

Asymmetric Consumption Smoothing

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

January 10, 2020

Abstract

Using data from an account aggregator, we document that households increase consumption when they receive (expected) tax refunds, *as if* they are liquidity constrained. However, this behavior is not due to liquidity constraints or hand-to-mouth behavior. These same households smooth consumption when making payments in other years, primarily by transferring funds among liquid accounts. Further, even households carrying credit card debt smooth consumption when making payments, and even high-liquidity households spend out of refunds. Thus the households we study follow a heuristic of spending out of increases in liquidity, while at the same time acting in anticipation of payments to maintain stable consumption.

Keywords: Consumption, household finance, liquidity constraints, hand-to-mouth, mental accounting, tax refunds

JEL Classification: G4, D12, E21, H31

*Baugh is with the College of Business, University of Nebraska-Lincoln. Ben-David is with the Fisher College of Business, The Ohio State University, and the National Bureau of Economic Research (NBER). Park is with Nanyang Technological University. Parker is with Sloan School of Management, Massachusetts Institute of Technology, and NBER. Early versions of this paper were titled “A Test of Consumption Smoothing and Liquidity Constraints: Spending Responses to Paying Taxes and Receiving Refunds.” We thank the online account aggregator for providing the data set. For helpful comments, we thank Manuel Adelino, Sumit Agarwal, Andrew Chen, Peter Ganong, Uri Gneezy, Damon Jones, Olivia Kim, Emi Nakamura, Scott Nelson, Pascal Noel, Michaela Pagel, Michael Palumbo, Joel Slemrod, René Stulz, and particularly our discussant, Annette Vissing Jorgensen, as well as participants in seminars at the Federal Reserve Bank of Cleveland, The Ohio State University, and the Philadelphia Federal Reserve Bank Payments Conference, the Annual Conference of the WFA 2018, the JPMorgan Chase Institute 2018 Academic Conference, and the 2015 AEA winter meetings. We are grateful for support from a NBER Household Finance Grant, the Ohio Supercomputer Center, and the Dice Center at the Fisher College of Business. Emails: bbaugh2@unl.edu, ben-david.1@osu.edu, hoonsukpark@ntu.edu.sg, JAParker@mit.edu.

1 Introduction

In theory, if people optimize and have diminishing marginal utility of consumption, they stabilize their consumption over time and insulate their standard of living from fluctuations in liquidity. However, studies show that in practice people increase their consumption spending significantly in response to increases in liquidity, such as predictable, moderately-sized increases in income.¹ There are two explanations of this behavior. First, households could face liquidity constraints: they lack the *means* to stabilize their standard of living due to financial frictions that make borrowing or accessing illiquid assets costly or impossible. Second, households could *choose* to consume from current income.

Using transaction-level data on households that use an account aggregator, we document an asymmetry in consumption behavior. The households in our sample increase consumption in response to the arrival of expected tax refunds *as if* they are either liquidity constrained or living hand-to-mouth. In fact, however, they are neither. In other years, these same households do not decrease consumption in response to expected tax payments. Instead, they increase transfers of funds prior to making payments, primarily between (liquid) savings and checking accounts. Even quite liquid households spend out of refunds and even less-liquid households smooth consumption when making payments. We conclude that excess sensitivity of consumption is at least partly driven by factors other than liquidity constraints and hand-to-mouth behavior. Households in our sample use liquidity to shield consumption from decreases in income when making payments, as rational consumption theories predict, but at the same time follow a heuristic to consume out of transitory increases in income.

We reach these conclusions by comparing the consumption responses of a sample of households as they file their U.S. federal income tax returns, receive refunds, and make payments in different years.² This setting has the nice feature that the increases and decreases in liquidity arise from the same source so that we measure the effect of the sign of the cash flow rather than the effect of the source of the change in liquidity. Further, we focus on a subsample of households that receive tax refunds in some years and make payments in other

¹See, e.g., Bodkin (1959), Zeldes (1989), Parker (1999), Hsieh (2003), Johnson, Parker, and Souleles (2006), and the review in Hassan and Fuchs-Schundeln (2016).

²In using variation in tax refunds, we follow Souleles (1999), who studies the consumption response to the arrival of tax refunds, as well as the contemporaneous work of Farrell, Greig, and Hamoudi (2019).

years so that our results are not driven by differences in the types of households that receive refunds or make payments.

Our sample and data come from a large administrative account-level dataset from an account aggregator that contains every transaction into or out of linked checking, savings, and credit card accounts at a daily frequency from 2011 through 2015.³ For a subset of more than 300,000 household-years, we can observe a reasonably complete financial picture of the households, and identify the date of tax filing and the date and amount of tax refund or payment in the current and previous year. While our administrative data have many advantages, they come with two caveats. First, administrative data track transactions, not consumption. Consequently, we measure only a subset of consumption spending that we can clearly identify as spending on consumer goods and services (retail, restaurants, etc.). The second caveat is that our sample is not a random sample of the U.S. population, primarily because households have to select into the account aggregation service and file taxes in a way that creates an observable transaction in the data. Our sample appears to be more liquid and may be more financially sophisticated than the typical American household. However, as we discuss at the end of the Introduction, this particular sample has some advantages for our analysis.

We measure the consumption response to expected income changes by regressing consumption spending onto distributed leads and lags of refunds and tax payments, and distributed leads and lags of an indicator of tax filing interacted with news about tax status learned during return preparation.⁴ The leads and lags of news serve two purposes. First, and particularly important for payments because they sometimes occur contemporaneously with filing, they control for the impact of information arrival so that we can cleanly measure the effect of an anticipated tax payment or refund arrival. Second, the leads and lags of news 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.

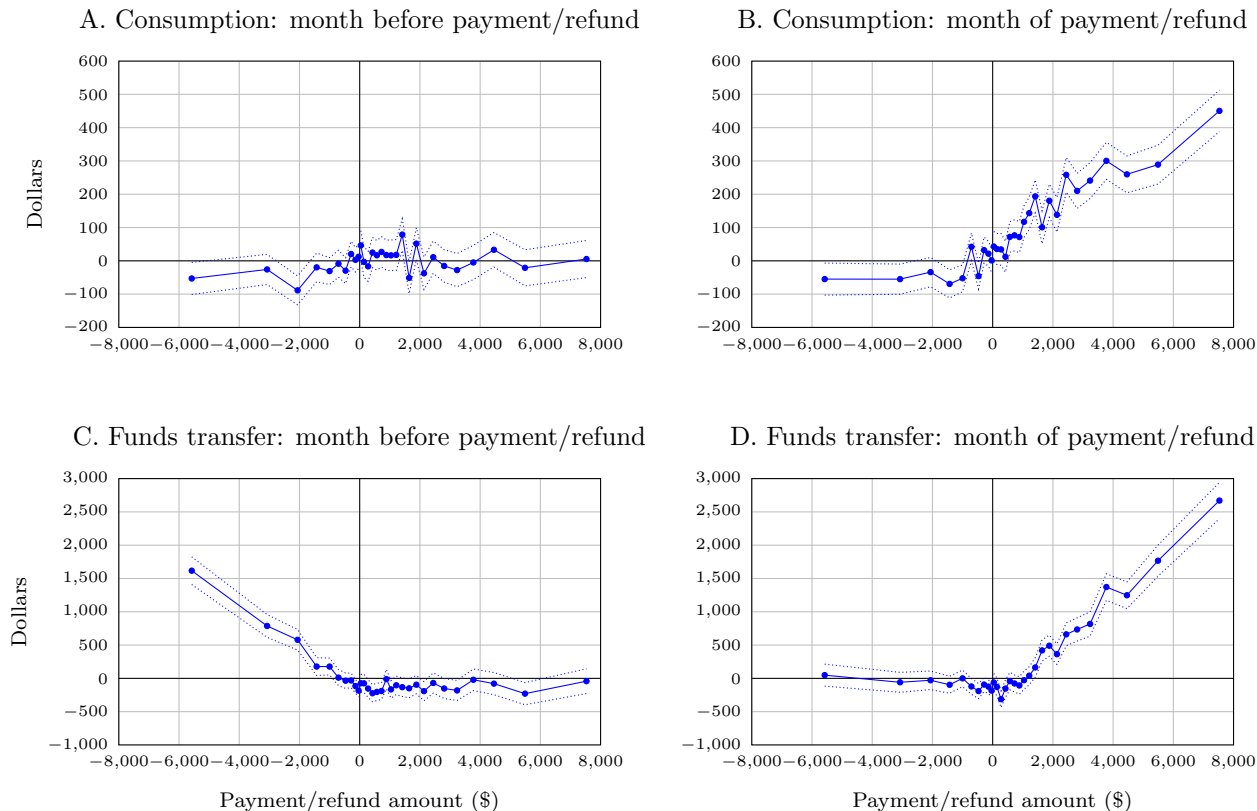
Our main findings are conveyed by Figure 1, which simply plots the abnormal consump-

³We build on research on consumption using high-frequency data (e.g., Stephens, 2003, 2006; Parker, 2017) and account-level data (e.g., Gelman, Kariv, Shapiro, Silverman, and Tadelis, 2014; Agarwal and Qian, 2014; Olafsson and Pagel, 2018; Baker, 2018; Aydin, 2019; Olafsson and Pagel, 2019).

⁴As we describe in Section 5, we measure news as the difference between the actual refund or payment and a prediction based on the household's past tax refund or payment.

Figure 1. Consumption and Fund Transfers Around Tax Payment or Refund

The figure presents consumption and fund transfers around tax payment and refund dates, as functions of payment or refund amounts. Panel A shows the abnormal consumption in the month prior to the tax refund or payment dates. Panel B shows the abnormal consumption in the month of the tax refund or payment dates. Panel C shows the abnormal fund transfers into the observed accounts in the month prior to the tax refund or payment date. Panel D shows the abnormal transfers out in the month following the tax refund or payment date. The markers show averages at every 5% of the data for those who received refunds, and every 10% of the data for those who made payments. The dotted lines represent two standard errors confidence intervals.



tion spending and fund transfers around tax time. We de-mean the data by household-year and day, and then compute average values in the month before or after the refund or payment date for different payment and refund amounts.

First, we observe a striking asymmetry between the consumption response to refunds and the response to payments (Figure 1, panel B). Consistent with a large literature, households spend a significant part of their tax refund on consumption (including on nondurable consumption, such as restaurants) in the month following receipt. In contrast, in years when they make tax payments, these same households do not reduce consumption spending. This result is novel, in part because few papers have studied the consumption response to

decreases in income.⁵ Our results hold for both small and large refunds and payments, indicating that there is a fundamental difference between positive liquidity shocks and negative ones, regardless of amount. Moreover, the result is almost scale-invariant: households spend a nearly constant fraction of their tax refund, regardless of its size.

According to our regression estimates, people spend over 7% of their refunds on consumption during the month following refund receipt, an amount that increases slowly to 15% over several months. These responses are lower bounds on the true increase in consumption because we cannot definitively classify some of the transactions, such as checks, which could be consumption, saving, debt payment, etc. We use daily variation in inflows and outflows to identify the effect and show that this spending increase starts the *day of* refund arrival and that there is no increase in spending related to the timing of the refund arrival prior to the day of arrival. In contrast, the consumption response to payments is economically and statistically insignificant. We also find no evidence that making a payment has an effect on consumption around the exact day of payment.

Second, households do not increase or reduce consumption shortly in advance of cash flow events. Panel A of Figure 1 shows that households do not adjust consumption down in the month prior to making a payment or up in the month before getting a refund. We also find no response in the previous two months, and, in a regression analysis, no change in consumption in response to the information about future cash flows that arrives as tax returns are prepared and filed.

Third, households actively manage liquidity to smooth consumption ahead of payments but not refunds. Panel C of Figure 1 shows that households smooth consumption through tax payments by increasing transfers among accounts in the month before making the payment

⁵Notably, Gelman, Kariv, Shapiro, Silverman, and Tadelis (2020) studies transitory unexpected declines in income and conclude that “even workers with surprisingly low liquid assets can smooth consumption using low-cost methods to shift the timing of payments for committed forms of expenditure.” In contrast, both Christelis, Georgarakos, Jappelli, Pistaferri, and van Rooij (2019) and Fuster, Kaplan, and Zafar (2018) find that people respond in surveys that they would cut consumption more in response to transitory negative than positive income shocks. The few other studies of consumption responses to decreases in income have almost exclusively focused on permanent or highly-persistent decreases. Stephens (2001), Ganong and Noel (2019), and Jorring (2018) find substantial declines in spending in the event of permanent worker displacement, when unemployment benefits expire, and when mortgage payments rise. Conversely, Souleles (2000) and Aguiar and Hurst (2005) find that consumption is well smoothed when college expenses start and when people retire, respectively.

by about a third of their anticipated payment. These transfers are mostly between checking and savings accounts rather than from outside (possibly less-liquid) accounts. We also find that households increase transfers in response to the news learned about upcoming payments or refunds prior to and at filing. We also find evidence that households in our sample have substantial liquidity on average, and that they manage liquidity over longer horizons: in years when they make payments, households have accumulated somewhat higher balances in their core accounts.

In contrast, these same households do not similarly transfer funds and raise consumption ahead of expected refunds. Instead, households getting refunds manage liquidity only after receipt (Figure 1, panel D).

We provide evidence that these findings are not driven by either the endogeneity of the timing of the filing of tax returns (results hold separately for households filing in each month) or heterogeneity across households (results hold for households expecting small refunds or payments). Payment and refund status are also endogenous, determined in part by past income not subject to withholding. But Gelman, Kariv, Shapiro, and Silverman (2019) provides evidence that this endogeneity does not cause our finding of consumption asymmetry. The paper develops a quantitative model of endogenous tax status that matches the distribution of tax refunds. It finds that there is little heterogeneity in marginal propensity to spend across different tax amounts and no sharp difference between refunds and payments. We discuss these issues and others in more detail in Sections 6 and 8.

We find that this consumption asymmetry occurs throughout the distribution of liquidity which is further evidence that liquidity constraints are not driving this behavior. We split the sample by different measures of liquidity and find that households across the distribution of liquidity both smooth consumption through payments and choose to wait to spend out of expected increases in income. We find that for household-years in the bottom third of the ex-ante distribution of liquidity, households have large propensities to consume out of refunds.⁶ But these same households do not cut consumption prior to or when making

⁶This finding is consistent with the literature comparing consumption responses of households with different levels of liquidity that began Zeldes (1989), followed by Jappelli, Pischke, and Souleles (1998), Agarwal, Liu, and Souleles (2007), and Aaronson, Agarwal, and French (2012). See also Kaplan, Violante, and Weidner (2014).

payments. Because even during their least liquid years, households appear to have the funds to smooth consumption over payments; their spending response to refunds appears to be a choice that is not driven by constraints.

