

MIT Open Access Articles

Asymmetric Consumption Smoothing

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

As Published: 10.1257/aer.20181735

Publisher: American Economic Association

Persistent URL: <https://hdl.handle.net/1721.1/134039>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Asymmetric Consumption Smoothing[†]

By BRIAN BAUGH, ITZHAK BEN-DAVID, HOONSUK PARK,
AND JONATHAN A. PARKER*

Analyzing account-level data from an account aggregator, we find that households increase consumption when they receive expected tax refunds, as if they face liquidity constraints. However, these same households smooth consumption when making payments in other years, primarily by transferring funds among liquid accounts. Even households carrying credit card debt smooth consumption when making payments, and even highly liquid households spend out of refunds. This behavior is inconsistent with pure liquidity constraints or hand-to-mouth behavior and is most consistent with a mental accounting life-cycle model. (JEL D12, E21, G51, H24, H31)

A central feature of consumer behavior is that many consumers delay spending out of expected income until it arrives. The leading theoretical explanation, exemplified by the buffer stock model, is that some people are financially constrained and so delay consumption until income arrives. Another leading model assumes that some households are hand-to-mouth consumers, spending income as it arrives. These theories of the consumer are central components of modern New Keynesian macroeconomics.

Our study presents a novel and robust consumption pattern that fits neither prototype model. We use high-frequency account-level data to document that when faced with anticipated income, households, across the liquidity spectrum, increase consumption on the date of cash flow arrival. In contrast, when faced with anticipated payments, the same households keep their consumption level intact and use liquid reserves or credit to fund payments. Among the models that we review, this asymmetric consumption pattern is best explained by a model of mental accounting

*Baugh: College of Business, University of Nebraska-Lincoln (email: bbaugh2@unl.edu); Ben-David: Fisher College of Business, The Ohio State University, and NBER (email: ben-david.1@osu.edu); Park: Faculty of Business and Economics, University of Melbourne (email: hoonsuk.park@unimelb.edu.au); Parker: Sloan School of Management, Massachusetts Institute of Technology, and NBER (email: JAParker@mit.edu). Emi Nakamura was the coeditor for this article. 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 dataset. For helpful comments, we thank Manuel Adelino, Sumit Agarwal, Michael Boutros, Andrew Chen, Peter Ganong, Uri Gneezy, Damon Jones, Olivia Kim, Scott Nelson, Pascal Noel, Michaela Pagel, Michael Palumbo, Joel Slemrod, René Stulz, and particularly our discussant, Annette Vissing-Jørgensen, 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 Western Finance Association 2018, and the JPMorgan Chase Institute 2018 Academic Conference. 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.

[†]Go to <https://doi.org/10.1257/aer.20181735> to visit the article page for additional materials and author disclosure statements.

that would lead to different economic dynamics and policy prescriptions from leading macroeconomic models.

Specifically, we study the consumption responses of households as they file their US federal income tax returns, receive refunds, and make payments in different years.¹ Households in our select sample, and so to some extent in the population, exhibit an asymmetry in consumption behavior in response to negative and positive cash flows. When faced with a tax payment, households do not cut consumption but instead make account transfers in advance of payments. Even low-liquidity households in our sample maintain stable consumption. Yet, these same households, in other years, increase consumption in response to the arrival of expected tax refunds. Even the most liquid household-years do so, albeit at lower rates. The latter observation has been interpreted as evidence of liquidity constraints.² Thus, households in our sample shield consumption from decreases in income when making transitory payments but at the same time follow a heuristic to consume out of transitory increases in income.

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 can identify the date of tax filing and the date and amount of tax refund or payment in the current and previous year. We focus on a subsample of households that receive tax refunds in some years and make payments in other years so that our results are not driven by differences in the types of households that receive refunds or make payments.

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 US population, both because households have to select into the account aggregation service and because they have to file taxes in a way that creates an observable transaction in the data. In particular, our sample omits households that postpone payment by filing an extension and those that take up refund anticipation loans instead of receiving refunds (issues we discuss in detail in Sections I and VI). Thus, our sample is likely to be more liquid than the typical American household, and indeed our measures of liquidity confirm this conjecture. Further, our select sample may also be more financially sophisticated. However, this sample has some advantages for our analysis. From an *a priori* standpoint, our finding of excess sensitivity of consumption to predictable refunds is less likely to be due to liquidity constraints than it would be in the population.

¹In 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).

²See, e.g., Bodkin (1959); Zeldes (1989); Parker (1999); Hsieh (2003); Johnson, Parker, and Souleles (2006); and the survey articles by Hassan and Fuchs-Schündeln (2016) and Gomes, Haliassos, and Ramadorai (forthcoming).

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

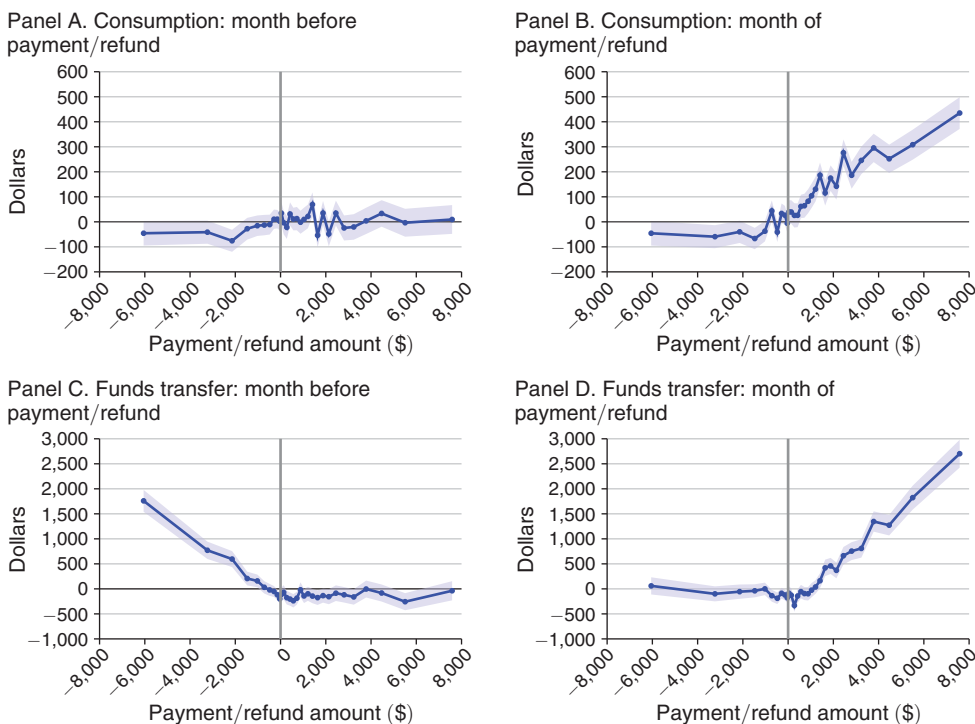


FIGURE 1. CONSUMPTION AND FUND TRANSFERS AROUND TAX PAYMENT OR REFUND

Notes: 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 percent of the data for those who received refunds and at every 10 percent of the data for those who made payments. The shaded region represents two standard errors confidence intervals.

We measure the consumption response to expected income changes by regressing consumption spending on 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. The leads and lags of filing and 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, they 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 plots average consumption spending and fund transfers after de-meaning by both household-year and calendar day for different payment and refund amounts and in the month before and in the month after the refund or payment date.

First, we observe a striking asymmetry between the consumption response to refunds and the response to payments (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 receiving the refund. In contrast, in years when they make tax payments, these same households do not reduce consumption spending following making the payment. This result is novel, in part because few papers have studied the consumption response to *decreases* in income.⁴

According to our regression estimates, people spend over 7 percent of their refunds on consumption during the month following refund receipt, an amount that increases slowly to 15 percent over several months. These responses underestimate the true increase in consumption (by roughly a factor of 2) because we cannot definitively classify some account transactions. This spending increase starts the *day of* refund arrival, and 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, and unrelated to the precise day of payment.

Second, 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 our regression analysis as well as 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 (panels C and D). In years when households receive refunds, they make transfers only after receipt. In contrast, households smooth consumption through tax payments by increasing transfers among accounts in the month before making the payment by about one-third of their anticipated payment. These transfers are mostly among observed checking and savings accounts rather than from outside (likely nontransaction) accounts. We also find that households increase transfers in response to the news learned about upcoming payments or refunds prior to and at filing. Finally, households in our sample have substantial liquidity on average and they manage liquidity over longer horizons: in years when they make payments, households have accumulated somewhat higher balances in their core accounts.

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. While payment and refund status are also endogenous, determined in part by past income not subject to withholding, Gelman et al. (2019) provides evidence that this endogeneity does not cause consumption asymmetry: the paper finds no sharp difference between

⁴Notably, Gelman et al. (2020, p. 2) studies transitory unexpected declines in income and concludes 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 et al. (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. Shea (1995b), Stephens (2001), Ganong and Noel (2019), and Jorring (2018) find substantial declines in spending in the event of wage cuts in union contracts, permanent worker displacement, when unemployment benefits expire, and when mortgage payments rise, respectively. Conversely, Souleles (2000) and Aguiar and Hurst (2005) find that consumption is well smoothed when college expenses start and when people retire.

the propensity to spend in response to refunds versus payments in a quantitative life-cycle model, as we discuss in Sections IV and VI.

We show that this consumption asymmetry occurs throughout the distribution of liquidity, which is further evidence that spending response to refunds for these households is an unconstrained choice. Households in the bottom tercile of the ex ante distribution of liquidity have large propensities to consume out of refunds.⁵ But in this same sample of households with low liquidity, households do not cut consumption prior to or when making payments, inconsistent with financial constraints driving the response to refunds.

