net interest earnings as interest earned on accounts less interest paid on credit cards. We drop accounts for which we do not see any interest earned or paid.

Figure 8 displays the heterogeneity in spending response by this measure of ex-ante account liquidity for our baseline sample. Households in the bottom two-thirds of the distribution of liquidity show similar and large consumption responses to refunds (significant), and no measurable response to making payments. However, households in the top third of the liquidity distribution are almost unresponsive to both receiving refunds and making payments. In sum, the spending response is asymmetric for the bottom two-thirds of the liquidity distribution, and the spending response to refunds is decreasing in liquidity.

A commonly used proxy for liquidity constraints is low income. Splitting households by income during the three months prior to February of each year shows stronger spending responses to refunds for people in the bottom two-thirds of the income distribution (Figure 9). We see no measurable declines in consumption spending when making a payment for any of our three income groups. Again, both findings are consistent with the presence of liquid constraints.

A slightly different way to address a similar question is to ask whether households with credit cards, and therefore with access to credit, respond less strongly to arrival than those without credit cards. This measure is an imperfect one, in that households with cards may be at or near their borrowing limits, and households with and without cards might be different in other ways. Nonetheless, Figure 10 shows three things. First, the asymmetry occurs for both households with credit cards (16% of the sample) and those without (84% of the sample). Both types of households do not cut spending around the time of a payment, and both types of households increase spending in response to receiving a refund. Second, both types of households spend cumulatively roughly the same amount from their refunds. But, third, a slight difference exists in the pattern of spending. Households with credit cards spend a few percent of the refund in the days before it arrives, whereas households without cards do not.

An additional prediction of the rational model is that people who are short of liquidity and expect a refund should file earlier. As noted in section 6.3, the timing of filing is not exogenous and can provide information about which households are likely constrained. The prediction of asymmetric spending responses that we are focusing on comes not only from the presence of low liquidity and borrowing constraints, but also from the optimization of people trying to keep consumption stable. Thus, optimizing behavior in the face of credit constraints predicts larger and potentially more immediate spending responses among people who file earlier. Figure 6 shows this pattern. Consistent with this theory, households that file in February have higher propensities to spend from refunds than those that file in March, which in turn have higher propensities to spend than those that file in April.

Similarly, households with little liquidity should file earlier in case they owe taxes, so that they have time to accumulate the liquidity before they have to make the payment at the tax deadline. If households failed to take this liquidity consideration into account, we would expect to see (larger) spending reductions for households that file near the deadline and have to make payments. Figures 6.b, 6.d, and 6.f show no evidence that the propensity to cut spending in response to making a payment increases as households file closer to the deadline.

In sum, these findings are consistent with household optimization in the face of liquidity constraints.

We now discuss two predictions of other models and find household behavior is not consistent with these alternative theories. First, the differences in spending responses by month constitute evidence against a behavioral theory in which some people have self-control problems that lead them to both procrastinate filing and accumulate little liquidity. Filing at the deadline is not associated with greater spending from refunds or cutting back more in response to payment.[21] In fact, people who file in February spend the most, and, although not statistically significant, are

[21] Another related theory is that households that have time-consistency problems are sophisticated about these problems, i.e., understand their bias and act to correct it. In this case, households with time-consistency problems value the commitment of filing later (rather than simply always intending to file tomorrow and failing to do so until the deadline). The prediction for naïfs or sophisticates is the same: People who file later are those most likely to spend when a refund arrives.

estimated to cut back on spending when they make their payment (whereas people filing near the deadline seem to smooth better through their payment).

Second, Kueng (forthcoming) shows that households in Alaska do not smooth predictable payments from the Alaska sovereign wealth fund, and that larger payments are better smoothed, consistent with near-rational behavior. Even if households were naively not smoothing refunds, larger payments should have lower spending responses given the concavity of the consumption function. We find no evidence of either behavior. Figure 11 divides households by their expected refunds and shows that spending out of refunds is independent of the expected size of the refund across groups.[22]

9 The Limited Spending Response to Filing and Information about Taxes

Our dataset allows us to examine also the behavior of households around the tax filing date, when the news about the refund or payment was revealed. In most of our previous analysis, we have controlled with distributed leads and lags of (i) a filing indicator, (ii) the amount of news about taxes if positive, and (iii) the amount of news if negative. Now we present the estimated coefficients on these controls, and so characterize how the household spending on consumption reacts to the information uncovered prior to and at filing.

Forward-looking households should on average consume more in response to good news about refund or payment and consume less in response to bad news about refund or payment. As for the responses to refund arrival or payment, the average estimated responses are changed if some households face (possibly) binding liquidity constraints. But liquidity constraints do not cause as significant an asymmetric response to news. First, a household facing a tightly binding constraint before and after news cannot spend more or less as they learn about the size of this refund, and so has a symmetric non-response. A continuously unconstrained household can respond equally to good and bad news, which is also symmetric. Second, the response is also presumably small if the unconstrained household follows the permanent-income hypothesis.

Finally, the responses may exhibit asymmetry for news that causes a household to move between constrained and unconstrained status, or that causes changes in the likelihood of constrained status in the near future. Households that are constrained today can cut spending in response to large enough negative news but cannot respond to positive news. Households that have binding (or probabilistically binding)

[22] Because payments are on average only about one-third as large as refunds, our main asymmetry is also evidence against larger responses for larger changes in liquidity. But, due to the concavity of utility, the utility costs of not smoothing declines in spending are also larger than the costs of not smoothing increases, hence does not serve as clear evidence prima facie.

constraints when they make their payments can respond to good and bad news about their payments. The response to large positive news is limited by the relaxation of the constraint, whereas the response to negative news is always complete. Households that are unconstrained can become constrained in response to large enough negative news, which therefore amplifies the reduction in their spending. No amount of good news can lead to a larger response. In sum, we expect a stronger reaction of spending to bad news than to good, which is the reverse of the reaction to liquidity.

Our first set of results, based on estimation of equation (1), finds evidence consistent with households being either too constrained to respond or too unconstrained to react much to news. Figures 12.a to 12.i show all impulse responses estimated using our baseline sample. The first column shows estimates from our baseline specification. The second shows a similar analysis but without the distributed lag of the day of filing indicator. The two exercises reveal similar estimates for the responses of interest.

We find a general tendency for consumption to increase before, but also after, filing. Panel a shows that consumption spending increases on average by about $100 over the month before filing and continues to increase steadily at a slightly higher rate than $300/month the four months after filing. Given that on average households receive neither good nor bad news about their tax status over this period, this pattern is not a response to outcomes that are better or worse than expected. If the effect were only prior to filing, an average increase in consumption could be due to precautionary saving. As uncertainty is resolved, spending would on average rise. But the fact the increase continues well after filing undermines this interpretation. Moreover, this finding is robust across many variations in specification. But note that we have confirmed all of our results are robust to whether or not this filing indicator is included in the regression as a control. The robustness of the results is evident in the second column of Figure 12.

Figures 12.b, 12.c, 12.d, and 12.e show no economically significant change in spending in the period before filing related to the size of the news uncovered during the preparation of taxes prior to filing. Some reaction occurs after filing, but not in the direction predicted by theory. Figures 12.b and 12.c show a small *reduction* in spending following *good news* about refund less payment due. Note the vertical range of the figures is the same as that for refund and payment – 60% of the news uncovered during filing – but news has about half the variance as refund minus payment. Figures 12.f through 12.i show our main results (as in Figure 3).

Could these results be due to biased expectations on the part of households? An arbitrary pattern of bias could lead to arbitrary bias in the effect of news and filing on spending. However, if the bias has a central tendency, this average bias would lead to a spending response to filing. Pessimism, like precautionary saving, would appear as an average increase in spending around filing as households get good news that they are receiving more money than expected. Such potential pessimism could explain the average increase in spending around the time of filing. But evidence suggests households have reasonably accurate and unbiased estimates of taxes (Smeeding, Phillips, and O'Connor, 2000; Jones, 2012; Porto and Collins, 2017; Caldwell, Nelson, and Waldinger, 2018).

We next focus in on households in the bottom quintile of expected refund less payment (whose responses to refunds and payments are displayed in Figure 4), whom we expect to be more likely to be able to respond to news about taxes. Figure 13.a finds no noticeable response to good news, but 13.b displays a 5% decrease in response to bad news, as predicted by the theory, but not close to statistically significant and somewhat after the news is revealed (prior to and at the date of filing).[23]

The second row of Figure 13 shows the response to news for households that file in April. These households have little time to save to make payments, and so we might expect them to make larger spending adjustments in response to bad news. We find no evidence that these households increase spending in response to good news, and no evidence that they cut spending in response to bad, at least for the first month after filing (in April). Of course, this subsample represents a select group of households. Only liquid households should wait to file until April, because they would have little time to save to make a payment should they receive bad news and owed taxes. Another possibility is that households with better access to credit adjust their spending in response to news about refund less payment. Figures 13.e and 13.f show the response among households with credit cards, which again shows a small decrease in spending ahead of filing in response to good news, which grows after filing, and a very similar but statistically insignificant response to bad news.

We find similar weak evidence of responses and little evidence for any asymmetry among other subgroups of households where responses to news could be particularly large. These subgroups include groups of more illiquid households – those with low incomes, those filing in February, those without credit cards, those expecting large refunds – and groups of more liquid households – those with high incomes, those in the top third of the distribution of net interest, those not expecting large refunds. The spending response of households in the bottom third of the distribution of net interest appears similar to the spending of households with credit cards.

[23] We also find a more precisely estimated non-response to good news and a slight increase in spending before and (only) shortly after filing in the larger sample of accounts that allows unlinked credit cards but requires that households receive refunds in some years and have taxes due in other (as in panels c and d in Figure 5).

We conclude that the lack of spending responses by households to news about their refund or taxes is generally consistent with the rational model with liquidity constraints. However, the evidence we are missing is a spending response to news by households making payments. Households making payments, particularly those with little liquidity or with little time to react, should adjust spending in response to news. The missing evidence may be due to the inability to detect responses in the subset of households that are (likely) constrained at the date of payment. Less than 10% of our sample pays taxes, the average payment is one-third the size of the average refund, and our measures of the amount and timing of news are less precise than our measures of cash flow.

This last factor – mismeasurement – raises concerns about our measurement of the spending response to expected cash flows. But, as noted, these results are robust to a variety of checks that allow for much more flexible and extensive controls for the arrival of information, such as including distributed lag and lead polynomials that interact news amounts with the date of cash flow and that interact cash-flow amounts with the date of filing. We also note that the institutional setting is one where refunds arrive with significant delay after filing, so that our daily analysis, and the precise effect of the day of refund arrival, strongly suggest we are correctly measuring the effect of cash flow and not confounding the effects of news with cash flow or cash flow with news.

10 Final Discussion

We conclude that strong evidence suggests household behavior around tax refunds and payments is well described by intertemporal optimization and the presence of liquidity constraints. Households spend from expected income and do not cut back on spending when they determine that they must make expected payments. This behavior is not consistent with near-rationality or mental accounts unless mental accounts are assumed to differ by the direction of cash flow.

References

Aaronson, Daniel, Sumit Agarwal, and Erik French, 2012, Spending and Debt Response to Minimum Wage Hikes, *American Economic Review* 102(7), 3111–3139.

Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles, 2007, The Reaction of Consumer Spending and Debt to Tax Rebates, *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), 4205-4230.

Aguiar, Mark and Erik Hurst. 2005, Consumption versus Expenditure, *Journal of Political Economy* 113(5), 919-948.

Baker, Scott, forthcoming, Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data, *Journal of Political Economy*.

Bodkin, Ronald, 1959, Windfall Income and Consumption, *American Economic Review* 49(4), 602–614.

Caldwell, Sydnee, Scott Nelson and Daniel Waldinger, 2018, Tax Refund Expectations and Financial Behavior, Working Paper, MIT.

Campbell, John Y., and Gregory N. Mankiw, 1989, Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence, in NBER Macroeconomics Annual 1989, Volume 4, Blanchard and Fischer, eds.

DiMaggio, Marco, Amir Kirmani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao, 2017, Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging, *American Economic Review* 107(11), 3550-88.

Ganong, Peter and Pascal Noel, Consumer Spending During Unemployment: Positive and Normative Implications, *American Economic Review*, forthcoming.

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, 212-215.

Hsieh, Chang-Tai, 2003, Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund, *American Economic Review* 93(1), 397–405.

Jappelli, Tulio, Jörn-Steffen Pischke, and Nicholas S. Souleles, 1998, Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources, *Review of Economics and Statistics* 80(2), 251–262.

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.

Jones, Damon, 2012, Inertia and Overwithholding: Explaining the Prevalence of Income Tax Refunds, *American Economic Journal: Economic Policy* 4(1), 158–85.

Jorring, Adam T., The Cost of Financial Mistakes: Evidence from U.S. Consumers, University of Chicago Manuscript, 2018.

Kueng, Lorenz, forthcoming, Excess Sensitivity of High-Income Consumers, *Quarterly Journal of Economics*.

Olafsson, Arna and Michaela Pagel, forthcoming, The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software, Review of Financial Studies.

Park, Hoonsuk, 2017, Loss Aversion and the Consumption Response to Adjustable-Rate Mortgage Payments, Working Paper.

Parker, Jonathan A., 1999, The Reaction of Household Consumption to Predictable Changes in Social Security Taxes, *American Economic Review* 89(4), 959–973.

Porto, Nilton and J. Michael Collins, 2017, The Role of Refund Expectations in Savings: Evidence from Volunteer Income Tax Preparation Programs in the United States, *The Journal of Consumer Affairs* 51(1), 183-199.

Reis, Ricardo, 2006, Inattentive Consumers, *Journal of Monetary Economics*, 53(8), 1761-80.