We also find that for household-years in the top tercile of the distribution of liquidity, households do not cut consumption when making payments. But they still increase spending when refunds arrive, albeit at a lower rate than low-liquidity households. This spending behavior among liquid household-years also suggests that spending following refunds is a choice or “splurge.”⁷

In sum, households in our sample increase their consumption when their refunds arrive *as if* they are liquidity constrained, but they appear, in fact, to be choosing not to draw on existing liquid savings. Thus, they are increasing consumption as if they are what Olafsson and Pagel (2018) refers to as the ‘liquid hand-to-mouth’ (See also Parker, 1999; Kueng, 2018). However, inconsistent with a simple hand-to-mouth heuristic, these same households do not reduce consumption when making payments but instead accumulate and transfer liquid funds prior to making a payment. Thus, households in our sample appear to follow a type of mental-accounting heuristic: they access liquid savings for making payments but not for smoothing consumption prior to refund arrival, instead treating refunds as windfalls and choosing to splurge when refunds arrive (Thaler, 1985; Arkes, Joyner, Pezzo, Nash, Siegel-Jacobs, and Stone, 1994).

Our findings imply that a propensity to consume out of predictable increases in income is not necessarily indicative of liquidity constraints. For many households, liquidity constraints surely determine behavior. However, we find, in our non-representative sample, households can smooth consumption, but choose not to. Hence, excess sensitivity of consumption to income is not purely driven by constraints.

⁷Consistent with impulsive consumption behavior documented in Agarwal, Chomsisengphet, Meier, and Zou (2019) and Ben-David and Bos (2017).

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 flows and the timing and amount of cash flows using the structure of the U.S. individual income tax system. The U.S. individual income tax covers all sources of household income in each calendar year. For most labor income, employers withhold income taxes from worker's pay during the calendar year, typically following Internal Revenue Service (IRS) guidelines based on pay and family structure. The employer remits these funds to the IRS during the year.⁸ 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. They then must fill out and submit (file) tax returns—some variant of the IRS 1040 tax form and additional schedules—to calculate their total taxes owed.⁹ 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.¹⁰ 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. If taxes owed are less than the withholding, the IRS remits the difference to the taxpayer as a tax refund, typically a few weeks after the taxpayer files his or her return.¹¹

Most households receive refunds, a pattern we expect 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. 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

⁸No corresponding system exists for most capital income. Instead, 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.

⁹People with low incomes or no taxes due do not have to file. Married individuals can file taxes jointly.

¹⁰Taxpayers can file for extensions, which can delay the legal requirement to file until October 15. However, people are responsible for interest charges from April 15, as well as additional penalties if they have significantly under-paid.

¹¹IRS Publication 2043 indicates that 90% of refunds were processed within 21 days of filing. For example, see for 2013: <https://www.irs.gov/pub/irs-prior/i1040--2013.pdf>. See also Slemrod, Christian, London, and Parker (1997).

significant underpayment and the associated penalties and interest. Jones (2012) shows that inaction is a dominant feature of withholding behavior. Gelman et al. (2019) shows that a rational model with uncertainty over income and penalties for under-withholding can also explain the share of households that over-withhold. Finally, the earned income tax credit (EITC) leads many low-income households to have a negative tax liability for the calendar year and so to receive a refund. Our sample likely contains few such households.¹²

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 when they fill out their tax forms before filing. Thus, information about future cash flows 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. 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 costs of missing the April 15 deadline.

3 Theory: Refunds, Tax Payments, and Consumption

Lifecycle Model In lifecycle 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 one percent of the change in amount revealed, spent over the remaining lifetime.

¹²Although households cannot choose negative withholding or estimated tax payments, households with children who qualify for the EITC can file a W-5 form with their employer and receive up to 60% of the EITC credit early.

Buffer-Stock Model In models with liquidity constraints, these responses may be asymmetrically constrained.¹³ Households can always increase savings to prepare to make a tax payment. But households with limited liquidity cannot always borrow to increase spending in anticipation of a tax refund.¹⁴ Thus, a model with liquidity constraints predicts different responses to news about refunds versus 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 about 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 have plenty of 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.

Mental Accounting In models of mental accounts Thaler (1985) and Arkes et al. (1994), people treat some accounts as current accounts that serve for current consumption, while others are reserved for future use (e.g., retirement). These descriptive theories can predict that a refund is consumed because it typically flows into the checking account which is designated for current consumption. And these mental rules prohibit consuming from a saving account ahead of refund arrival, but accessing a saving account to make tax payments that are not considered consumption.

¹³Shea (1995) investigates this asymmetry in aggregate data and shows that aggregate consumption responds more to predictable decreases in income than to predictable increases.

¹⁴Or 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 on unsecured borrowing, or choose not to take on the costs associated with turning less liquid assets into consumption.

Hand-to-Mouth and Models of Inattention Other theories of consumer behavior have quite different implications and predict symmetric spending responses. If households are living hand-to-mouth (as in Campbell and Mankiw, 1989) or behave as target savers as 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 (e.g., Caballero, 1995).

4 Data and Variable Construction

4.1 Data Source

The data we use were provided by a free online account aggregator. This service allows households to view their financial information such as account balances and spending by category. The service also assists with financial management, such as offering alerts for upcoming bills or approaching credit limits. Users sign up and provide access to account information for accounts across different financial institutions. Once someone signs up, the aggregator has access to the household's account information and there is low attrition in our sample.

The raw data cover daily transactions for 2.7 million households from July 2010 to May 2015, and include all checking, savings, debit card, and credit card transactions for any account once linked to the service by the household. We observe permanent household identifiers, and the date, amount, and description of every transaction, in a form such as is typically found on monthly bank or credit card statements.

4.2 Variable Construction

We use the text of each description in conjunction with the data provider’s categorization of each transaction to map financial transactions into economic concepts as follows.

First, we 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.¹⁵ To remove unusual tax activity such as that occurring through business owners, we exclude any household-year containing more than one such tax refund or payment. We further remove any household that has ever incurred a tax payment or tax refund of over \$10,000, in large part because the very few households who made payments greater than \$10,000 are quite different from most of our sample along many dimensions.

Similarly, 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 designate the filing date of the household as the transaction date of the tax-preparation software. We exclude households that have tax-preparation transactions on multiple days (as would be the case for a family filing separately on different days). We also implicitly exclude households that elect to deduct the preparation charges directly from their refunds (since we would not observe a payment associated with filing).

We construct a measure of consumption spending that consists only of outflows that we are highly confident represent spending on consumption. Consumption is defined as the sum of outflows on the following categories: gas, restaurants, retail, groceries, cash, entertainment, health care, travel, utilities, miscellaneous bills (e.g., gym memberships), and insurance. To ensure that consumption is not mechanically related to taxes, we exclude from it any filing fees or tax payments. These outflows are primarily those that the household makes via debit cards and online bill payments as well as those on linked credit cards. Im-

¹⁵Conventional 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 a 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 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 “sbtprg,” “tax products p,” “block bank des hrbb,” “block bank hrbb,” and “republic trs.”

portantly, we are unable to categorize outflows made by check. Our measure of consumption spending, thus, understates actual consumption spending not only because of categories that we omit from our measure out of caution, but also all those made by checks.

Our data allow us to observe consumption by households as long as the household consumes via a linked debit card or linked credit card.¹⁶ We do not observe spending that occurs on an unlinked account such as an unlinked credit card account. To overcome this issue, we create a scaling factor that scales up observed credit card spending by the household-specific ratio of annual payments to all credit cards (linked and unlinked) divided by the payments to linked credit cards (separately for each household-year). We drop all household-years with scaling factors above 2 since, for these households, we observe less than half of their credit card spending on consumption. The scaling factor has a median of 1.083 and an average of 1.167. Our estimated consumption responses are only slightly larger due to the scaling-up of credit card spending, consistent with the 16.7% implied by the average scaling factor.

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 minus 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 definitively categorize into either consumption or savings. This variable is equal to the sum of the following categories: checks and otherwise uncategorized outflows. Given that we omit miscellaneous payments from consumption, we are likely understating true consumption. We measure income based on direct deposits, which captures labor income, pension income, and Social Security distributions.

Our data do not directly contain account balances so we construct two measures from the information we do have. First, we use ‘net interest earned’ as a proxy for each household’s liquidity, which we construct as the interest earned on checking and savings accounts less the interest paid for balances carried on credit card accounts. Second, some transaction lines for interest earned include the account balance in the text. We extract account balances for

¹⁶We observe all transactions on linked cards, but only payments to unlinked cards not spending on these cards.

about 5 percent of household for which we observe balances in this way for all bank linked accounts.

4.3 Sample Construction

The central question that we investigate is whether households react differently to cash payments versus cash receipts. To answer this question in our empirical setting, we need to keep household characteristics as constant as possible. Doing so ensures that the effects that we document result from neither within-household variation nor across-household differences. Therefore, we focus on the subsample of households who we observe both making tax payments and receiving refunds across different years in our sample period.

We arrange our data into household-years running from November 1 to October 31.¹⁷ We drop household-years for which all necessary variables are not present, as described in the previous subsection. Further, to ensure we have active users rather than dormant account holders, we impose a simple activity filter for households. We require that households have at least 25 transactions of at least one dollar each month.


We next require that the tax filing date precedes the tax refund or tax payment date. However, 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 if an individual pays a tax-preparation fee with a debit card that takes two days to post while paying their taxes with a credit card that posts immediately.. On average, refunds are received 10.2 days after filing, and payments are made 8.6 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. Finally, in addition to observing the refund or payment for a given tax year, we require that we also observe a refund or payment in the year prior for reasons explained in Section 5.1.

The resulting distribution of tax payments and refunds is shown in Figure 2. Panel A shows the distribution of tax payments/refunds in our broad sample. Panel B shows the distribution for households that make a payment in at least one year and receive a refund

¹⁷Each year consists of 365 days, with the exception of 2015, when our sample ends, which consists of 237 days.

Figure 2. Distribution of Payments and Refunds

This figure presents histograms of payment/refund amounts. Panel A shows the histogram for the broad sample, while panel B shows this same histogram for the final sample, households with a payment in at least one year and a refund in another year.



in at least one year, which is our main sample. The distribution of refunds and payments is more symmetric in our main sample than in the original broader sample.

4.4 Sample Statistics

After applying all filters except that each household must receive at least one refund and make one tax payment in our sample, our baseline sample contains 312,212 household-years from 199,077 unique households, which leads to a dataset with 104 million household-day observations. The summary statistics for this broad-sample population are presented in Appendix Table A.I. Our original dataset consists of households that have selected to use the aggregator. It appears to be broadly representative of the population with some exceptions.¹⁸ From this population, we use the subset of households that are observed to have at least one refund and one payment (in different years). The summary statistics for this, our main sample of households, are presented in Table 1.

As shown in Table 1, the average daily income is \$186.14, corresponding to an average annual salary of \$67,941. The distribution of salary is shown in Figure 3, which illustrates

¹⁸Relative to U.S. Census, households in our sample are well dispersed geographically, though we have high concentrations of households in California, New York, and Texas. See a detailed distribution of the sample in Baugh, Ben-David, and Park (2018), who use the same broad sample.

Table 1. Summary Statistics

This table presents summary statistics for our final sample, which includes only households that make a tax payment in one year and receive a tax refund in another. *Refund – Payment* (\$) is the average refund (less payment) amount across household years. *Predicted Refund* (\$) shows the average predicted refund size. *News Amount* (\$) is the difference between the realized refund and the predicted refund. To ensure that the prediction is unbiased, we use households that only make payments or only have refunds to make this prediction, so the mean is nonzero. *Distance Filing Refund* (days) is the number of days between filing and the subsequent payment/refund. *I(Linked Credit Card)* is an indicator variable that equals one if the household has a credit card linked to the account aggregator, and zero otherwise. *I(Unlinked Credit Card)* is an indicator variable that equals one if the household has an unlinked credit card, and zero otherwise. *I(Any Credit Card)* is an indicator that equals one if the household has either linked or unlinked credit cards, and zero otherwise. *Net Flow* (\$) is the difference between inflows and outflows to the core accounts. *Consumption* (\$) is observed consumption. *Scaled Consumption* (\$) is the *Consumption* variable scaled up to compensate for the presence of unlinked credit cards. *Savings and Loans* (\$) shows payments to savings accounts (net investing outflows, net transfers out) and net loan payments (mortgage, auto loan, and net decrease in credit card debt). *Misc Payments* (\$) is the sum of checks and uncategorized outflows. *Net Interest* (\$), *Interest Expense* (\$), and *Interest Earned* (\$) are net interest, interest expense, and interest earnings, respectively. Finally, *Net Credit Card (CC) Charge* (\$) is all linked credit card expenditures after excluding tax-related transactions (such as filing fees and tax payments made on credit cards).