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 is consistent with a growing body of evidence that even highly liquid households have substantial spending responses to predictable increases in liquidity (Parker 1999, Kueng 2018, McDowall 2019).⁶ These findings imply that a propensity to consume out of predictable increases in income that decreases with liquidity is not necessarily indicative of liquidity constraints. For many households in the population, liquidity constraints very likely determine behavior. However, in our nonrepresentative sample, households can smooth consumption but choose not to. Hence, the well-documented excess sensitivity of consumption to increases in expected income in the population is not purely driven by constraints.

We discuss these findings in light of the theories that are commonly used to explain household behavior. Among the theories, the models most consistent with our results are those of mental accounting (Thaler 1999). Our empirical evidence matches the predictions of the behavioral life-cycle model of Shefrin and Thaler (1988), in which households maintain three mental accounts: current income, current assets, and future income. Specifically, an anticipated tax refund is considered future income (which a rule-of-thumb prevents households from consuming). When the refund is received, it has a similar status as a “bonus from work,” part of which is consumed upon arrival and part saved. Payments, on the other hand, are viewed as unrelated to the income-consumption process and therefore are funded from less liquid savings accounts and do not impact the consumption process. In general, because models of mental accounting are based on observations of behavior, their predictions are generally limited, which has led to limited empirical testing. Shefrin and Thaler (1988) also proposes that households divert funds out of their sight to savings accounts (as opposed to transaction accounts) as a mental commitment device. Under the assumption that savings accounts can provide commitment, the hyperbolic discounting model of Angeletos et al. (2001) also fits most, but not all, of the behavior we find in our sample and also provides more testable predictions.

⁵This finding is consistent with the literature comparing consumption responses of households with different levels of liquidity starting with Zeldes (1989), and followed by a large literature including Jappelli, Pischke, and Souleles (1998); Agarwal, Liu, and Souleles (2007); Aaronson, Agarwal, and French (2012); and Kaplan, Violante, and Weidner (2014).

⁶And with papers that find that the propensity to consume is persistent rather than purely driven by liquidity (see Parker 2017; Aguiar, Bils, and Boar 2020).

In sum, households in our sample consume more when their refunds arrive as if many are liquidity constrained. But they appear, in fact, to be choosing not to draw on existing liquid savings, as if they are what Olafsson and Pagel (2018) refers to as the “liquid hand-to-mouth.” But inconsistent with a simple hand-to-mouth heuristic, these same households do not reduce consumption when making payments. These consumers instead follow a different heuristic, accessing liquid savings for making payments but not for smoothing consumption prior to refund arrival. Rather, they consume or even impulsively splurge when refunds arrive (as in Agarwal et al. 2019, Ben-David and Bos forthcoming).

I. US Individual Income Tax Returns

This section describes the US tax system and how it generates cash flows with two key features. First, the income tax system causes both expected inflows and outflows of funds, allowing us to measure the effect of the sign of the cash flow rather than the effect of the source of the change in liquidity. Second, we can identify news about future cash flows and the timing and amount of cash flows using the structure of the US individual income tax system. We can, therefore, separate the effect of the arrival of news about future after-tax income from the arrival of the change in income.

The US individual income tax covers all sources of household income in each calendar year. For most labor income, employers withhold income taxes from workers’ 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.⁷ 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 fill out tax return forms, some variant of the IRS 1040 tax form and additional schedules, to calculate their total taxes owed and submit (file) their returns by the April 15 deadline.⁸

If the withholding (and estimated taxes paid) exceed taxes due, the household receives a refund of the difference, typically a few weeks after filing.⁹ If the taxes due exceed the amount previously paid, then the household must remit payment by the April deadline. However, any household can also file for an extension that pushes the deadline back to October 15 but also leads to possible interest and penalty charges. If the household has paid at least 90 percent of its total taxes by the April 15 deadline, then the unpaid balance accrues interest at roughly 3.5 percent during the years in our sample. If the household has not paid at least 90 percent of taxes by April 15, then it must pay an additional penalty of 0.5 percent per month, for an annual rate close to 10 percent. Ultimately, long-term nonpayment leads to legal penalties that can include wage garnishment and/or incarceration.

⁷No 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.

⁸People with low incomes or no taxes due do not have to file. Married individuals can file taxes jointly. The actual deadline is delayed when April 15 falls on a Sunday.

⁹IRS Publication 2043 indicates that 90 percent of refunds were processed within 21 days of filing. For example, for 2013, see <https://www.irs.gov/pub/irs-prior/i1040--2013.pdf>. See also Slemrod et al. (1997).

To measure when households file returns, we take advantage of the fact that many households use online tax-preparation companies such as TurboTax to help them fill out and file their tax forms. These tax preparation companies also offer customers the ability, for a fee, to receive their refunds immediately from the company instead of with delay from the federal government.¹⁰ While these fees are not large, on the order of \$25–\$35, because the refunds are expected within weeks, the implicit annualized interest rates on these short-term loans are large.

As we describe in the next section, because of the way we measure tax filing, payments, and refunds, our sample omits both people who postpone payment until after April 15 and those who choose to receive their refunds immediately from their tax preparer. The implication of this sample selection is that our sample underrepresents people who are significantly liquidity constrained, which in turn bolsters our conclusion that the consumption responses to refunds that we find are not solely driven by liquidity constraints.

In our data and in the population, 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 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.¹¹

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 after filing (or occasionally at filing for some payments). The timing of the arrival of information is based on a (constrained) 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.

¹⁰In this case, the actual refund is paid to the company instead of to the taxpayer.

¹¹Although 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 percent of the EITC credit early.

II. Data and Variable Construction

A. 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 until the household actively discontinues the service by requesting account deletion. Therefore, 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 *bank* 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.

B. 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, with further details provided in the online Appendix.

First, we identify federal tax refunds and payments by querying the transaction descriptions such as "us treasury des tax" and "irs treas tax," among other terms.¹² 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 \$10,000.¹³

Second, we identify tax-preparation transactions by querying for payments to TurboTax, H&R Block, TaxAct, or TaxSlayer. The filing date assigned to each household is this transaction date. We exclude household-years that have tax-preparation transactions on multiple days (as would be the case for a family filing separately on different days). This process also implicitly excludes household-years with multiple returns or in which tax preparation charges are deducted directly from a refund (since we would not observe a filing payment).

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

¹²When we predict a refund, we also use refunds paid directly to households by tax preparers as described in the online Appendix.

¹³The very few households that made payments greater than \$10,000 are quite different from most of our sample along many dimensions. We confirmed that our results are very similar if we instead truncate at \$20,000.

(e.g., gym memberships), and insurance.¹⁴ The outflows that we can categorize are primarily those that the household makes via debit cards and online bill payments as well as those on linked credit cards. Similarly, we construct a measure of savings and debt payment and a measure of labor and pension income. We construct a proxy for account balances based on interest earned, and a direct measure from account balances reported in “interest earned” transactions.

We also cannot classify spending on unlinked cards. To overcome this issue, we scale 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).¹⁵ We drop all household-years with scaling factors above 2 since, for these households, we observe less than one-half of their credit card spending. Across all household-years, the scaling factor has a median of 1.086 and an average of 1.175. Consequently, our estimated consumption responses are only about 17 percent larger due to the scaling-up of credit card spending.

Despite this adjustment, our measure of consumption understates the true consumption amount, and so the true consumption responses, for two reasons. First, we omit consumption spending done with counterparties that we cannot clearly assign to one of the categories above. Second, we are unable to categorize outflows made by check. To evaluate the quantitative importance of these omissions, we construct a measure of miscellaneous payments that consists of payments that we cannot definitively categorize into either consumption or savings. This variable is equal to the sum of checks and otherwise uncategorized outflows.

C. 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. Therefore, we focus on the subsample of households that 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 October 1 to September 30 of each year.¹⁶ We drop household-years for which any necessary variable is not present, as well as those for which there are not at least 25 transactions of at least \$1 each month (to ensure we have active users rather than dormant account holders). We focus on people filing (and paying) roughly on time, so restrict our attention to household-years with a filing date and refund or payment date both before June 1. We require our inferred filing date to weakly precede payment or refund receipt, as described in the online Appendix. Finally, we require that we also observe a refund or payment in the year prior for reasons explained in Section IIIA.

¹⁴To ensure that consumption is not mechanically related to taxes, we exclude from it any filing fees or tax payments.

¹⁵This process does not bias our results if households randomized which card was used for any given transaction or, more reasonably, if the choice of card is independent of the timing and amount of the tax variables in our estimating equation (3). We confirm that our results are similar without this adjustment, as reported in online Appendix Table A.II.

¹⁶The exception is 2015, when our sample ends, which consists of 237 days.




FIGURE 2. DISTRIBUTION OF PAYMENTS AND REFUNDS

Notes: Panel A shows the histogram for the broad sample. 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.


After applying these filters but before requiring that each household must receive at least one refund and make one tax payment in our sample, our dataset contains 307,702 household-years from 196,565 unique households (102 million household-day observations). The summary statistics for this broad-sample population are presented in online Appendix Table A.I. While our data contain only households who have selected to use the aggregator, the sample appears to be broadly representative of the population with some exceptions (discussed below and in the online Appendix).

From this population, we focus our analysis on the subset of households that are observed to have at least one refund and one payment (in different years). This restriction not only means that we measure the response to refunds and payment for the same households, but also makes the distribution of refund and payment amounts more symmetric. The distribution of tax payments and refunds in our broad sample is shown in panel A of Figure 2; panel B shows the distribution for our final sample of households that make a payment in at least one year and receive a refund in at least one year.

