For Online Publication: Appendices to Consumption Behavior Across the Distribution of Liquid Assets

Robert A. McDowall[1]

New York University

Data Appendices

Appendix A. Transaction Categorization

I categorize expenditures in accordance with the 2019 NIPA Handbook.[2] Expenditures include credit, debit, and deposit outflows categorized according to Merchant Category Codes (MCCs), attributed to the time of purchase. Within the total expenditure category I attribute payments to unobservable credit accounts as contemporaneous expenditure. Where appropriate I test the robustness of this assumption by attributing these outflows to debt repayment instead, or restricting to a subset of the population that does not make any payments to unobservable credit card accounts. Throughout the paper expenditures are classified as follows:

- *Total Expenditure* (e): All account (credit and deposit) outflows, excluding account transfers and credit card balance payments for which card purchases are observable.
- *Non-Durables* ($e_{\mathcal{ND}}$): Groceries, entertainment, fuel, discount and drug stores, direct market catalogs, or services such as utilities, telecommunications, insurance, health expenses, other bills, food services, travel services and other personal and professional services.

Additionally, I construct a taxonomy decomposing *Total Expenditure* =

[1]Department of Economics, New York University, 19 West 4th Street, New York, NY 10012 (e-mail: robertmcdowall@nyu.edu).

[2] https://www.bea.gov/sites/default/files/methodologies/nipa-handbook-all-chapters.pdf

+ *Non-Durables Goods* ($e_{\mathcal{ND}_{\mathcal{G}}}$): Groceries, entertainment, fuel, discount and drug stores, direct market catalogs.

+ *Services* ($e_{\mathcal{S}}$): Education, healthcare, travel, telecommunications, utilities, housing, rent, other bills, financial services, personal or professional services, and food services.

+ *Durables* ($e_{\mathcal{D}}$): Auto purchases, repairs, and parts; healthcare equipment; home improvement goods and appliances.

+ *Illiquid Debt Payments* ($e_{\mathcal{B}}$): student loans, auto loans.

The remainder of total expenditure not attributed to one of the subcategories above includes unclassified paper checks, cash outflows, and payments to unobserved credit card accounts. This unclassified proportion makes up roughly one third of average monthly total expenditure.

I develop an imputation procedure to assign outflows in the forms of cash, paper checks, and payments to unobserved credit card accounts to expenditure categories. I detail this procedure below. Throughout the paper income is classified in the following manner:

- *Categorized Income*: Labor income (direct deposit and payroll), tax refunds, social security payments, unemployment insurance, investment income.
- *Total Income*: Categorized income plus paper checks and cash deposits.

Roughly 73% of income in the transaction data is categorized by source (ie. payroll, social security, unemployment insurance, etc.) while the rest comes in the form of paper checks (11%), cash (2.5%), and ACH deposits and miscellaneous inflows. Balance sheet variables are defined in the following manner:

- *Transaction Account Balances*: The sum of checking and savings account balances.
- *Total Liquid Balances*: Transaction account balances plus observable brokerage, money market, and certificates of deposit.

Furthermore, the taxonomy of inflows and outflows from transaction accounts is completed by defining transfers to illiquid savings accounts and unobserved demand deposit accounts. At the daily frequency, while checking and savings accounts are perfectly fungible (money can be transferred immediately between these accounts within the bank), money market and brokerage account transfers operate on some delay (usually one to two business days), while liquidations of retirement accounts and certificates of deposit often entail some cost.

Appendix B. External Validation

Here I report supplemental external validation measures. I compare transaction and account data to survey micro data (SCF, CEX, SCPC) and per-capita macro data (PCE, USDA, IRS) for the year 2016. Each of the data sources overlap in this year. In a manner similar to Baker (2015), I also compare the distributions of observables within the transaction data to those of the SCF. The purpose of these benchmarking exercises is twofold: (1) to evaluate the directions of possible bias in the results that follow, and (2) in the spirit of (1) to alleviate concerns that the data's lens on household's overall financial activity is limited.

Table B.1: Expenditure Comparison, Monthly Averages, 2016

Source	Total	Non-Dur.	Durables	Svcs.	Food Svcs.	Groceries
CEX	4775.92	980.92[3]	633.67[4]	2386.83[5]	337.42[6]	303.17
PCE	8455.85[7]	1754.98	891.89	-	-	-
USDA	-	-	-	-	548.84	627.98
BANK	5347.84	1059.18	168.40	1252.30	306.41	220.49

I benchmark household income measures to the CEX and Survey of Consumer Finances (SCF), and tax return outcomes to those reported by the Internal Rev-

[3] CEX measure includes: Food at home, laundry and cleaning, postage/stationery, apparel, motor oil/gasoline, entertainment, smoking supplies, and drugs.

[4] Housekeeping and other household supplies, furnishings, and equipment; reading; medical supplies; auto repairs; and vehicle purchases.

[5] Food away from home, alcoholic beverages, transportation, insurance, education, housing services, personal services, telecommunications, and other bills.

[6] Food away from home and alcoholic beverages.

[7] Calculated from U.S. BEA annual Personal Consumption Expenditures and U.S. Census Bureau Total Household data.

enue Service (IRS). Total monthly take-home income tracks above the CEX measure, driven in part by the exclusion of unbanked households in the transaction data. While the CEX and transaction data measure take-home (post-tax) income, the SCF survey requests that households report gross (pre-tax) annual income. I adjust SCF income to post-tax levels using Congressional Budget Office (CBO) data on average tax rates within pre-tax income brackets. This procedure is described below.

Table B.2: Income Comparison, Monthly Averages, 2016

Source	Monthly Income	Federal Tax Refund	State/Local Tax Refund	Federal Tax Payment
CEX	5347.92	-	-	-
IRS	-	2860	1622	5422
BANK	5949.35	2844.61	1218.37	1591.59

I compare the transaction data to the SCF and the Survey of Consumer Payment Choice (SCPC). Both surveys are designed to be broadly representative. While the SCF measures the total value of accounts held by households, the SCPC asks respondents to exclude accounts exclusively held by their spouse or partner and represents the sum of primary and secondary checking accounts. The transaction data substantially understates total liquid balances available to households, but tracks transaction accounts (checking and savings) quite well. This understatement appears to be more pronounced at higher levels of liquid assets.

Table B.3: Liquid Asset Comparison, Quantiles, 2016

	Checking			Savings			Liquid Balances[8]			Credit Card Debt		
Source	25th	Median	75th	25th	Median	75th	25th	Median	75th	25th	Median	75th
SCF	300	1700	5600	0	10	5000	800	3800	16000	0	0	2100
SCPC	200	1000	3500	-	-	-	-	-	-	0	0	2000
BANK	338	1251	3687	0	0	300	459	1796	6182	0	0	0

Additionally, the SCPC surveys how many checking and savings accounts consumers manage. Conditioning on at least one checking account, I compare the

[8]Liquid balance measures include checking, savings, money market, brokerage accounts, and certificates of deposit (retirement account balances are excluded).

frequencies of holding additional accounts across the datasets (Table 6). The transaction data understates the number of checking accounts available to consumers. However, according to the SCPC survey the median balance in secondary checking accounts is $0, and the 75th percentile of secondary checking account balances is just $100. This provides suggestive evidence that primary accounts are largely representative of day-to-day financial activity.

Table B.4: Number of Accounts, Frequencies, 2016

	Checking			Savings			
Source	1	2	3+	0	1	2	3+
SCPC	0.665	0.247	0.087	0.208	0.471	0.205	0.116
BANK	0.879	0.107	0.014	0.61	0.345	0.039	0.012

Appendix B.1. SCF Post-Tax Adjustment

In order to obtain a measure of take-home income from the pre-tax figures reported by SCF respondents I utilize the CBO's 2016 Distribution of Household Income report. Specifically, I utilize the income thresholds reports by the CBO for two-person households across quintiles and for the top one percent of earners. The adjustment includes federal taxes only - namely personal income, payroll, excise, and corporate income taxes. Table $A.24$ reports the adjustment made within each income bracket.

Table B.5: Average Taxes by Bracket, 2016. Source: CBO

Bracket	Q1	Q2	Q3	Q4	Q5	99th
Bounds	<32.5	32.5-54.8	54.8-81.8	81.8-126.1	126.1-546.8	>546.8
Average Tax Rate	0.017	0.094	0.139	0.179	0.265	0.333

In addition to the distribution of annual income reported in the main text, here I report moments of the SCF income distribution (before and after adjustment), as well as moments of annual take-home income observed in the transaction data.