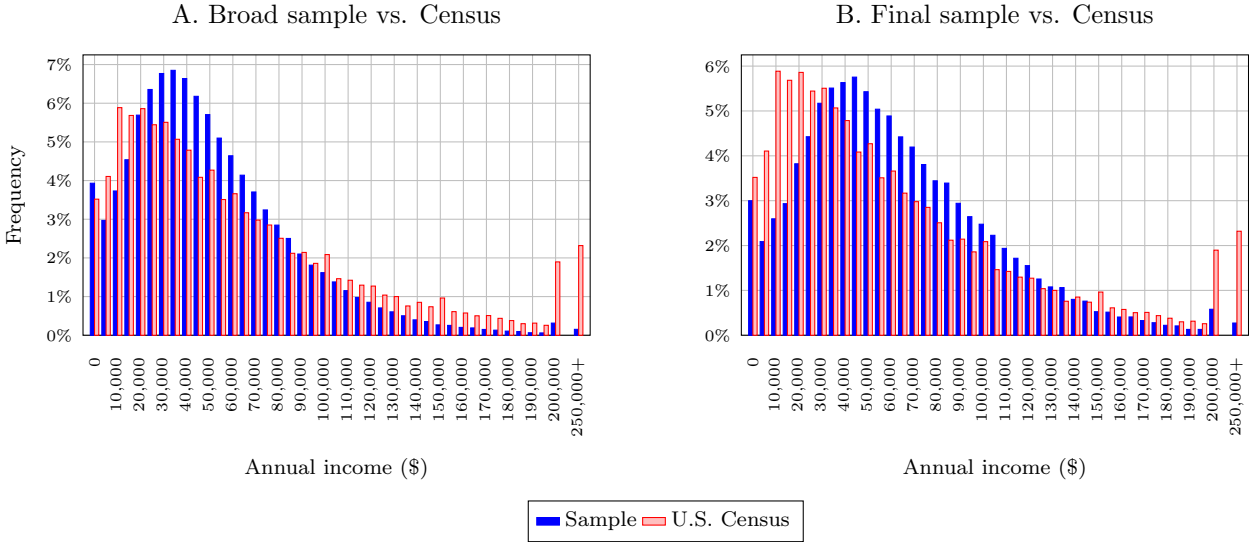
| Variable | Count | Mean | SD | p1 | p10 | p25 | p50 | p75 | p90 | p99 |
|--------------------------|------------|----------|----------|-----------|-----------|-----------|----------|----------|----------|----------|
| Households | 37,581 | | | | | | | | | |
| Household-Years | 62,953 | | | | | | | | | |
| Filing Date | 62,953 | Mar 16 | 27.51 | Jan 21 | Feb 04 | Feb 20 | Mar 23 | Apr 12 | Apr 15 | Apr 19 |
| Refund Date | 62,953 | Mar 26 | 25.63 | Feb 01 | Feb 14 | Mar 04 | Apr 06 | Apr 16 | Apr 21 | May 03 |
| I(Positive Refund) | 62,953 | 0.57 | 0.50 | 0.00 | 0.00 | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| Refund – Payment (\$) | 62,953 | 564.37 | 2,577.07 | -6,621.00 | -2,210.00 | -630.00 | 271.00 | 1,827.00 | 3,785.00 | 7,826.00 |
| Predicted Refund | 62,953 | 1,355.76 | 1,592.05 | -2,032.92 | -304.43 | 200.80 | 1,162.92 | 2,150.22 | 3,479.34 | 6,182.52 |
| News Amount (\$) | 62,953 | -791.39 | 2,610.03 | -8,054.27 | -3,951.50 | -2,079.45 | -701.54 | 525.06 | 2,196.77 | 6,326.82 |
| Filing to Ref/Pay (days) | 62,953 | 9.53 | 12.53 | -2.00 | 0.00 | 2.00 | 6.00 | 11.00 | 21.00 | 67.00 |
| I(Linked Credit Cards) | 62,953 | 0.81 | 0.40 | 0.00 | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| I(Unlinked Credit Cards) | 62,953 | 0.83 | 0.38 | 0.00 | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| I(Any Credit Cards) | 62,953 | 0.96 | 0.18 | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| Household-Days | 20,953,219 | | | | | | | | | |
| Net Flow (\$) | 20,953,219 | -54.58 | 4,063.04 | -3,902.98 | -677.49 | -221.36 | -52.72 | 0.00 | 364.00 | 4,470.35 |
| Consumption (\$) | 20,953,219 | 99.73 | 176.16 | 0.00 | 0.00 | 0.00 | 36.58 | 120.62 | 265.58 | 911.89 |
| Scaled Consumption (\$) | 20,953,219 | 116.26 | 207.87 | 0.00 | 0.00 | 0.00 | 41.66 | 139.31 | 309.72 | 1,081.11 |
| Savings and Loans (\$) | 20,953,219 | 38.83 | 509.89 | -995.88 | -150.70 | -37.75 | 0.00 | 0.00 | 75.56 | 2,556.93 |
| Misc Payments (\$) | 20,953,219 | 21.93 | 544.21 | -2,000.00 | 0.00 | 0.00 | 0.00 | 30.00 | 168.47 | 1,799.53 |
| Income (\$) | 20,953,219 | 186.14 | 815.13 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 3,990.73 |
| Net Interest (\$) | 20,953,219 | -0.69 | 9.10 | -18.93 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 1.37 |
| Interest Expense (\$) | 20,953,219 | 0.81 | 9.55 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 18.99 |
| Interest Earned (\$) | 20,953,219 | 0.10 | 1.49 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 1.40 |
| Net CC Charge (\$) | 20,953,219 | 57.64 | 171.31 | 0.00 | 0.00 | 0.00 | 0.00 | 47.22 | 158.92 | 798.26 |

the income distribution of our initial broad sample (panel A) and final sample (panel B) relative to U.S. Census data.¹⁹ As shown, households in our sample vary widely in income. We also have a disproportionate number of very low-income accounts, perhaps because we do not identify all income correctly. More generally, while the distribution of our initial sample

¹⁹2013 Current Population Survey from the U.S. Census (HINC01).

Figure 3. Distribution of Income

This figure presents the distribution of income in our samples and from the Census for the same years. Panel A shows the histogram for the broad sample of households, while panel B shows this same histogram for our final sample, households with a payment in at least one year and a refund in another year.



corresponds to the U.S. distribution reasonably well, our final sample of households tends to be higher income. To some extent this is to be expected because the income that we observe is net of taxes and benefits such as 401k contributions, health care premiums, etc.

In terms of the consumption, household consumption averages \$99.73 per day (corresponding to \$36,401 annually) in directly observed spending and \$116.26 per day (corresponding to \$42,435 annually) when scaled up to account for spending on unlinked cards. Nearly 81% of households have credit cards linked to the aggregator, 83% of households have credit cards that are not linked to the aggregator, and 96% of all households have at least one credit card that is either linked or unlinked.

In our sample, 57% of household-years receive a refund. The average value of *Refund* – *Payment* is \$2,351.38 for the broad sample and \$564.37 in our final sample. In untabulated results, we find that, conditional on receiving a refund, the average refund is \$2,119. Thus, the average refund is quite large since it constitutes 37% of an average household’s monthly income (= \$2,119 / (\$67,941 / 12)). Conditional on making a payment, the average payment is \$1,494.

Table 1 also shows that the average household files on March 16, and that there is

Table 2. Payment and Refunds Statistics by Month of Filing

This table shows the distribution of filing statistics for our final sample, households that make a payment in at least one year and receive a refund in another year. Panel A shows the distribution of filing month across households, while panel B shows the distance between filing and payment/refund. In both panels, the first column shows the distribution for all households, the second column shows the distribution for households making a payment, and the third column shows the distribution for households receiving a refund.

| Month of Filing | Panel A: Filing Month | | | Panel B: Days from filing to payment/refund | | |
|-----------------|-----------------------|------------|-----------|---------------------------------------------|------------|-----------|
| | All | I(Payment) | I(Refund) | All | I(Payment) | I(Refund) |
| January | 6% | 3% | 9% | 19.3 | 36.7 | 15.4 |
| February | 26% | 15% | 34% | 14.1 | 23.2 | 11.0 |
| March | 25% | 24% | 25% | 10.5 | 12.1 | 9.4 |
| April | 43% | 58% | 32% | 4.8 | 2.1 | 8.5 |

significant variation: the standard deviation of filing date is 27.5 days. The average household makes a payment on March 16 or receives a refund on March 26, and again there is large variation across households (standard deviations of 26 and 28 days for filing and refund dates, respectively). The average distance between filing and refund or payment is 10.2 days, with a still substantial standard deviation of 8.7 days. Table 2 provides more details and shows that there is substantial variation in refund and payment dates across months, and in the number of days between filing and refund or payment.

Table 2, panel A, shows the distribution of filing dates throughout tax season. The overall distribution has a slight bimodal tendency, driven by the relatively higher propensity of households with refunds to file early and households with payments to file near the deadline. Table 2, panel B, shows the delay in days between filing and refund receipt and filing and tax payment separately. 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 pay when they file or choose instead to pay right before the deadline, although many payments fit neither scenario.

5 Estimation Method

5.1 Information Acquired During Tax Preparation and Filing

We measure the news about the future amount of tax refund or payment as the difference between the actual amount paid or received and the expected amount. We compute the expected amount using information on the previous year’s amount and take the residual from the equation below as a measure of the information revealed by tax preparation.

To set the notation, let $Refund_{h,y}$ be the amount of any refund received in year y and be zero if a payment is made that year, and analogously for $Payment_{h,y}$. We run the following regression:

$$Refund_{h,y} - Payment_{h,y} = \beta_0 + \beta_1 Refund_{h,y} + \beta_2 Payment_{h,y} + \beta_3 \mathbf{1}[Refund_{h,y} > 0] + \eta_{h,y} \quad (1)$$

where $\mathbf{1}[\cdot]$ denotes the indicator function. We run this regression on the broad sample of households rather than the final sample since the final sample is selected on the outcome of refunds and payments. The predictive regression has a fit goodness (R^2) of 50%.²⁰ Our measure of information about tax information uncovered during filing, or “news,” is the residual in this regression, which we denote by

$$News\ Amount_{h,y} = Refund_{h,y,t} - Payment_{h,y} - E_{y-1}[Refund_{h,y} - Payment_{h,y}]. \quad (2)$$

The distribution of *News Amount* is shown in Table 1 for our final sample, and in Appendix Table A.I for the broader sample. Whereas the average *News Amount* in the broader sample is close to zero, the average *News Amount* in the final sample is $-\$791.39$. This difference reflects our selection of households with both payments and refunds, who then have worse outcomes on average because payments are more prevalent than in the larger sample used for estimating *News Amount*.

This empirical model is identified from cross-sectional variation and has a short time

²⁰Adding 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.

dimension; consequently, it effectively endows agents with knowledge of the increase in average refund over the few years we study. This assumption is supported by the fact that this was a period without a federal tax reform and with few changes in tax law more generally. Consistent with this stability, according to the IRS, average refunds declined reasonably steadily by \$82 per year from their peak in 2010.²¹ 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 Responses to Cash Flows and Information

We summarize household behavior in two ways. First, we present transparent plots of the data as in Figure 1 in the Introduction. We de-mean the average value of $Y_{h,y,t}$, both by calendar day and by subtracting the average value of $Y_{h,y,t}$ for the household over the seven months that exclude the two months before and three months after the refund or payment.²² We create 10 equally-sized bins of tax payments and 20 equally-sized bins of tax refunds and compute the average value of tax payment or refund along the x -axis in the 30 days before or after the payment or refund.

Second, we estimate and display the impulse responses 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 with a different slope for refunds than payments (linear with a kink at zero). We also estimate (and so control for) the arrival of information by estimating in the same regression the impulse response to the news uncovered prior to and at filing, allowing the spending response to be affine in amount of news, but with different slopes for positive news and for negative news (that is, affine with different coefficients on either side of zero).

To be precise, let $Refund_{h,y,t}$ be the amount of refund received on day t in year y and be zero on all other days of that year or if a payment is made that year. Define $Payment_{h,y,t}$ analogously and let $File_{h,y,t}$ be an indicator variable for the day t of year y on which house-

²¹IRS Statistics of Income, Tax Stats, <https://www.irs.gov/statistics/soi-tax-stats-amount-of-refunds-issued-including-interest-by-state-irs-data-book-table-8>.

²²The resultant de-meaned value of $Y_{h,y,t}$ is interpreted as abnormal $Y_{h,y,t}$ with seasonality and household-year effects stripped away.

hold h files its tax return. Let $News\ Amount_{h,y,t} = File_{h,y,t} * News\ Amount_{h,y}$, so that $News\ Amount_{h,y,t}$ is the amount of news only on the day of filing, similarly to the way that $Refund_{h,y,t}$ and $Payment_{h,y,t}$ are the amounts only on the day a refund is received and a payment made, respectively. Finally, let $PosNews_{h,y,t} = \max[News\ Amount_{h,y,t}, 0]$ and $NegNews_{h,y,t} = \max[-News\ Amount_{h,y,t}, 0]$.

Our main estimating equation is

$$\begin{aligned}
Y_{h,y,t} = & \sum_{k=-28}^K \beta_k^+ PosNews_{h,y,t+k} + \sum_{k=-28}^K \beta_k^- NegNews_{h,y,t+k} + \sum_{k=-28}^K \phi_k File_{h,y,t+k} \quad (3) \\
& + \sum_{k=-28}^K \gamma_k^+ Refund_{h,y,t+k} + \sum_{k=-28}^K \gamma_k^- Payment_{h,y,t+k} + \alpha_{h,y} + \tau_t + u_{h,y,t}
\end{aligned}$$

where Y is an account inflow or outflow measure $\alpha_{h,y}$ is a household-year-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 capture the prior, contemporaneous, and lagged response of the dependent variable to news about a refund or payment ($\hat{\beta}_k$), to the date of filing ($\hat{\phi}_k$), and to getting a refund or making a payment ($\hat{\gamma}_k$). The responses to refunds, payments, and news are measured as a share of the refund, payment, and news, respectively. Thus, for example, when Y is consumption spending, $\hat{\gamma}_k^+$ measures the increase in consumption caused by the arrival as a share of the refund amount k days after the refund arrives, and $\hat{\beta}_k^-$ measures the increase in spending caused by the arrival of information as a percentage of the news uncovered during filing k days after filing. We refer to these coefficients as the marginal propensity to consume out of refunds and out of negative news.

Most of the identification of $\hat{\beta}_k$ and $\hat{\phi}_k$ comes from the “event time,” which is relative to the day of filing. Similarly, for identification of $\hat{\gamma}_k$, event time is relative to the day a refund arrives or the payment is made. The day fixed effects (τ_t) control for 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. Finally, the household-year fixed effects ($\alpha_{h,y}$) absorb any correlation between the average household-specific level of outflows and refund or payment amount. Such differences in our data arise not only due to differences in economic circumstances such as standard of living but also due to possible differences in

the scope of our measurement, such as the share of actual consumption spending that we identify as such.

We smooth the daily impulse responses by imposing that the daily coefficients are constant within weeks from $k = -28$ to -15 days, and for $k > 14$ days. Standard errors allow for arbitrary heteroscedasticity, within-day correlation, and within-household-year correlation in $u_{h,t}$.

We report cumulative impulse responses of our regression estimates from equation 3, which are simply the sums of estimated coefficients over k up to some κ . For example, for consumption and γ^+ , the impulse response at κ , $\sum_{k=-28}^{\kappa} \gamma_k^+$, is the total marginal propensities to consume out of the refund starting 28 days before receiving the refund up to κ days after, controlling for seasonal factors, individual differences, and the effect of news gathered during and shortly before filing. Standard errors for the cumulated daily total are calculated for the endpoint of each discrete interval (using the variance-covariance matrix of the coefficients).