Figure 3 shows the distribution of direct deposit income in our initial broad sample (panel A) and final sample (panel B) relative to US Census data.¹⁷ Income in our sample is net of withheld taxes and benefits such as 401(k) contributions, health care premiums, etc., while the Census reports income before deductions. The average annual income in our final sample is \$68,543, and the distribution of households in our data are clearly higher income than in the population as measured by the Census.

¹⁷ 2013 Current Population Survey from the US Census (HINC01).

Panel A. Broad sample versus census



Panel B. Final sample versus census




FIGURE 3. DISTRIBUTION OF INCOME

Note: 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.

D. Sample Statistics

Table 1 presents summary statistics for our final sample. The average income is \$187.79 per day (corresponding to \$68,543 annually). Household consumption averages \$100.24 per day (corresponding to \$36,588 annually) in directly observed spending and \$116.97 per day (corresponding to \$42,694 annually) when scaled up to account for spending on unlinked cards. The difference between income and consumption highlights how conservative our method of categorizing transactions as consumption is. Of the households in our sample, 81 percent have credit cards linked to the aggregator, 83 percent have credit cards that are not linked to the aggregator, and 97 percent have at least one credit card that is either linked or unlinked.

TABLE 1—SUMMARY STATISTICS

Variable	Count	Mean	SD	p1	p25	p50	p75	p99
Households	37,180							
Household days	20,829,451							
Net flow (\$)	20,829,451	−55.03	4,089.89	−3,935.51	−222.42	−53.00	0.00	4,493.31
Consumption (\$)	20,829,451	100.24	177.98	0.00	0.00	36.77	121.03	915.21
Scaled consumption (\$)	20,829,451	116.97	210.58	0.00	0.00	41.87	139.84	1,084.95
Savings and loans (\$)	20,829,451	39.14	521.57	−1,000.00	−38.28	0.00	0.00	2,573.67
Misc payments (\$)	20,829,451	100.01	399.97	0.00	0.00	0.00	40.00	1,922.00
Income (\$)	20,829,451	187.79	828.34	0.00	0.00	0.00	0.00	4,008.48
Net interest (\$)	20,829,451	−0.69	9.15	−18.96	0.00	0.00	0.00	1.40
Interest expense (\$)	20,829,451	0.81	9.61	0.00	0.00	0.00	0.00	19.02
Interest earned (\$)	20,829,451	0.11	1.55	0.00	0.00	0.00	0.00	1.43
Net CC charge (\$)	20,829,451	58.23	174.14	0.00	0.00	0.00	47.86	802.50
Household-years	62,404							
Filing date	62,404	Mar 17	27.53	Jan 21	Feb 20	Mar 24	Apr 12	Apr 19
Refund/payment date	62,404	Mar 26	25.97	Feb 01	Mar 04	Apr 06	Apr 16	May 08
I(positive refund)	62,404	0.56	0.50	0.00	0.00	1.00	1.00	1.00
Refund − Payment (\$)	62,404	509.00	2,651.44	−7,263.00	−662.00	240.50	1,798.00	7,869.00
Predicted refund (\$)	62,404	1,364.30	1,606.58	−2,048.24	195.90	1,166.06	2,164.74	6,248.64
News amount (\$)	62,404	−855.29	2,748.89	−8,997.48	−2,124.96	−718.50	512.20	6,433.61
Filing to ref/pay (days)	62,404	9.72	12.98	−2.00	2.00	6.00	11.00	69.00
I(linked credit cards)	62,404	0.81	0.39	0.00	1.00	1.00	1.00	1.00
I(unlinked credit cards)	62,404	0.83	0.38	0.00	1.00	1.00	1.00	1.00
I(any credit cards)	62,404	0.97	0.18	0.00	1.00	1.00	1.00	1.00

Note: 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. (See online Appendix for variable definitions.)

While in our broad sample, 90 percent of household-years receive a refund, in our final sample, 56 percent of household-years receive a refund. The average value of *Refund − Payment* is \$2,339 for the broad sample and \$509 in our final sample. In untabulated results, we find that, conditional on receiving a refund, the average refund is \$2,805 in our broad sample and \$2,121 in our final sample, which are more than one-third of average monthly income. Conditional on making a payment, the average payment is \$1,687 in our broad sample and \$1,574 in our final sample.

Table 1 also shows that the average household files on March 17 and that there is significant variation: the standard deviation of filing date is 27.5 days. In untabulated results, we find that the average household makes a payment on April 4 or receives a refund on March 19, and again there is large variation across households (standard deviations of 28 and 26 days for filing and refund dates, respectively). The average distance between filing and refund or payment is 9.7 days, with a still substantial standard deviation of 13 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. Panel A shows that the distribution of filing dates 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. 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.

TABLE 2—PAYMENT AND REFUNDS STATISTICS BY MONTH OF FILING

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	19.6	37.4	15.5
February	26	15	34	14.2	23.3	11.1
March	25	24	25	10.6	12.2	9.3
April	43	58	32	5.2	2.5	8.8

Notes: Panel A shows the distribution of filing month. Panel B shows the distance between filing and payment/refund. In both panels, the second and third columns show the distributions only for households making a payment or receiving a refund, respectively, in that year. The sample is our final sample of households that make a payment in at least one year and receive a refund in another year.

III. Estimation Method

A. 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 0 if a payment is made that year, and analogously for $Payment_{h,y}$. We run the following regression:

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

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 percent.¹⁸ Our measure of information about tax information uncovered during filing, or “news,” is the residual in this regression, which we denote by

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

The distribution of $NewsAmount$ is shown in Table 1 for our final sample, and in online Appendix Table A.I for the broader sample. Whereas the average $NewsAmount$ in the broader sample is 0, the average $NewsAmount$ in the final sample is $-\$855$. This difference reflects our selection of households with both payments and

¹⁸ Adding the previous year's income and its interaction with the previous year's indicator variable leads to only a trivial increase in fit. Adding two years' prior income as well leads to a slightly greater increase in fit (about 1 percent) but a large decline in sample size.

refunds, who then have worse outcomes on average because payments are more prevalent than in the larger sample used for estimating *NewsAmount*.

This empirical model is identified from cross-sectional variation and has a short time 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 2010.¹⁹ 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.

B. 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.²⁰ 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 for 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 the amount of news, but with different slopes for positive and negative news about the tax amount (that is, affine with different coefficients on good and bad news).

To be precise, let $Refund_{h,y,t}$ be the amount of refund received on day t in year y and be 0 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 household h files its tax return. Let $NewsAmount_{h,y,t} = File_{h,y,t} \times NewsAmount_{h,y}$, so that $NewsAmount_{h,y,t}$ 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 $PosNews_{h,y,t} = \max[NewsAmount_{h,y,t}, 0]$ and $NegNews_{h,y,t} = \max[-NewsAmount_{h,y,t}, 0]$.

¹⁹IRS 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>.

²⁰The 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.

Our main estimating equation is

$$\begin{aligned}
 (3) \quad Y_{h,y,t} = & \sum_{k=-28}^K \beta_k^+ \text{PosNews}_{h,y,t+k} + \sum_{k=-28}^K \beta_k^- \text{NegNews}_{h,y,t+k} \\
 & + \sum_{k=-28}^K \phi_k \text{File}_{h,y,t+k} + \sum_{k=-28}^K \gamma_k^+ \text{Refund}_{h,y,t+k} \\
 & + \sum_{k=-28}^K \gamma_k^- \text{Payment}_{h,y,t+k} + \alpha_{h,y} + \tau_t + u_{h,y,t},
 \end{aligned}$$

where Y is an account inflow or outflow measure, $\alpha_{h,y}$ is a household-year-specific intercept, and τ_t is a day-specific intercept; K is set to the maximum identifiable lag. The β_k , γ_k , and ϕ_k coefficients, respectively, 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.²¹

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 heteroskedasticity, within-day correlation, and within-household-year correlation in $u_{h,t}$.

In the sections that follow, 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

²¹ We include the leads and lag *File* (in addition to its interactions with the amount of news) for two reasons. First, although presumably a quantitatively small effect, as uncertainty is resolved, precautionary saving may decline. Second, the household may learn information unrelated to the amount of news about refund or payment during return preparation, as we discuss in Section VII. In practice, omitting this term makes very little difference to any of the other coefficients.

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). We show the standard errors by shading the region in the figures surrounding the lines that represent coefficient estimates. The shaded regions in all figures represent two standard error confidence intervals.

IV. 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 V 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 VI shows that households have substantial liquidity and use liquidity management to smooth consumption through payments but not refunds. Section VII 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.

A. 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 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 (and the responses for unscaled consumption are reported in online Appendix Table A.II).

Panel B of Figure 4 shows that people consume about 8 percent of their refunds over the month following receipt, a number which rises to 15 percent 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).