Shea, John, 1995, Myopia, Liquidity Constraints and Aggregate Consumption: a Simple Test, *Journal of Money, Credit and Banking* 27(3), 798-805.

Slemrod, Joel B., Charles Christian, Rebecca London, and Jonathan A. Parker, , 1997, April 15 Syndrome, *Economic Inquiry* 35(4), 695-709.

Smeeding, Timothy M., Katherin Ross Phillips, and Michael O'Connor, 2000, The EITC: Expectation, Knowledge, Use, and Economic and Social Mobility, *National Tax Journal*, 53(4), 1187-210.

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

Souleles, Nicholas S., 2000, College Tuition and Household Savings and Consumption, *Journal of Public Economics*, 77, 185–207.

Stephens, Melvin, 2003, "3rd of tha Month": Do Social Security Recipients Smooth Consumption between Checks? *American Economic Review* 93(1), 406–422.

Stephens, Melvin, 2006, Paycheque Receipt and the Timing of Consumption, *Economic Journal* 116(July), 680–701.

Thaler, Richard H., Anomalies: Saving, Fungibility, and Mental Accounts, *Journal of Economic Perspectives* 4(1), 193–205.

Zeldes, Stephen P., 1989, Consumption and Liquidity Constraints: An Empirical Investigation, *Journal of Political Economy* 97(2), 305–346.

Table 1. Geographic Distribution of the Sample

This table shows the geographic distribution of the households in the sample relative to the 2010 US Census.

State	% Households Residing			State	% Households Residing		
	Data	U.S. Census	Data - Census		Data	U.S. Census	Data - Census
Alabama	0.4%	1.5%	-1.2%	Montana	0.1%	0.3%	-0.2%
Alaska	0.2%	0.2%	0.0%	Nebraska	0.1%	0.6%	0.6%
Arizona	1.6%	2.1%	-0.4%	Nevada	1.7%	0.9%	0.9%
Arkansas	0.3%	0.9%	-0.6%	New Hampshire	0.2%	0.4%	0.4%
California	21.6%	12.1%	9.5%	New Jersey	1.7%	2.8%	2.8%
Colorado	0.5%	1.6%	-1.1%	New Mexico	0.4%	0.7%	0.7%
Connecticut	1.0%	1.2%	-0.1%	New York	16.6%	6.3%	6.3%
Delaware	0.1%	0.3%	-0.1%	North Carolina	2.3%	3.1%	3.1%
District of Columbia	0.2%	0.2%	0.0%	North Dakota	0.1%	0.2%	0.2%
Florida	8.7%	6.1%	2.6%	Ohio	0.5%	3.7%	3.7%
Georgia	3.4%	3.1%	0.3%	Oklahoma	0.7%	1.2%	1.2%
Hawaii	0.3%	0.4%	-0.1%	Oregon	0.6%	1.2%	1.2%
Idaho	0.1%	0.5%	-0.4%	Pennsylvania	1.2%	4.1%	4.1%
Illinois	5.5%	4.2%	1.3%	Rhode Island	0.3%	0.3%	0.3%
Indiana	0.4%	2.1%	-1.7%	South Carolina	1.2%	1.5%	1.5%
Iowa	0.2%	1.0%	-0.8%	South Dakota	0.1%	0.3%	0.3%
Kansas	0.7%	0.9%	-0.2%	Tennessee	1.1%	2.1%	2.1%
Kentucky	0.3%	1.4%	-1.1%	Texas	14.6%	8.1%	8.1%
Louisiana	0.4%	1.5%	-1.0%	Utah	0.2%	0.9%	0.9%
Maine	0.1%	0.4%	-0.3%	Vermont	0.0%	0.2%	0.2%
Maryland	1.9%	1.9%	0.0%	Virginia	2.5%	2.6%	2.6%
Massachusetts	1.9%	2.1%	-0.2%	Washington	1.0%	2.2%	2.2%
Michigan	1.0%	3.2%	-2.2%	West Virginia	0.1%	0.6%	0.6%
Minnesota	0.2%	1.7%	-1.5%	Wisconsin	0.2%	1.8%	1.8%
Mississippi	0.2%	1.0%	-0.7%	Wyoming	0.0%	0.2%	0.2%
Missouri	0.9%	1.9%	-1.1%				

TABLE 2. SUMMARY STATICS

This table shows the basic summary statistics for the variables in our various samples. *Baseline* denotes our baseline sample in which the filing date is required, whereas *Unrestricted* denotes our expanded sample, which does not require the observation of filing date. *Households* is the number of households in the given sample, whereas *Household-Years* is the number of household-years in the sample. *Filing Date* and *Refund Date* are the dates of filing and refund, respectively. *Pos Refund* is an indicator that takes the value of 1 if the refund is positive, and zero otherwise. *Refund Amount* is the size of the annual refund or payment, with refunds taking a positive sign and payments taking a negative sign. *Lag Refund Amount* is defined similarly. *Predicted Refund* is our prediction of the current year□s refund amount, which we arrive at by pooling all observations and regressing refund amount on lagged refund amount, an interaction term of lagged refund amount with an indicator for a positive refund, and an interaction term of lagged refund amount with an indicator for a tax payment. *Surprise* is defined as the actual refund amount minus the expected refund amount. *Disatance Filing Refund* is defined as the distance from filing date to refund or payment date. *Linked CC* is an indicator that takes the value of 1 if the household has linked credit cards, whereas *No CC* is an indicator that takes the value of 1 if the household has zero linked credit cards. *Avg Monthly Income* is the average of monthly income in the first three months of each household year (November, December, January), conditional on non-zero values. *Avg Monthly Net Interest* is the average of monthly net interest received in the first three months of each household year (November, December, January), conditional on non-zero values. *Net Flow* is the net daily inflow received across all accounts. *Consumption* is the observed daily consumption across all accounts and includes categories such as gas, restaurant, retail, groceries, cash, entertainment, healthcare, travel, utilities, miscellaneous bills, and insurance. *Savings and Loans* is the observed daily flows to the following categories: mortgages, auto loans, net investing (flows to investing accounts □ flows from investing accounts), credit card repayments (credit card payments minus net credit card expenditures), and other loan repayments (e.g., student loans). *Miscellaneous Payments* is the observed daily values for miscellaneous payments not clearly categorized into either consumption or savings. This variable is equal to the sum of checks and net uncategorized transactions (uncategorized inflows □ uncategorized outflows). Lastly, *Income* denotes daily flows of income. *Panel A* shows our baseline sample, *Panel B* shows our unrestricted sample, *Panel C* shows the subsample of our baseline sample without credit cards, *Panel D* shows the subsample of our baseline sample with linked credit cards, *Panel E* shows the subsample of our baseline sample that receives a refund, and *Panel F* shows the subsample of our baseline sample that makes a tax payment.

Panel A: Baseline sample with observed filing date

Variable	count	mean	sd	p1	p10	p25	p50	p75	p90	p99
Households	11,138	-	-	-	-	-	-	-	-	-
Household-Years	15,456	-	-	-	-	-	-	-	-	-
Filing Date	15,456	Feb 27	29.89	Jan 13	Jan 26	Feb 03	Feb 19	Mar 30	Apr 15	Apr 20
Refund Date	15,456	Mar 10	29.02	Jan 29	Feb 05	Feb 12	Feb 28	Apr 09	Apr 22	May 06
Pos Refund	15,456	0.94	0.24	0.00	1.00	1.00	1.00	1.00	1.00	1.00
Refund Amount	15,456	2,169.51	2,447.64	-1,937.00	121.00	628.00	1,350.50	3,214.03	5,664.10	9,450.00
Lag Refund Amount	15,456	2,212.90	2,450.07	-1,479.00	192.59	634.00	1,334.05	3,266.02	5,697.00	9,700.00
Predicted Refund	15,456	2,327.42	1,768.10	-204.22	839.41	1,163.25	1,676.85	3,094.25	4,877.75	7,814.58
Surprise	15,456	-157.92	1,676.34	-4,934.51	-1,591.49	-782.59	-337.41	416.54	1,639.14	4,908.22
Distance Filing Refund	15,456	10.29	9.68	-1.00	4.00	6.00	8.00	12.00	18.00	58.00
Linked Cc	15,456	0.16	0.37	0.00	0.00	0.00	0.00	0.00	1.00	1.00
No Cc	15,456	0.84	0.37	0.00	0.00	1.00	1.00	1.00	1.00	1.00
Avg Monthly Income	11,961	4,648.29	3,388.40	126.84	1,180.06	2,425.61	3,977.25	6,026.06	8,776.47	16,682.64
Avg Monthly Net Interest	7,259	1.89	13.09	-20.74	0.01	0.02	0.07	0.50	4.01	50.75
Net Flow	5,124,640	2.17	547.42	-1,283.10	-225.00	-70.45	-4.45	0.00	70.10	2,020.49
Consumption	5,124,640	64.35	156.23	0.00	0.00	0.00	9.82	64.95	176.39	688.26
Savings and Loans	5,124,640	9.23	152.50	-37.30	0.00	0.00	0.00	0.00	0.00	300.00
Misc Payments	5,124,640	-1.44	347.66	-919.38	-0.40	0.00	0.00	0.00	35.36	847.02
Income	5,124,640	67.91	375.31	0.00	0.00	0.00	0.00	0.00	0.00	1,781.14

Table 2. Summary Statics (Cont.)

Panel B: Unrestricted sample without observed filing date

Variable	count	mean	sd	p1	p10	p25	p50	p75	p90	p99
Households	80,747	-	-	-	-	-	-	-	-	-
Household-Years	154,507	-	-	-	-	-	-	-	-	-
Filing Date	15,456	Feb 27	29.89	Jan 13	Jan 26	Feb 03	Feb 19	Mar 30	Apr 15	Apr 20
Refund Date	154,507	Mar 07	27.64	Jan 29	Feb 05	Feb 12	Feb 27	Mar 29	Apr 20	May 12
Pos Refund	154,507	0.97	0.16	0.00	1.00	1.00	1.00	1.00	1.00	1.00
Refund Amount	154,507	2,660.71	2,582.76	-830.00	333.00	836.06	1,743.00	4,166.00	6,451.79	9,899.00
Lag Refund Amount	154,507	2,667.50	2,621.25	-747.00	313.00	815.00	1,719.00	4,172.75	6,479.48	10,191.00
Predicted Refund	154,507	2,659.38	1,903.52	139.74	927.75	1,296.05	1,959.27	3,759.48	5,451.83	8,174.80
Surprise	154,507	1.33	1,741.91	-4,735.85	-1,535.23	-721.99	-261.62	665.00	2,027.10	5,212.75
Distance Filing Refund	15,456	10.29	9.68	-1.00	4.00	6.00	8.00	12.00	18.00	58.00
Linked Cc	154,507	0.16	0.37	0.00	0.00	0.00	0.00	0.00	1.00	1.00
No Cc	154,507	0.84	0.37	0.00	0.00	1.00	1.00	1.00	1.00	1.00
Avg Monthly Income	117,150	4,185.97	3,106.06	96.45	976.85	2,152.47	3,595.20	5,458.68	7,871.64	15,045.65
Avg Monthly Net Interest	67,268	1.79	12.93	-20.28	0.01	0.02	0.06	0.45	3.82	48.77
Net Flow	51,000,000	1.66	517.65	-1,184.63	-208.72	-58.85	0.00	0.00	52.09	1,850.04
Consumption	51,000,000	61.12	159.58	0.00	0.00	0.00	3.19	56.07	169.54	707.64
Savings and Loans	51,000,000	8.04	125.55	0.00	0.00	0.00	0.00	0.00	0.00	230.89
Misc Payments	51,000,000	-2.51	340.39	-880.00	-0.01	0.00	0.00	0.00	32.37	750.00
Income	51,000,000	60.25	336.65	0.00	0.00	0.00	0.00	0.00	0.00	1,619.82

Panel C: Baseline sample with observed filing date and no credit cards

Variable	count	mean	sd	p1	p10	p25	p50	p75	p90	p99
Households	9,387	-	-	-	-	-	-	-	-	-
Household-Years	12,906	-	-	-	-	-	-	-	-	-
Filing Date	12,906	Feb 27	29.79	Jan 13	Jan 26	Feb 03	Feb 18	Mar 27	Apr 15	Apr 21
Refund Date	12,906	Mar 09	28.91	Jan 29	Feb 05	Feb 12	Feb 28	Apr 08	Apr 22	May 06
Pos Refund	12,906	0.94	0.23	0.00	1.00	1.00	1.00	1.00	1.00	1.00
Refund Amount	12,906	2,250.07	2,426.57	-1,430.00	162.10	651.00	1,407.50	3,337.00	5,793.00	9,519.00
Lag Refund Amount	12,906	2,294.62	2,434.50	-956.00	223.00	659.78	1,396.50	3,408.55	5,854.15	9,700.00
Predicted Refund	12,906	2,384.40	1,768.55	41.53	861.72	1,182.17	1,722.67	3,198.82	4,993.05	7,814.58
Surprise	12,906	-134.34	1,640.93	-4,635.27	-1,588.73	-777.36	-326.09	456.30	1,663.29	4,882.08
Distance Filing Refund	12,906	10.29	9.59	-1.00	4.00	6.00	8.00	12.00	18.00	57.00
Linked Cc	12,906	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
No Cc	12,906	1.00	0.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Avg Monthly Income	10,087	4,512.57	3,122.43	129.60	1,165.86	2,420.52	3,938.61	5,920.15	8,401.87	14,875.54
Avg Monthly Net Interest	5,953	1.28	10.83	-19.16	0.01	0.02	0.06	0.36	2.39	39.04
Net Flow	4,272,130	1.41	523.35	-1,223.30	-227.51	-71.71	-4.52	0.00	74.00	1,932.95
Consumption	4,272,130	65.22	155.89	0.00	0.00	0.00	10.00	66.27	179.87	692.69
Savings and Loans	4,272,130	8.97	112.13	0.00	0.00	0.00	0.00	0.00	0.00	260.50
Misc Payments	4,272,130	-2.03	338.49	-925.00	-0.60	0.00	0.00	0.00	37.62	830.00
Income	4,272,130	66.87	361.89	0.00	0.00	0.00	0.00	0.00	0.00	1,735.67

Table 2. Summary Statics (Cont.)