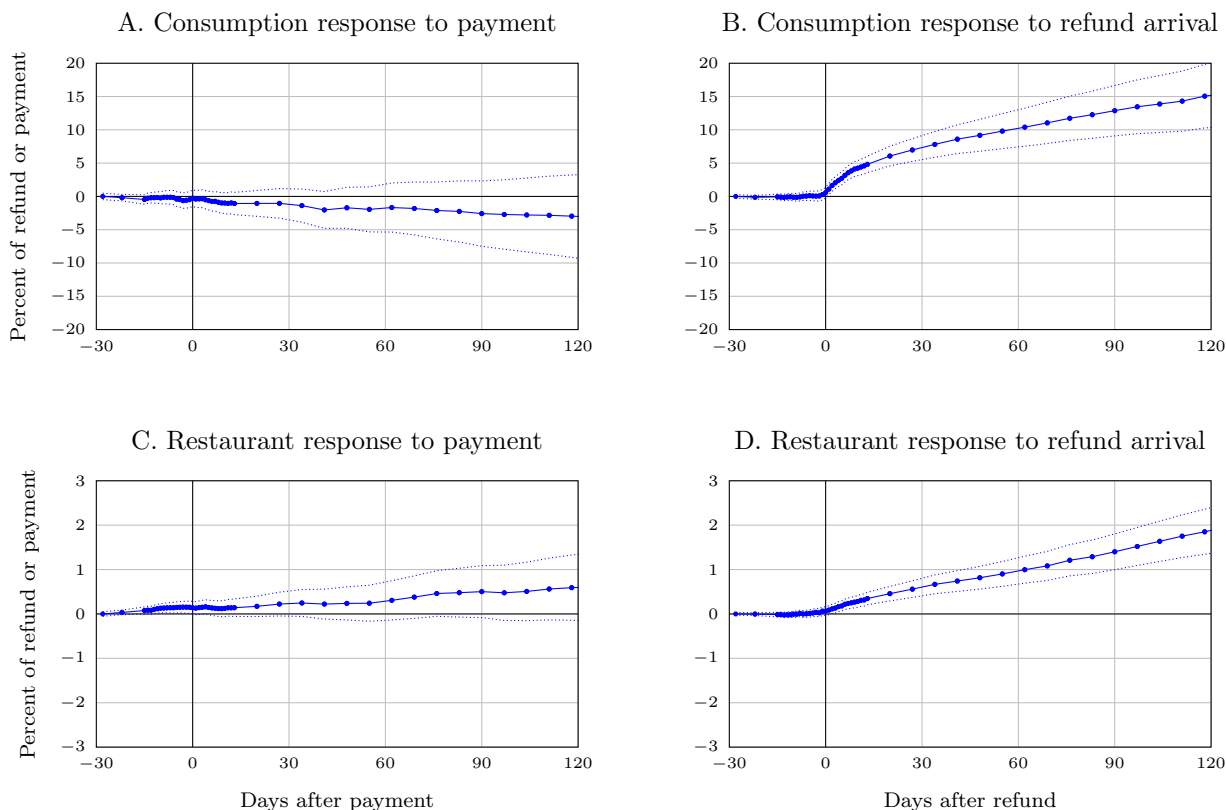
6 Asymmetric Consumption Responses

In this section, we present our main finding: in our sample, households have different consumption responses to expected increases and expected decreases in income. The following sections investigate why. Section 7 shows that the asymmetric response is not driven purely by liquidity because it occurs for the most and least liquid households in our sample. Section 8 shows that households have substantial liquidity and use liquidity management to smooth consumption through payments but not refunds. Section 9 shows that, despite making transfers across accounts in response to bad news about refunds and payments, households do not adjust consumption in response to such news.

Consumption Response to Tax Refunds Our first result is that households increase consumption spending when they receive refunds but do not decrease spending when they make tax payments. Specifically, Figure 4 shows estimates of cumulated coefficients, $\sum_{k=-28}^{\kappa} \gamma_k^+$ and $\sum_{k=-28}^{\kappa} \gamma_k^-$, for different horizons κ from the estimation of equation 3 on our measure of consumption spending. Each impulse response shows the cumulative increase in spending

Figure 4. Consumption Response to Payment of Taxes and Arrival of Refunds

The figure shows the consumption response to making tax payments (panels A and C) and the arrival of tax refunds (panels B and D). Panels A and B show the response of all transactions classified as consumption. Panels C and D show the subset of transactions classified as restaurants. The x -axis represents the number of days after the tax payment or receipt of refund. The y -axis shows the propensity to spend out of payment or refund. These responses are computed from equation 3. The dotted lines represent two standard errors confidence intervals. Standard errors are doubled-clustered at the household-year and at the calendar date levels.



as a percentage of refund and as a percentage of payment starting 28 days before the refund arrived or payment was made. These cumulative spending responses are also reported in the first row of Table 3.

Figure 4, panel B, shows that people consume about 8% of their refunds over the month following receiving a refund, a number which rises to 15% of their refund over the following three months. The spending increase starts the *exact day* on which the refund arrives. We find no evidence of increases in consumption spending prior to the day of arrival (related to the timing of refund arrival rather than filing).

Consumption Response to Tax Payments In contrast to the consumption response to refunds, the response to payments is small and statistically indistinguishable from zero. Panel A of Figure 4 shows the change in spending around the time when households make a tax payment, controlling for the arrival of information about payment. There is a slight decline in consumption that accumulates to (an insignificant in both senses) 4% of payment after four months. We also find no evidence of any decline in spending around the day of payment as we might have expected given the strong response to refunds on the day of arrival. These results focus on the response to cash flow but confirm the temporal pattern of the spending asymmetry displayed in Figure 1, panel B, which showed that month-after spending increases linearly with refund amount and does not decrease with payment amount.

Because our measurement of consumption is conservative, our estimated propensities to consume out of refunds are surely underestimates. The response of outflows that we cannot characterize—which primarily consist of checks and so is probably mostly consumption—has a pattern and magnitude very similar to that of consumption. The final row of Table 3 shows that miscellaneous payments exhibit the same asymmetry as our baseline consumption measure. If miscellaneous payments were consumption, then the propensity to consume out of refunds would be 17% after a month and 28% after four months, which are more than twice as large as our baseline, conservative estimates.²³

Nondurable Consumption Response to Tax Refunds and Payments Panels C and D of Figure 4 show the same response of consumption to payments and refunds, but only for a narrow and easily-identifiable type of consumption: spending at restaurants.²⁴ These figures show that spending on restaurants exhibits the same clear asymmetry, consistent with consumption rising substantially in response to refund receipt but being stable around the time tax payments are made. Panels C and D also rule out the interpretation that only spending on durable goods increases, which could then just represent an increase in

²³Around tax payments, miscellaneous payments rise instead of falling as they would if tax payments were reducing consumption. This increase possibly reflects payment of state and local taxes. Note that our results for our baseline consumption measure hold almost identically using only households in states without state income taxes.

²⁴We use the classification of restaurant transactions provided by the debit card or credit card providers, so this includes everything from fast food chains to gourmet dining experiences.

Table 3. Cumulative Consumption Response to Payments and Refunds

This table shows the cumulative response (in percent) of account outflows to expected payments and refunds. The cumulative response is calculated from day -29 , i.e., one month prior to the payment or refund. The cumulative response is calculated as $\sum_{k=-29}^{\kappa} \gamma_k^+$ and $\sum_{k=-29}^{\kappa} \gamma_k^-$, for different horizons κ from the estimation of equation 3 on the measure of consumption spending. Standard errors, shown in parentheses, are clustered by the household-year and calendar day.

| Variable: | Panel A: Percent of Payment | | | | Panel B: Percent of Refund | | | |
|---------------------------|-----------------------------|-----------------|-----------------|-----------------|----------------------------|-----------------|-----------------|-----------------|
| | Days after Payment | | | | Days after Refund | | | |
| | 0 | 28 | 56 | 84 | 0 | 28 | 56 | 84 |
| Consumption | -0.49 (0.62) | -1.04 (1.13) | -1.95 (1.73) | -2.25 (2.34) | 0.27 (0.41) | 6.98 (0.87) | 9.80 (1.34) | 12.27 (1.81) |
| Restaurant | 0.15 (0.07) | 0.22 (0.14) | 0.24 (0.21) | 0.48 (0.28) | 0.06 (0.05) | 0.56 (0.10) | 0.90 (0.14) | 1.29 (0.19) |
| Paid Using Credit Card | -1.09 (0.63) | -0.07 (1.15) | 0.24 (1.70) | 1.31 (2.30) | 0.43 (0.37) | 2.25 (0.71) | 3.57 (1.06) | 5.06 (1.42) |
| Miscellaneous Payments | -0.59 (1.59) | 7.96 (3.10) | 6.78 (4.45) | 6.84 (5.90) | 0.11 (0.84) | 10.14 (1.67) | 13.26 (2.46) | 15.25 (3.24) |

savings/investment. Thus, our estimated spending responses are neither purely driven by spending on durable goods nor spuriously driven by miscategorized or misinterpreted account outflows.

We again emphasize that our results are not driven by differences across households that receive refunds and those that make payments because all households in our sample both make payments and receive refunds (in different years).²⁵ However, households can differ across years, and in particular, they have different amounts of liquidity. We address this issue in detail in Sections 7 and Section 8. Before doing so, we first address two possible concerns with our results so far.

The Timing of Refunds and Payments Could the consumption asymmetry be driven by the endogeneity of the timing of filing? People can always postpone payment until

²⁵It is also true that persistent differences across households within the canonical model, such as from differences in impatience, would not produce our main consumption asymmetry. More impatient and so more illiquid households should have larger sensitivity of consumption to cash flows, should withhold less, and thus should be more likely to make payments. Thus, differences across households in impatience or liquidity would lead households with lower spending reactions to be more likely to get refunds.

the deadline in mid-April.²⁶ A bias from this source seems unlikely because we find the asymmetric consumption response for many different, more homogeneous subgroups of our sample.

For example, consider the possibility that the spending response to refunds is driven by years in which households have little liquidity and receive large refunds so that they cannot smooth consumption in these years. Panels A and B of Figure 5 show that there is a large asymmetry in consumption response among only those households expecting either payments or small refunds. This subsample contains only household-years in the lowest quintile of the expected value of $Refund - Payment$. In this subsample, $Refund - Payment$ averages $-\$62.40$ (a payment, relative to a refund of $\$564$ in the whole sample) with a standard deviation of $\$2,966$.

Focusing more closely on the issue of timing, the consumption asymmetry is observed for people independent of when they file their taxes. Panels C and D of Figure 5 show that we find the same asymmetry for households filing in each month of the year. Households filing in April are partly those with substantial liquidity who are not constrained, and yet we find substantial spending out of refunds. Alternatively, households who file and pay taxes late have the least time to save and prepare to smooth consumption. Yet, these households also both smooth consumption over payments and spend out of refunds. This lack of differences in spending responses by month constitute evidence against a specific 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.²⁷

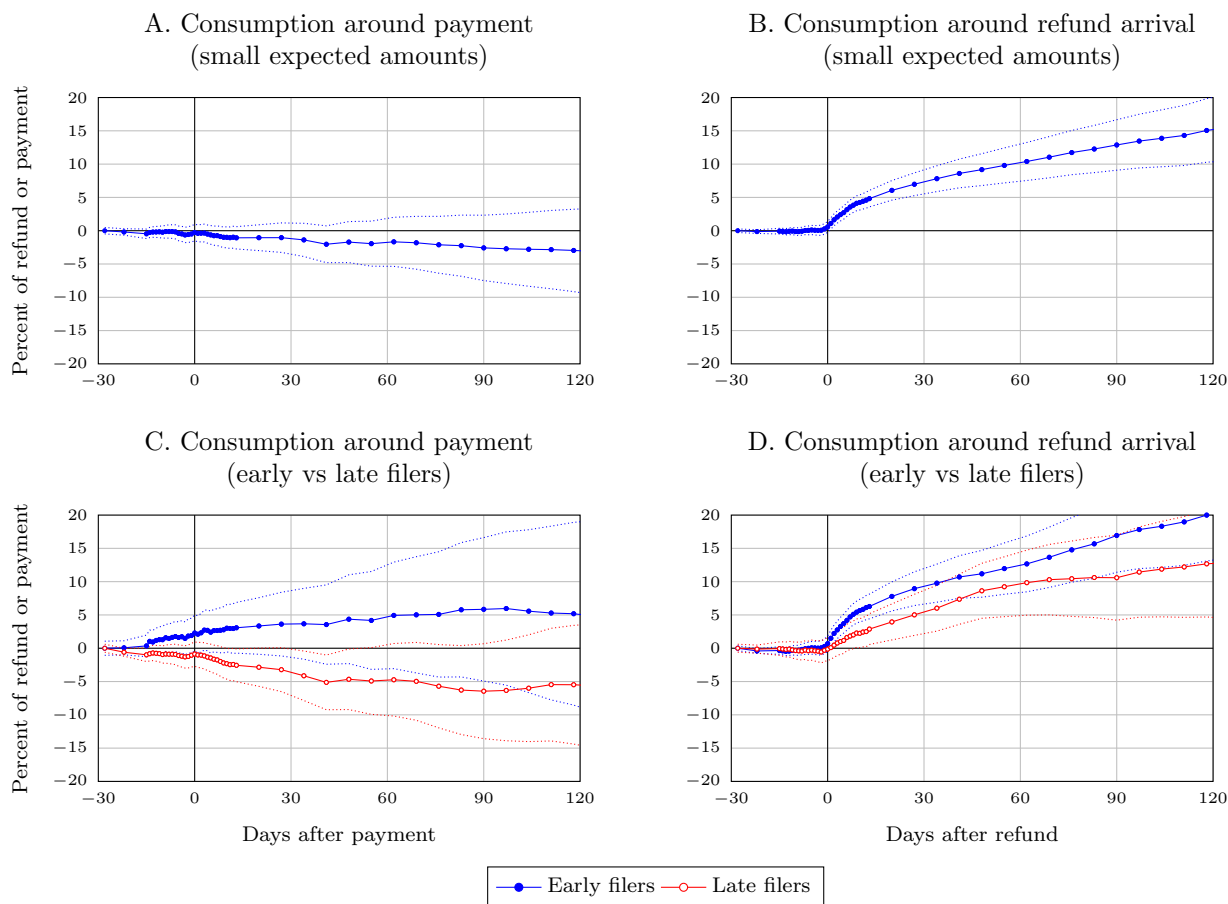
Mismeasurement and Robustness Could mismeasurement generate our finding of an asymmetric consumption response? Both the timing and amount of news about refund

²⁶Of households that file in February and owe taxes, 33% pay in April, and the average time between filing and payment is 23 days. Of those that file in March, 47% pay in April, and the mean time between filing and payment is 12 days.