B. 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 percent 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 panel B of Figure 1, which

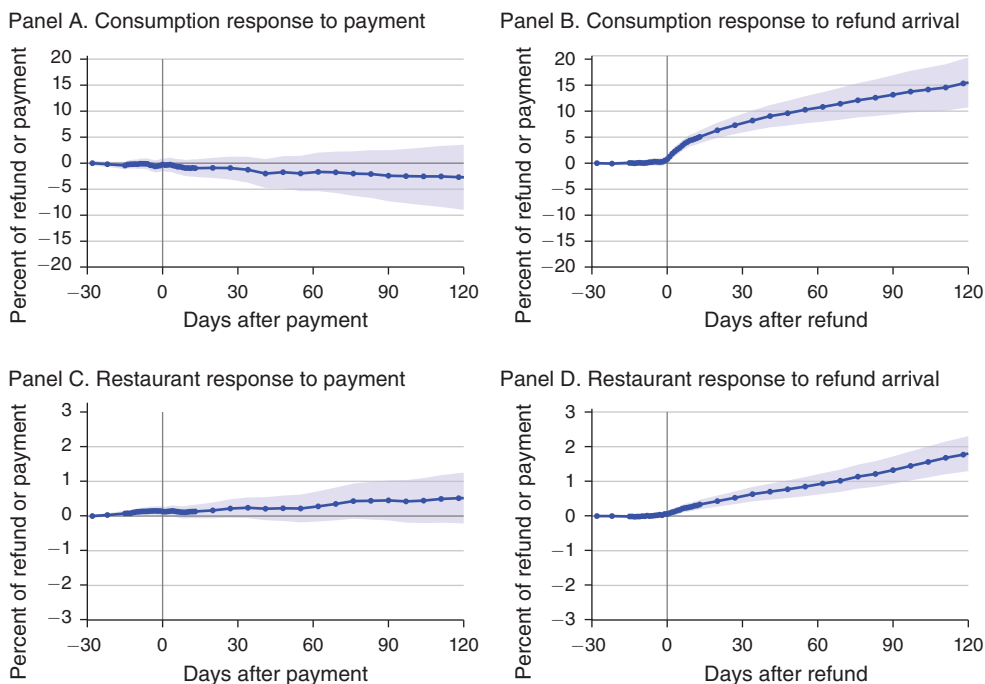


FIGURE 4. CONSUMPTION RESPONSE TO PAYMENT OF TAXES AND ARRIVAL OF REFUNDS

Notes: 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 change in consumption as a percentage of the payment amount (> 0) or refund amount.

TABLE 3—CUMULATIVE CONSUMPTION RESPONSE TO PAYMENTS AND REFUNDS

	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	0.06 (0.60)	-0.15 (1.10)	-0.71 (1.69)	-0.39 (2.30)	0.53 (0.41)	7.32 (0.87)	10.39 (1.33)	12.71 (1.80)
Restaurant	0.14 (0.07)	0.21 (0.13)	0.22 (0.19)	0.47 (0.26)	0.07 (0.05)	0.53 (0.10)	0.85 (0.14)	1.20 (0.19)
Paid using CC	-0.86 (0.60)	0.12 (1.11)	0.23 (1.65)	1.19 (2.24)	0.54 (0.38)	2.41 (0.73)	3.73 (1.08)	5.18 (1.45)
Misc payments	-1.37 (1.55)	6.20 (3.04)	3.66 (4.43)	2.88 (5.88)	0.07 (0.85)	9.89 (1.67)	12.52 (2.45)	14.50 (3.23)

Notes: 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.

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 are 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 percent after a month and 28 percent after four months, which are more than twice as large as our baseline, conservative estimates.²²

C. Strictly 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.²³ These panels 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 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 unlikely to be 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). However, one possibility is that the asymmetric response is due to households that slip into our sample by making a small payment or getting a small refund in just one year. Online Appendix Figures A.I and A.II show that this is not the case: the asymmetry remains when we drop household-years with amounts less than \$2,000 and, more severely, when we require households to have a refund and a payment each greater than \$2,000. Pursuing this point further, in theory we could identify the asymmetry purely from differences within the same household by interacting the day-of-the-year effects with the individual effects, but doing so makes inference very imprecise. Instead, we confirm that we find the same results when we interact the individual effects with indicators for day of the week, month of the year, the first three days of the month, the last three days of the month, and the 14th to the 16th of the month (a total of 30 time fixed effects; see online Appendix Figure A.III).²⁴

However, it is important to note that households can differ across years, and in particular, they do have different amounts of liquidity. We address this issue in detail in Sections V and VI. Before doing so, we first address two final possible concerns with our results so far.

²²Around 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.

²³We use the classification of restaurant transactions provided by the debit card and credit card providers, so this includes everything from fast food chains to gourmet dining experiences.

²⁴It 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 a stronger 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.

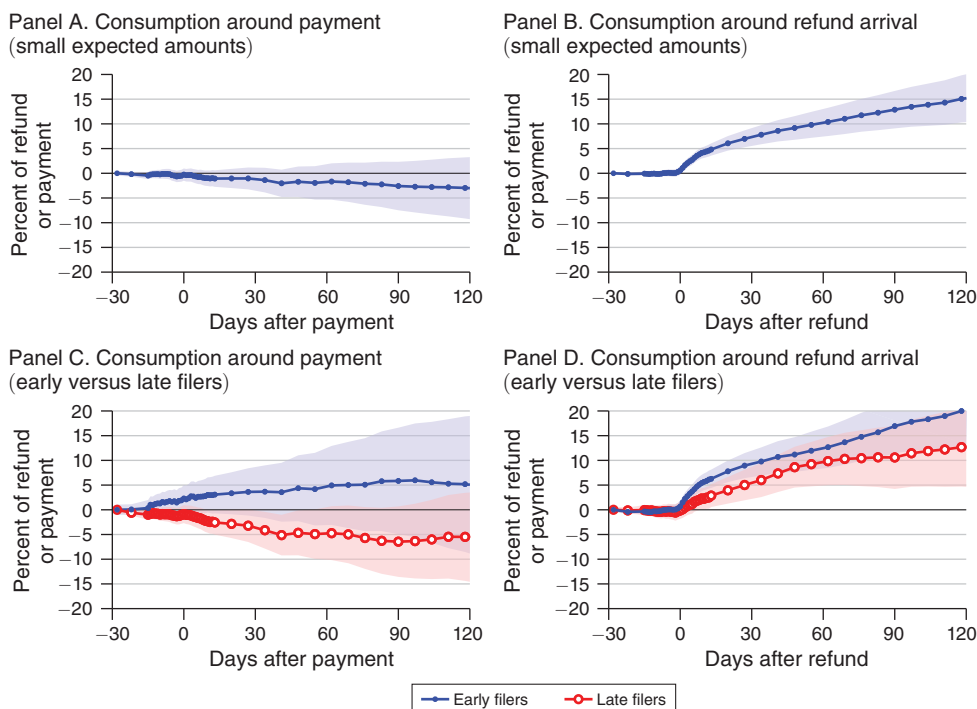


FIGURE 5. CONSUMPTION FOR SMALL EXPECTED AMOUNTS AND EARLY/LATE FILERS

Notes: 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.

D. 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 the deadline in mid-April.²⁵ A bias from this source seems unlikely because we find the asymmetric consumption response for many different, more homogeneous subgroups of our sample.²⁶

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 $\text{Refund} - \text{Payment}$. In this subsample, $\text{Refund} - \text{Payment}$ averages $-\$62.40$ (a payment, relative to a refund of $\$564$ in the whole sample) with a standard deviation of $\$2,966$.

²⁵ Of households who file in February and owe taxes, 33 percent pay in April, and the average time between filing and payment is 23 days. Of those who file in March, 47 percent pay in April, and the mean time between filing and payment is 12 days.

²⁶ As we discuss in both the introduction and Section VI, ours is a select sample that excludes households that choose to postpone payments past the April 15 deadline.

Focusing more closely on the issue of timing, the consumption asymmetry is observed for households 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. For refunds, households choosing to file in April are likely those who do not need liquidity and so are not constrained. Yet these households spend substantial amounts of their refunds. For payments, households filing in April have the least time to save and prepare to smooth consumption. Yet, these households also smooth consumption over payments. This lack of difference in spending responses by month constitutes 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.²⁷ The lack of differential spending responses to refund arrival also suggests that liquidity constraints are not the main reason for the high propensity to consume out of refunds on arrival. Optimizing households that are short on liquidity should file earlier to better smooth consumption. Thus, if liquidity constraints were driving spending from refunds, we would expect much higher spending responses among households filing earlier (and revealing that they are liquidity constrained). In fact, panel D shows only small differences.

E. Mismeasurement and Robustness

Could mismeasurement generate our finding of an asymmetric consumption response? Both the timing and amount of news about refund 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 two or more weeks after filing. Our statistical procedure thus might not cleanly separate the consumption response to filing and news from the response to payment.²⁸

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. Therefore, 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 observe. Further evidence is provided by the differential effect on the exact day of payment and refund.²⁹ Thus, this type of mismeasurement cannot account for our asymmetry.

²⁷ Another related theory is that households that have time-consistency problems are sophisticated about these problems, i.e., understand their bias and act to correct it. In this case, households with time-consistency problems value the commitment of filing later (rather than simply always intending to file tomorrow and failing to do so until the deadline). The prediction for naïfs or sophisticates is the same: people who file later are those most likely to spend when a refund arrives.

²⁸ 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 refunds.

²⁹ To 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 percent response to refund arrival and a –100 percent 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.

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. Finally, this result does not just apply to the somewhat atypical sample of households that regularly both make payments and receive refunds. As we show in online Appendix Figure A.IV, we find the same asymmetric consumption response in our broad sample of all households regardless of whether they ever make payments or receive refunds.

We conclude that our main finding is unlikely to be 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.

V. 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 consumption. 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.

Two different properties of liquidity are useful for interpreting these results. First, in the US population, low liquidity is persistent in both survey and account-level data. Thus, low or high liquidity is largely a characteristic of a household not a transitory state. Second, however, liquid wealth does vary some over time, particularly for wealthier households. Section VI documents that in our data, liquidity in checking and savings accounts is endogenous and related to impending payments: people adjust their liquidity around tax time and have slightly more liquidity in years with payments.