It is important to note that the SCF adjusted measure is not adjusted for state income taxes. Accounting for this component adds substantial complication. In addition to substantial variation in personal income tax rates across states, there is substantial variation across households within states. A number of states have no

Table B.6: Annual Income, 2016

Source	25th	Median	75th
SCF	30,379	58,733	104,302
SCF, Adjusted	29,863	50,569	85,632
BANK	31,754	50,568	82,484

personal income tax (including Florida and Texas) while others levy substantial personal income taxes (California at 13.3% and New York at 8.82% among the highest[9]). According to the Tax Policy Center, 30 percent of taxpayers itemized their deductions in 2016. Those who itemize can deduct state income taxes from their federal return. According to the Tax Foundation the choice to itemize is overwhelmingly weighted towards those facing high federal tax rates. Amongst those earning gross income of one hundred to two hundred thousand 76% choose to itemize, and of those earning over two hundred thousand 93.4% itemize.

Appendix C. Tax Refund Responses

Appendix C.1. Event Identification

Tax refunds are identified from transactions to which either a state treasury or the U.S. Department of the Treasury is the counter-party. Since counter-party identification is necessary for transaction identification, those reconciling their taxes via paper check are not included in the main analysis. For the population receiving refunds, this does not appear to be overly restrictive. The IRS reports that roughly 80% of refunds are paid via direct deposit. Direct deposit refunds are about 26% larger, however ($\$2,995$ on average, versus $\$2,370$ for refunds issued via paper check in 2016[10]). Differences in the form of refund receipt are likely driven by age, income, and whether a households is banked. I address these biases further in the external validation section above.

Tax filing dates are identified from the first payment of the calendar year that a household makes to either a brick and mortar or online tax service provider. These filing dates are identified for roughly 18% of households with observed tax refund

[9] Source: Turbotax
[10] Source: IRS

activity. Below I show that the distribution of tax reconciliation dates for this subset largely aligns with the broader population and that the empirical results are similar across these populations; alleviating concerns of selection bias in filing date identification[11].

Roughly 55% of refunds are received by the end of week 10 (early March), while an additional 13% are clustered in the two weeks around the filing deadline[12]. Tax refunds receipt represents a significant cash flow event, equating to 39% of average monthly income. The timing of refund arrival is driven by variation in processing times within and across counter-parties, whereas variation in payments is driven by individual selection and the externally imposed filing deadlines.

Appendix C.2. Imputation Procedure

As described in the main body of the text, I perform an imputation to assign an appropriate portion of cash outflows, unclassified checks, and payments to unobserved credit card accounts to non-durable expenditures. The purpose of these procedure is to overcome a significant difficulty faced by users of administrative transaction data - the categorization of unclassified transactions.

The procedure makes two broad assumptions: 1.) that the proportion of cash, unclassified checks, and payments to unobserved credit card accounts that the household expends on non-durables is commensurate with that of the identifiable portion of total expenditure, and 2.) that the excess response of these categories at income receipt scales proportionally. In the following section (Appendix F.3) I show that shifts in the composition of expenditure around refund receipt are small, and so (2) is likely a reasonable approximation. For assumption (1), the SCPC provides some suggestive evidence. In 2018 roughly 40% of cash and paper check transactions represented purchases of retail goods. Likewise, roughly 36% of payment card (credit, debit, pre-paid) transactions were toward retail goods.

The pre-imputation non-durable response for the total population is reported here:

[11] It is likely that identification of filing dates restricted to brick and mortar and online tax service providers trims two tails of the income distribution: those who self-prepare, and those who employ private accountants to prepare their taxes.

[12] Those filing early likely seek liquidity, while those filing later likely have more complicated returns or prefer to delay the task.

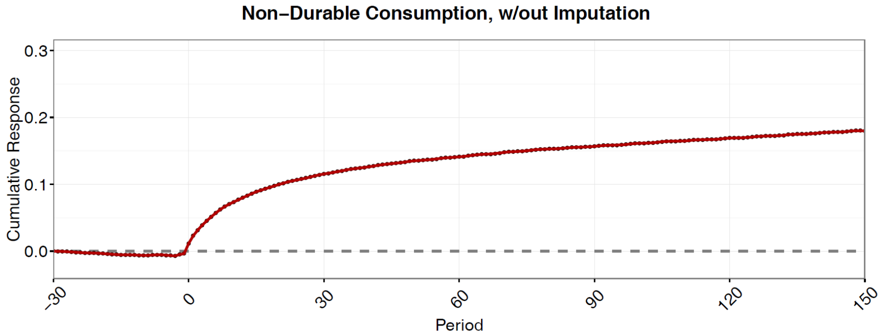

I report the cross-sectional non-durable consumption responses obtained before the imputation procedure (Figure C.17).

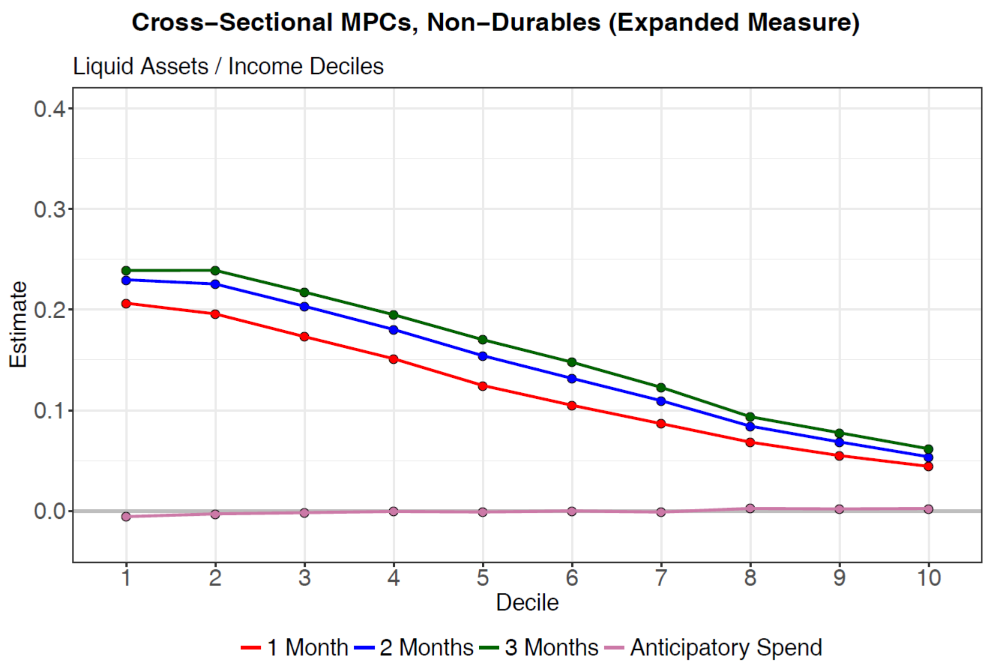

Figure C.1: Non-Durable Responses in the Cross-Section, Expanded Measure Before Imputation

Define the mean observable proportion assigned to non-durables for individuals in population q, $\xi^q \equiv \frac{1}{N}\sum_{i=1}^{N} \frac{e^{i,q}_{\mathcal{ND}_\mathcal{E}}}{e^{i,q}-e^{i,q}_{\mathcal{C}}}$, where $e^{i,q}$ denotes total expenditure, $e^{i,q}_{\mathcal{ND}_\mathcal{E}}$ non-durables under the measure including services, and $e^{i,q}_{\mathcal{C}}$ the unidentified cash, check, and credit expenditures. The imputed non-durable consumption responses for cross-sectional quantile q at lag j are then obtained via $\delta^{\mathcal{ND}_\mathcal{I},q}_{t-l} = \delta^{\mathcal{ND}_\mathcal{E},q}_{j} + \xi^q \cdot \delta^{\mathcal{C},q}_{j}$. Expenditures used to compute the expenditure share ξ are taken from the month prior to tax refund receipt.

As reported above, the average MPCs pre-imputation are 0.121 after one month and 0.16 after one quarter. At the median, MPCs are roughly one-third lower, in line with the one third of expenditure that is unclassified. Qualitatively, the results are robust to this upward shift. As reported in the main text, post-imputation non-durable expenditures accounts for 47% of the total expenditure response. Figure C.18 plots the total expenditure results.

