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

Robert A. McDowall[1]

New York University

Data Appendix

This document contains extended appendices for the paper.

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 and attribute these outflows to debt repayment instead, or restrict 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* =

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

[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

- ☐ *Groceries* ($e_{\mathcal{ND}_\mathcal{F}}$): Grocery, drug, and liquor stores (ie. food at home).

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

  - ☐ *Food Services* ($e_{\mathcal{S}_\mathcal{F}}$): Restaurants and bars (ie. food away from home).

+ *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 in the main text. 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	Expenditure	Non-Durables	Durables	Services	Food Services	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 Revenue 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.

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 trans-

[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.

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

action 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 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

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

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.

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

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 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.

[9]Source: Turbotax

Appendix C. Identification

Identification relies on variation in the timing of inflows with respect to calendar time, each other, and to the date of filing in the case of tax refunds and payments. This section outlines these arguments.

Appendix C.1. One Event

In the case of one event identification relies on calendar variation in the timing of individuals' shocks. Consider specification

$$y_{i,t} = \alpha_i + \lambda_t + \delta^1_{t+s} I^1_{i,t+s} + \delta^1_{t+s-1} I^1_{i,t+s-1} + \epsilon_{i,t}$$

Taking expectations, assuming $\mathbb{E}[\epsilon_{i,t} - \epsilon_{i,t-1} | I_{i,t+s}, I_{i,t+s-1}] = 0$, and differencing across time t and individuals i and j

$$\mathbb{E}[\Delta y_{it} - \Delta y_{jt}] = \mathbb{E}[\delta^1_{t+s}(I^1_{i,t+s} - I^1_{j,t+s}) + \delta^1_{t+s-2}(I^1_{j,t+s-2} - I^1_{i,t+s-2})]$$

Variation in the calendar timing of treatment (simplest case in which individual i receives $I^1 > 0$ at time $t+s$ and individual j receives $I^1 > 0$ at time $t+s+k$, $k \geq 1$) yields

$$\mathbb{E}[\Delta y_{it} - \Delta y_{jt}] = \mathbb{E}[\delta^1_t I^1_{i,t}]$$

Which is the parameter of interest. Identification of the anticipatory response follows in a similar manner. Below I plot the observed calendar variation in dates of filing and tax refund receipt.

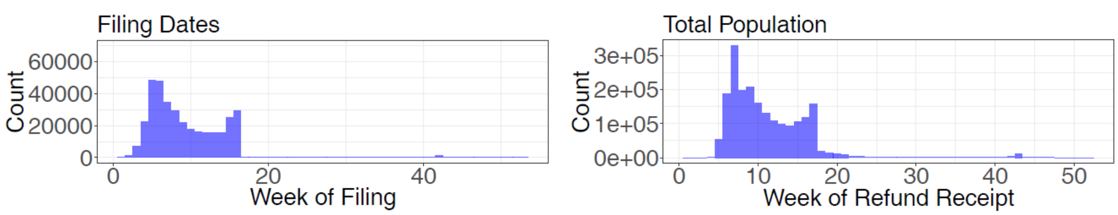

Figure C.1: Timing of Tax Refunds

Appendix C.2. Multiple Events

For the case of multiple events, consider specification

$$y_{i,t} = \alpha_i + \lambda_t + \delta^1_t I^1_{i,t} + \delta^1_{t-1} I^1_{i,t-1} + \delta^2_{t+s} I^2_{i,t+s} + \delta^2_{t+s-1} I^2_{i,t+s} + \epsilon_{i,t}$$

Assuming $\mathbb{E}[\epsilon_{i,t} - \epsilon_{i,t-1} | I^1_{i,t}, I^1_{i,t-1}, I^2_{i,t+s}, R^2_{i,t+s-1}] = 0$, and differencing as above

$$\begin{aligned}\mathbb{E}[\Delta y_{i,t} - \Delta y_{j,t}] = \mathbb{E}[&\delta^1_t(I^1_{i,t} - I^1_{j,t}) + \delta^1_{t-2}(I^1_{j,t-2} - I^1_{i,t-2}) \\ &+\delta^2_{t+s}(I^2_{i,t+s} - I^2_{j,t+s}) + \delta^2_{t+s-2}(I^2_{j,t+s-2} - I^2_{i,t+s-2})]\end{aligned}$$

Variation in the timing of treatment (simplest case in which individual i receives $I^1 > 0$ at time t and individual j receives $I^1 > 0$ at time $t + k$ $(1 \leq k < s)$, individual i receives $I^2 > 0$ at time $t + s$ and individual j receives $I^2 > 0$ at time $t + s + k$))

$$\mathbb{E}[\Delta y_{i,t} - \Delta y_{j,t}] = \mathbb{E}[\delta^1_t(I^1_{i,t}) + \delta^2_{t+s}(I^2_{i,t+s})]$$

In addition to variation in calendar time, as in the one shock case, identification here further relies on variation in timing between the two shocks, I^1 and I^2 (simplest case, individual h that receives $I^1 > 0$ at time t and $I^2 > 0$ at time $t + s + l$, $l \neq k$). Differencing again yields

$$\mathbb{E}[(\Delta y_{i,t} - \Delta y_{j,t}) - (\Delta y_{i,t} - \Delta y_{h,t})] = \mathbb{E}[\delta^1_t I^1_{h,t}]$$

Which is the parameter of interest. Below I plot the observed variation in the time between filing and refund receipt as well as the days between receipt for those receiving federal and state refunds.