Panel D: Baseline sample with observed filing date and linked credit cards

Variable	count	mean	sd	p1	p10	p25	p50	p75	p90	p99
Households	1,751	-	-	-	-	-	-	-	-	-
Household-Years	2,550	-	-	-	-	-	-	-	-	-
Filing Date	2,550	Mar 03	30.15	Jan 13	Jan 27	Feb 04	Feb 25	Apr 04	Apr 15	Apr 20
Refund Date	2,550	Mar 13	29.37	Jan 30	Feb 06	Feb 14	Mar 08	Apr 13	Apr 22	May 10
Pos Refund	2,550	0.91	0.29	0.00	1.00	1.00	1.00	1.00	1.00	1.00
Refund Amount	2,550	1,761.80	2,512.79	-5,947.00	12.00	547.00	1,153.50	2,618.00	5,156.00	8,993.00
Lag Refund Amount	2,550	1,799.31	2,487.04	-3,908.00	53.29	546.10	1,146.00	2,576.00	5,153.00	9,572.00
Predicted Refund	2,550	2,039.05	1,737.69	-1,345.58	737.21	1,098.77	1,538.89	2,588.01	4,478.65	7,720.67
Surprise	2,550	-277.25	1,840.85	-6,553.55	-1,624.46	-795.85	-373.27	194.58	1,469.27	5,555.87
Distance Filing Refund	2,550	10.29	10.09	-1.00	3.00	6.00	8.00	12.00	18.00	62.00
Linked Cc	2,550	1.00	0.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
No Cc	2,550	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Avg Monthly Income	1,874	5,378.86	4,492.11	113.04	1,245.74	2,436.32	4,181.22	6,761.40	11,129.89	22,398.66
Avg Monthly Net Interest	1,306	4.69	20.21	-42.78	0.01	0.03	0.24	3.09	14.74	87.98
Net Flow	852,510	5.95	654.81	-1,650.00	-212.00	-64.79	-4.03	0.00	51.01	2,539.03
Consumption	852,510	60.01	157.88	0.00	0.00	0.00	7.99	59.07	159.39	664.12
Savings and Loans	852,510	10.51	277.10	-252.12	0.00	0.00	0.00	0.00	0.00	575.64
Misc Payments	852,510	1.50	390.36	-893.06	0.00	0.00	0.00	0.00	28.00	929.00
Income	852,510	73.10	436.34	0.00	0.00	0.00	0.00	0.00	0.00	2,105.52

Panel E: Baseline sample with observed filing date and receiving refund

Variable	count	mean	sd	p1	p10	p25	p50	p75	p90	p99
Households	10,484	-	-	-	-	-	-	-	-	-
Household-Years	14,487	-	-	-	-	-	-	-	-	-
Filing Date	14,487	Feb 25	29.16	Jan 13	Jan 25	Feb 03	Feb 17	Mar 24	Apr 14	Apr 20
Refund Date	14,487	Mar 08	28.66	Jan 29	Feb 05	Feb 11	Feb 27	Apr 03	Apr 22	May 07
Pos Refund	14,487	1.00	0.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Refund Amount	14,487	2,394.72	2,294.33	27.00	363.00	770.00	1,482.00	3,424.00	5,825.00	9,572.00
Lag Refund Amount	14,487	2,349.33	2,384.43	-498.00	298.00	719.00	1,433.00	3,435.47	5,858.00	9,764.00
Predicted Refund	14,487	2,422.73	1,741.28	256.74	916.75	1,225.62	1,749.45	3,218.57	4,995.87	7,861.53
Surprise	14,487	-28.02	1,537.00	-3,904.63	-1,344.02	-691.28	-288.81	491.74	1,737.34	4,983.95
Distance Filing Refund	14,487	10.55	9.21	0.00	5.00	6.00	8.00	12.00	18.00	56.00
Linked Cc	14,487	0.16	0.37	0.00	0.00	0.00	0.00	0.00	1.00	1.00
No Cc	14,487	0.84	0.37	0.00	0.00	1.00	1.00	1.00	1.00	1.00
Avg Monthly Income	11,267	4,524.77	3,216.91	120.05	1,163.90	2,390.33	3,914.80	5,899.29	8,504.62	15,579.17
Avg Monthly Net Interest	6,645	1.63	12.83	-21.29	0.01	0.02	0.07	0.44	3.39	47.20
Net Flow	4,812,039	2.08	532.27	-1,244.07	-222.38	-69.44	-3.99	0.00	71.63	1,965.93
Consumption	4,812,039	63.98	154.96	0.00	0.00	0.00	9.53	64.37	175.90	685.32
Savings and Loans	4,812,039	9.09	141.68	-23.00	0.00	0.00	0.00	0.00	0.00	294.74
Misc Payments	4,812,039	-1.53	337.26	-900.00	-0.50	0.00	0.00	0.00	35.00	820.00
Income	4,812,039	66.56	363.95	0.00	0.00	0.00	0.00	0.00	0.00	1,745.91

Table 2. Summary Statics (Cont.)

Panel F: Baseline sample with observed filing date and making payment

Variable	count	mean	sd	p1	p10	p25	p50	p75	p90	p99
Households	654	-	-	-	-	-	-	-	-	-
Household-Years	969	-	-	-	-	-	-	-	-	-
Filing Date	969	Mar 30	23.73	Jan 26	Feb 18	Mar 18	Apr 11	Apr 16	Apr 17	Apr 21
Refund Date	969	Apr 05	20.43	Jan 31	Feb 28	Apr 03	Apr 15	Apr 17	Apr 18	May 05
Pos Refund	969	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Refund Amount	969	-1,197.44	2,183.93	-11,328.00	-3,090.00	-1,161.00	-420.00	-118.00	-31.00	-2.00
Lag Refund Amount	969	173.18	2,512.93	-9,934.00	-1,469.00	-279.00	225.10	917.00	2,136.00	7,272.80
Predicted Refund	969	902.52	1,539.22	-4,177.13	-199.52	359.65	863.26	1,370.88	2,265.21	6,033.85
Surprise	969	-2,099.96	2,342.09	-12,700.10	-4,375.24	-2,463.24	-1,418.65	-860.32	-527.57	917.54
Distance Filing Refund	969	6.38	14.43	-3.00	-1.00	0.00	1.00	4.00	22.00	70.00
Linked Cc	969	0.25	0.43	0.00	0.00	0.00	0.00	0.00	1.00	1.00
No Cc	969	0.75	0.43	0.00	0.00	1.00	1.00	1.00	1.00	1.00
Avg Monthly Income	694	6,653.73	5,063.30	150.74	1,532.87	3,393.76	5,505.56	8,449.28	13,447.57	26,662.37
Avg Monthly Net Interest	614	4.77	15.33	-8.59	0.02	0.06	0.30	2.65	13.63	83.53
Net Flow	312,601	3.47	742.62	-1,965.99	-266.73	-87.91	-11.11	0.00	44.29	2,953.01
Consumption	312,601	70.16	174.57	0.00	0.00	0.00	14.18	73.20	183.04	731.83
Savings and Loans	312,601	11.35	268.78	-200.00	0.00	0.00	0.00	0.00	0.00	491.04
Misc Payments	312,601	-0.09	480.07	-1,286.00	0.00	0.00	0.00	0.00	43.50	1,200.00
Income	312,601	88.75	519.29	0.00	0.00	0.00	0.00	0.00	0.00	2,518.43

Figure 1: Histograms of Annualized Income

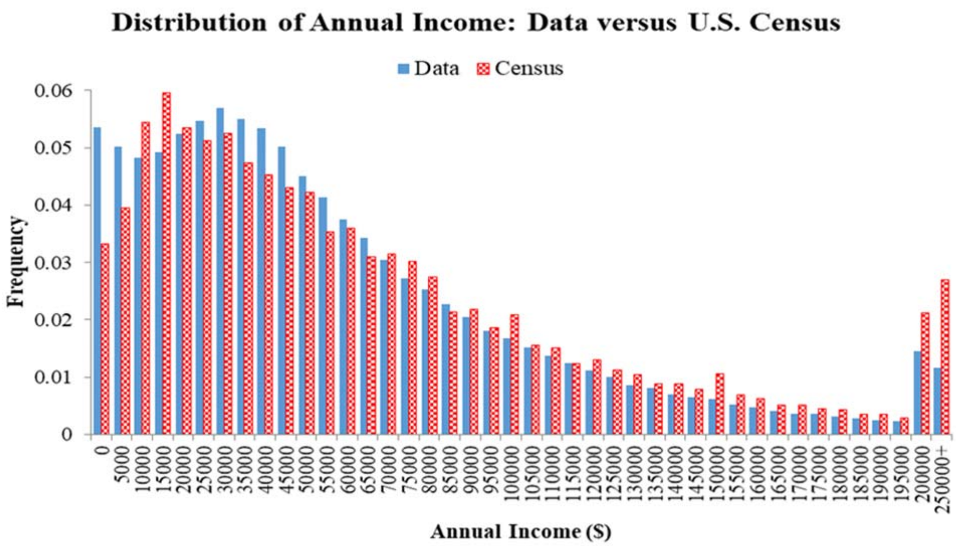

Figure 2: Dates and Amounts of Filing, Refunds, and Payments

a. Density of filing dates

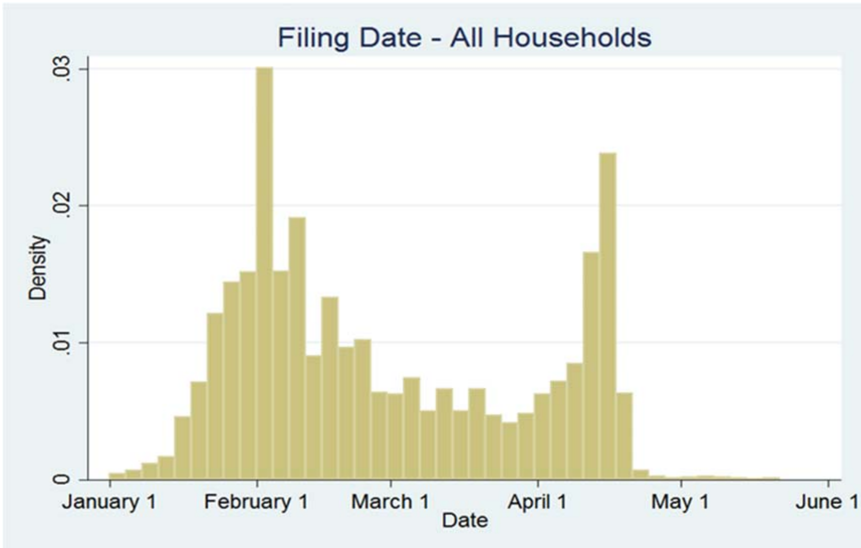

b. Density of filing dates for accounts with taxes due

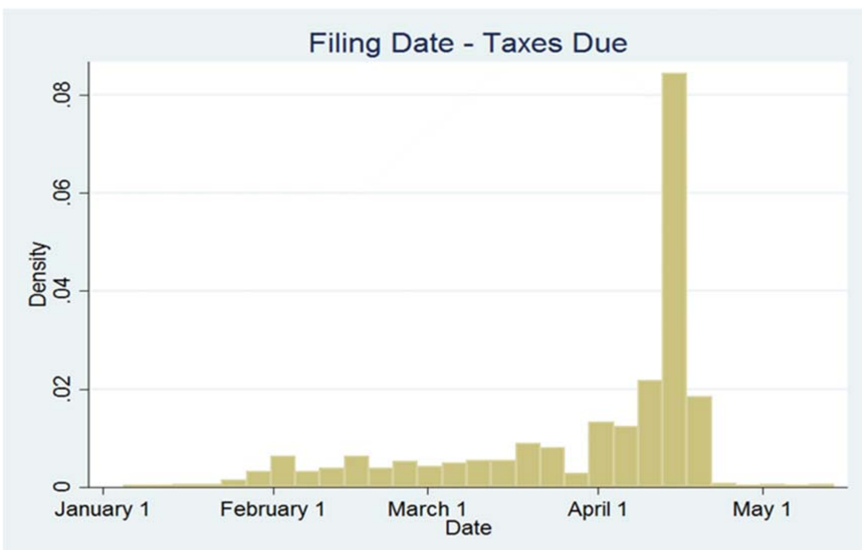

Figure 2: Timing of Filing, Refunds, and Payments (Cont.)