²⁷Another 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.

Figure 5. Consumption for Small Expected Amounts and Early/Late Filers

The figure shows the consumption response to making tax payments (panels A and C) and to the arrival of tax refunds (panels B and D). Panels A and B show the response among the subsample of household-years with small expected amounts, defined as the bottom quintile of expected amount (refund less payment). Panels C and D compare the consumption response of early versus late filers. Early filers are defined as those who filed their returns before March. Late filers are those who filed their returns in April. The dotted lines represent two standard errors confidence intervals. Standard errors are doubled-clustered at the household-year and at the calendar-date levels.



or payment are measured with error. Measurement error is particularly a concern for the consumption response to payments because most payments are made within three days of filing, whereas no refunds arrive on the day of or day after filing, and most arrive after a two or more weeks delay. Our statistical procedure thus might not cleanly separate the consumption response to filing and news from the response to payment.

However, payments are on average associated with bad news, and households should respond to bad news by decreasing consumption or at least not increasing it. Thus, if

our procedure were not cleanly separating the effect of making a payment from the effect of filing and news about payment, we would be biased toward finding a *larger* decline in spending in response to making payments, not the insignificantly small changes we actually find. Further evidence is provided by the differential effect on the exact day of payment and refund.²⁸ Thus, this type of mismeasurement cannot account for our asymmetry.

Because there is always a temporal delay between filing and refund, our methodology has much more power to separately identify the response to news and the response to payments or refunds.

Moreover, our main asymmetry is a very robust result. We find the asymmetry when the equation is estimated with a log dependent variable and indicators of refund, payment, positive news, and negative news. We find similar results if we instead *i*) omit the controls for the arrival of news about tax status, *ii*) interact the amount of the refund or payment with the timing of filing as a control (instead of interacting the amount of news about future taxes), and *iii*) control for the news that arrives but with the timing related to the cash flow rather than filing.

We conclude that our main finding is unlikely due to a variety of possible biases. Therefore 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 Asymmetric Consumption Smoothing by Liquidity

In this section, we show that liquidity constraints have a limited role in driving the asymmetric consumption response. Households exhibit the consumption asymmetry across the liquidity distribution. First, households do not reduce consumption in response to payments even in the household-years in which they are in the lowest tercile of liquidity. Thus, even households with low liquidity have sufficient funds and debt capacity to stabilize consump-

²⁸To validate our estimation strategy, we confirmed that it measures the effect of the news and cash flows on the tax-induced cash flows with near perfect accuracy (which we found required daily data rather than data collapsed to the weekly level). For *Refund – Payment*, there is a 100% response to refund arrival and –100% response to payments, and no response to news. 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 estimated to decay over time.

tion. Second, these same households wait until arrival to increase spending in response to refunds even in the household-years in which they have substantial liquidity, although the increase in spending is more modest. To be clear, we are not arguing that liquidity is exogenous. In fact, in Section 8 we document that people adjust their liquidity around tax time and do have slightly more liquidity in years with payments.

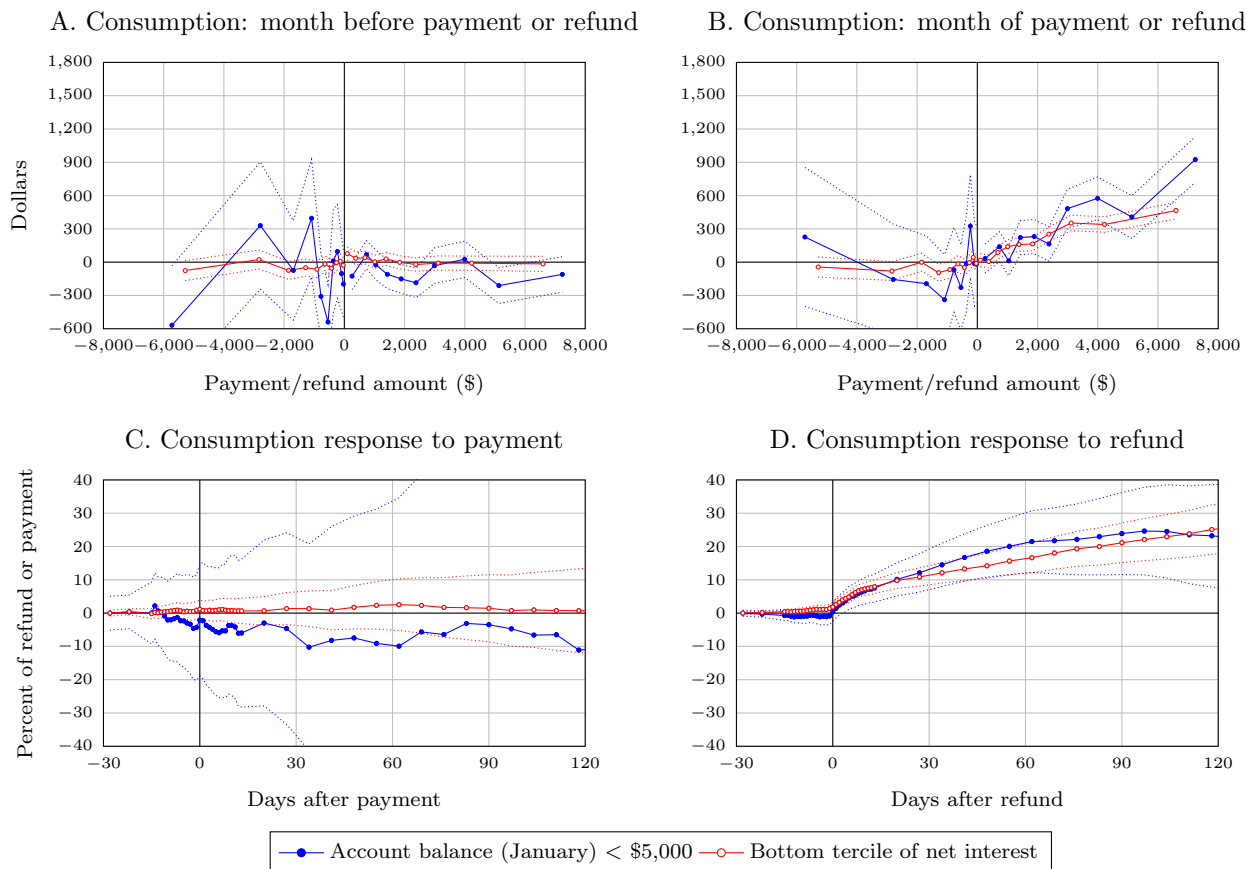
Households with Low liquidity We construct a sample containing the lowest tercile of households by liquidity using two different measures of liquidity. First, we select the lowest tercile of net interest earned during November, December, and January, which we construct as the interest earned on checking and savings accounts less the interest paid on linked credit card accounts. To be included in our low liquidity sample requires paying at least \$8.07 in net interest expense per month. Since interest rates on checking/savings are close to zero over this time period, this implies a revolving credit card balance of roughly \$500. Our second measure of low liquidity is based on the small subsample of accounts for which we can observe the account balance from a text line in the account.²⁹ We start with our broad sample (as described in Table A.I) and keep only accounts for which we ‘observe’ balances in this way for all core accounts in January, which is less than 5% of our sample, and then among these, we keep only households with total balances less than \$5,000.

Figure 6 displays the spending response for households with low ex-ante liquidity according to these measures. Panels A and B show that households with little liquidity do not cut consumption prior to making payments, and that they increase consumption after receiving refunds. Panels C and D show that for these household-years in the bottom of the distribution of liquidity, payments do not cause households to reduce consumption prior to or after payments, but refunds do cause these households to increase consumption following the arrival of refunds. This pattern is strong evidence that low liquidity is not driving the spending response to refunds since it is not hindering the smoothing of consumption through tax payments.

²⁹E.g., “\$0.16 interest earned for average daily balance of \$3,810.72.”

Figure 6. Consumption Response for Households with Low Liquidity

Panels A and B show the abnormal consumption around tax payment and refund dates, as a function of payment and refund amounts for household with low liquidity, defined in two different ways. The markers denote averages at every 10% of the data for those who received refunds, and every 10% of the data for those who made payments. Panels C and D show the cumulative response of consumption to making tax payments and receiving refunds respectively. The horizontal axes measure days since payment or refund arrival. The dotted lines represent two standard errors confidence intervals. Standard errors are doubled-clustered at the household-year and at the calendar date levels.



Households with High liquidity We focus on households with highly liquid using similar measures of net interest and balances. For our measure based on net interest, we drop all households without at least one credit card since a lack of credit cards suggest a lack of credit and debt capacity. Then, in this sub-sample, we keep all household-years with above zero net interest earned, which effectively removes any household paying interest on a credit card balance. This leaves us with 15,861 households and 26,024 household-years. Our second measure is all households from the broad sample that have January balances exceeding \$20,000, leaving us with 10,689 households and 15,821 household-years.

Figure 7. Consumption Response for Households with High Liquidity

Panels A and B show the abnormal consumption around tax payment and refund dates, as a function of payment and refund amounts for households with high liquidity, defined in two different ways. The markers denote averages at every 10% of the data for those who received refunds, and every 10% of the data for those who made payments. Panels C and D show the cumulative response of external saving and debt payment to making tax payments and receiving refunds respectively. The horizontal axes measure days since payment or refund arrival. The dotted lines represent two standard errors confidence intervals. Standard errors are doubled-clustered at the household-year and at the calendar date levels.

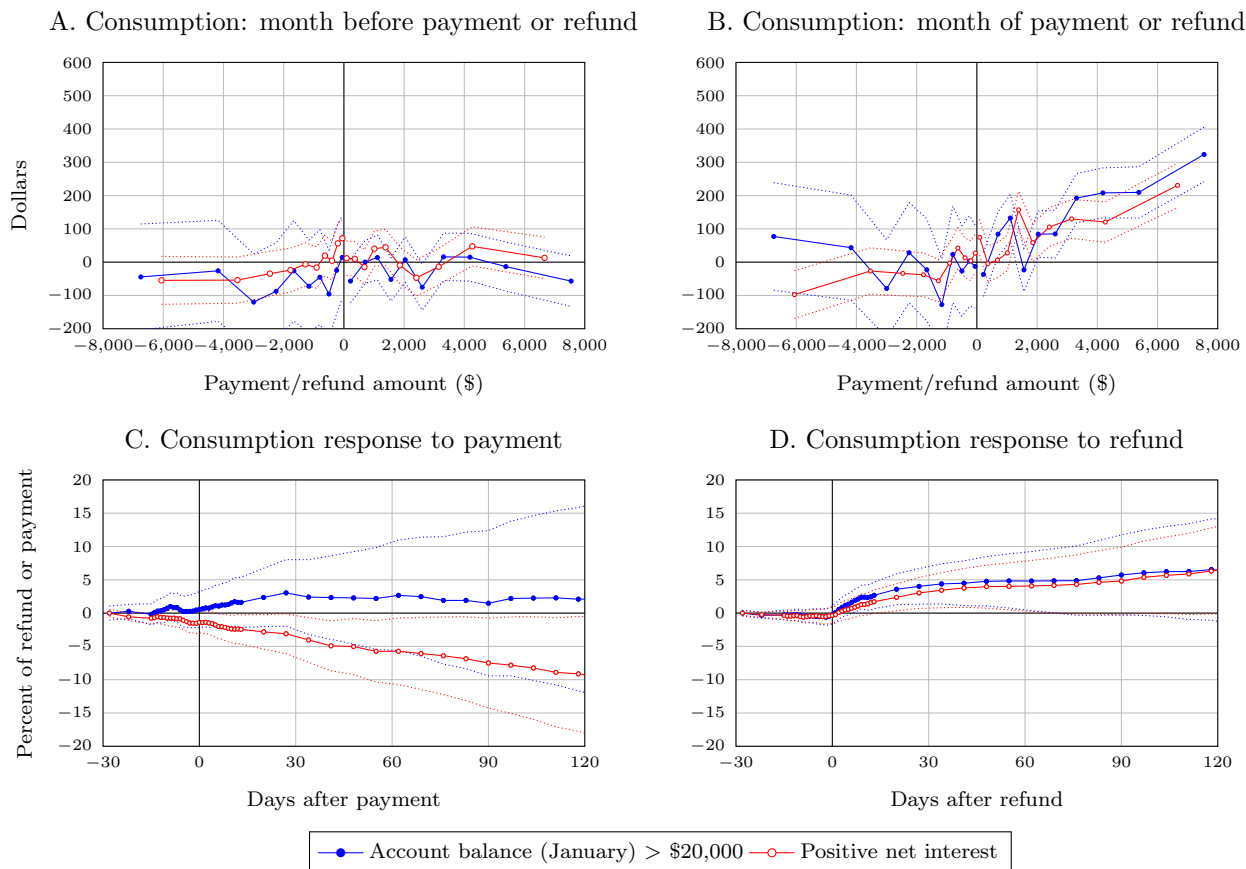


Figure 7 displays the spending responses for households with high ex-ante liquidity. Panel A shows that households with substantial liquidity do not cut consumption prior to making payments according to both measures of liquidity. More importantly, panel B shows that they do increase consumption after receiving refunds, a result that is both economically and statistically significant. Panel D shows that refunds cause these households to increase consumption spending in the three months following the arrival of a refund by 5-6% of the refund. While this result is only borderline statistically significant after three months (and again, likely an underestimate), the response begins on the day of refund arrival, as we

found for less liquid households.

Finally, panel C shows that payments do not cause households with substantial liquidity to reduce consumption prior to making payments. However, households in the top of the distribution of net interest earned do show some decline in consumption following a payment. This decline echos the similar but insignificant decline for all households (panel A of Figure 4), for those with small payments and refunds (panel A of Figure 5), and for those households filing in April (panel C of Figure 5). This evidence is inconsistent and not as large or significant as the increase in consumption following refunds. We do not even see hints of this behavior for households with low liquidity. Thus, despite being statistically weak, this evidence raises the possibility that liquid households also choose to reduce consumption some following payments (and seemingly by less and with a different time pattern).