A. Households with Low Liquidity

Our primary measure of liquidity is based on the small subsample of accounts for which we can observe the account balance from a text line in the account (e.g., “\$0.16 interest earned for average daily balance of \$3,810.72”). Because we “observe” balances in this way for all core accounts in January for just under 5 percent of our sample, we present results for both our final sample of households that receive refunds in some years and make payments in other years and our broad sample that does not impose this restriction (as described in online Appendix Table A.I).

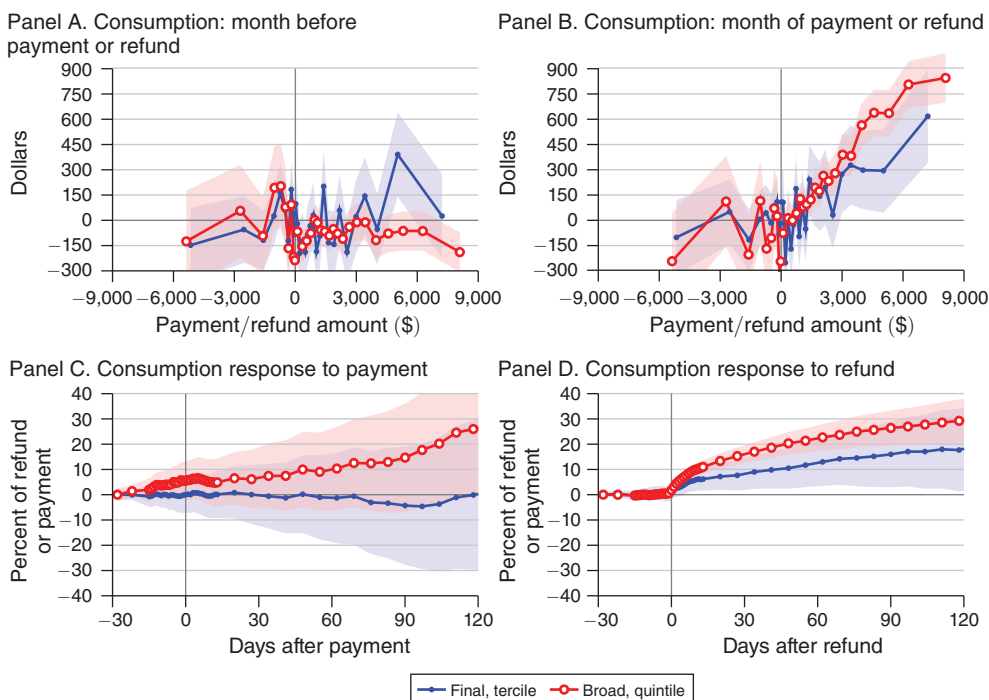


FIGURE 6. CONSUMPTION RESPONSE FOR HOUSEHOLDS WITH LOW ACCOUNT BALANCES

Notes: Panels A and B show the abnormal consumption around tax payment and refund dates for household-years with low liquidity. For households in our final sample, we take the bottom tercile of observed account balances. For households in our broader sample, we take the bottom quintile of observed account balances. The markers denote averages at every 5 percent of the data for those who received refunds, and at every 10 percent of the data for those who made payments. Panels C and D show the cumulative response of external savings and debt payments to making tax payments and receiving refunds, respectively.

Figure 6 presents our main results for households with low ex ante liquidity according to their January balance for the bottom tercile of our final sample and the bottom quintile of our broad sample.³⁰ Panels A and B show the excess consumption spending around payments and refunds (as in panels A and B in Figure 1 but for larger vertical axes). Households with low liquidity smooth consumption when making payments, yet these same households nevertheless increase consumption when receiving refunds in other years when they also have low liquidity.

Panels C and D show the same results from our estimation of equation (3) (as in panels C and D of Figure 4 but for larger vertical axes). Controlling for seasonal patterns and the timing of filing and information about taxes confirms this result: households with low liquidity spend an even greater fraction of their refunds when they arrive than the typical household, yet they still smooth consumption through payments. The propensity to spend from rebates is more than 1.5 times the average. This pattern is strong evidence that low liquidity is not driving

³⁰The broad sample does not require that we observe the household both receiving a refund and making a payment. There are about 3,935 household-years in each tercile of our final sample and about 8,323 household-years in each quintile of our broad sample.

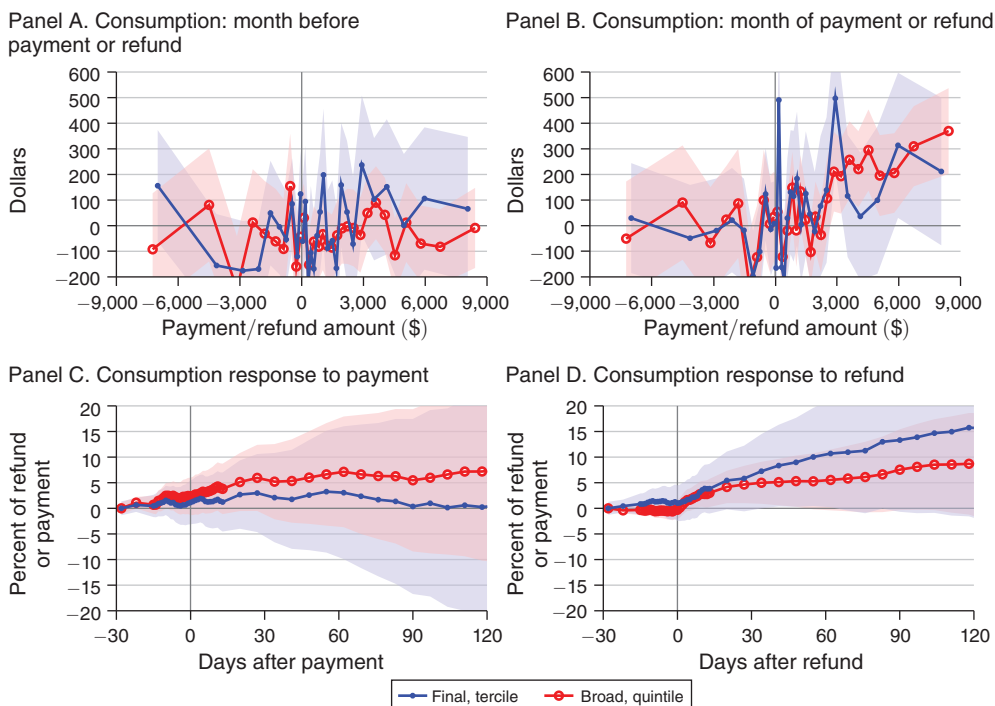


FIGURE 7. CONSUMPTION RESPONSE FOR HOUSEHOLDS WITH HIGH ACCOUNT BALANCES

Notes: Panels A and B show the abnormal consumption around tax payment and refund dates for household-years with high liquidity. For households in our final sample, we take the top tercile of observed account balances. For households in our broader sample, we take the top quintile of observed account balances. The markers denote averages at every 5 percent of the data for those who received refunds, and at every 10 percent of the data for those who made payments. Panels C and D show the cumulative response of external savings and debt payments to making tax payments and receiving refunds, respectively.

the spending response to refunds since it is not hindering the smoothing of consumption through tax payments. These results are quite robust across different measures of liquidity, as we note just below, including looking at households that pay interest on credit card balances.

We emphasize again that these results are from a sample of households that neither take up refund anticipation loans nor postpone payment until the later tax deadline, and so are less likely to be liquidity constrained than the general population. In the population and for typical payments, illiquid households may reduce consumption around making substantial payments.

B. Households with High Liquidity

Figure 7 displays the spending responses for households with high ex ante liquidity as measured by January account balances. Panels A and B show little evidence that households have lower consumption in the month before making payments and some evidence that households with high liquidity have higher consumption in the month following the receipt of a refund. The evidence is statistically weak in our final sample, due to losing 95 percent of our sample, but quite

strong for the broad sample. Panels C and D show the same results for our regression analysis. For our final sample (consisting only of households that we observe both making a payment and receiving a refund), the point estimate of the spending response to refunds is only just slightly smaller than that of the entire sample but is only on the edge of statistical significance. For the broad sample, the evidence for a substantial spending response is much more statistically significant (but requires stronger identifying assumptions to be interpreted as causal).

Optimizing households that both have liquidity and try to stabilize their standard of living should not cut consumption prior to or after making tax payments or receiving tax refunds. We confirm that prediction for tax payments.

These results are quite robust across different measures of liquidity. For example, households filing in February, who might be less liquid if they are impatient to get refunds, have stronger spending responses (as shown earlier in panels C and D of Figure 5). Further, we verify these results for a larger subsample of our data by measuring liquidity based on net interest earned during November, December, and January. Net interest earned is interest earned on all checking and savings accounts less interest paid on linked credit card accounts. While the sample is larger, this measure is less precise as a measure of liquidity due to different interest rates. Online Appendix Figures A.V and A.VI show that we find very similar results for low and high liquidity households, defined as those in the lowest and highest terciles (quintiles) of our final (broad) sample. Finally, we also confirmed these findings for other proxies for liquidity, such as splitting households by income during the three months prior to February of each year, and classifying households expecting small refunds or payments as less likely to be constrained.

C. Implications

Even less-liquid households smooth consumption through making tax payments, which shows that these households are not actually constrained and that their spending out of refunds even when more liquid is unlikely to be driven by liquidity constraints or financial frictions. The finding that households with lower liquidity have larger consumption responses to expected increases in liquidity than households with higher liquidity is consistent with many previous findings in the literature, which, as we have 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 liquidity, implying that households can and do manage their finances so as to smooth consumption.