c. Density of days between filing and refund receipt

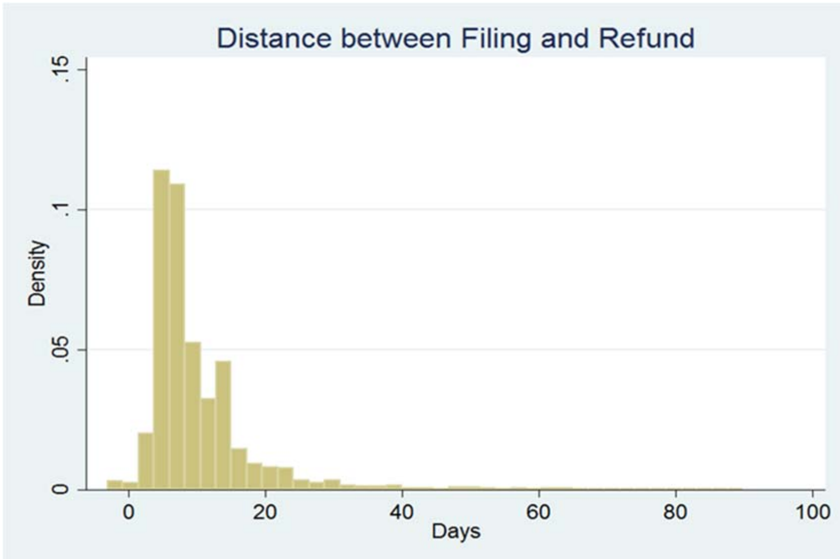

d. Density of days between filing and taxes paid

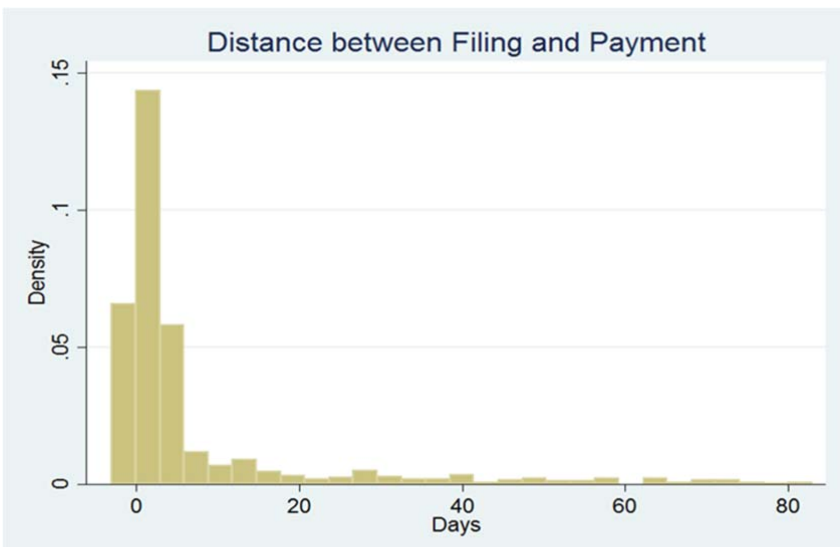

Figure 2: Timing of Filing, Refunds, and Payments (Cont.)

e. Density of refund payment

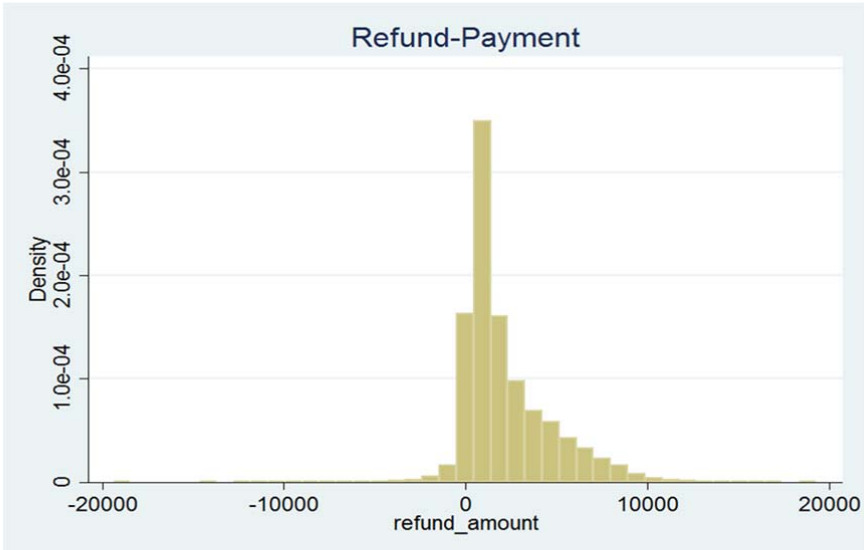

Figure 3: The response of consumption spending to refund and payment

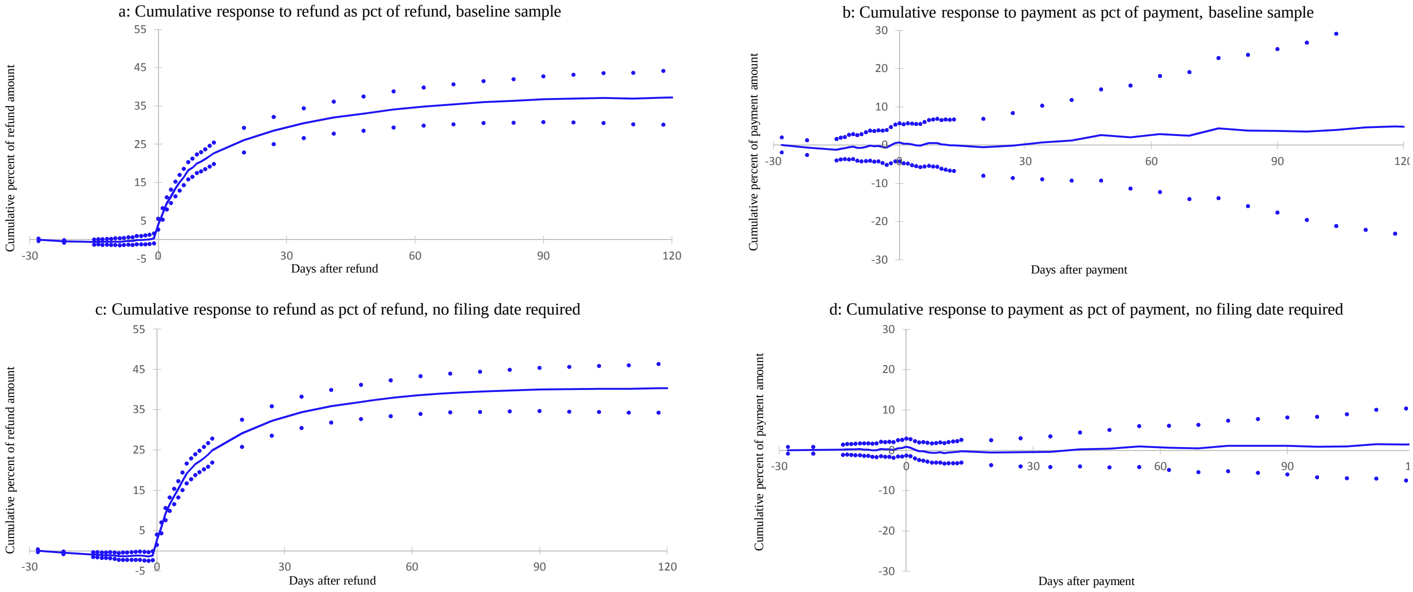

Note: The figures plot the sum of coefficients from estimation of equation (1) with consumption spending as the dependent variable, which measures the cumulative change in consumption spending from 28 days before refund or payment to 120 days afterwards as a percent of refund or payment. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure 4: The response of consumption spending for households expecting payments or small refunds

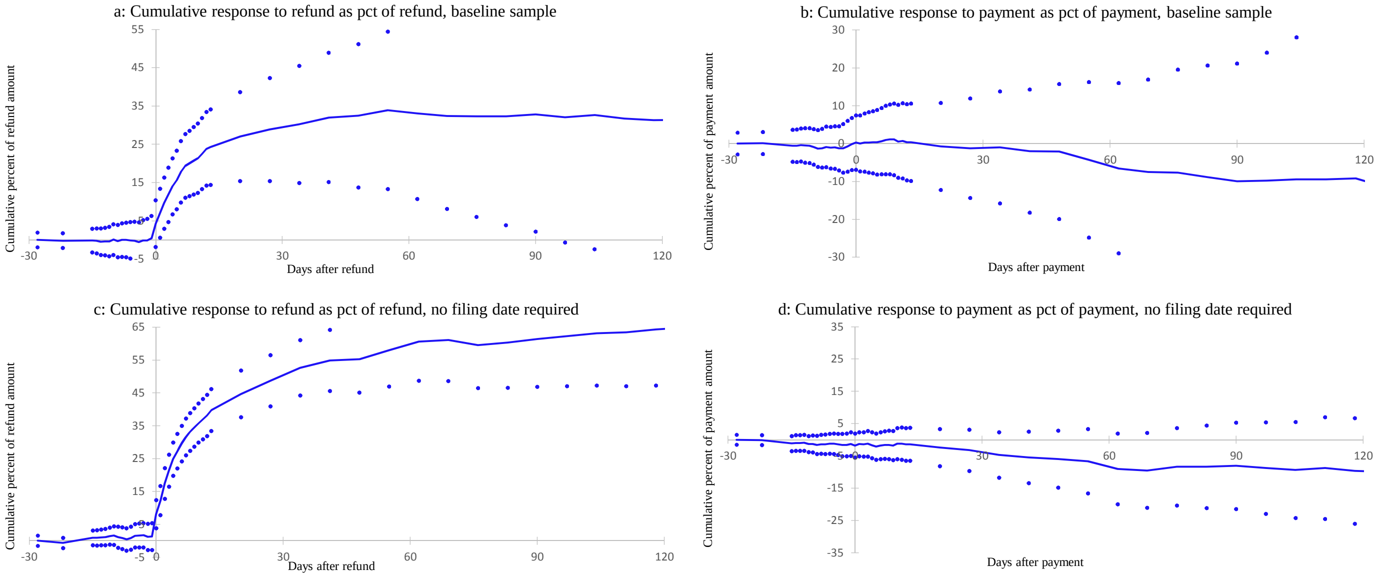

Note: The figures plot the sum of coefficients from estimation of equation (1) with consumption spending as the dependent variable, which measures the cumulative change in consumption spending from 28 days before refund or payment to 120 days afterwards as a percent of refund or payment. Dots show pointwise 95% confidence intervals at kink points of the impulse responses. The sample is the bottom 20% of households by E[refunds].

Figure 5: The consumption responses for households that make payments in some years and receive refunds in other years