This pattern of households with lower liquidity having larger consumption responses to expected increases in liquidity is typical of many previous findings in the literature, which, as noted, focuses almost entirely on expected increases in income or liquidity. The problem with interpreting this pattern as driven solely by liquidity constraints is that there is no consumption response to payments, even for households with lower ex-ante liquidity, implying that households can and do manage their finances so as to smooth consumption when they want to. One interpretation is that these households have a mental rule that prohibits them from tapping liquid savings for spending in anticipation of refunds but not in anticipation of payments. To be clear, liquidity surely plays a role in household consumption behavior in general, and the management of liquidity plays a role in the smoothing of consumption through payments as we show in the next section. However, in our sample, the negative correlation between liquidity and spending out of refunds appears to be driven by a correlation between the behavioral propensity to spend from refunds and liquidity.

We confirm these findings with other proxies for liquidity. 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 third of the income distribution (panels A and B of Figure A.I). We see no measurable declines in consumption spending when making payments among any of our three income groups. Households filing in February, who might be less liquid if they are impatient to get refunds, have stronger spending responses (as shown earlier

in Figure 5, panels C and D). Finally, households expecting small refunds or payments, who are less likely to be constrained, have spending responses to refunds. In all these cases, less liquid households spend out of refunds and thus appear possibly constrained, but smooth consumption through making tax payments, which shows that these households are not actually constrained.

The Endogeneity of Liquidity As we discuss in the next section, we find that in years when households make payments, they have slightly higher account balances in January. While this difference is small, and so cannot account on its own for the different responses to refunds and payments, it is consistent with households taking action to optimally smooth consumption over tax payments and not taking action to smooth consumption through refund arrival.

However, from a tax structure perspective, endogenous heterogeneity in liquidity seems unlikely to generate our findings. Households that receive a lot of non-labor income (that is not subject to automatic withholding) are likely to have more wealth and are more likely to make payments. But Gelman et al. (2019) shows that this effect is quantitatively small. Specifically, the authors present a model in which households with stochastic non-labor income face an approximation of the U.S. tax system. The model fits the observed distribution of payments and refunds and generates an average propensity to spend out of refunds of 30%. Most importantly for our purposes, the model implies relatively little variation in consumption reaction across the range of payments and refunds. The propensity rises by only 6% of payment/refund from a payment of \$4,000 to a refund of \$4,000. The rise is smooth, without even a kink at zero.

In sum, because there is no consumption response to payment, even for the least liquid household-years, we conclude that even these households have plenty of liquidity to stabilize consumption. Further, the most liquid household-years wait until refund arrival to increase spending, although the increase is more modest for this most liquid group. Thus households appear to be *choosing* to increase consumption only once refunds arrive but perfectly *smoothing* consumption through payments, mostly, as we now show, by tapping into liquid funds.

8 Asymmetric Liquidity Management

Our evidence so far has shown that households smooth payments but not refunds. One possibility is that households have little liquidity on average and that there is a fixed cost associated with accessing liquidity, such as due to complexity or an actual cost (e.g., taxation of capital gains). If this were the case then households could be willing to pay these costs to avoid large declines in consumption when making payments but not to increase consumption slightly ahead of refund arrival.

In this section we show that households in our sample appear to be quite liquid on average across the distribution of refunds and payments. They make payments by accumulating slightly more wealth in core accounts by January in years when they owe more taxes. And then they make transfers, almost entirely between observed, liquid accounts—e.g., from saving to checking account—both in response to an upcoming payment and in response to news uncovered during tax preparation and filing about higher payments or lower refunds. Households do not adjust earnings ahead of refunds or payments. Because these internal transfers are costless to make, this behavior arises due to mental costs or rules as discussed in Section 3.

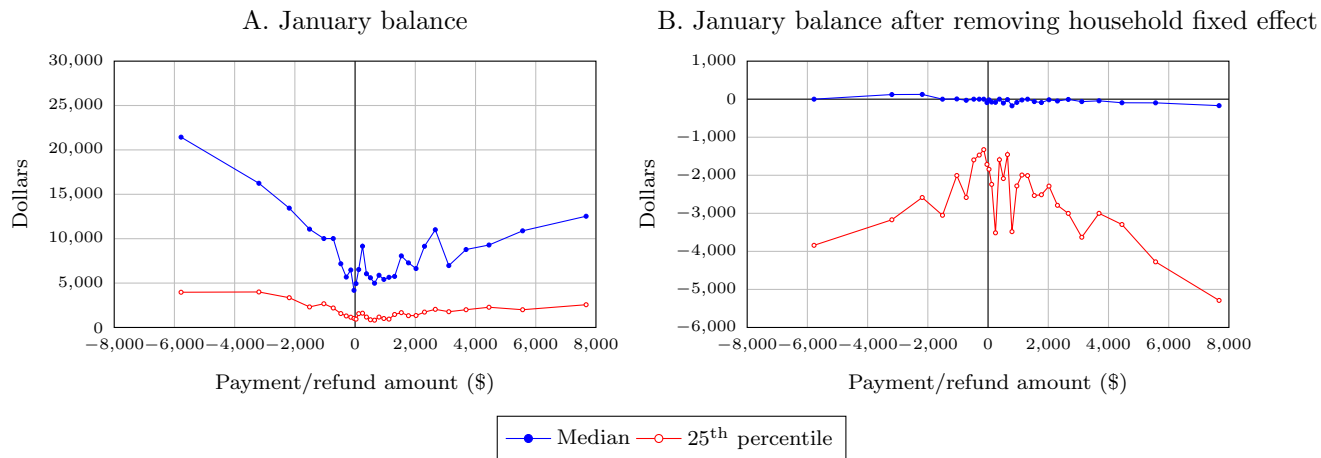
Account Balances As described in Section 7, we can observe account balances for a subset of our households by examining the text of interest transactions. While these balances ignore debt capacity on credit cards and liquid funds outside of the observed accounts, the balances that we do observe suggest that our sample of households has substantial liquidity, particularly relative to the amount of the typical payments and refunds (see Figure 2).

Panel A of Figure 8 plots the median and 25th percentile of the distribution of balances in January in each range of payments and refunds. These households have a significant amount of liquidity throughout the distribution of refunds and payments. The distribution is also somewhat skewed. Across this same distribution, households have on average \$15,000 to \$24,000 in liquid funds, far above the typical payment and refund amounts.

The median displayed in panel A reveals both slightly higher balances in years of payments and a slight V-shape. To some extent, these are both mechanical. Given our sample, a household that will make payment is more likely to have received a refund in the previous

Figure 8. January Account Balances Around Tax Payment or Refund

The figure presents the median and 25th percentile of the distribution of account balances in January for 10-equally sized bins of payments and 20-equally sized bins of refunds (same number of observations). Panel A is based on raw account balances, while panel B is based on account balances after removing household fixed effects (only for the 4,300 households with at least two years of balance information, leading to 10,778 household-years).



year, which, at least *ceteris paribus*, would raise the account balance. The V-shape arises because higher-income households tend to make larger payments and receive larger refunds, and have more wealth and a higher transaction demand for liquidity.

Panel B of Figure 8 plots the median and 25th percentile of the distribution of abnormal account balances, defined as the balance relative to the household mean January balance across years. The median difference in the account balance in a year when the household will make a payment is no different than in a year when the household will receive a refund. The 25th percentile of this distribution shows, if anything, that there are more households with relatively low account balances in years when they receive refunds than in years when they make payments. The inverted V-shape of the 25th percentile is a natural consequence of higher-income households having both greater year-to-year volatility of account balances and higher payments or refunds.

Account Transfers The pattern of transfers that households in our sample make reveals that both they accumulate liquid wealth following refunds and that they have liquid wealth and draw it down prior to making tax payments. Panel C of Figure 1 in the Introduction

Table 4. Cumulative Changes as a Percentage of Refund or Payment

This table shows the cumulative response (in percentage of refund or payment) of various variables to expected payments and refunds. The cumulative response is calculated from day -29 , i.e., one month prior to the payment or refund. The cumulative response is calculated as $\sum_{k=-29}^{\kappa} \gamma_k^+$ and $\sum_{k=-29}^{\kappa} \gamma_k^-$, for different horizons κ from the estimation of equation 3 on the measure of consumption spending. Standard errors, shown in parentheses, are clustered by the household-year and calendar day.

| Variable: | Panel A: Percent of Payment | | | | Panel B: Percent of Refund | | | |
|---------------------------------------|-----------------------------|-------------------|--------------------|---------------------|----------------------------|-------------------|-------------------|-------------------|
| | Days after Payment | | | | Days after Refund | | | |
| | 0 | 28 | 56 | 84 | 0 | 28 | 56 | 84 |
| Transfers In | 25.18 (4.38) | 25.52 (6.10) | 18.95 (7.87) | 13.86 (10.01) | 2.40 (1.28) | 32.80 (2.64) | 38.04 (3.81) | 42.11 (4.92) |
| External Savings and Loan Payments | -4.85 (2.66) | -5.25 (5.01) | -4.59 (7.32) | -3.91 (9.55) | 0.48 (0.99) | 13.90 (2.13) | 13.17 (3.26) | 13.01 (4.51) |
| Change in Account Balances | 19.24 (30.47) | -99.86 (59.72) | -114.23 (89.06) | -138.83 (117.57) | 21.62 (9.15) | 92.73 (18.19) | 99.56 (26.58) | 112.57 (35.18) |
| Probability of Overdraft | -0.04% (0.02%) | 0.01% (0.04%) | 0.03% (0.07%) | 0.00% (0.09%) | -0.02% (0.02%) | -0.17% (0.04%) | -0.24% (0.05%) | -0.32% (0.07%) |
| Income | -5.09 (7.40) | -17.40 (14.73) | -27.21 (21.72) | -33.95 (28.52) | 0.05 (2.75) | -1.70 (5.58) | -1.96 (9.32) | -1.04 (13.11) |

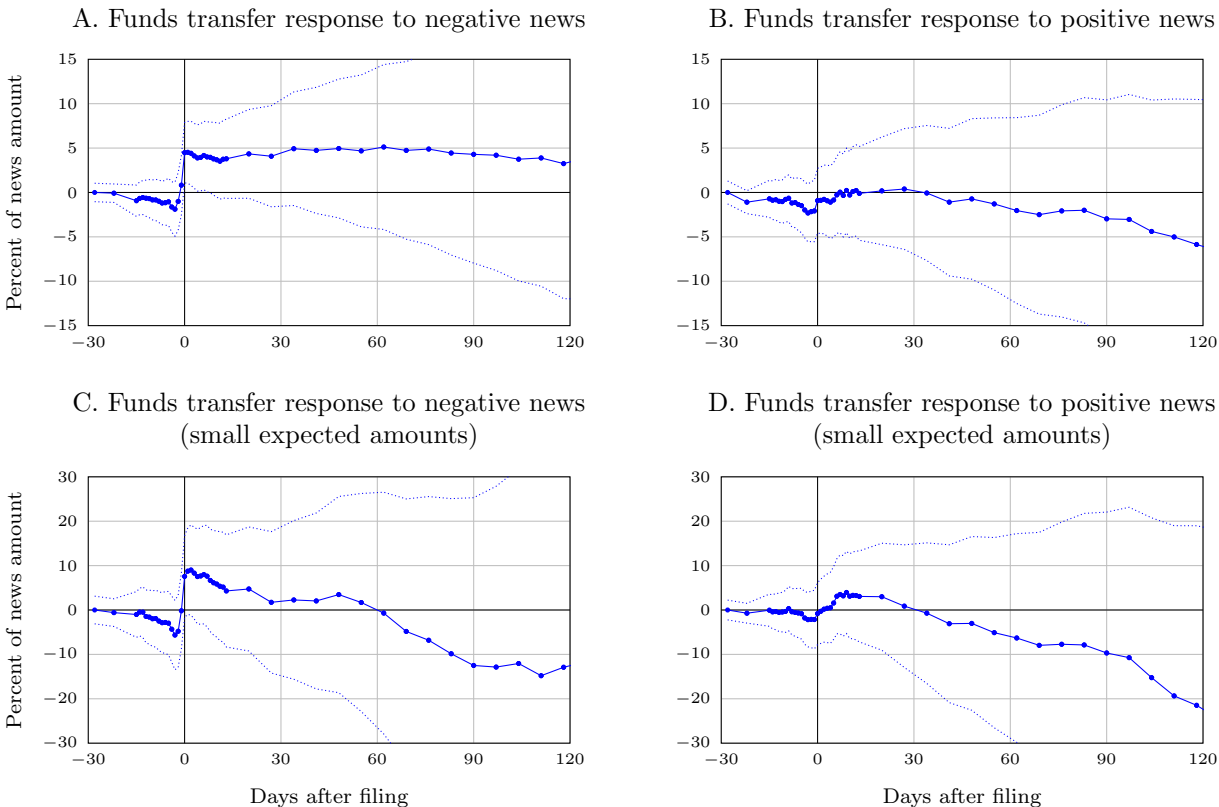
showed that households smooth consumption through tax payments by increasing transfers into either observed account (including transfers from one account to the other) in the month before making a payment. These transfers are equivalent in amount to about a quarter of the anticipated payment.

The first number in Table 4 shows an almost identical result from our estimation of equation 3, which measures the cumulative response to payment from 29 days before making, controlling for the arrival of information about the payment. Upcoming payments cause households to increase transfers into the observed accounts (including internal transfers) by 25% of the upcoming payment. The arrival of a refund increases transfers, but only following arrival, and presumably from checking accounts into savings accounts. More than 40% of the refund amount is transferred in the four months following the arrival (first row of Table 4, panel B). And this liquidity is presumably then available for the years in which households make payments.

We also find that households manage liquidity and make transfers between accounts in

Figure 9. Funds Transfers in Response to News about Tax Amount

The figure shows the fund transfer response around negative and positive news. Panels A and C show the response around negative news. Panels B and D show the response around positive news. Panels A and B show the response for all household-years. Panels C and D show the response of household-years with small amounts of expected payments or refunds, defined as being the bottom quintile of absolute expected refunds or payments. The x -axis represents the number of days after households filed their tax returns. The y -axis shows the dollars response per \$100 payment or refund. The response is computed from equation 3. The dotted lines represent two standard errors confidence intervals. Standard errors are doubled-clustered at the household-year and at the calendar date levels.



response to *news* about payments or refunds. Panel A of Figure 9 shows that households also increase transfers into their accounts when they learn that they will either owe more taxes than expected or receive lower than expected refunds. There is no corresponding reaction to good news. Thus, households react to bad news by moving funds across accounts to be able to make payments, but they do not move funds to increase consumption in response to goods news (a result confirmed by the analysis of the consumption response to news in the next section).