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.

D. 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, there is another form of endogenous liquidity that arises from the structure of the US tax system. Households that receive a lot of nonlabor 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 in a model of liquidity constraints, this feature has a quantitatively small effect on the propensity to adjust spending in response to tax refunds and payments. Specifically, in that paper, households with stochastic nonlabor income face an approximation of the US tax system. The model is calibrated to match the observed distribution of payments and refunds and generates an average propensity to spend out of refunds of 30 percent. Most importantly for our purposes, the model implies relatively little variation in consumption reaction across the range of payments and refunds. The propensity to spend rises by only 6 percent 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. Thus, this source of endogenous liquidity is unlikely to be the cause of our results.

In sum, because there is no consumption response to payments, 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. 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.

VI. 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 they then make transfers, almost entirely among observed, liquid accounts, e.g., from a savings to a 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 nearly costless to

make in terms of time or money, asymmetric consumption behavior appears to be due to mental costs or heuristics.

A. Account Balances

As described in Section V, 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 twenty-fifth 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 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 twenty-fifth 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 twenty-fifth 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 twenty-fifth percentile is a natural consequence of higher-income households having both greater year-to-year volatility of account balances and higher payments or refunds.

B. Account Transfers

The pattern of transfers that households in our sample make reveals both that 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 showed that households smooth consumption through tax payments by increasing transfers into an observed account (including transfers from one account to another) in the month before making a payment. These transfers are equivalent in amount to about one-quarter of the anticipated payment.

We find an almost identical result from our estimation of equation (3), which measures the cumulative response to payment from 29 days before it is made, controlling for the arrival of information about the payment. The first number in

Panel A. January balance



Panel B. January balance after removing household fixed effects




FIGURE 8. JANUARY ACCOUNT BALANCES AROUND TAX PAYMENT OR REFUND

Note: Panel A is based on raw account balances, and panel B is based on account balances after removing household fixed effects (only for the 2,511 households with at least two years of balance information, leading to 6,085 household-years).

Table 4 shows that upcoming payments cause households to increase transfers into the observed accounts (including internal transfers) by 25 percent of the upcoming payment. Similarly confirming Figure 1, the arrival of a refund increases transfers, but only following arrival, and presumably from checking accounts into savings accounts. More than 40 percent of the refund amount is transferred into the observed accounts in the four months following the arrival (first row of panel B in Table 4). Note that these transfers do not include transfers out of observed accounts into presumably less-liquid saving accounts, as we discuss after this analysis of transfers into observed accounts. Both sorts of transfers “save” funds that are then presumably available for the years in which the household makes payments.

TABLE 4—CUMULATIVE CHANGES AS A PERCENTAGE OF REFUND OR PAYMENT

	Panel A. Percent of payment Days after payment				Panel B. Percent of refund Days after refund			
	0	28	56	84	0	28	56	84
Transfer in	24.44 (4.21)	24.45 (5.93)	18.07 (7.80)	12.95 (9.95)	2.58 (1.36)	32.94 (2.70)	37.85 (3.87)	41.95 (5.02)
External savings and loan payments	-4.82 (2.52)	-4.74 (4.78)	-3.83 (6.98)	-3.25 (9.12)	0.33 (1.00)	14.28 (2.14)	13.57 (3.28)	13.15 (4.51)
Changes in account balances	19.54 (27.70)	-99.27 (54.16)	-114.98 (80.79)	-138.11 (106.54)	22.96 (9.28)	100.05 (18.25)	108.72 (26.78)	127.25 (36.02)
Probability of overdraft	-0.04 (0.02)	0.03 (0.04)	0.06 (0.06)	0.07 (0.09)	-0.01 (0.02)	-0.16 (0.03)	-0.23 (0.05)	-0.29 (0.07)
Income	-4.72 (7.02)	-16.72 (13.84)	-26.79 (20.54)	-33.72 (26.97)	0.10 (2.75)	-1.54 (5.57)	-1.78 (9.25)	-0.58 (12.95)

Notes: This table shows the cumulative response (in percentage of refund or payment) of different account measures 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 $\sum_{k=-29}^{\kappa} \gamma_k^p$ and $\sum_{k=-29}^{\kappa} \gamma_k^r$, 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. *Transfer in* is the sum of all transfers into any observed account (including among accounts). *External savings and loan payments* is the sum of outflows to major financial services companies, clear debt payments, and credit cards.

We also find that households manage liquidity and make transfers in 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. We find 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 nonresponse of households to good news about taxes is because households expecting large refunds are constrained. That is, households that are saving to make payments or that are only expecting to receive small refunds can save less and consume a little more in response to news that they will owe less tax this year. Panels C and D of Figure 9 rule out this hypothesis, using the broad sample of all households (for reasons of statistical power). Households expecting payments or small refunds also react to bad news about the tax amount but do not respond to goods news.³¹

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.

³¹ We find the same result with less precision for panels C and D for our baseline sample of households, as shown in online Appendix Figure A.VII.

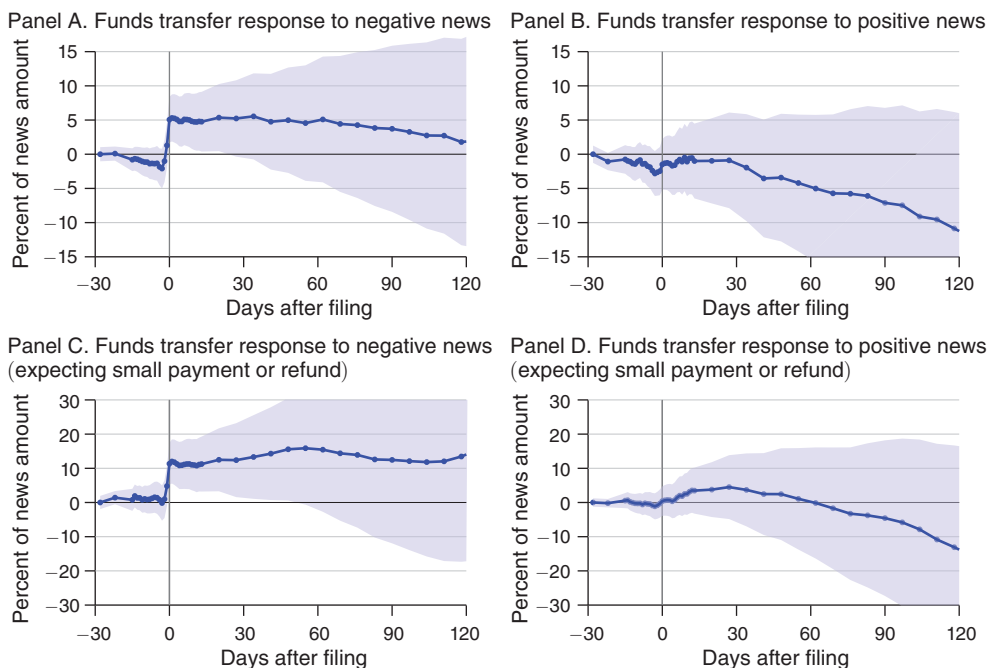


FIGURE 9. FUNDS TRANSFERS IN RESPONSE TO NEWS ABOUT TAX AMOUNT

Notes: 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 using our final sample. Panels C and D show the response of household-years in our broad sample with small amounts of expected payments or refunds, defined as the bottom quintile of absolute expected refunds or payments.

C. Debt Payments and Transfers from External Accounts

Only a small share of these transfers in advance of payments come from accounts other than the observed 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 savings in response to impending payments. But this net dissaving accounts for only 5 percent of the upcoming payment. Comparing rows 1 and 2 of Table 4, we see that this accounts for only one-fifth of the 25 percent 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 in net savings). Households accumulate less than \$500 from these external sources. Panel C of Figure 1, in comparison, shows a large increase in transfers among 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

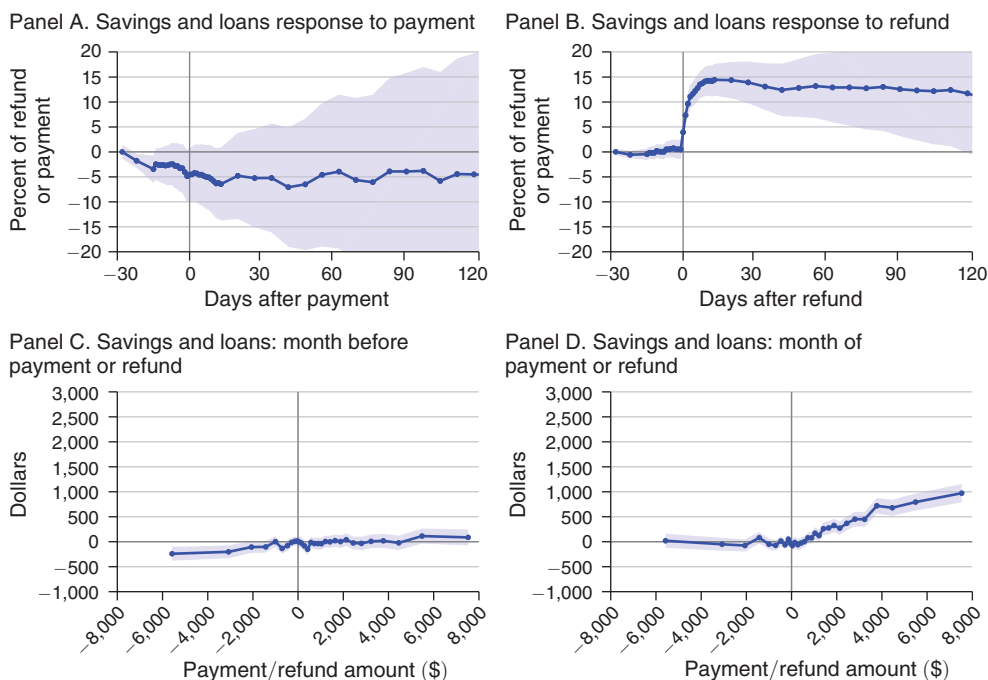


FIGURE 10. RESPONSE OF SAVINGS AND LOANS TO TAX PAYMENTS AND REFUNDS

Notes: Panels A and B show the cumulative response of external savings and debt payments to making tax payments and receiving refunds, respectively. Panels C and D show the abnormal external savings and debt payments around tax payment and refund dates, with respect to payment or refund amounts. The markers denote averages at every 5 percent of the data for those who received refunds, and every 10 percent of the data for those who made payments.

and debt payment. Taking panel B of Figure 10 for example, the cumulative increase in transfers to noncore accounts rises to 13 percent of the refund three months following arrival. Comparing rows 1 and 2 in panel B of Table 4 shows that this is still only one-third of total transfers observed, meaning that most transfers triggered by tax refunds also occur among core accounts.