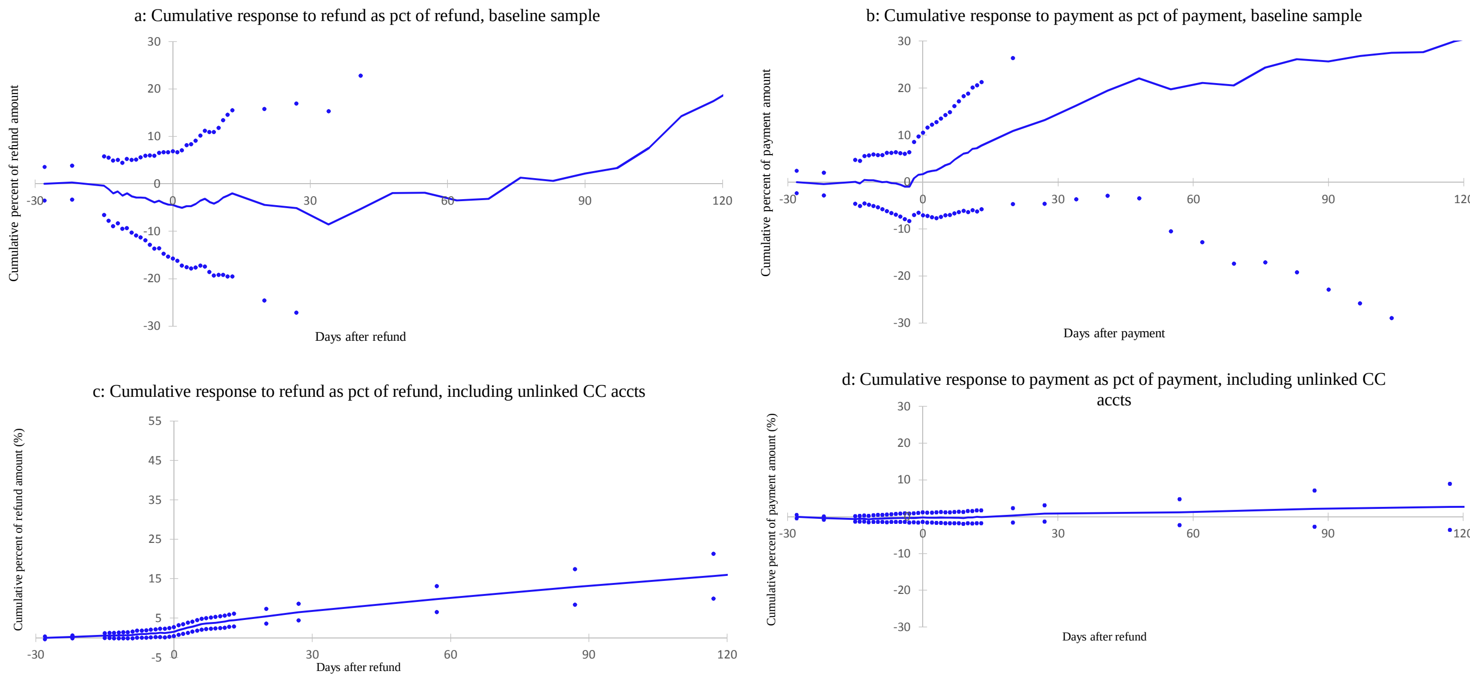

Figure 6: The response of consumption expenditures by month

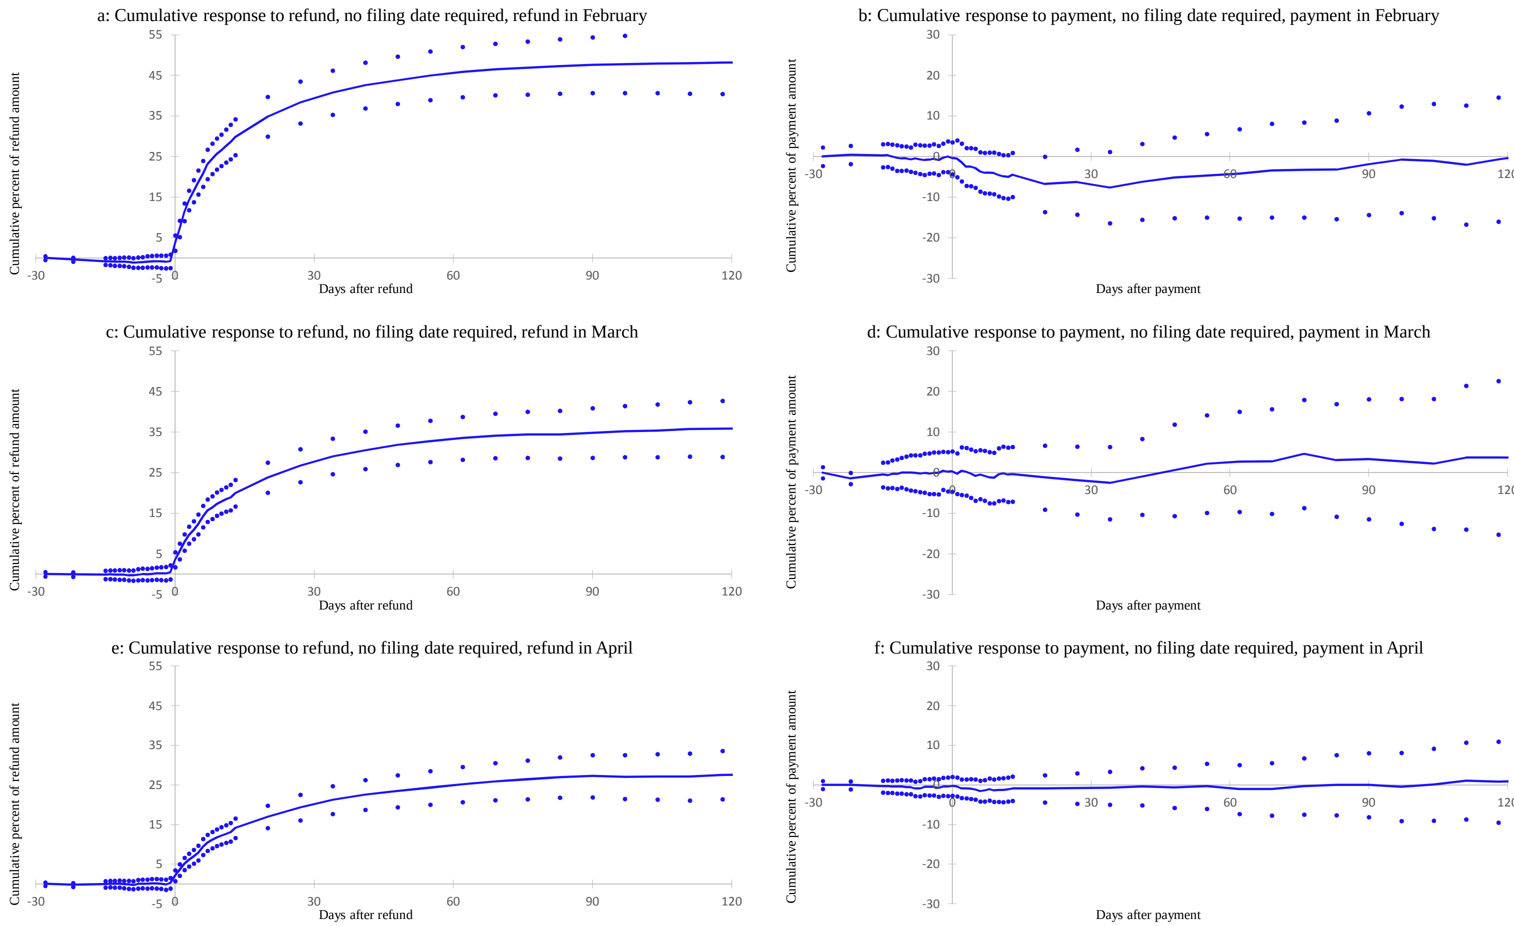

Note: Each row shows the sum of coefficients from estimation of equation (1) with a different sample. The first column displays the cumulative response to a refund, the second column the cumulative response to a payment, in each case from 28 days before refund or cash flow to 120 days after. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure 7: The response of account flows to refunds and payments

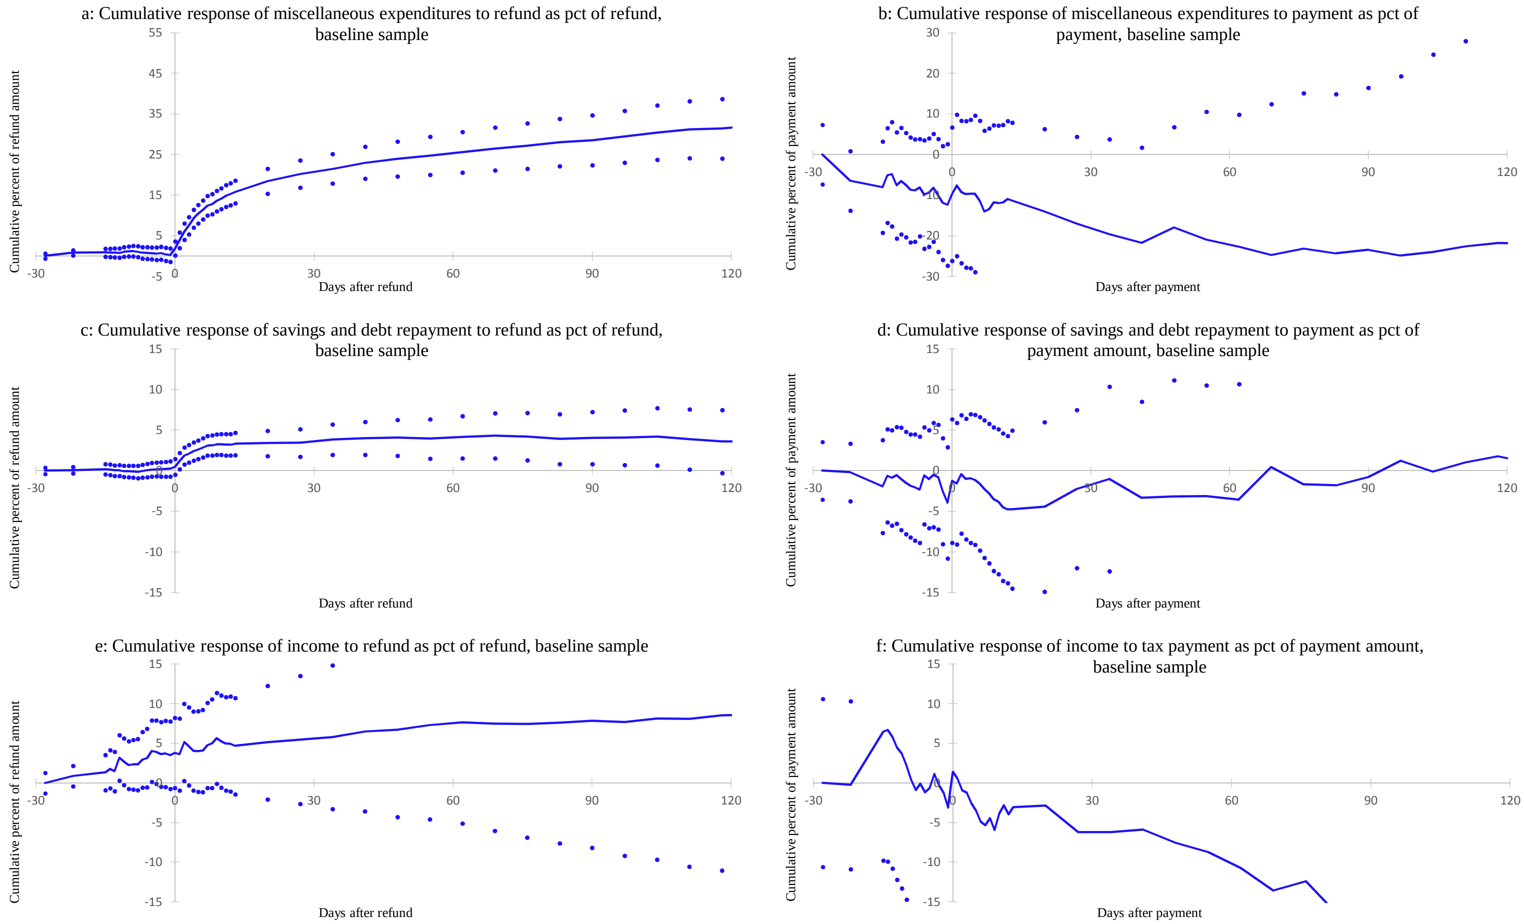

Note: Each row shows the sum of coefficients from estimation of equation (1) with a different dependent variable. The first column displays the cumulative response to a refund, the second column the cumulative response to a payment, in each case from 28 days before refund or cash flow to 120 days after. Note different scales. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure 8: The response of consumption expenditures by net interest, baseline sample

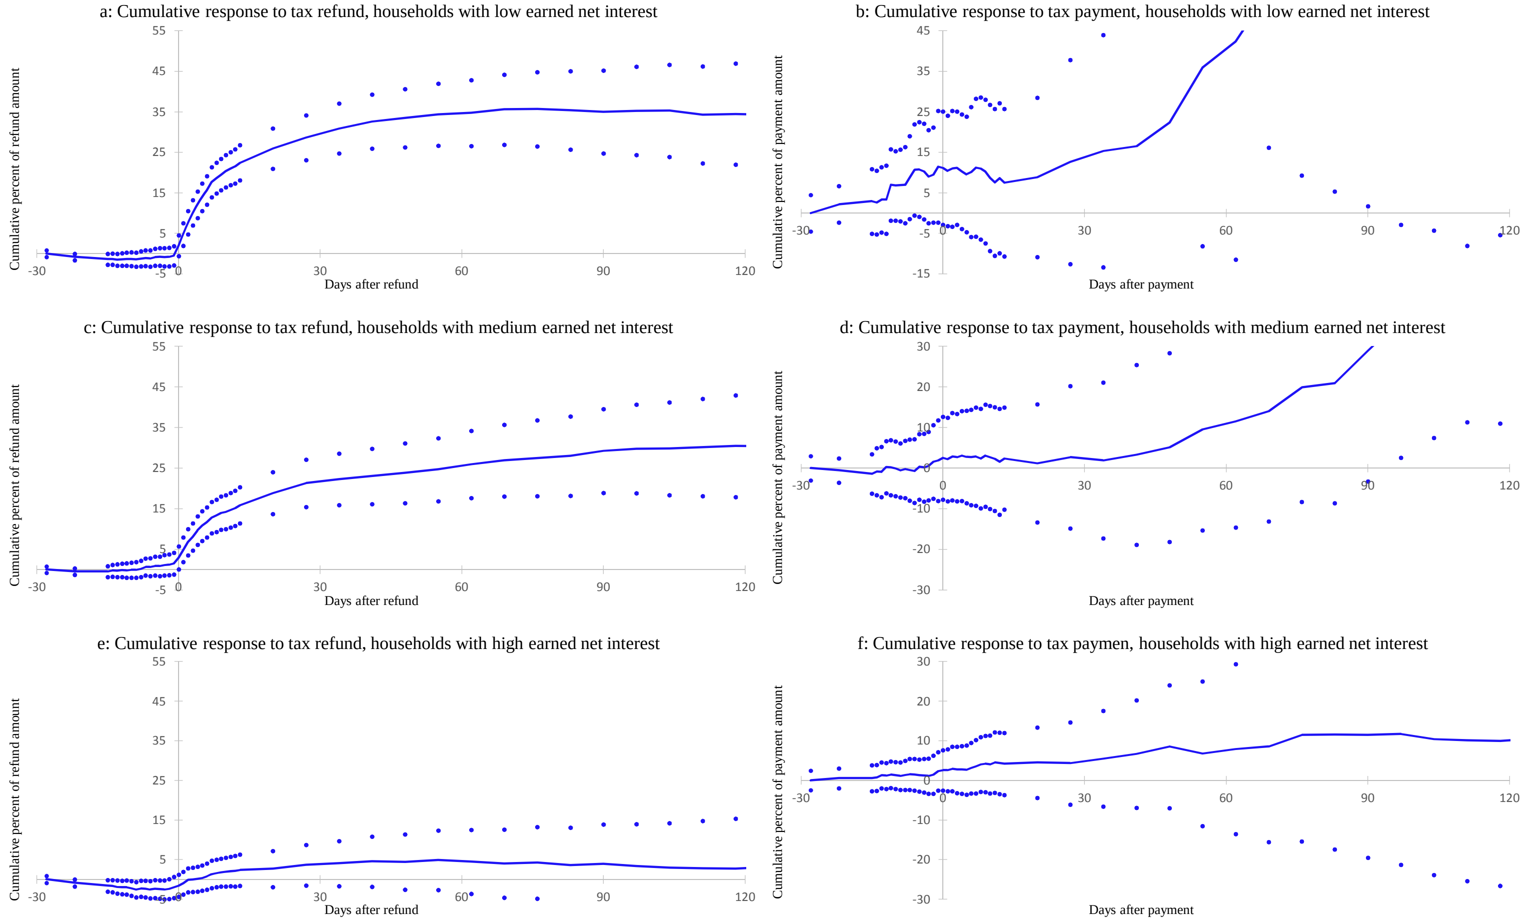

Note: Each row shows the sum of coefficients from estimation of equation (1) with a different sample. The first column displays the cumulative response to a refund, the second column the cumulative response to a payment, in each case from 28 days before refund or cash flow to 120 days after. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure 9: The response of consumption expenditures by income level, baseline sample