Perhaps the non-response of households to good news about taxes is because households expecting large refunds are constrained? That is, households who are saving to make pay-

ments or who are only expecting to receive small refunds can save less and consume a little more in response to news that they will owe less taxes this year. Panels C and D of Figure 9 rule out this hypothesis. Households expecting payments or small refunds also react to bad news about the tax amount and do not respond to goods news.

The response of transfers to upcoming payments, to the arrival of refunds, and to bad news, all confirm what we see in the simple plots in Figure 1. Households actively prepare to make payments. They transfer funds to smooth consumption. But they do not transfer funds to raise consumption in advance of refunds.

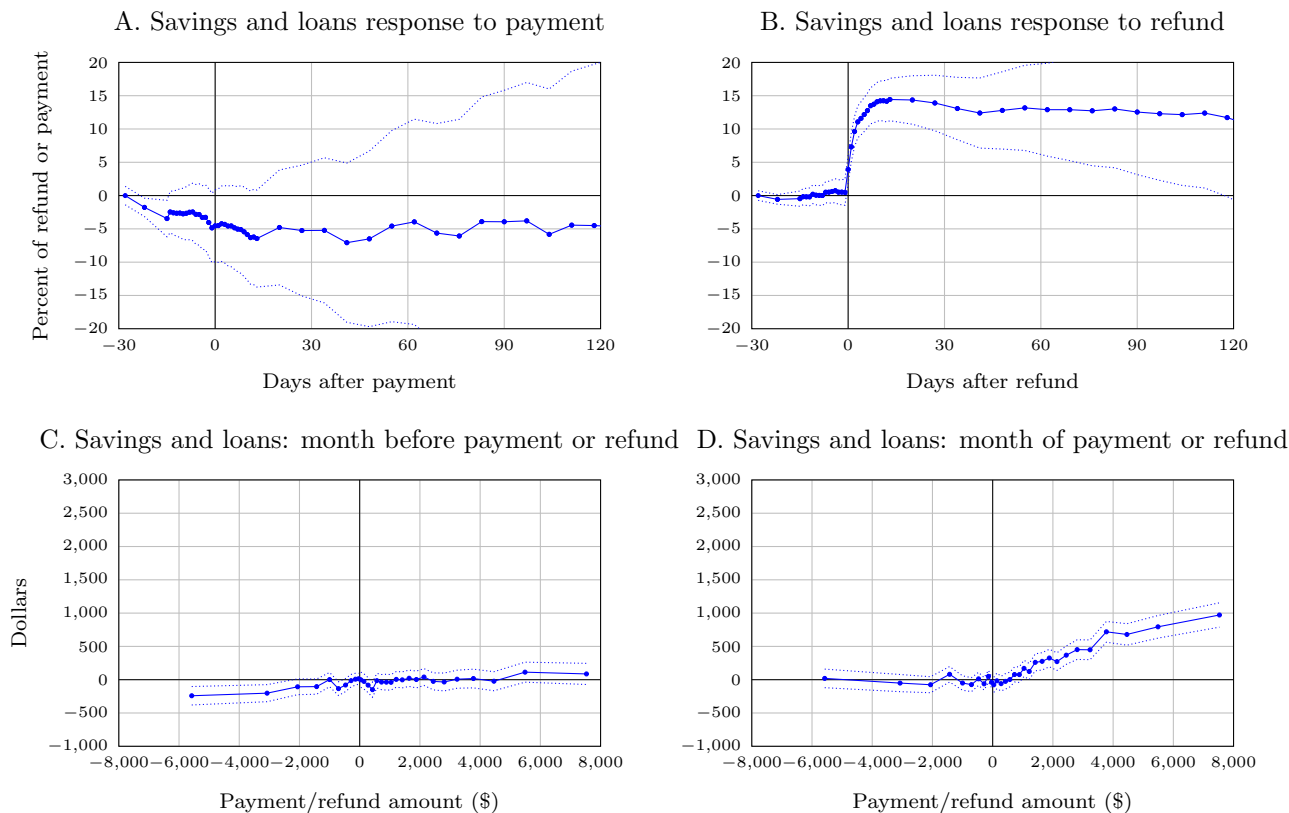
Debt Payments and Transfers from External Accounts Only a small share of these transfers in advance of payments come from accounts other than the primary checking and savings accounts. We aggregate all transfers to financial accounts outside of these core accounts and all debt payments (including to credit card accounts) to measure net external savings (or reduction in debt). Panel A of Figure 10 shows that households reduce debt payments and external saving in response to impending payments. But this net dissaving accounts for only 5% of the upcoming payment. Comparing rows 1 and 2 of Table 4, we see that this accounts for only one-fifth of the 25% of payment amount transferred into the account in the 28 days prior to payment. Thus, most transfers we observe are from internal rather than external accounts.

Panel C of Figure 10 confirms this result by simply plotting the abnormal dissaving from external accounts (reduced loan payment, decreases net savings). Households accumulate less than \$500 from these external sources. Panel C of Figure 1 in comparison shows a large increase in transfers between core accounts. We conclude that while transfers rise prior to payment, tax payments are largely made using funds in observed, liquid accounts.

In response to refunds, panels B and D of Figure 10, and panel B of Table 4 all show that the arrival of a refund causes a subsequent increase in external savings and debt payment. Taking panel B of Figure 10 for example, the cumulative increase in transfers to non-core accounts rises to 13% of the refund three months following arrival. Comparing rows 1 and 2 of panel B of Table 4 shows that this is still only a third of total transfers observed, so that most transfers triggered by tax refunds also occur between core accounts.

Figure 10. Response of Saving and Loans to Tax Payments and Refunds

Panels A and B show the cumulative response of external saving and debt payment to making tax payments and receiving refunds respectively. The horizontal axes measure days since payment or refund arrival. Panels C and D show the abnormal external saving and debt payment around tax payment and refund dates, with respect to payment or refund amounts. The markers denote averages at every 5% of the data for those who received refunds, and every 10% of the data for those who made payments. The dotted lines represent two standard errors confidence intervals. Standard errors are doubled-clustered at the household-year and at the calendar date levels.



The Dynamics of Account Balances Given the modest increases in consumption, miscellaneous outflows, and external savings and debt repayment following refunds, we try to measure changes in core account balances over time in response to refunds (and decreases in response to making payments). We create synthetic measures of account balances by cumulating inflows and outflows over time. Unfortunately, because the evolution of account balances is highly variable across households and over time, our estimate of the impact of tax payments on account balances has very low power; nevertheless, (statistically insignificant) points estimates suggest that balances rise by 20% of payments in advance of payment and then a month after the payment have fallen by more than 100% of the payment (row 3 of

Table 4). In response to refunds, our results are also noisy, but now statistically significantly different from zero. We estimate that core account balances rise during the month ahead of refunds by 22% of the refund, and points estimates suggest that account balances have risen by the entire refund amount by two months after arrival, but 95% confidence intervals contain more reasonable estimates (including roughly 50% of the refund amount).

To address the issue of low power, row 4 of Table 4 shows the response of a proxy for core account balances, the probability of an overdraft. The decline in overdrafts before payments confirms that balances rise prior to payments. Similarly, the overdraft probability confirms that balances fall after payments and rise following refunds (all consistent with the balances in core accounts).

Earned Income and Consumption Smoothing Households could, in theory, increase earnings to increase liquidity and make tax payments, but we do not find any increases in income inflows to their accounts ahead of an expected payment (final row of Table 4).³⁰

Delay of Tax Payments and Refund Anticipation Loans A final way that people might manage liquidity is by delaying payment. Could our main asymmetry be driven by this asymmetry in the tax system? That is, households do not actually have to pay taxes due at the April 15 deadline. Instead, they can postpone payment and borrow at a significant interest rate and potentially face penalties from the government. Ultimately, of course, taxpayers who fail to pay their taxes face the possibility of incarceration, so this strategy is not without risk, but borrowing from the government is an option in the short term.

We note that the institutional structure is actually roughly symmetric. Households can delay payment and borrow from the government, but they can also borrow in advance of a refund. That is, once taxes are filed, the payment from the government provides sufficient collateral that most tax preparers and preparation programs will advance the refund to the taxpayer upon filing (e.g., TurboTax Refund Advance in some states). Thus, for both refunds and payments, the taxpayer has a high-interest-rate option that would allow them to smooth

³⁰We also find no evidence that income responds to any news about the tax amount, or that it falls before or after refund receipt.

consumption.³¹

We conclude that, in practice, people have a roughly symmetric opportunity to postpone payment even further (and suffer interest penalties) or to borrow against refunds (and pay high interest rates). This type of liquidity management occurs but is not the source of our asymmetry.

In sum, households smooth consumption through payments by accumulating more wealth, by reacting to bad news when it arrives, and by transferring funds (primarily) among liquid accounts. They also manage funds to accumulate slightly more liquid assets by January in years when they owe more taxes. And yet, they do not take these actions to smooth consumption in response to refunds but instead increase consumption following refund arrival.

9 Lack of Consumption Response to News

This section shows that households do not adjust consumption prior to payments or refunds as news arrives. Nor do they adjust consumption in response to impending payments or refunds, or when making a payment.

In theory, households with a lot of liquid wealth should increase consumption in response to good news about refund or payment and should reduce consumption in response to bad news about refund or payment symmetrically. But these responses should be trivial because tax refunds and payments are small fractions of wealth and lifetime income.

Alternatively, if households are forward-looking, smooth consumption, and face occasionally binding liquidity constraints, then theory predicts that the reaction of consumption to negative news should be larger than the reaction to positive news—the reverse of the asymmetry we observe for positive and negative expected cash flows. Households that have few funds and face tightly binding liquidity constraints are unable to adjust consumption in response to positive or negative news about future cash flows. Quantitatively significant and asymmetric consumption responses arise only for households that are “weakly” constrained or close to constrained. Households that are weakly constrained do not respond to good news

³¹People who delay payment to October 15 (or later) do not appear in our sample. Nor do people who take out refund anticipation loans. Thus, the asymmetry we observe is for the sample of households that follow the “normal” route and receive a refund or make payment on time.

but do cut spending in response to negative news that is large enough to (probabilistically) relax their constraints. Households that are close-to-constrained will respond to bad news but will not increase spending fully in response to good news that (probabilistically) imposes a future binding constraint.

We find that households do not decrease consumption in response to bad news about their refund or payment amount, nor do they lower consumption in advance of making tax payments. First, as shown in panel A of Figure 1 in the Introduction, consumption does not decline in the month before making a payment (that is, relative to the timing of “cash flow” time rather than filing) or rise once the payment is made.

Second, and more importantly, there is 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. Panels A and B of Figure 11 show these (lack of) consumption responses to good and bad news, respectively, relative to the event time of tax filing (and controlling for the consumption response to making a payment or receiving a refund).

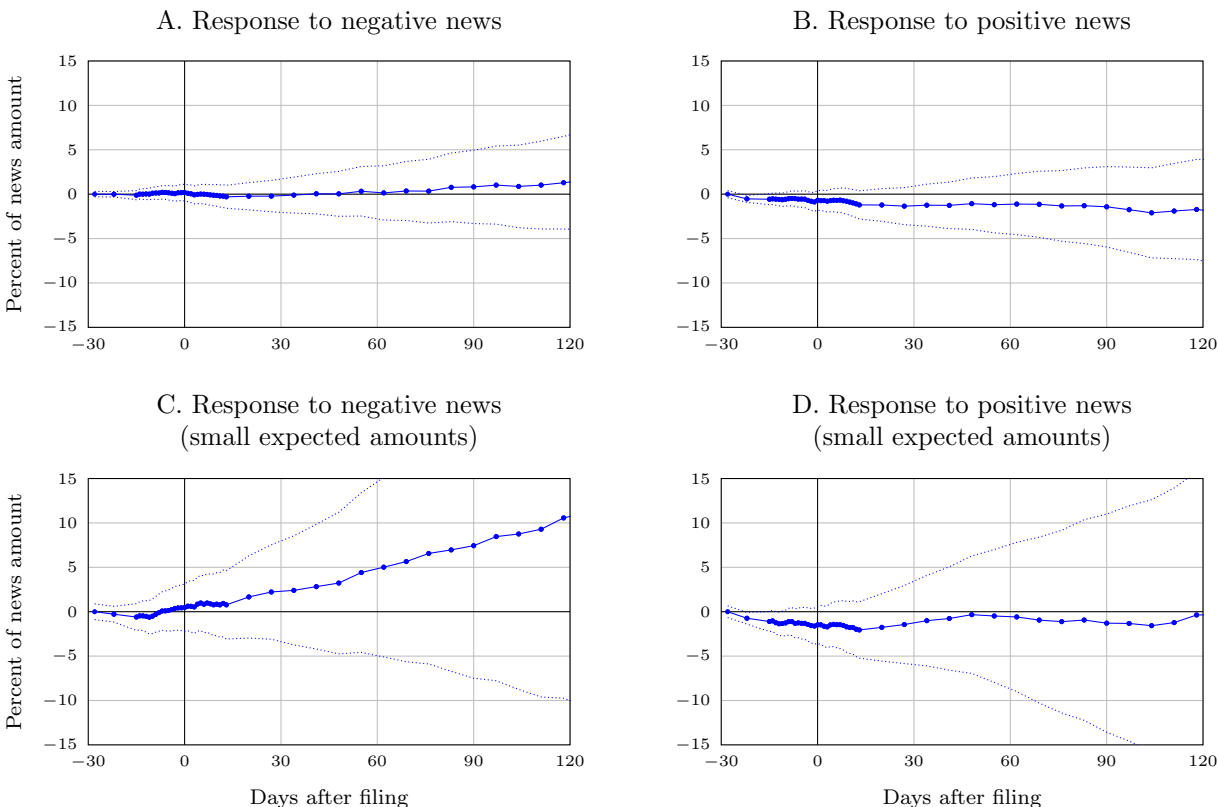
Finally, we again examine the subsample of households that are likely to make payments or receive small refunds. Households that are liquidity constrained by a future payment may adjust consumption nearly completely in response to news about the value of that payment. Those expecting small refunds may also be able to respond, while those expecting large refunds may simply be too constrained to respond to news. Panels C and D of Figure 11 shows that we find no statistically significant response of consumption to good or bad news about future payment or refund.

Do these results show that households do not adjust their consumption as news arrives, or could our findings be due to mismeasurement of news about taxes uncovered during filing? This could occur in two different ways.