D. 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).³²

E. 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

³²We also find no evidence that income responds to any news about the tax amount, or that it falls before or after refund receipt.

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 percent of payments in advance of payment and then a month after the payment have fallen by more than 100 percent 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 percent of the refund, and points estimates suggest that account balances have risen by the entire refund amount by two months after arrival, but 95 percent confidence intervals contain more reasonable estimates (including roughly 50 percent of the refund amount).³³

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).

F. 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 reasonable interest rates 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 consumption.³⁴

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

³³Note that online Appendix Table A.II displays the cumulative response of all categories of account inflows and outflows in one place.

³⁴People 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 their payment on time.

(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.

VII. 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 a refund or payment and should reduce consumption in response to bad news about a 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 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.

This discussion assumes that return preparation provides information about payment or refund and not about future nontax income or current wealth. That is, we assume that while the household may learn about the previous year’s income, this information is only relevant for tax status condition on their (known) current financial wealth, current after-tax take-home pay, and expected future non-tax income. If this assumption were incorrect, then the estimated responses to news would also include the response to information about current wealth or future income.

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, consumption does not decline in the month before making a payment (that is, relative to the timing of the “cash flow” 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).

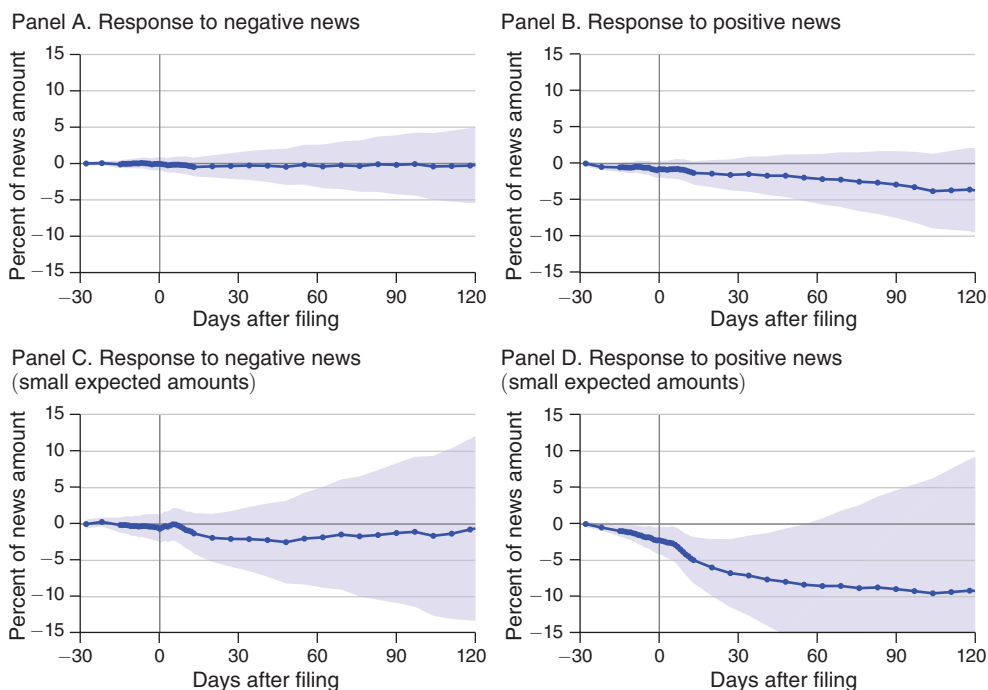


FIGURE 11. CONSUMPTION RESPONSE AROUND NEWS DURING TAX PREPARATION

Notes: 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 using our final sample. Panels C and D show the response of household-years in our broad sample with small amounts of expected payments or refunds, defined as being the bottom quintile of absolute expected refunds or payments.

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 report results based on our broad sample and show no statistically significant response of consumption to bad news about future payment or refund. However, we find some decline in consumption in response to positive news, which suggests that, at least for this subsample of households that do not all receive both payments and refunds, positive news about payments may be negative news about resources for consumption, such as investment wealth. Online Appendix Figure A.VIII shows no significant effects for either positive or negative news for our baseline sample (and point estimates suggest that positive news increases consumption and negative news decreases it).

While panels A and B of Figure 11 show that households do not adjust their consumption as news arrives, could our findings be due to mismeasurement of news about taxes uncovered during filing?

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 are receiving more money than expected. 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 2020).

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 households make transfers in response to our measure of news.

VIII. Theory: Refunds, Tax Payments, and Consumption

This section contrasts our results with the prediction of the main theories of household behavior that explain the previously observed high marginal propensity to consume out of transitory income.

A. Buffer-Stock Model

In models with liquidity constraints, spending responses to expected refunds and payments may be asymmetrically constrained.³⁵ 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.³⁶ Thus, a model with liquidity constraints predicts different responses to news about refunds versus payments, as we find. However, in contrast to our findings, the buffer-stock model predicts (i) no spending increase in response to refund arrival for liquid households, (ii) a decline in consumption prior to payment, (iii) increases in spending shortly before refund arrival, and (iv) large changes in response to news about payments (or small refunds) for the least liquid households (see Figure 5 and discussion).

While we also find no spending responses to news while people prepare their taxes, this is not a strong rejection of the buffer-stock model for two reasons. First, for liquid households, tax refunds and payments are small relative to lifetime income; as such, the consumption responses to news are predicted to be small (relative to the power of our tests). Second, for relatively less-liquid households that are expecting to make payments and thus should have stronger responses, the power of our tests is weak, and we cannot rule out some response to information.

³⁵ Shea (1995a) investigates this asymmetry in aggregate data and shows that aggregate consumption responds more to predictable decreases in income than to predictable increases.

³⁶ Or if they can borrow, they may choose not to make the effort or pay the fixed cost to obtain credit, or they may choose not to pay the higher interest rate on unsecured borrowing, or they may choose not to take on the costs associated with turning less liquid assets into consumption.

B. Mental Accounting

In models of mental accounts (Thaler 1985, Shefrin and Thaler 1988, Arkes et al. 1994), households use rules-of-thumb to determine their income and consumption paths. The theory is driven by empirical observations (Thaler 1999) and hence is limited in its predictions. In a specific framework, Shefrin and Thaler (1988) proposes a theory of consumption and saving based on a behavioral life cycle. According to the theory, households maintain three mental accounts: current income, current assets, and future income. The *current income* mental account contains regular income and funds consumption. Hence, consumption tracks current income, and households' consumption is hypersensitive to income. *Current assets* are less liquid and are saved for future use (e.g., retirement). The *future income* account is reserved for anticipated income (e.g., future inheritance), and households follow a rule-of-thumb of not consuming out of it until arrival (even if the arrival is guaranteed).

In the context of our empirical setting, the behavioral life-cycle theory of Shefrin and Thaler (1988) predicts our main findings. Anticipated tax refunds are classified as future income until their actual arrival, and hence are not consumed ahead of time. Upon arrival, refunds join the same pool as income from work and consumption, but because refunds are not regular income, they are viewed like bonuses and so are partly consumed and partly saved. McDowall (2019) contains an elegant formulation of this idea that matches the spending response of households out of refunds. Shefrin and Thaler (1988) further argues that households use less-liquid accounts to store current assets as a commitment mechanism against frivolous consumption. Thaler (1994) proposes that a higher saving rate could be achieved simply by changing the mental classification of tax refunds: by depositing refunds into retirement accounts (e.g., IRAs).³⁷

Shefrin and Thaler (1988) does not explicitly discuss tax payments. However, it is plausible that tax payments would be considered as external to the income-consumption cycle that takes place in the current income account. As such, households would use their liquid or less-liquid assets to fund payments without lowering consumption. An interesting possibility is that people mentally associate the tax payment with the source of the income that caused the additional taxes. Thus, people who have lots of dividends or capital gains (realized or just distributed by mutual funds) and therefore owe taxes feel that it is appropriate to pay from these investment accounts. While this may be the case, this does not explain much of the behavior we document. Transfers from non-core accounts such as brokerage accounts make up very few of the transfers we observe prior to payment.

The predictions of a model of hyperbolic discounting with a commitment saving device are quite similar (Angeletos et al. 2001). A rational hyperbolic saving model predicts some anticipatory spending ahead of refund arrival. In our empirical setting, however, the economic costs of transferring funds from a savings account to a checking account are negligible.