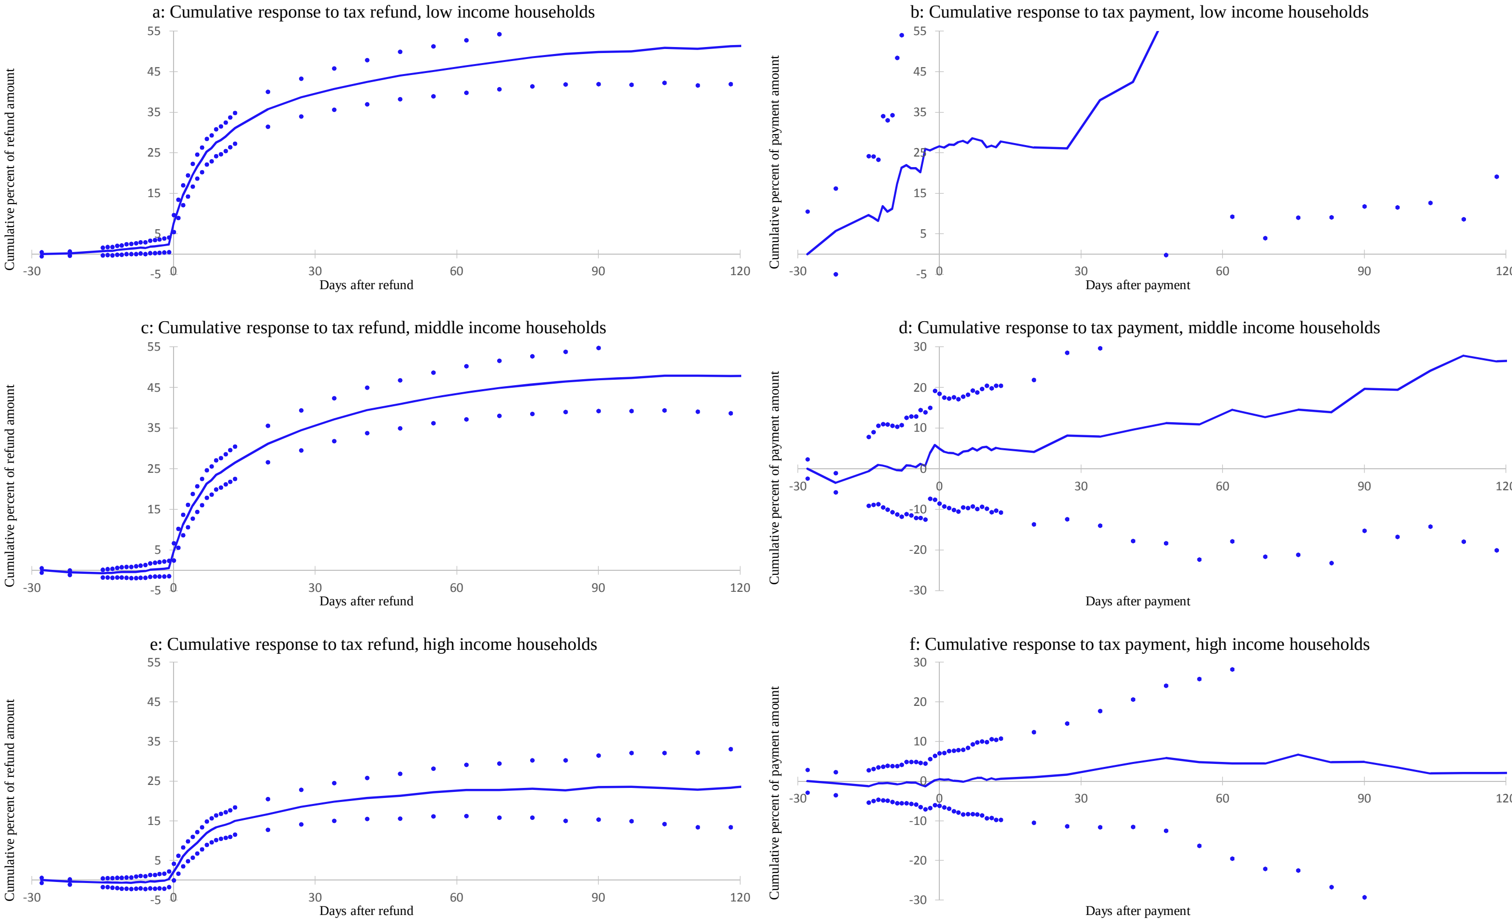

Note: Each row shows the sum of coefficients from estimation of equation (1) with a different sample. The first column displays the cumulative response to a refund, the second column the cumulative response to a payment, in each case from 28 days before refund or cash flow to 120 days after. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure 10: The response of consumption expenditures and credit cards, baseline sample

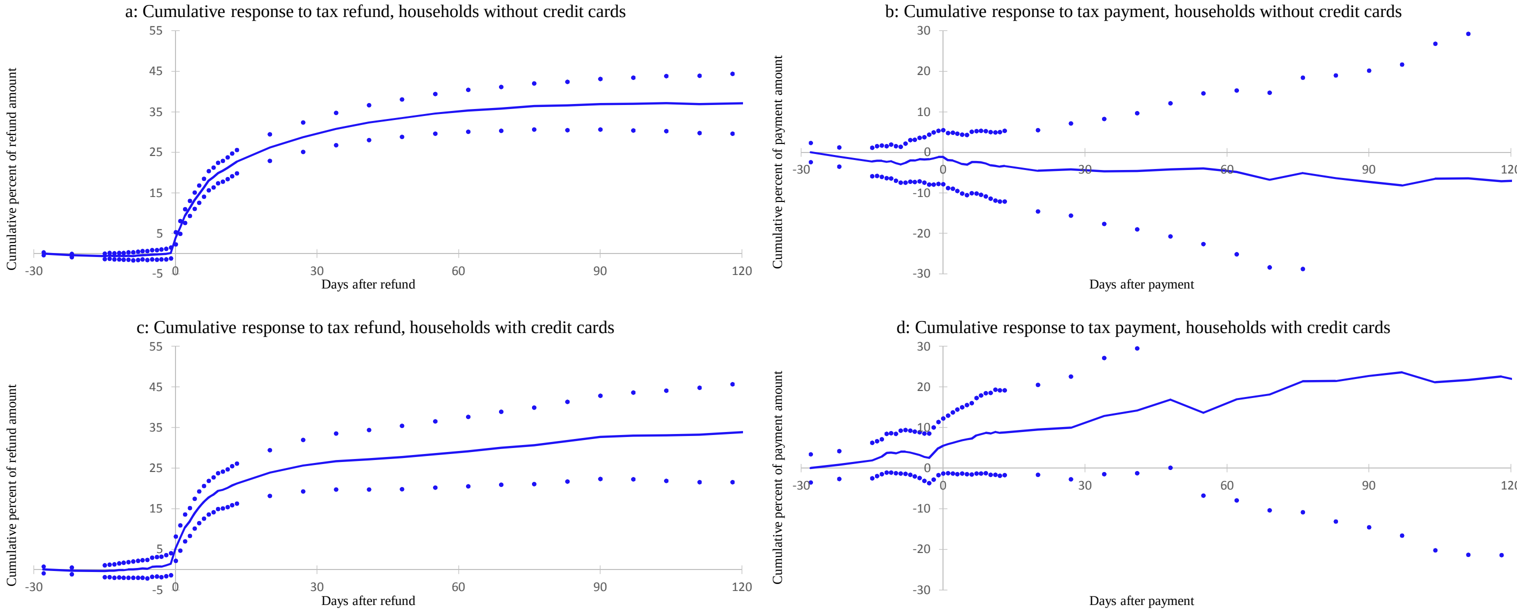

Note: Each row shows the sum of coefficients from estimation of equation (1) with a different sample. The first column displays the cumulative response to a refund, the second column the cumulative response to a payment, in each case from 28 days before refund or cash flow to 120 days after. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure 11: The response of consumption by predicted refund less payment, baseline sample

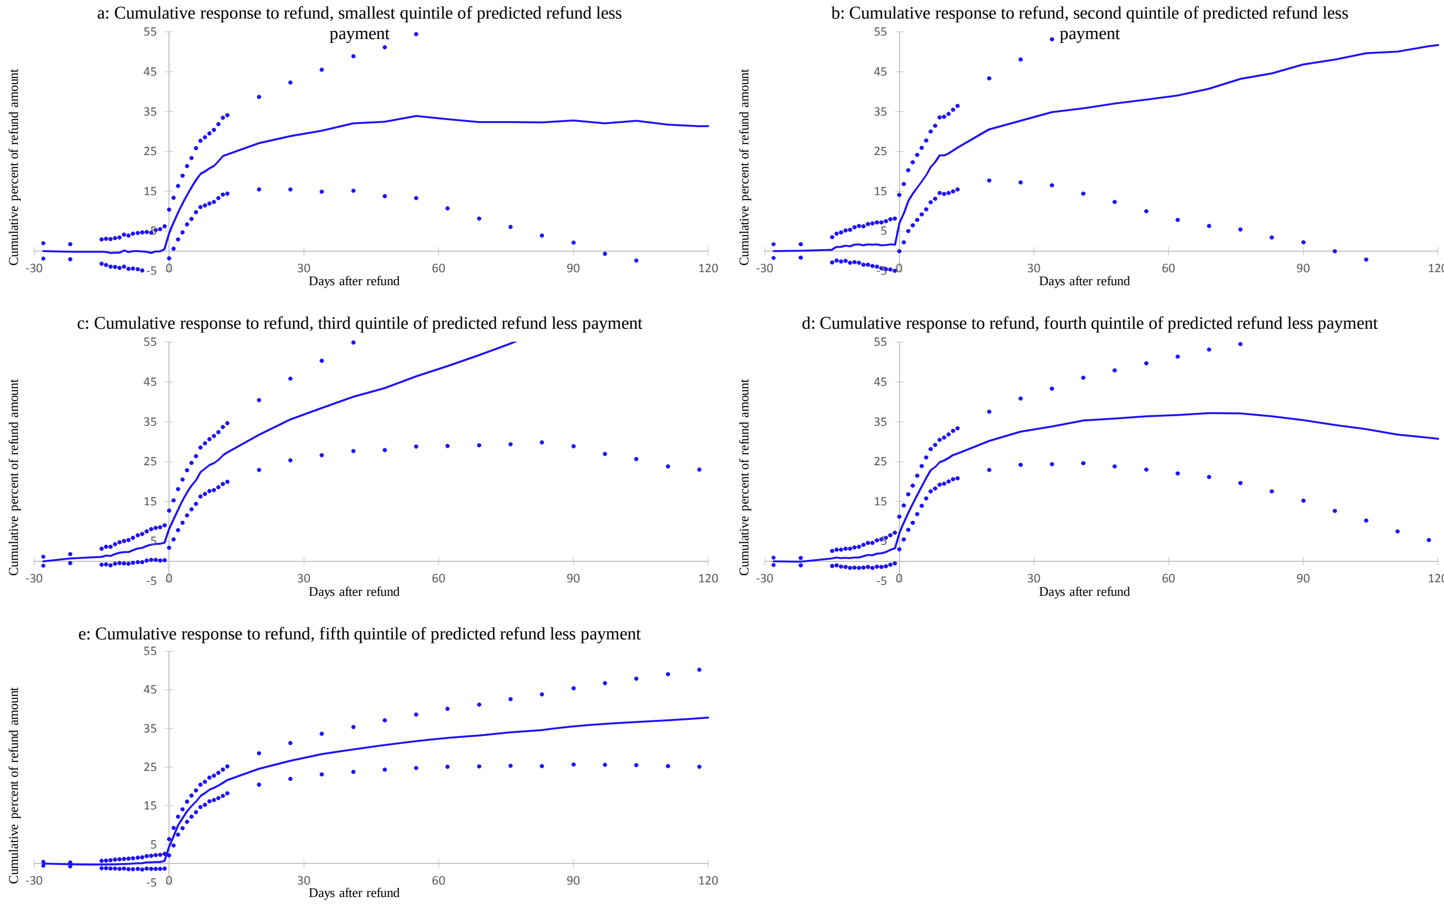

Note: Each panel shows the cumulative response to a refund from estimation of equation (1) with a different sample from 28 days before refund or cash flow to 120 days after. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure 12: The response of consumption to filing and information, baseline sample