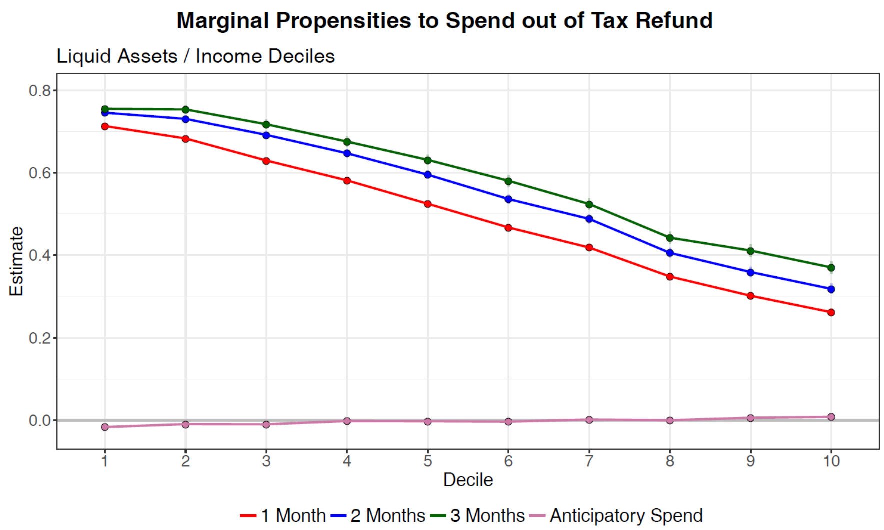

Figure C.2: Total Expenditure Responses in the Cross-Section

Appendix C.3. Expenditure Composition around Refund Receipt

Here I report the average proportion of expenditure by consumption category around tax refund receipt. On the intensive margin average expenditures on durables increases by 93% in the week of refund receipt with respect to the week prior (from \$48 to \$93). Additionally, durable expenditures increase from 3.6% to 4.3% of the expenditure basket (a 19.4% increase). On the extensive margin, the proportion of households observed making durable purchases of more than \$100 increases from 6.9% to 8.6% to 14.8% in the month before, week before, and week of refund receipt, respectively (the increase is from 0.4% to 0.5% to 1.5% of households for purchases in excess of \$1000).

However it is notable that the expenditure response is not simply story of timing large durable purchases to refund receipt. Average expenditure towards

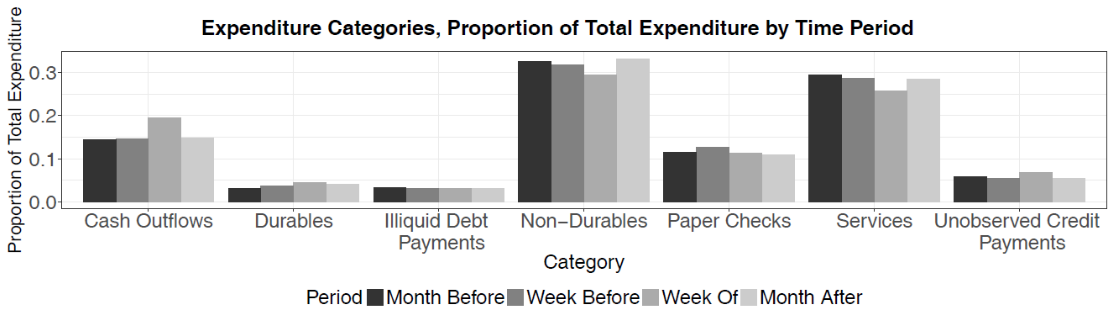

Figure C.3: Weekly Consumption Basket Around Tax Refund

non-durable goods and services increase by 47% and 28% respectively, and their contributions to total expenditure as measured by proportions of the weekly consumption basket are largely stable - crowded out mainly by durables and cash outflows in the week of refund receipt.

The proportion of expenditure in the form of cash outflows increases from 14.6% to 19.4% in the week of refund receipt (an increase of 32.8%). For context, according to the 2016 SCPC, U.S. households reported that 27.4% of their monthly expenditure was in the form of cash. Figure C.19 plots the proportions of consumption basket around refund receipt. Below I report the expenditure composition around refund receipt for the first and fifth quintiles of liquidity:

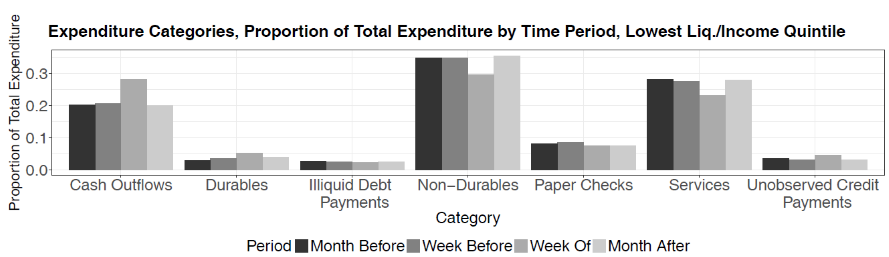

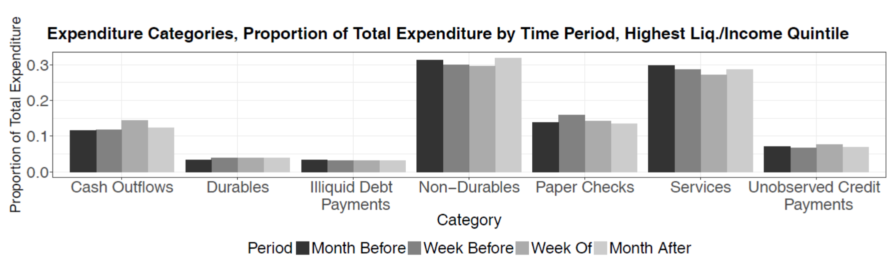

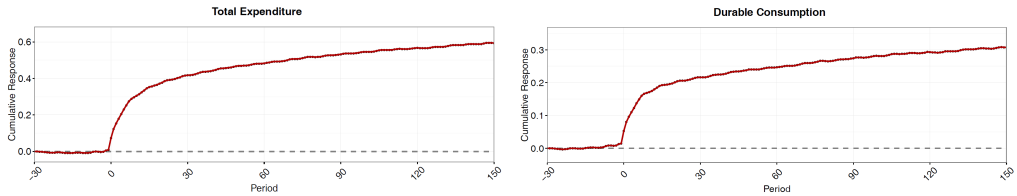

Figure C.4: Total and Durable Expenditure Responses

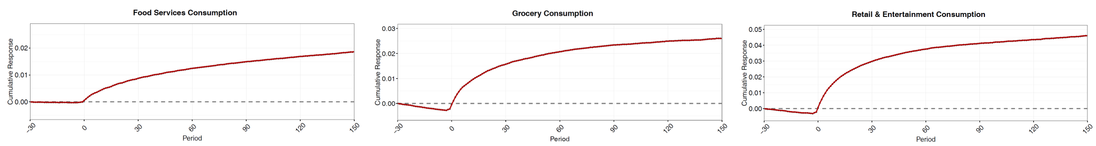

Figure C.5: Additional Expenditure Categories

Appendix C.4. Tax Refund Expenditure Responses

In this section I provide a view of the household balance sheet response to tax refund receipt. On average, the cumulative proportion of a tax refund expended in the month prior to receipt is 0.009 - less than one cent of every dollar received. Tax refund responses to receipt are immediate. On the day of receipt 7.4 cents of every dollar received are expended, 41.7 cents are expended within 30 days, and 59.4 within 150 days. On average, 73% of the 150-day total expenditure response to tax refunds occurs within 30 days of receipt. This includes 66% of non-durable expenditure and 47% of food services consumption responses occurring within the first 30 days. The residual total expenditure not spent towards non-durables or debt payments are classified as durables. I report this category below.

The findings are qualitatively robust to the category of expenditure considered. Crucially, they hold for food services expenditures (restaurants and bars), which can be largely attributed to contemporaneous consumption.

Below I show that credit card balance accumulation in advance of receipt is negligible with a less than one to two cents of every refund dollar spent towards debt repayment.

Appendix C.5. Receipt of Multiple Tax Refunds

In this section I analyze consumption and expenditure responses at the date of tax filing and subsequent refund receipt for those households receiving state and

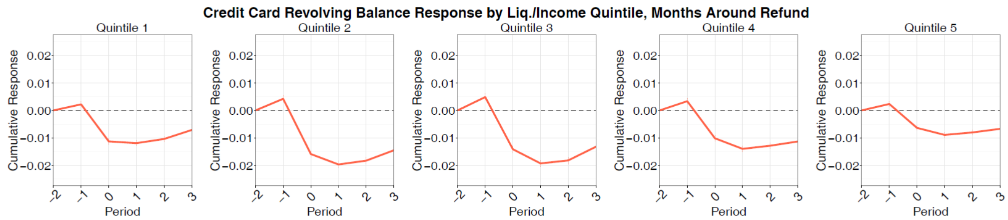

Figure C.6: Credit Card Balances around Refund Receipt

federal refunds in the same year. Notably, excess sensitivity is restricted to receipt (as opposed to filing), and anticipatory spending out of a second refund is muted and excess sensitivity is significant. This, even though receipt of the first occurs shortly beforehand, and both are sizable.