First, households might have biased expectations about their refunds and so our measures of news could be incorrect. 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

Figure 11. Consumption Response Around News During Tax Preparation

The figure shows the consumption response around negative and positive news. Panels A and C show the response around negative news. Panels B and D show the response around positive news. Panels A and B show the response for all household-years. Panels C and D show the response of household-years with small amounts of expected payments or refunds, defined as being the bottom quintile of the absolute expected refunds or payments. The x -axis represents the number of days after households filed their tax returns. The y -axis shows the dollars response per \$100 payment or refund. The response is computed from Equation 3. The dotted lines represent two standard errors confidence intervals. Standard errors are doubled-clustered at the household-year and at the calendar date levels.



are receiving more money than expected.³² But prior evidence suggests that 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).

Second, it could simply be that households learn information about their future tax payment or refund far earlier than our statistical model suggests. If our measure of news and/or its timing were unrelated to the information uncovered during filing, then we would find no response of consumption or any other variable to news. In fact, our measure of news does contain information, and households react to this information. Figure 9 shows that

³²We find some smooth, slow increase in spending related to the time of filing.

households make transfers in response to our measure of news.

10 Final Discussion

We observe a specific sample of households that increase spending when they receive an anticipated tax refund, consistent with prior research measuring the presence of probabilistically-binding liquidity constraints or hand-to-mouth behavior (Zeldes, 1989; Olafsson and Pagel, 2018). This effect is stronger for less liquid households, as again has been interpreted as evidence of liquidity constraints. However, this behavior does not appear to be driven by a lack of liquidity: these same households completely smooth consumption when making anticipated tax payments, implying that they also have the liquidity to smooth consumption through refunds. They smooth consumption primarily by moving funds around between savings and checking accounts prior to making payments as well as when bad news about taxes due or refunds arrives. People bring in a small fraction of their payments from outside accounts. They also accumulate slightly higher account balances in January in years when they face higher payments or lower refunds.

Thus, in our sample, people spend out of tax refunds largely by choice rather than due to liquidity constraints. They have the liquidity to increase consumption prior to arrival, but they choose not to access it. This behavior is consistent, for example, with a heuristic in which the savings account is reserved for lumpy or necessary costs like tax payments and not used for discretionary current consumption. Not only do these households delay consumption until arrival, but they also choose to splurge, treating refunds as windfall gains and increasing consumption by more than dictated by consumption smoothing (e.g. if they were surprised by arrival). At the same time, the behavior of people facing tax payments is reasonably well-described by a textbook model of an inter-temporally-optimizing consumer. Therefore, we reject simple hand-to-mouth behavior in favor of one-sided deviations from optimal behavior, as in an asymmetric, mental-accounting heuristic.

References

- Aaronson, Daniel, Sumit Agarwal, and Eric French, 2012, The spending and debt response to minimum wage hikes, *American Economic Review* 102, 3111–3139.
- Agarwal, Sumit, Souphala Chomsisengphet, Stephan Meier, and Xin Zou, 2019, In the mood to consume: Effect of sunshine on credit card spending, Working paper, Columbia Business School Research Paper.
- Agarwal, Sumit, Chunlin Liu, and Nicholas S Souleles, 2007, The reaction of consumer spending and debt to tax rebates: Evidence from consumer credit data, *Journal of Political Economy* 115, 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, 4205–4230.
- Aguiar, Mark, and Erik Hurst, 2005, Consumption versus expenditure, *Journal of Political Economy* 113, 919–948.
- Arkes, Hal R, Cynthia A Joyner, Mark V Pezzo, Jane Gradwohl Nash, Karen Siegel-Jacobs, and Eric Stone, 1994, The psychology of windfall gains, *Organizational Behavior and Human Decision Processes* 59, 331–347.
- Aydin, Deniz, 2019, Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines, Working paper, University of Washington St Louis.
- Baker, Scott R, 2018, Debt and the response to household income shocks: Validation and application of linked financial account data, *Journal of Political Economy* 126, 1504–1557.
- Baugh, Brian, Itzhak Ben-David, and Hoonsuk Park, 2018, Can taxes shape an industry? Evidence from the implementation of the “Amazon Tax”, *Journal of Finance* 73, 1819–1855.
- Ben-David, Itzhak, and Marieke Bos, 2017, Impulsive consumption and financial wellbeing: Evidence from an increase in the availability of alcohol, Technical report, National Bureau of Economic Research.
- Bodkin, Ronald, 1959, Windfall income and consumption, *American Economic Review* 49, 602–614.
- Caballero, Ricardo J, 1995, Near rationality, heterogeneity, and aggregate consumption, *Journal of Money, Credit, and Banking* 27, 29–48.
- Caldwell, Sydnee, Scott Nelson, and Daniel Waldinger, 2018, Tax refund expectations and financial behavior, Working paper, MIT.
- Campbell, John Y, and N Gregory Mankiw, 1989, Consumption, income, and interest rates: Reinterpreting the time series evidence, *NBER Macroeconomics Annual* 4, 185–216.

- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten van Rooij, 2019, Asymmetric consumption effects of transitory income shocks, *Economic Journal* 129, 2322–2341.
- Farrell, Diana, Fiona Greig, and Amar Hamoudi, 2019, Tax time: How families manage tax refunds and payments, Technical report, JPMorgan Chase Institute.
- Fuster, Andreas, Greg Kaplan, and Basit Zafar, 2018, What would you do with \$500? Spending responses to gains, losses, news and loans, Working paper, National Bureau of Economic Research.
- Ganong, Peter, and Pascal Noel, 2019, Consumer spending during unemployment: Positive and normative implications, *American Economic Review* 109, 2383–2424.
- Gelman, Michael, Shachar Kariv, Matthew D Shapiro, and Dan Silverman, 2019, Rational illiquidity and consumption: Theory and evidence from income tax withholding and refunds, Working paper, National Bureau of Economic Research.
- 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.
- Gelman, Michael, Shachar Kariv, Matthew D Shapiro, Dan Silverman, and Steven Tadelis, 2020, How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown, *Journal of Public Economics* forthcoming.
- Hassan, Tarak, and Nicola Fuchs-Schundeln, 2016, Natural experiments in macroeconomics, in Ernst Berndt, and Charles Hulten, eds., *Handbook of Macroeconomics*, volume 2, 923–1012 (North-Holland).
- Hsieh, Chang-Tai, 2003, Do consumers react to anticipated income changes? Evidence from the Alaska permanent fund, *American Economic Review* 93, 397–405.
- Jappelli, Tullio, 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, 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, 1589–1610.
- Jones, Damon, 2012, Inertia and overwithholding: Explaining the prevalence of income tax refunds, *American Economic Journal: Economic Policy* 4, 158–185.
- Jorring, Adam T, 2018, The costs of financial mistakes: Evidence from U.S. consumers, Working paper, Boston College.
- Kaplan, Greg, Giovanni L Violante, and Justin Weidner, 2014, The wealthy hand-to-mouth, *Brookings Papers on Economic Activity* 1, 77–153.

- Kueng, Lorenz, 2018, Excess sensitivity of high-income consumers, *Quarterly Journal of Economics* 133, 1693–1751.
- Olafsson, Arna, and Michaela Pagel, 2018, The liquid hand-to-mouth: Evidence from personal finance management software, *Review of Financial Studies* 31, 4398–4446.
- Olafsson, Arna, and Michaela Pagel, 2019, Borrowing in response to windfalls, Working paper, Columbia University.
- Parker, Jonathan A, 1999, The reaction of household consumption to predictable changes in Social Security taxes, *American Economic Review* 89, 959–973.
- Parker, Jonathan A, 2017, Why don't households smooth consumption? Evidence from a 25 million dollar experiment, *American Economic Journal: Macroeconomics* 9, 153–183.
- 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, *Journal of Consumer Affairs* 51, 183–199.
- Reis, Ricardo, 2006, Inattentive consumers, *Journal of Monetary Economics* 53, 1761–1800.
- Shea, John, 1995, Myopia, liquidity constraints, and aggregate consumption: A simple test, *Journal of Money, Credit and Banking* 27, 798–805.
- Slemrod, Joel, Charles Christian, Rebecca London, and Jonathan A Parker, 1997, April 15 syndrome, *Economic Inquiry* 35, 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, 1187–1209.
- Souleles, Nicholas S, 1999, The response of household consumption to income tax refunds, *American Economic Review* 89, 947–958.
- Souleles, Nicholas S, 2000, College tuition and household savings and consumption, *Journal of Public Economics* 77, 185–207.
- Stephens, Melvin, 2001, The long-run consumption effects of earnings shocks, *Review of Economics and Statistics* 83, 28–36.
- Stephens, Melvin, 2003, “3rd of the month”: Do Social Security recipients smooth consumption between checks?, *American Economic Review* 93, 406–422.
- Stephens, Melvin, 2006, Paycheque receipt and the timing of consumption, *Economic Journal* 116, 680–701.
- Thaler, Richard, 1985, Mental accounting and consumer choice, *Marketing science* 4, 199–214.
- Zeldes, Stephen P, 1989, Consumption and liquidity constraints: An empirical investigation, *Journal of Economic Perspectives* 97, 305–346.

Table A.I. Summary Statistics: Broader Population

This table presents summary statistics for our final sample, which includes only households that make a tax payment in one year and receive a tax refund in another. *Refund – Payment* (\$) is the average refund (less payment) amount across household years. *Predicted Refund* (\$) shows the average predicted refund size. *News Amount* (\$) is the difference between the realized refund and the predicted refund. To ensure that the prediction is unbiased, we use households that only make payments or only have refunds to make this prediction, so the mean is nonzero. *Distance Filing Refund* (days) is the number of days between filing and the subsequent payment/refund. *I(Linked Credit Card)* is an indicator variable that equals one if the household has a credit card linked to the account aggregator, and zero otherwise. *I(Unlinked Credit Card)* is an indicator variable that equals one if the household has an unlinked credit card, and zero otherwise. *I(Any Credit Card)* is an indicator that equals one if the household has either linked or unlinked credit cards, and zero otherwise. *Net Flow* (\$) is the difference between inflows and outflows to the core accounts. *Consumption* (\$) is observed consumption. *Scaled Consumption* (\$) is the *Consumption* variable scaled up to compensate for the presence of unlinked credit cards. *Savings and Loans* (\$) shows payments to savings accounts (net investing outflows, net transfers out) and net loan payments (mortgage, auto loan, and net decrease in credit card debt). *Misc Payments* (\$) is the sum of checks and uncategorized outflows. *Net Interest* (\$), *Interest Expense* (\$), and *Interest Earned* (\$) are net interest, interest expense, and interest earnings, respectively. Finally, *Net Credit Card (CC) Charge* (\$) is all linked credit card expenditures after excluding tax-related transactions (such as filing fees and tax payments made on credit cards).

| Variable | Count | Mean | SD | p1 | p10 | p25 | p50 | p75 | p90 | p99 |
|--------------------------|-------------|----------|-----------|-----------|-----------|----------|----------|----------|----------|----------|
| Households | 199,077 | | | | | | | | | |
| Household-Years | 312,212 | | | | | | | | | |
| Filing Date | 312,212 | Mar 03 | 28.75 | Jan 15 | Jan 28 | Feb 06 | Feb 25 | Apr 02 | Apr 14 | Apr 19 |
| Refund Date | 312,212 | Mar 13 | 27.92 | Jan 30 | Feb 07 | Feb 18 | Mar 09 | Apr 11 | Apr 21 | May 04 |
| I(Positive Refund) | 312,212 | 0.90 | 0.30 | 0.00 | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| Refund – Payment (\$) | 312,212 | 2,351.38 | 2,541.05 | -4,199.00 | -10.00 | 753.00 | 1,905.08 | 3,825.00 | 5,893.00 | 8,959.00 |
| Predicted Refund (\$) | 312,212 | 2,338.91 | 1,638.32 | -1,050.33 | 751.65 | 1,241.90 | 1,986.13 | 3,300.95 | 4,696.11 | 6,788.52 |
| News Amount (\$) | 312,212 | 12.47 | 1,927.22 | -5,730.66 | -1,990.24 | -840.77 | -115.35 | 925.34 | 2,283.68 | 5,396.51 |
| Filing to Ref/Pay (days) | 312,212 | 10.08 | 9.80 | -1.00 | 3.00 | 6.00 | 7.00 | 11.00 | 18.00 | 57.00 |
| I(Linked Credit Card) | 312,212 | 0.72 | 0.45 | 0.00 | 0.00 | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| I(Unlinked Credit Card) | 312,212 | 0.83 | 0.37 | 0.00 | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| I(Any Credit Card) | 312,212 | 0.94 | 0.23 | 0.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| Household-Days | 104,055,302 | | | | | | | | | |
| Net Flow (\$) | 104,055,302 | -26.51 | 52,936.51 | -3,096.53 | -547.47 | -180.87 | -37.18 | 0.00 | 249.95 | 3,786.44 |
| Consumption (\$) | 104,055,302 | 89.17 | 165.70 | 0.00 | 0.00 | 0.00 | 27.58 | 105.48 | 243.00 | 835.08 |
| Scaled Consumption (\$) | 104,055,302 | 103.40 | 194.40 | 0.00 | 0.00 | 0.00 | 31.37 | 121.24 | 281.75 | 984.98 |
| Savings and Loans (\$) | 104,055,302 | 31.57 | 422.79 | -748.21 | -106.00 | -12.99 | 0.00 | 0.00 | 30.60 | 1,997.19 |
| Misc Payments (\$) | 104,055,302 | 14.46 | 481.51 | -1,664.88 | 0.00 | 0.00 | 0.00 | 20.00 | 132.00 | 1,475.00 |
| Income (\$) | 104,055,302 | 150.31 | 689.53 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 3,386.73 |
| Net Interest (\$) | 104,055,302 | -0.67 | 8.49 | -18.56 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.49 |
| Interest Expense (\$) | 104,055,302 | 0.76 | 8.89 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 18.62 |
| Interest Earned (\$) | 104,055,302 | 0.07 | 1.18 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.51 |
| Net CC Charge (\$) | 104,055,302 | 41.33 | 144.42 | 0.00 | 0.00 | 0.00 | 0.00 | 20.63 | 112.77 | 625.74 |

Figure A.I. Heterogeneity in Consumption Response Across Liquidity Terciles

The figure shows the consumption response to making tax payment dates (panel A) and to the arrival of tax refunds (panel B) across income terciles. The consumption response is measured for those household-years in the bottom third (blue) and top third (red) of income. The dotted lines represent two standard errors confidence intervals. Standard errors are doubled-clustered at the household-year and at the calendar date levels.