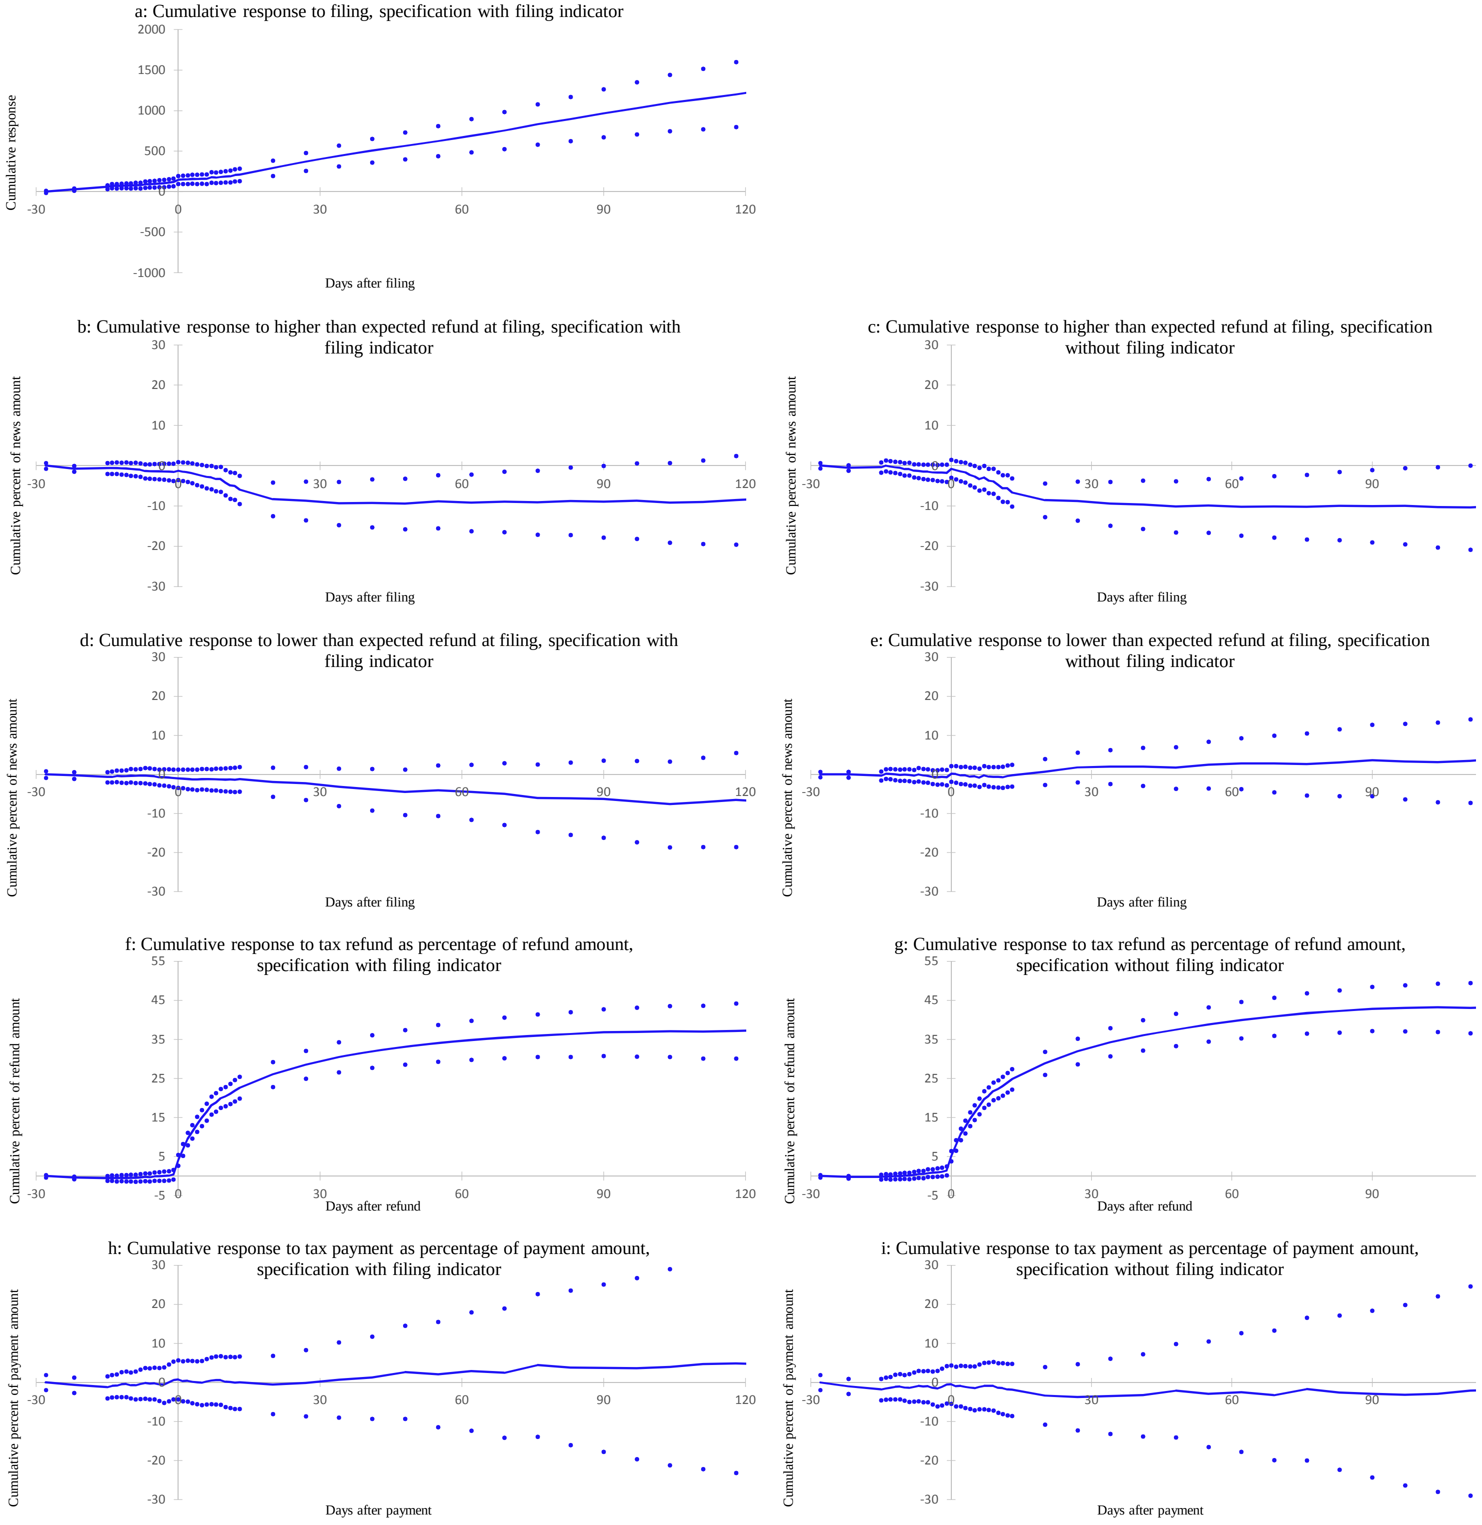

Note: The first columns shows the sum of coefficients from different lagged polynomials in equation (1). The second column shows the same witl coefficients on the date of filing set to zero. All cumulated impulse responses run from 28 days before filing or cash flow to 120 days after. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure 13: The response of consumption spending to refund for different subsamples

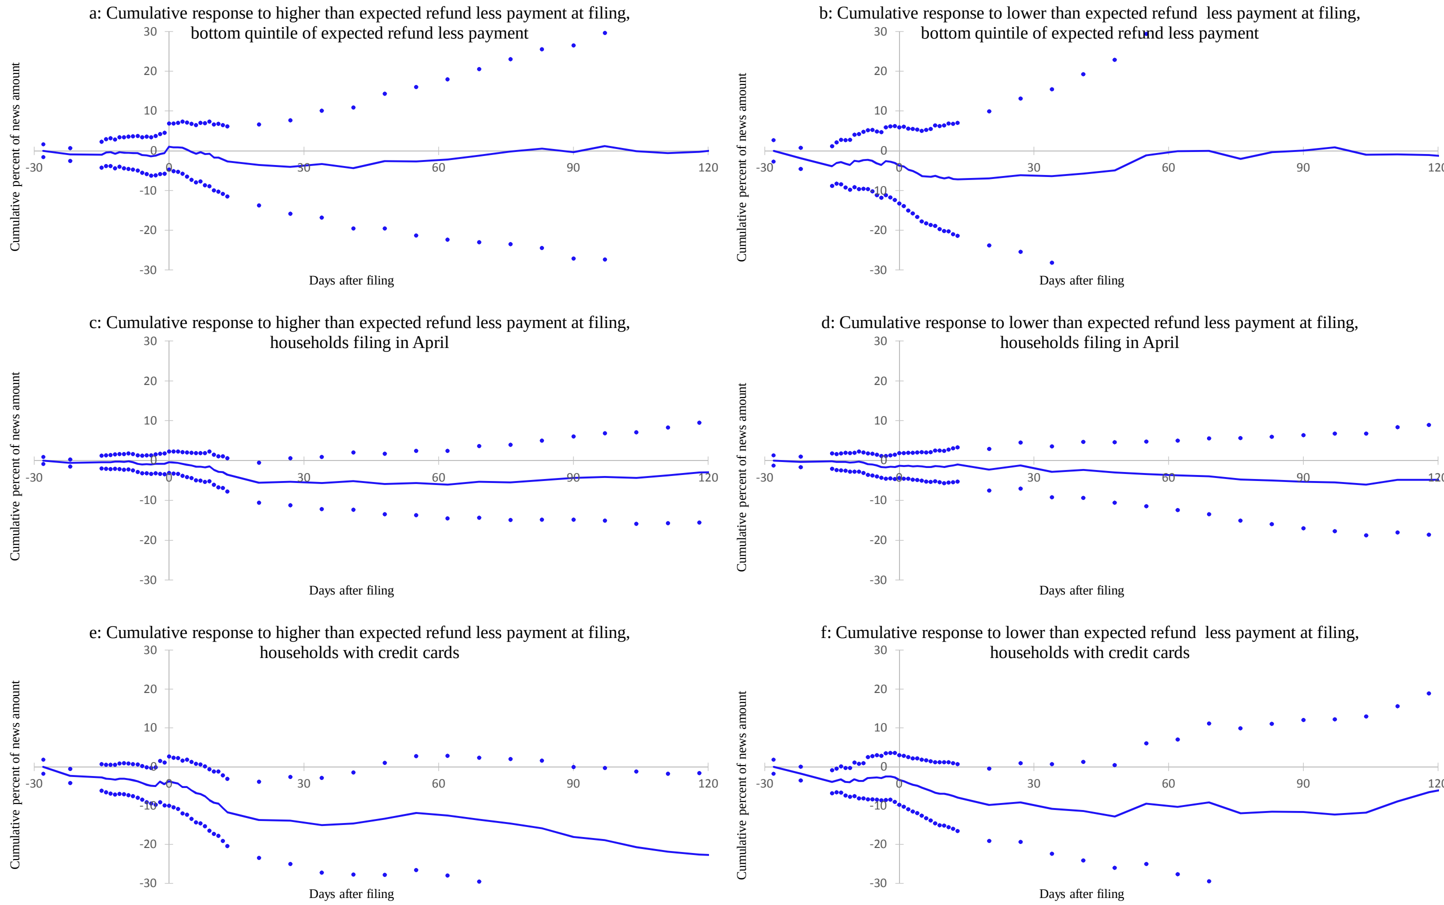

Note: Each figure uses the baseline sample and them imposes its sample restriction. Each row shows the sum of coefficients from estimation of equation (1) with a different sample. The first column displays the cumulative response to good news about taxes due and the second column displays the response to bad news, in each case from 28 days before filing to 120 days after. Dots show pointwise 95% confidence intervals at kink points of the impulse responses

Figure A.1: The response of consumption expenditures by filing month

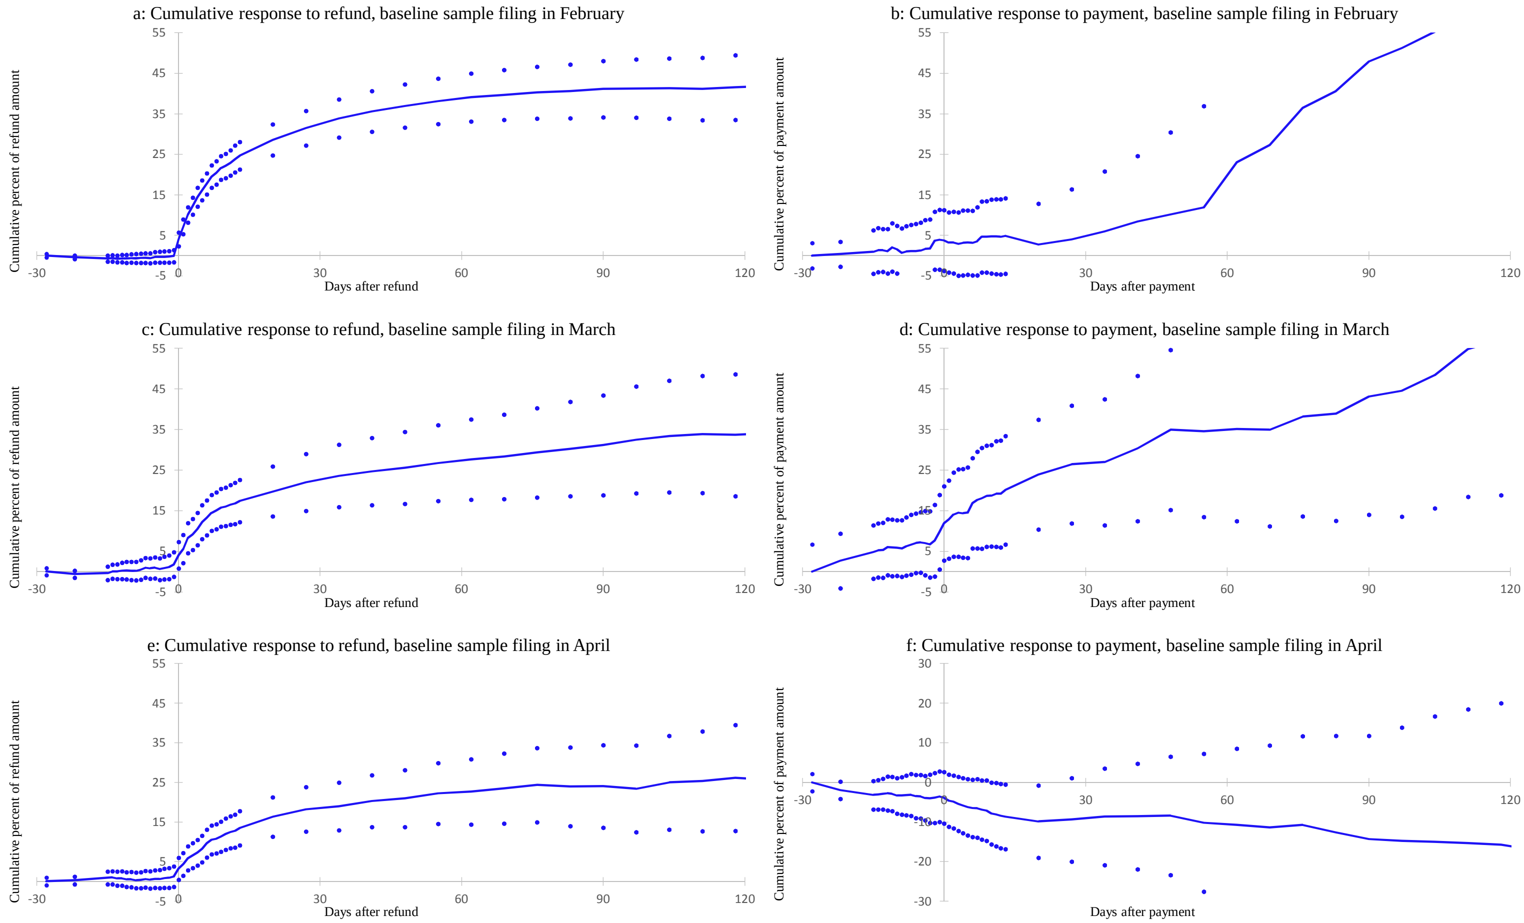

Note: Each row shows the sum of coefficients from estimation of equation (1) with a different sample. The first column displays the cumulative response to a refund, the second column the cumulative response to a payment, in each case from 28 days before refund or cash flow to 120 days after. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure A.2: The estimated cumulative responses of tax refund-payment and filing fee