I begin by estimating specification (1), where I_i^1 is the first tax refund received in the calendar year received by household i, and I_i^2 is the second. Identification relies on both variation in calendar of refund arrival, and variation in the timing between receipt of the first and second refunds. The former (latter) is driven by variation within (between) state and federal refund processing times. More formal arguments for identification are outlined above. Both refunds are of a significant magnitude - the average state refund is \$1218, whereas the average federal refund is \$2845. On average, these two refunds arrive within ten days of one another (See Figure 2). Figure C.23 plots the estimated cumulative total expenditure responses proportional to each refund.

I report the non-durable responses to each refund receipt below:

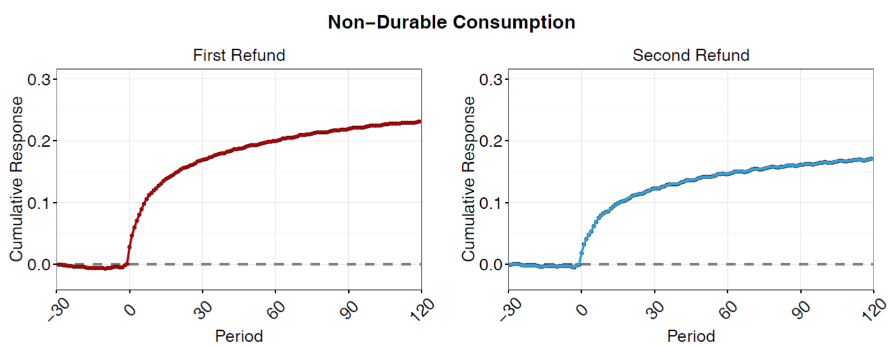

Figure C.7: Tax Refund Non-Durable Consumption Responses, Multiple Refunds

Appendix C.6. Tax Filing

In this section I analyze the consumption response at tax filing. Embedded within the tax filing event is a combination of a resolution of uncertainty and the allocation of household attention to future income receipt. While the date of tax refund receipt is the source of some uncertainty given variation in processing times and the risk of errors in a household's return, the date of tax filing is directly chosen by household. As described above, tax filing dates are identified from the first payment of the calendar year that a household makes to a brick-and-mortar or online tax services provider.

Filing dates are identified for 17.3% of the population receiving refunds. Below I show that this subpopulation is broadly similar to the general refund population in terms of observables, and that the distribution of refund arrival dates (largely driven by filing date self-selection) largely coincides across populations. The latter provides suggestive evidence that the results below are not driven by unobservables.

Table D.18 reports moments summarizing each population of refund recipients. The two groups are broadly similar in terms of liquid wealth and income, with the filing date population appearing to be slight more homogenous. Refunds for the filing date identified subgroup are slightly smaller than the broader population.

Table C.7: Summary Statistics, Tax Refund Recipients, broad population versus filing dates identified subset

	Broad Population				**Filing Dates**			
	Mean	25th	Median	75th	Mean	25th	Median	75th
Liquid Balances	7279	581	1828	5699	7458	815	2292	6366
Income	5259	2425	3868	6245	5329	2624	4088	6389
Tax Refund, First	2072	360	1120	2993	1904	338	1038	2661
Tax Refund, Second	2057	317	937	2666	1753	292	829	2157

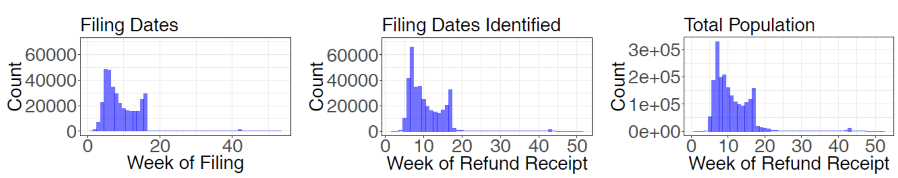

Figure C.8: Filing Date Sub-Population Comparison

I estimate specification (1), where I_i^1 is the first tax refund received in the calendar year received by household i, and I_i^0 is the return at the date of filing. Consumption responses are estimated with respect to a household's total observed tax return (ie. $I_{i,0}^0 = (I_{i,0}^1 + I_{i,0}^2)$ for a household receiving state and federal returns).[13]

Figures C.25 reports the cumulative total expenditure (less payments to tax service providers at the filing date) around the dates of filing and refund receipt. Relative to 30 days prior to filing 0.018 cents of every tax return dollar is expended. An additional 0.57 cents are expended at the date of filing, with a 30 day response of 1.28% of the refund. In comparison, the excess sensitivity of total expenditure on the day of refund receipt is 6.20 cents relative to the day before, with 42.5% of the refund expended over 30 days.

The above result holds for total expenditure, as well as more discretionary non-durable purchases. This dichotomy between receipt and filing, suggests consumption responses are unlikely to be driven by myopia on the part of households, as filing entails a measurable allocation of resources both monetarily (the average cost to file is $82) and in terms of time taken to fill out the return. By itself, the lack of response at filing aligns with an environment in which expectations regarding the size of returns are correct, on average. But systematic inattention to the inflow in advance of filing would surely produce a large degree of excess sensitivity on this day.

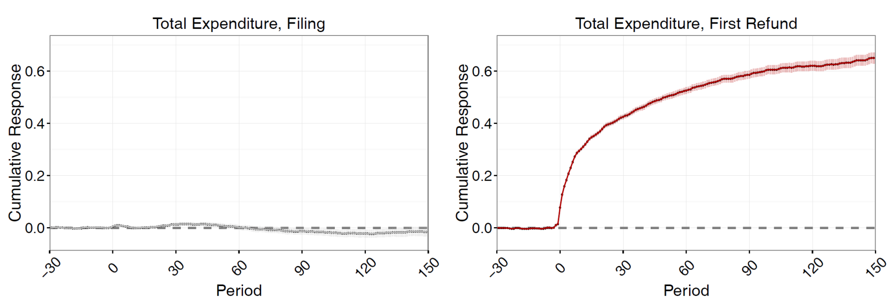

Figure C.9: Expenditure Response at Tax Filing

These consumption responses are consistent with a model in which households

[13]Results are broadly identical when measuring with respect to only the first refund, and are qualitatively similar when the filing date is demarcated by an indicator (ie. $I_i^0 = \mathbb{I}_i^0$), in which case $\{\delta_j^0\}_{t-l}^{t+L}$ measure the response in dollar terms, as opposed to as a proportion of the return.

face externally imposed liquidity constraints. However, this result holds for households with substantial liquid wealth (Figure C.26) who receive multiple large refunds and credit card holders (according to the SCPC, roughly 75% of U.S. households held credit cards in 2016.) who could conceivably borrow interest free for the short interval in advance of refund receipt. Further, this unwillingness to consume in advance of receipt is unlikely to be driven by expectations of delays in refund arrival given that the median refund arrival time is just 8 days after filing, with 94% arriving within 30 days and 97% within 60 days.

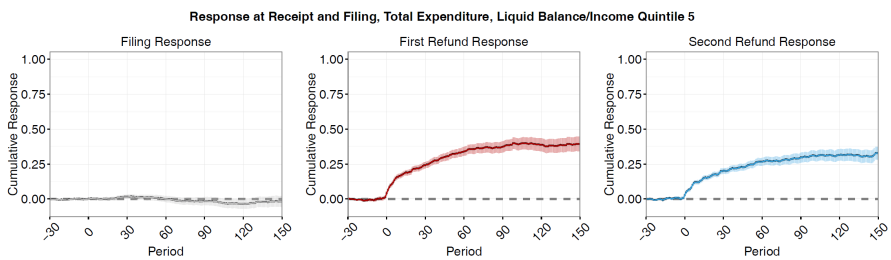

Figure C.10: Expenditure Response, Multiple Refunds and Filing, Highest $\frac{Liquid\ Asset}{Income}$ Quintile

Appendix C.7. Refund Receipt by Income & Cash-on-Hand

I this section I subdivide the population receiving refunds into low, middle, and high income households according to annual income in the calendar year prior to the tax event. These groups include households observed earning less than $\$40,000$, between $\$40,000$ and $\$120,000$, and greater than $\$120,000$ of take home income[14], respectively. Within the population of households receiving tax refunds, the proportion of households within each income group is 0.345, 0.56, and 0.095. I further subdivide income groups according to low, middle, and high levels of cash on hand. These subgroups include households with below median, between the median and $75th$ percentile, and above the $75th$ percentile of $\frac{Liquid\ Account}{Total\ Income}$ for their income group, respectively.[15]

Expenditure responses across these subpopulations share the qualitative characteristics of the wider population (Figure C.27). Responses in levels predictably

[14]The $\$120,000$ cutoff is quite extreme. According to the 2014 ACS, less than 6% of American household's *pre-tax* income exceeded this figure.