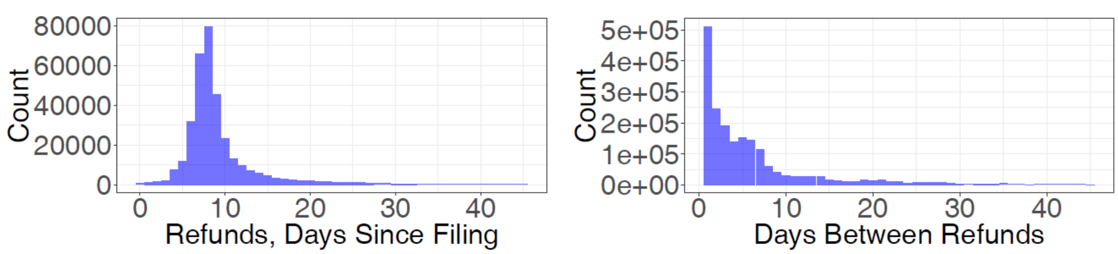

Figure C.2: Time Between Events

Appendix D. Obtaining Cumulative Coefficients

In order to interpret the coefficients as the cumulative change in $y_{i,t}$ as a proportion of the inflow, I^1, I normalize the regression. Consider the simplest case

$$\begin{aligned}y_{i,t} =& \alpha_i + \lambda_t + \delta_0 I_{i,0} + \delta_1 I_{i,1} + \epsilon_{i,t} \\ =& \alpha_i + \lambda_t + \delta_0 I_{i,0} + \delta_1 I_{i,1} + \delta_0 I_{i,1} - \delta_0 I_{i,1} + \epsilon_{i,t} \\ =& \alpha_i + \lambda_t + \delta_0 (I_{i,0} - I_{i,1}) + (\delta_1 + \delta_0) I_{i,1} + \epsilon_{i,t} \\ =& \alpha_i + \lambda_t + \beta_0 \Delta I_{i,0} + \beta_1 I_{i,1} + \epsilon_{i,t}\end{aligned}$$

In a more general case such as

$$y_{i,t} = \alpha_i + \lambda_t + \sum_{j=t-l}^{t+L} \delta_j^1 I_{i,j}^1 + \epsilon_{i,t}$$

this procedure proceeds recursively

$$y_{i,t} = \alpha_i + \lambda_t + \beta_{t-l}^1 \Delta I_{i,t-l}^1 + ... + \beta_{t+L-1}^1 \Delta I_{i,t+L-1}^1 + \beta_{t+L}^1 I_{i,t+L}^1 + \epsilon_{i,t}$$

Where $\beta_{t-l}^1 = \delta_{t-l}^1$, $\beta_{t-l+1}^1 = \delta_{t-l}^1 + \delta_{t-l+1}^1$ ... $\beta_{t+L}^1 = \sum_{j=t-l}^{t+L} \delta_j^1$. A large number of leads and lags can be produced efficiently in `R` via the following code, with 100 lags and leads chosen as an example

```
%Create last lag, t+L = 100
  data= data  %>% group_by(household)
              %>% mutate(inflow_lag_100 = lag(inflow, 100))
              %>% ungroup()

%Make dataframe compatible with data tables
  setDT(data, keep.rownames=TRUE, key=NULL, check.names=FALSE)

%Create rest of the lag columns, difference to create cumulative
  data = data[, sprintf("inflow_lag_%d", 0:99) :=
                mapply( `-`,
                    shift(inflow, 0:99,  fill = NA, type = 'lag'),
                    shift(inflow, 1:100, fill = NA, type = 'lag'), %0,
                    SIMPLIFY = FALSE),
                    by = household]

%Create lead columns
  data = data[, sprintf("inflow_lead_%d", 1:100) :=
              mapply( `-`,
                    shift(inflow, 1:100, fill = NA, type = 'lead'),
                    shift(inflow, 0:99,  fill = NA, type = 'lead'), %0,
                    SIMPLIFY = FALSE),
                    by = household]
```

For non-cumulative coefficients, the respective lines are replaced by the commented ‘`%0,`’.

Appendix E. Regressions on Large Panel Data

This section outlines how the above econometric methodology is executed in practice for very large panel data. The follow sections outline the implementation of fixed effects, and the execution of a regression on the resulting dataframe. This approach utilizes the `biglm` package, which I also provide an overview of below.

This approach allows for panel regressions to be run on very large data that cannot be held in memory. As such, household fixed effects are obtained first and regression coefficients are obtained by iterating through chunks of the dataset and updating the least squares estimator components. I describe each of these steps in turn.