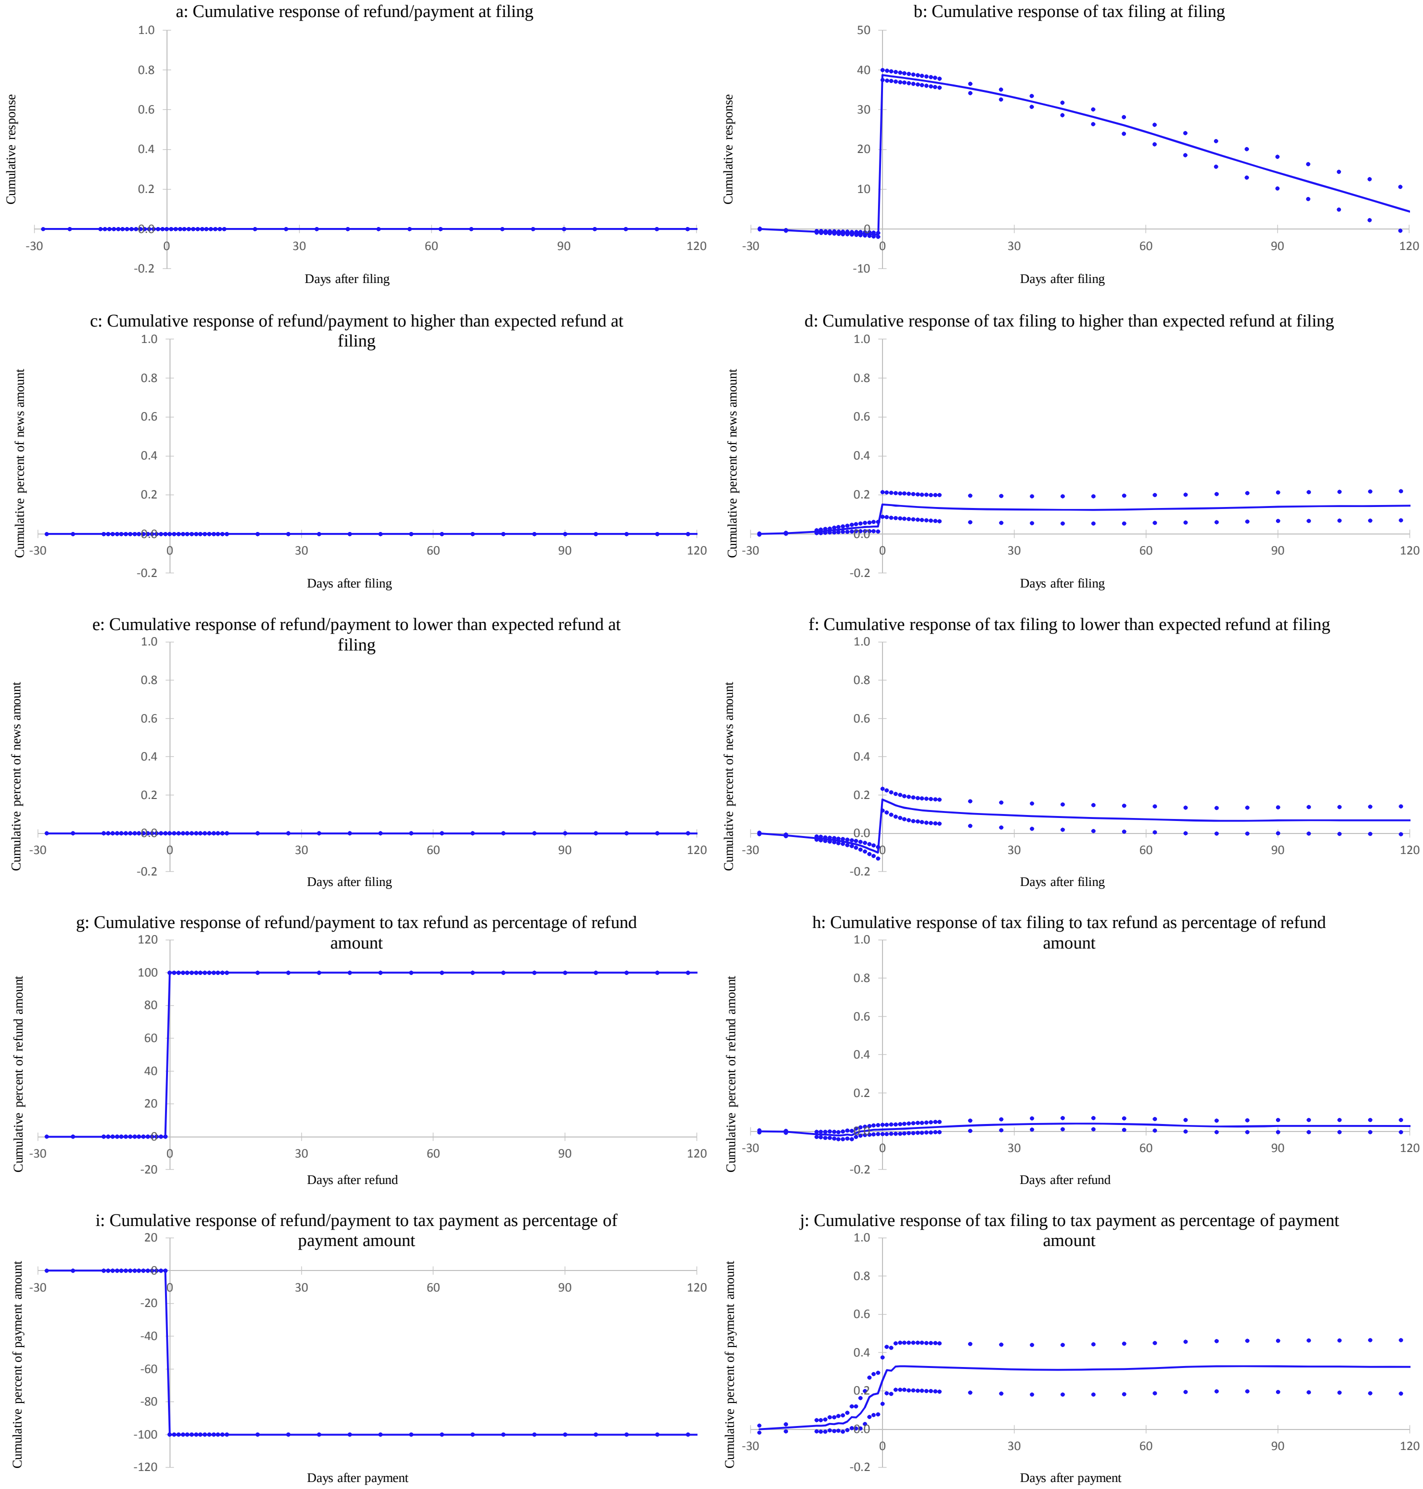

Note: Baseline sample. Each column of figures shows the sum of coefficients from estimation of equation (1) with a different dependent variable, which measures the cumulative change in the dependent variable from 28 days before filing to 120 days after either filing of date of refund or payment. Figures a) to c) plot the responses of refund minus payment to filing date and its interactions with positive and negative news; figures f) to h) plot the response of filing paymen to the same. Figures d) and e) plot the response of of refund minus payment to the arrival of a refund or the making of a payment as a percent of the amount; Figure i) and j) plot the response of filing payment to the same. Note different scales. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.

Figure A.3: The response of consumption expenditures to payments by predicted refund amount

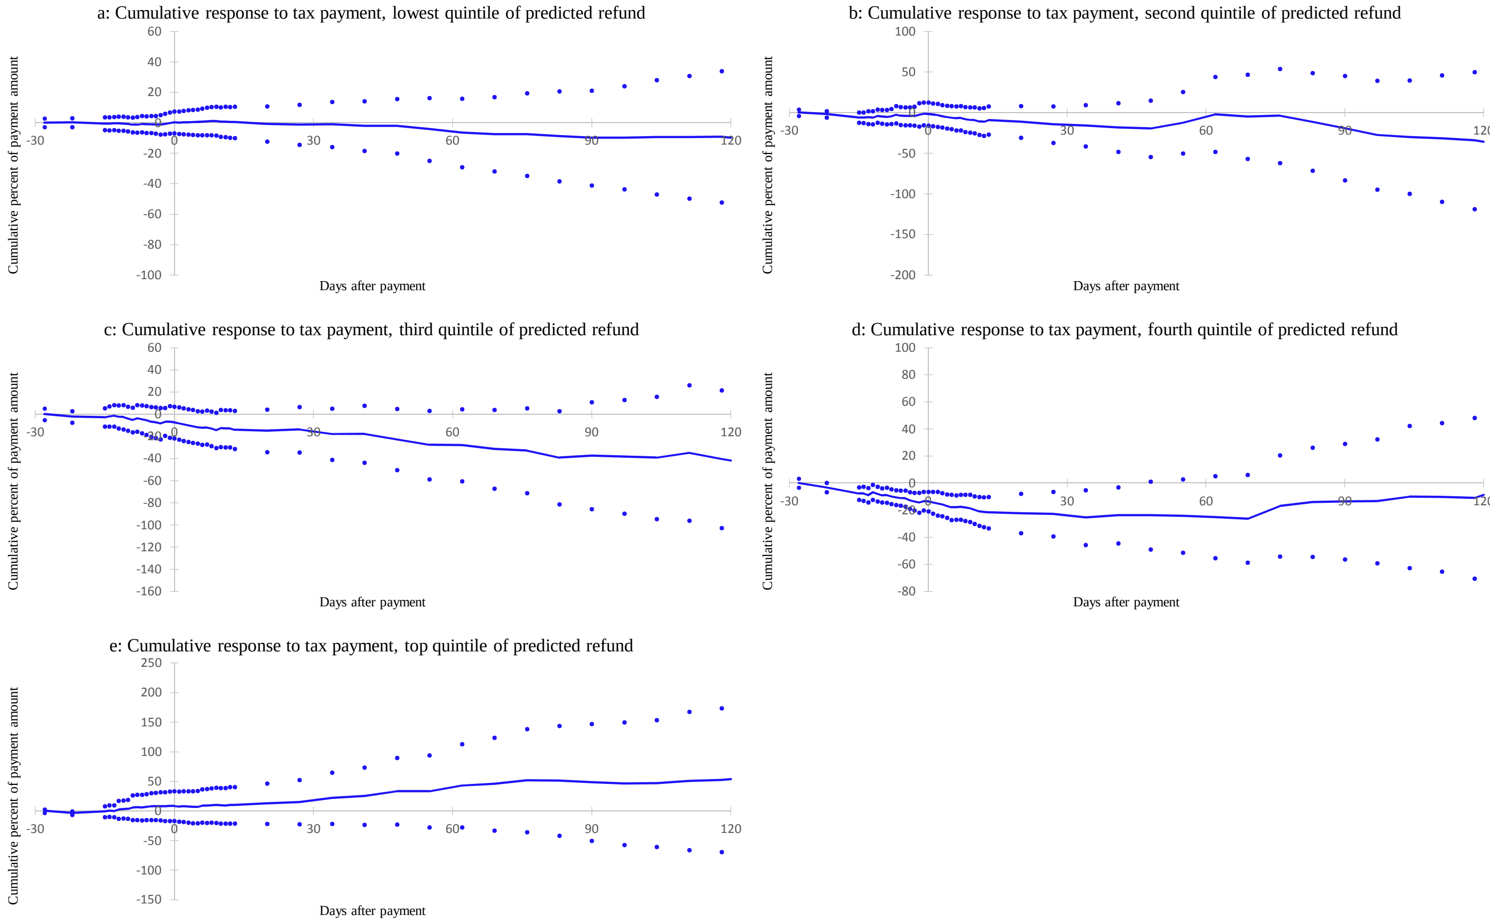

Note: Each panel shows the sum of coefficients on the distributed lag of tax payment amount from estimation of equation (1) with a different sample. Dots show pointwise 95% confidence intervals at kink points of the impulse responses.