[15]In a variance decomposition across observables, liquid balances and income account for over 75% of the explainable variation in MPCs.

Table C.8: Summary Statistics by Income Level

		Median			Mean
	$\frac{Income}{LiquidBalance}$	Income	Liquid Balances	Expenditure	Tax Refund
Low Income	Low	2608	499	2374	1694
($< 40k$)	Middle	2769	1564	2603	1673
	High	2827	5059	2769	1634
Middle Income	Low	5650	1546	5238	2344
($40k - 120k$)	Middle	5897	4763	5649	2388
	High	5738	13823	5646	2417
High Income	Low	13567	5854	12492	3540
($> 120k$)	Middle	13894	14601	13010	3887
	High	13801	35539	13141	4326

decline in cash balances within each income bracket. Amongst high income individuals in the middle cash-on hand group (approximately 1.05 months of income in liquid accounts), the expenditure response over 120 days is 0.41. Amongst the high income, high cash on hand group (over 10 weeks of income in liquid accounts), the 120 day response is 0.34. This particular subgroup represents roughly 2.4% of the overall refund-receiving population. It is only at these high levels of income and cash-on-hand that a notable degree of anticipatory spending is observed - roughly 3.5 cents of every refund dollar in the three weeks prior to receipt. The level of excess sensitivity in the (three) week(s) following receipt, however, is an additional 9.6 (17.6) cents of every refund dollar.

Income processes estimated for each of the nine groups suggest that, within income brackets, household incomes display similar levels of month-to-month persistence and volatility across levels of cash on hand. These income processes are estimated in a similar manner to that of the main text. This suggests that households holding higher levels of cash-on-hand is not driven simply by income volatility, but rather a degree of self-selection related to household choice among savings vehicles and alternate balance sheet structures (ie. stock market participation).

Appendix C.8. The Non-Hand-to-Mouth and Large Refunds

In this section I focus the analysis to a population of households deemed to be non-hand-to-mouth households. The criterion for this classification is chosen to be restrictive - households with ten or more weeks of income in their liquid accounts. I begin by describing this population of households and comparing observables to the population of more constrained households. In this section I present two

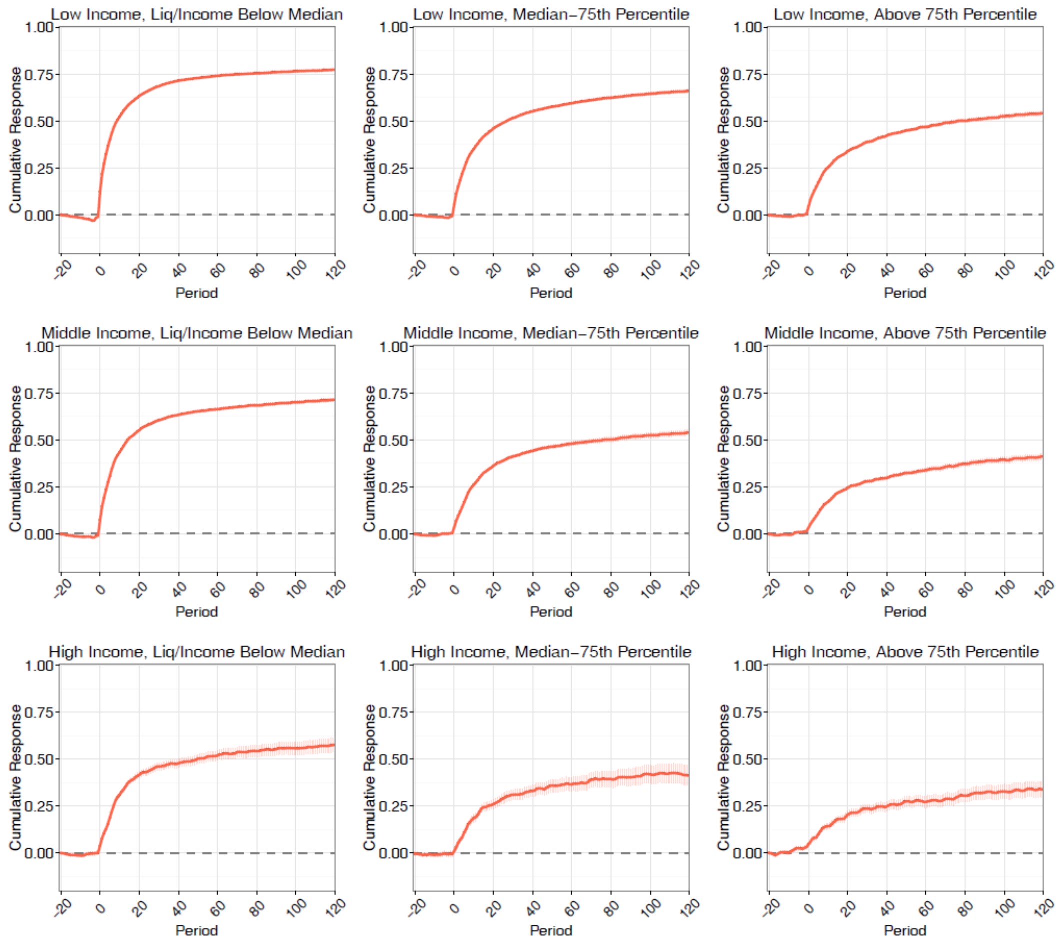

Figure C.11: Expenditure Responses, Income and Asset Distributions

main results. First, even amongst the population of decidedly non-hand-to-mouth households anticipatory expenditure is insignificant and excess sensitivity is great. Second, the results hold even for households receiving very large tax refunds, as measured relative to their average monthly income or expenditure. This result is notable - it suggests that significant promised rewards do not produce large anticipatory responses, even amongst those households with sufficiently liquidity to smooth through the cash-flow event.

In this section I restrict the analysis to the biggest tax refund of a household's calendar year in order to test whether these large salient events can generate significant anticipatory responses. Further, I expand the sample to include all of those households with average $\frac{Liquid\ Account}{Total\ Income} > 2.5$ in the year prior to the tax event for

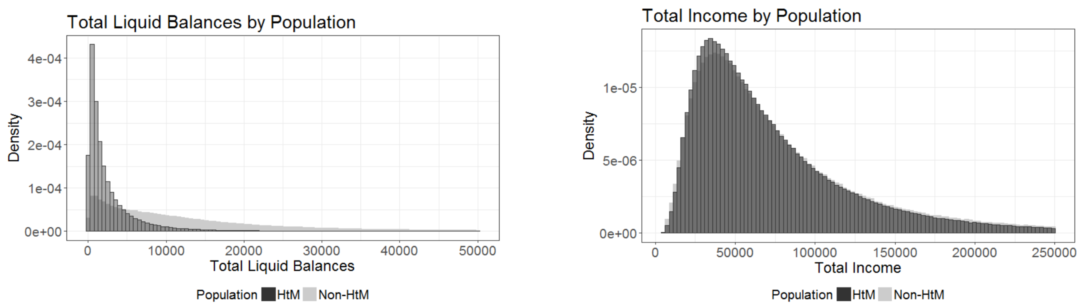

Figure C.12: Hand-to-Mouth and Non-Hand-to-Mouth Comparisons

all years in the sample 2014-17. The median non-hand-to-mouth household holds almost three months of income in liquid assets, while tax refunds represent roughly 2.5 weeks of monthly income (Table 11).

Table C.9: Summary Statistics, the Non-Hand-to-Mouth

	Mean	25th	Median	75th
Liquid Balances	22028	4466	11372	24172
Income	5587	2311	3864	6387
Tax Refund, Larger	3359	1090	2325	4605

Annual incomes amongst the non-hand-to-mouth largely align with those of the rest of the population (hereafter referred to as the hand-to-mouth). As shown in the previous section, levels of liquidity are not highly correlated with levels of income volatility within income brackets. Further, as reported above, these states are relatively persistent over time. Average ages of primary account holders across the non-hand-to-mouth and hand-to-mouth populations of refund recipients are 47 and 41, respectively. This suggests some component of liquidity may be driven by life-cycle liquid wealth accumulation.