³⁷Shefrin and Thaler (1988) predicts that a tax cut without a change in the withholding rate should result in a greater savings rate because tax refunds will increase over time and would be considered a "bonus" (partly consumed and partly saved). In contrast, higher net income would be allocated for consumption. Sahm, Shapiro, and Slemrod (2012) compares the response of consumption to rebates and reductions in withholding. Feldman (2010) shows that savings rates decline the year following a decrease in the tax withholding requirement.

C. *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).

IX. 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, a finding that 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 have the liquidity to smooth consumption through refunds. They smooth consumption primarily by moving funds around among 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 can weather temporary declines in income without cutting consumption; therefore, they spend out of tax refunds largely *by choice* rather than due to liquidity constraints. This behavior is consistent 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. Among the models describing household behavior, our results best match mental accounting behavior (Shefrin and Thaler 1988), and our study provides a relatively clean test of its predictions.

Our results have important implications for at least two types of policies. First, the behavior that we document would lead to quantitatively different dynamic responses to income shocks in economic models. Because these consumers choose not to spend out of future income, their responsiveness to information about future income or interest rates would be limited and not determined by the distribution of liquid wealth. Similarly, their response to changes in current income would differ by sign but not be as dependent on liquidity. In macroeconomic models such as heterogeneous-agent New Keynesian models, these features would alter the response of consumption demand both to economic shocks and to policy responses ranging from fiscal transfers to forward guidance.

Second, Shefrin and Thaler (1988) and Thaler (1994) argue that in models where people behave in the way that we find they do, savings would be higher if tax refunds were larger (i.e., higher withholding rate) and if they were deposited into households' savings accounts (e.g., retirement saving accounts) directly instead of into households' checking accounts.

REFERENCES

- Aaronson, Daniel, Sumit Agarwal, and Eric French. 2012. "The Spending and Debt Response to Minimum Wage Hikes." *American Economic Review* 102 (7): 3111–39.
- Agarwal, Sumit, Souphala Chomsisengphet, Stephan Meier, and Xin Zou. 2019. "In the Mood to Consume: Effect of Sunshine on Credit Card Spending." Unpublished.
- Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles. 2007. "The Reaction of Consumer Spending and Debt to Tax Rebates: Evidence from Consumer Credit Data." *Journal of Political Economy* 115 (6): 986–1019.
- Agarwal, Sumit, and Wenlan Qian. 2014. "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore." *American Economic Review* 104 (12): 4205–30.
- Aguiar, Mark A., Mark Bils, and Corina Boar. 2020. "Who Are the Hand-to-Mouth?" NBER Working Paper 26643.
- Aguiar, Mark, and Erik Hurst. 2005. "Consumption versus Expenditure." *Journal of Political Economy* 113 (5): 919–48.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives* 15 (3): 47–68.
- 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 (3): 331–47.
- Aydin, Deniz. 2019. "Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines." Unpublished.
- 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 (4): 1504–57.
- Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker. 2021. "Replication Data for: Asymmetric Consumption Smoothing." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E120690V1>.
- Ben-David, Itzhak, and Marieke Bos. Forthcoming. "Impulsive Consumption and Financial Well-Being: Evidence from an Increase in the Availability of Alcohol." *Review of Financial Studies*.
- Bodkin, Ronald. 1959. "Windfall Income and Consumption." *American Economic Review* 49 (4): 602–14.
- Broda, Christian, and Jonathan A. Parker. 2014. "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption." *Journal of Monetary Economics* 68: S20–S36.
- Caballero, Ricardo J. 1995. "Near-Rationality, Heterogeneity, and Aggregate Consumption." *Journal of Money, Credit, and Banking* 27 (1): 29–48.
- Caldwell, Sydnee, Scott Nelson, and Daniel Waldinger. 2020. "Tax Refund Uncertainty: Evidence and Welfare Implications." Unpublished.
- Campbell, John Y., and N. Gregory Mankiw. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." In *NBER Macroeconomics Annual 1989*, Vol. 4, edited by Olivier Jean Blanchard and Stanley Fischer, 185–216. Cambridge, MA: MIT Press.
- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten van Rooij. 2019. "Asymmetric Consumption Effects of Transitory Income Shocks." *Economic Journal* 129 (622): 2322–41.
- Farrell, Diana, Fiona Greig, and Amar Hamoudi. 2019. *Tax Time: How Families Manage Tax Refunds and Payments*. JPMorgan Chase Institute.
- Feldman, Naomi E. 2010. "Mental Accounting Effects of Income Tax Shifting." *Review of Economics and Statistics* 92 (1): 70–86.
- Fuster, Andreas, Greg Kaplan, and Basit Zafar. 2018. "What Would You Do with \$500? Spending Responses to Gains, Losses, News and Loans." NBER Working Paper 24386.
- Ganong, Peter, and Pascal Noel. 2019. "Consumer Spending during Unemployment: Positive and Normative Implications." *American Economic Review* 109 (7): 2383–424.

- Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, and Dan Silverman. 2019. "Rational Illiquidity and Consumption: Theory and Evidence from Income Tax Withholding and Refunds." NBER Working Paper 25757.
- Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2014. "Harnessing Naturally Occurring Data to Measure the Response of Spending to Income." *Science* 345 (6193): 212–15.
- 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* 189: 103917.
- Gomes, Francisco, Michael Haliassos, and Tarun Ramadorai. Forthcoming. "Household Finance." *Journal of Economic Literature*.
- Hassan, Tarak, and Nicola Fuchs-Schündeln. 2016. "Natural Experiments in Macroeconomics." In *Handbook of Macroeconomics*. Vol. 2A, edited by John B. Taylor and Harald Uhlig, 923–1012. Amsterdam: North-Holland.
- Hsieh, Chang-Tai. 2003. "Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund." *American Economic Review* 93 (1): 397–405.
- Jappelli, 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 (2): 251–62.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96 (5): 1589–610.
- Jones, Damon. 2012. "Inertia and Overwithholding: Explaining the Prevalence of Income Tax Refunds." *American Economic Journal: Economic Policy* 4 (1): 158–85.
- Jorring, Adam T. 2018. "The Costs of Financial Mistakes: Evidence from U.S. Consumers." Unpublished.
- Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity*: 77–138.
- Kueng, Lorenz. 2018. "Excess Sensitivity of High-Income Consumers." *Quarterly Journal of Economics* 133 (4): 1693–751.
- McDowall, Robert A. 2019. "Consumption Behavior across the Distribution of Liquid Assets." Unpublished.
- Olafsson, Arna, and Michaela Pagel. 2018. "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software." *Review of Financial Studies* 31 (11): 4398–446.
- Olafsson, Arna, and Michaela Pagel. 2019. "Borrowing in Response to Windfalls." Unpublished.
- Parker, Jonathan A. 1999. "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes." *American Economic Review* 89 (4): 959–73.
- Parker, Jonathan A. 2017. "Why Don't Households Smooth Consumption? Evidence from a \$25 Million Experiment." *American Economic Journal: Macroeconomics* 9 (4): 153–83.
- 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 (1): 183–99.
- Reis, Ricardo. 2006. "Inattentive Consumers." *Journal of Monetary Economics* 53 (8): 1761–800.
- Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod. 2012. "Check in the Mail or More in the Paycheck: Does the Effectiveness of Fiscal Stimulus Depend on How It Is Delivered?" *American Economic Journal: Economic Policy* 4 (3): 216–50.
- Shea, John. 1995a. "Myopia, Liquidity Constraints, and Aggregate Consumption: A Simple Test." *Journal of Money, Credit, and Banking* 27 (3): 798–805.
- Shea, John. 1995b. "Union Contracts and the Life-Cycle/Permanent-Income Hypothesis." *American Economic Review* 85 (1): 186–200.
- Shefrin, Hersh M., and Richard H. Thaler. 1988. "The Behavioral Life-Cycle Hypothesis." *Economic Inquiry* 26 (4): 609–43.
- Slemrod, Joel, Charles Christian, Rebecca London, and Jonathan A. Parker. 1997. "April 15 Syndrome." *Economic Inquiry* 35 (4): 695–709.
- Smeeding, Timothy M., Katherin Ross Phillips, and Michael O'Connor. 2000. "The EITC: Expectation, Knowledge, Use, and Economic and Social Mobility." *National Tax Journal* 53 (4): 1187–209.
- Souleles, Nicholas S. 1999. "The Response of Household Consumption to Income Tax Refunds." *American Economic Review* 89 (4): 947–58.
- Souleles, Nicholas S. 2000. "College Tuition and Household Savings and Consumption." *Journal of Public Economics* 77 (2): 185–207.

- Stephens, Melvin, Jr.** 2001. "The Long-Run Consumption Effects of Earnings Shocks." *Review of Economics and Statistics* 83 (1): 28–36.
- Stephens, Melvin, Jr.** 2003. "'3rd of Tha Month': Do Social Security Recipients Smooth Consumption between Checks?" *American Economic Review* 93 (1): 406–22.
- Stephens, Melvin, Jr.** 2006. "Paycheque Receipt and the Timing of Consumption." *Economic Journal* 116 (513): 680–701.
- Thaler, Richard.** 1985. "Mental Accounting and Consumer Choice." *Marketing Science* 4 (3): 199–214.
- Thaler, Richard H.** 1994. "Psychology and Savings Policies." *American Economic Review* 84 (2): 186–92.
- Thaler, Richard H.** 1999. "Mental Accounting Matters." *Journal of Behavioral Decision Making* 12 (3): 183–206.
- Zeldes, Stephen P.** 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97 (2): 305–46.