Non-hand-to-mouth households display average 150 day total expenditureresponses of 44.5 cents of every refund dollar, with 67.5% of the cumulative 150-day response occurring within the first 30 days. Roughly 2.7 cents are expended over the month in advance of receipt, with an additional 15 cents expended in just the first week.

In order to assess whether large refunds are associated with a greater degree of anticipatory spending or overall consumption smoothing I quantile the non-hand-to-mouth population according to $\frac{Refund}{Total\ Income}$. Table C.18 reports these quantiles. Above the 80th percentile, the average refund totals \$5782 and represents in excess of six weeks of household income. It should be noted that these results come with a caveat - that of a degree of self-selection in terms of refund size. I address this concern by testing household expenditure responses to another form of large expected income - annual bonuses. To address the concern that responses are driven by household expenditure on big ticket items (vacations, appliances, etc.) I confirm that the results persist for non-durables as well.

Table C.10: The Non-Hand-to-Mouth, Refund Size, Median Ratio by Quantile

Quantile	40th	60th	80th	90th
$\frac{Refund}{Total\ Income}$	0.75	1.13	1.73	2.95

Across quantiles of the $\frac{Refund}{Total\ Income}$ distribution excess sensitivity and muted anticipatory responses persist. Of the total 120-day expenditure responses the proportion coming in the first thirty days is 71%, 69%, 68%, and 73% across the 40th, 60th, 80th, and 90th quantiles, respectively. For non-durable consumption the results are similar - 74.5%, 71%, 66%, and 62%, respectively. Crucially, total anticipatory responses are insignificant - totaling less than 3.89 cents of the refund in the 3 weeks in advance of receipt amongst the 40th percentile group, and 22 cents in the next three weeks.

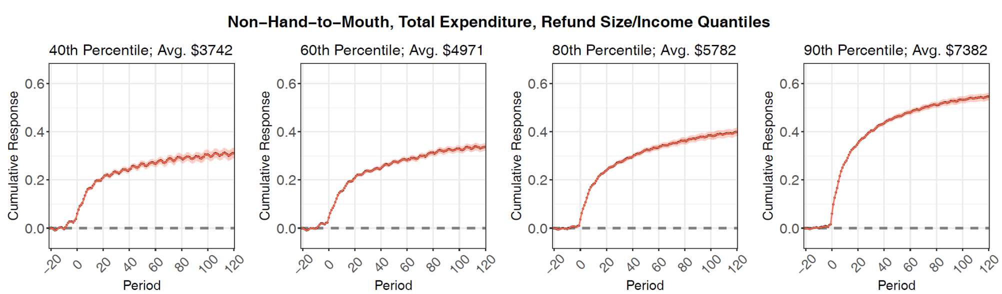

Figure C.13: Expenditure Responses of the Non-Hand-to-Mouth Across Refund Size Relative to Income

To put the above in monetary terms, a back of the envelope calculation implies households amongst the 90th percentile expend an additional \$2861 in the month of refund receipt[16]. In order to address the concern that the results are driven by a levels of expenditure on the part of the household in the baseline period, I instead quantile the non-hand-to-mouth population according to $\frac{Refund}{Total\ Expenditure}$. These results are robust to this alternative measure of liquidity.

Appendix D. Other Forms of Income

In this section I show that the paper's main results are robust to the form of income considered. Here I extend these analysis to regular paychecks and bonus checks. Within subsequent sections analyzing particular events I further detail the subpopulations of interest.

Regular paychecks include employer payroll and direct deposit inflows. Labor income accounts for around 75% of median monthly income, with this proportion decreasing in household total income[17]. The average household within the population receiving refunds receives 2.6 paychecks each month, totaling $1383 apiece. I estimate the consumption responses to regular paychecks within this subset in order to derive an internally valid comparison to tax refund responses.

Bonus checks are identified for a subset of the overall population by first establishing a household's regular pay cadence. Bonus pay is then defined as off-cadence checks from a household's employer that are at least one standard deviation and $100 larger than an average paycheck. I provide further exposition regarding this identification procedure below. The median bonus totals $5630, and the average check represents almost six weeks of household income.

Appendix D.1. Regular Paychecks

In this section I compare regular paycheck responses to tax refund responses for the tax refund population analyzed in the main text. Unlike tax refunds, the day of receipt for a regular paycheck is fully known in advance and occurs at a regular cadence (as opposed to once or twice each year). Whereas the magnitude of

[16]As reported in plot, the average refund above the 90th percentile is $7382, while median refund amongst this group is $6050.

[17]A measurable portion of those unclassified inflows coming in form of paper checks, cash, and unclassified ACH, are likely labor income. According to the National Automated Clearing House Association (NACHA), 82% of U.S. households received their pay via direct deposit/ACH in 2016.

a tax refund can chosen, to the extent described earlier; many workers (especially the salaried) face a fixed structure of compensation dictated by their individual employment contracts.

In comparison to the first tax refund households receive, which averages \$2072 paychecks are significantly smaller, at an average of \$1414. Households receive an average of 2.6 paychecks each month, a number that varies according to each worker's cadence of pay and the number of workers within each household depositing their paychecks to the primary account. In contrast with estimation of tax refund responses, three notable complications arise when estimating consumption responses to paychecks - calendar variation necessary for identification across households, correlation of receipt with calendar-driven expenditures within households, and overlapping response periods.

Taking each of the above in turn; the variation in pay schedules provided by weekly, bi-weekly, bi-monthly, and monthly pay frequencies as well as the alterations to these schedules provided by holidays, allows for disentangling day-of-week effects from paycheck responses. Paycheck receipt, however, is inevitably tied to calendar-driven expenditures (rent, mortgage, etc.), especially early in the month. To address this issue, it is necessary to restrict the analysis to discretionary categories (food services, non-durables) in order uncover the underlying household consumption behavior. Due to the high cadence of regular paychecks, the response periods between one paycheck and the next will overlap over long horizons, thereby biasing the results. In order to address this issue, I focus the analysis to a the week before and after a paycheck.

As with tax refunds, households exhibit excess sensitivity to receipt of regular paychecks across categories of expenditure (Figure D.30) and across quintiles of the liquid balance-to-income ratio (Figure D.31). Amongst the lowest liquidity quintiles, a dip in non-durable expenditure is observed, suggesting that expenditures among these subpopulations are particularly timed to paycheck receipt.

Table D.11: Cross-Section of Liquid Balance-to-Income, Regular Paychecks (Refund Population)

Quintile	1	2	3	4	5
Average Paycheck	1057	1301	1495	1600	1582

In dollar terms, the average household amongst the first quintile spends an

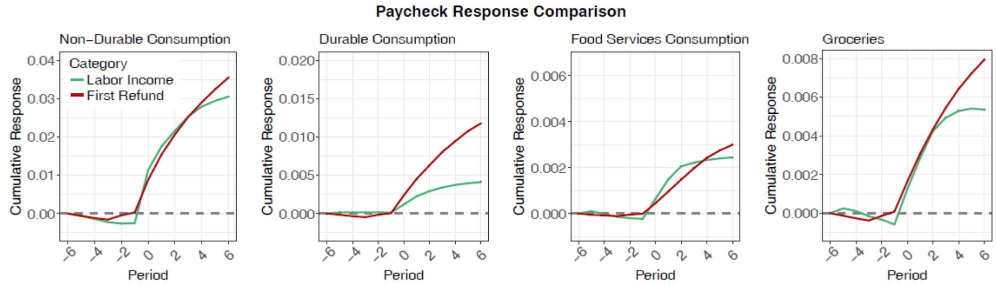

Figure D.14: Consumption Responses Around Regular Paychecks

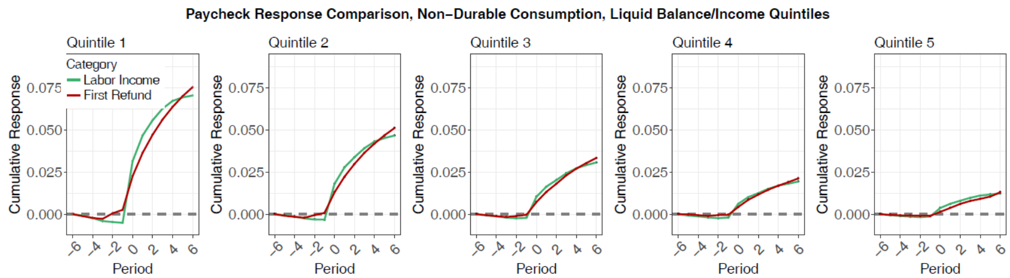

Figure D.15: Non-Durable Goods Responses Around Regular Paychecks, Cross-Section

additional \$123 on non-durable goods (including an extra \$46 in groceries) in the week following paycheck receipt, as compared to the week prior. Amongst this group average checking and savings account balances on the day in advance of paycheck receipt average just over \$500 (Figure D.32[18]).

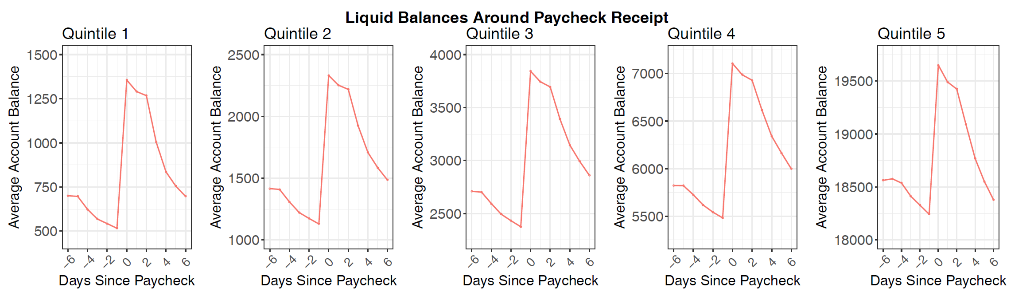

Figure D.16: Path of Liquid Balances Around Regular Paychecks

The consumption responses at payday are somewhat surprising, given the liq-

[18]Figure D.32 plots the average balances around paycheck receipt for each liquid balance quintile. A kink is visible on day 2 in each subplot - aligned with Sundays for most of the population (58% of all paychecks arrive on Fridays).

uid balances on hand that households could utilize to smooth day to day, even amongst the most constrained. Average daily non-durable expenditure in the week in advance of paycheck receipt amongst the first quintile is $24.59, meaning excess sensitivity to a paycheck amounts to an extra five days of baseline expenditure, on average, in the week following receipt. Households among the fourth and fifth quintiles spend an excess of $48 and $30, respectively, in the week after paycheck receipt. The deviations from perfect consumption smoothing observed amongst these populations of non-constrained households represent distortions of 10% and 7% of bi-weekly consumption.

These shifts in consumption are small relative to those induced by tax refund arrival. However, the kink in consumption responses at paycheck arrival is indicative of a preference to consume at receipt rather than at an earlier date, even for those with sufficiently liquidity. The cycles in liquid balances induced by paycheck receipt (Figure D.32) is further evidence of short-term budgeting via internally imposed constraints on the part of households.

Appendix D.2. Bonus Checks

In this section I study a population of 163,300 households observed receiving bonus paychecks. In what follows I describe the population of interest, and then estimate household expenditure responses to this cash flow event. I conclude by drawing internally valid comparisons between bonus and tax refund consumption responses for a subset of this population. Notably, the results of the main text are robust to this form of income.

Employers may aggregate bonuses and regular pay or process bonuses separately. Identifying bonus checks in transaction data requires first determining a worker's regular pay cadence, and then separately distinguishing large, off-cadence checks from their employer.[19] Workers with a regular cadence are identified as those having 90% of their paychecks come 13 to 17 days apart. Bonus pay is then defined as checks from a household's employer arriving outside this window that are at least one standard deviation and $100 larger than an average paycheck. The

[19] The tax code is agnostic about bonus checks arriving with regular paycheck or separately. In order to cleanly identify the magnitude of these inflows, I restrict the analysis to off-cadence checks.

analysis is restricted to those workers receiving at most two such checks during a calendar year.

Figure D.33 plots the timing of bonus checks by calendar week. The majority of bonuses arrive early in the year with a significant portion arriving shortly before March 15th (the federal deadline companies face for deducting bonuses from the previous year's earnings, otherwise known as 409A Day). Bonus checks are signif-

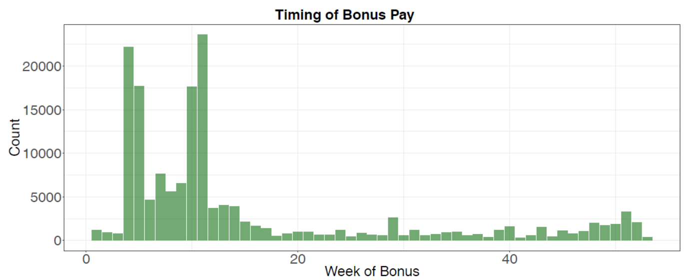

Figure D.17: Timing of Bonus Paychecks

icant, representing over a months income for the average bonus recipient (Table G.18). Also reported are the magnitudes of tax refunds for the subset of bonus recipients receiving them in the same calendar year.

Table D.12: Summary Statistics, Bonus Check Recipients

	Mean	25th	Median	75th
Liquid Balances	12690	2024	4930	13609
Income	8246	4207	6182	9396
Bonus Check	11445	3290	5733	10802
Tax Refund	2818	582	1170	3906

I estimate the expenditure responses of this population to bonus check receipt. Qualitatively, the results are similar to those of the refund population. This includes insignificant anticipatory spend and a large degree of excess sensitivity at receipt, with a majority of the response coming in the first thirty days.

In order to obtain an internally valid comparison across groups, I quintile this group according to the same $\frac{Liquid\ Asset}{Income}$ bounds as the refund population studied in the main text. Table G.22 summarizes.

I estimate the total non-durable responses to bonus receipt by quintile and report these a similar manner to the refund population (Figure D.34). The tight

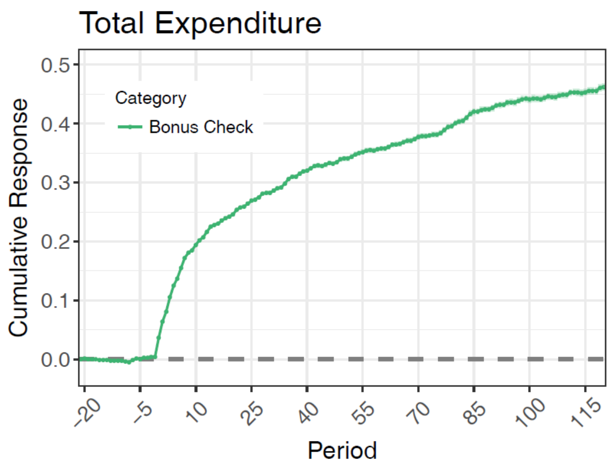

Figure D.18: Total Expenditure Response to Bonus Checks

Table D.13: Cross-Section of Liquid Balance-to-Income, Bonus Paychecks, Averages

Quintile	1	2	3	4	5
Total Income	6999	7916	8674	9078	8924
Bonus Check	6765	8746	11249	13673	16926

correlation between $\frac{Liquid\ Asset}{Income}$ and consumption responses, along with significant responses amongst the highly liquid, is evident here.

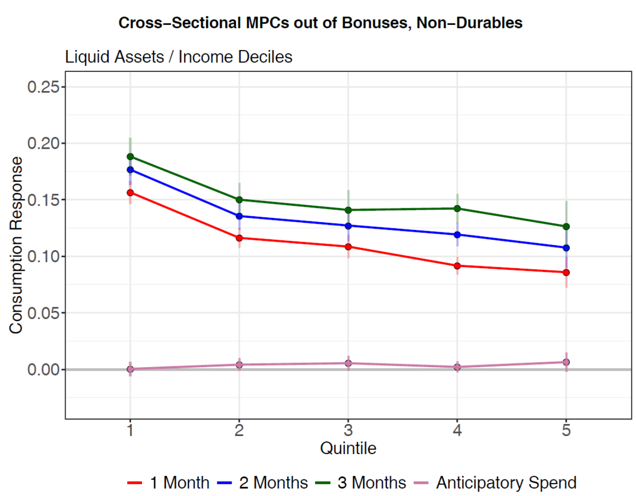

Figure D.19: Non-Durable Response to Bonus Checks

Appendix E. Comparison to Literature, Non-Durable Consumption Responses

Source	Income	Data	Frequency	Non-Durables (Quarter of Receipt)	Methodology	Notes
Souleles (1999)	Tax Refunds	CEX	Quarterly	0.03 (0.02)	OLS	Total response: 0.185 (0.077)
Parker et. al (2005)	2001 Stimulus	CEX	Quarterly	0.386 (0.135)	OLS, 2SLS	-
Parker et. al (2011)	2008 Stimulus	CEX	Quarterly	0.201 (0.067)	OLS	-
Kaplan & Violante (2011)	2001 Stimulus	CEX	Quarterly	0.219 (0.079)	2SLS	Trim 1.5% tails
Keung (2018)	Alaska Permanent Fund	Personal Financial Website	Monthly	0.22	OLS	-

Model Appendices

Appendix A. Structural Estimation

I estimate the components of the income process, $\{\rho, \sigma_{\epsilon}^2, \{\Gamma_t\}_{t=0}^T\}$, from a combination of BLS Consumer Expenditure Survey (CEX) data at the annual frequency, and bank income data at the monthly frequency. The deterministic growth path ($\{\Gamma_t\}_{t=0}^T$) is obtained from the 2000-2015 waves of the CEX public-use micro-data adjusted to 2016 dollars. The restrictions to the sampling criteria are standard (see, for example, Cagetti, 2003). The beta correction profile, $\{\hat{\beta}_t\}_{t=0}^T$, is obtained from Carroll (2012). These procedures are detailed in the accompanying Appendix.

I define total income as total income less taxes, pension contributions, education and health expenses, and asset income. The estimation is performed separately for working and retired adults. The former is restricted to individuals ages 24 to 64 who are married, have completed high school, report working at least 35 hours each week, and whose total earnings for the year exceed \$500. Retired individuals are allowed to be single and are restricted to being over 60 years of age. These restrictions yield samples of 6322 and 3328 working age and retired consumers, respectively. The deterministic profile, $\{\Gamma_t\}_{t=0}^T$, is obtained from the fitted values of a regression of log total income on a fourth degree polynomial in age, controlling for cohort effects, calculated separately for retired and working individuals. In order to obtain a monthly series I divide each annual series across twelve months and use a simple moving average to smooth the income path from year to year.

In order to estimate high-frequency movements in income I utilize the bank transaction data. I restrict to those households that are observed for at least 60 consecutive months, do not experience an observable unemployment spell (as measured by UI income), and whose primary account holder is male. To control for extreme outliers I trim households experiencing outcomes in the 1% tails of total income months and the 5% tails of monthly income growth. To control for the deterministic age component I obtain residuals from a first stage regression of log income on age, and then estimate the an $AR(1)$ process, $y_{i,a} = \rho y_{i,a-1} + \epsilon_{i,a}$. The results are reported in Table 7.

To focus the analysis to liquid asset accumulation, while retaining the life-cycle profile of earnings, I abstract from the drop in income at retirement. This modeling choice has a natural interpretation - a proportion of household income is placed into an illiquid savings account during each period of working life. This account subsequently pays out with certainty in each month of retirement. The beta correction profile, $\{\hat{\beta}_t\}_{t=0}^{T}$, is obtained from Carroll (2012).

Figure A.36 plots the resulting profiles of $\{\Gamma_t\}_{t=0}^{T}$ and from the first-stage estimation procedure as well as the the beta correction, $\{\hat{\beta}_t\}_{t=0}^{T}$, obtained from Carroll (2012).

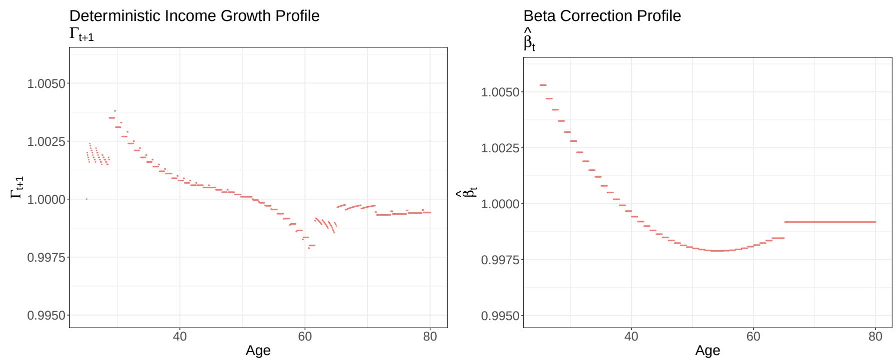

Figure A.20: Deterministic Profiles

Additionally, I obtain monthly job finding and separation rates from the CPS Labor Force Statistics data for the year 2016. I obtain unemployment insurance replacement rate data from the Office of Unemployment Insurance UI Replacement Rates Report for 2016.

To structurally estimate the model's four preference parameters $\{\beta, \gamma, \psi, \kappa\}$ I utilize a simulated method of moments approach. I estimate the model's preference parameters to fit both life-cycle liquid asset accumulation (the median liquid assets of 8 age groups) and the cross-section of consumption responses observed in the transaction data (the average thirty day consumption response for each decile of $\frac{LiquidAssets}{Income}$). In order to obtain estimates of liquid assets I utilize data from the Survey of Consumer Finances between 2001 to 2016. As in the other samples, I restrict to households of working age (24 to 64) who are married.

Appendix B. Parameter Identification

Estimate model to the objective satisfy

$$\min_{\beta,\gamma,0,\kappa} \Theta \sum_{i}^{N} \omega_i^a \mid d_{i,a}^{liq} - m_a^{liq}(\beta,\gamma,0,\kappa) \mid +(1-\Theta) \sum_{j}^{10} \mid d_j^{mpc} - m_j^{mpc}(\beta,\gamma,0,\kappa) \mid$$

Arguments for parameter identification are detailed in Section 3.2 of the main text. To further assess parameter identification, I report contours for each set of moments, plotting the median absolute distance against key model preference parameters (Figure B.37). For each plot I fix the absent parameters to a small window around their estimated values. Darker blue shades indicate smaller distances between the model generate moments and their empirical counterparts. Consistent with the previous literature, I find weak identification for β and γ in determination of life-cycle liquid asset accumulation.

Intuitively, increasing impatience can be traded off with an increased level of risk aversion to obtain similar liquid asset profiles. The consumption response contours indicate a similar trade-off between increasing patience (β) and increasing dissaving aversion (ψ). Jointly, the plots indicate that for a fixed level of γ, obtaining consumption responses closer to the data requires ascending the steep gradient away from the liquid asset minima.

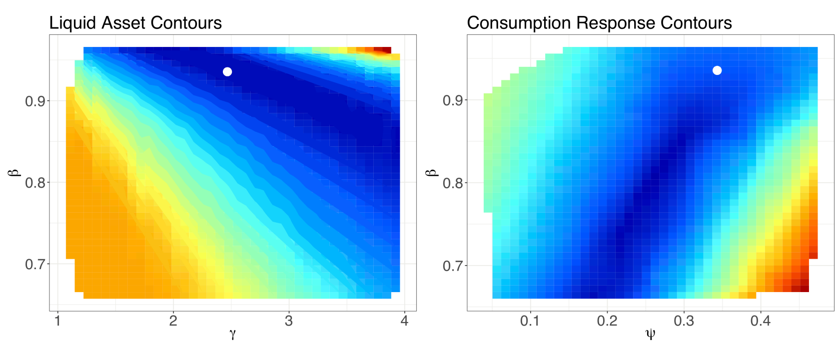

Figure B.21: Preference Parameter Contour Plots

Figure B.38 indicates the fundamental trade-off in the standard model between matching liquid asset balances and consumption responses.

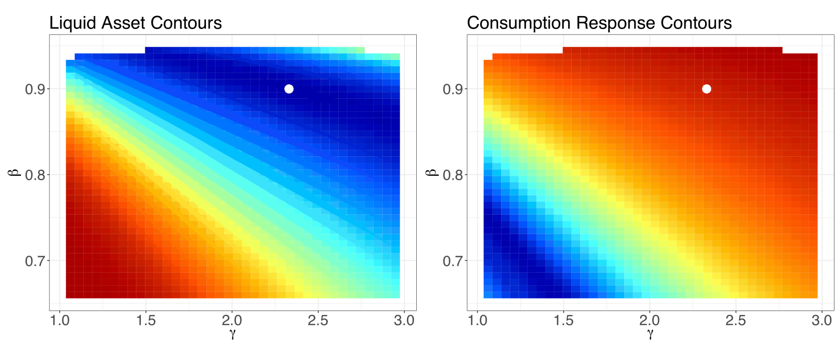

Figure B.22: Parameter Identification Contours, Standard Buffer-Stock Model