Appendix E.1. Create Dummy Variables for Fixed Effects

Because household fixed effects will be implemented before the regression is performed, it is useful to expand all time fixed effects columns and create individual dummy variables. Consider one fixed effect vector with three ($L = 3$) levels

```
%Convert time FE to factors
data$time_FE_var = as.factor(data$time_FE_var)

%Create an index column
data$indexer = 1:nrow(data)

#Expand matrix to replace each L level factor
#        with L-1 dummies
data = as.data.frame(indexer ~
                        household + y_var + inflow + time_FE_var, data)
```

The above will create a `data.frame` replacing the `time_FE_var` with $L-1$ dummies named `time_FE_var2` and `time_FE_var3` by default (the default ordering of factors is alphabetical, and this operation drops the first level).

Appendix E.2. Implementing Fixed Effects

In order to efficiently implement time (λ_t) and household fixed effects (α_i), demeaning of both the LHS and RHS variables is performed prior to the regression for the large number of household fixed effects. Consider the specification

$$c_{i,t} = \alpha_i + \lambda_t + \delta I_{i,t} + \epsilon_{i,t} \tag{E.1}$$

Taking within-household mean across time

$$\bar{c}_i = \alpha_i + \bar{\lambda} + \delta \bar{I}_i \tag{E.2}$$

Subtracting (E.2) from (E.1) extracts household fixed effects from the underlying data

$$c_{i,t} - \bar{c}_i = (\lambda_t - \bar{\lambda}) + \delta(I_{i,t} - \bar{I}_i) + \epsilon_{i,t} \tag{E.3}$$

The regression can now be run on the residualized values to obtain the coefficient of interest δ. Define $\tilde{x}_{i,t} \equiv x_{i,t} - \bar{x}_i$ and rename $\tilde{\lambda}_t \equiv \lambda_t - \bar{\lambda}$ so that $(E.3)$ collapses to

$$\tilde{c}_{i,t} = \tilde{\lambda}_t + \delta \tilde{I}_{i,t} + \tilde{\epsilon}_{i,t} \tag{E.4}$$

It is computationally efficient to perform this step prior to running the regression. Due to the number of individual households it is likely more efficient to read each individual household into memory before obtaining $\bar{c}_i$, $\bar{\lambda}$, and $\bar{I}_i$. This can be performed efficiently in `R` via `dplyr`:

```
%Residualize household FE
%Disclude household id column
%            and any non-demeaned variables in 'data' from this operation
  data= data  %>% group_by(household)
              %>% mutate_each( funs(. - mean(., na.rm = TRUE),
                                        -any_non_demeaned_vars )
              %>% ungroup()
```

Appendix E.3. Running the Regression

For very large panel datasets it is often still infeasible to read $\tilde{c}_{i,t}$, $\tilde{\lambda}_t$, and $\tilde{I}_{i,t}$ into memory order to perform regression specification (7). The solution is to read this data into memory in manageable chunks and update the coefficient matrix for each chunk. The `biglm` package supports this operation.

First, the large panel data must be split into chunks. As an example consider one hundred such chunks, where a set of households is assigned to a particular `chunk`, $i \in \{0, 1, ..., 99\}$. Each chunk is read into memory. On the first pass the `biglm` regression function is called, and on each subsequent pass this object is updated via the `update` function. This is performed as follows:

```
%Iterate through all chunks of data
      for( i == 0:99 ){

      %Read next chunk of data here, example:
              DataChunk = read.csv(paste0("data_chunk_",i,".csv"))
```

```
        %Perform additional operations here, example:
                %Create time FE dummies described in Appendix E.1

                %Lead and lag columns described in Appendix D

                %Household fixed effects described in Appendix E.2


  %If this is the first chunk (ie. chunk 0) then run regression
        if( i == 0 ){
            regression_obj <- biglm(formula = y_var ~ 0 +
                                      inflow_lag_2  +
                                      inflow_lag_1  +
                                      inflow_lag_0  +
                                      inflow_lead_1 +
                                      inflow_lead_2 +
                                      time_FE_var2  +
                                      time_FE_var3  , data = DataChunk)

  %Else, update regression object with next chunk of data
        }else{
            update(regression_obj, moredata = DataChunk)
        } %if
        } %for
```

Once the for loop completes the final regression output can be saved via the `saveRDS` command.

Appendix E.4. Adjusting Standard Errors

The `biglm` regression object returns unadjusted standard errors. Since household fixed effects are performed prior to the regression, it is important to adjust these error bounds for the degrees of freedom n_{α_i}. Consider a data set of size n. Here we consider household fixed effects, where the number of household intercept terms is denoted n_{α_i} where n_{α_i} may be very large. Consider the specification

$$c_{i,t} = \alpha_i + \delta I_{i,t} + \epsilon_{i,t} \tag{E.5}$$

There is one slope term plus n_{α_i} intercept terms so the standard error of the parameter estimate is

$$s_{\hat{\delta}} = \sqrt{\frac{1}{n - n_{\alpha_i} - 1} \frac{\sum_{i,t} \epsilon_{i,t}^2}{\sum_{i,t}(I_{i,t} - \bar{I}_{i,t})^2}}$$

Fixed effects are implemented before running the regression. Taking the within-household mean across time

$$\bar{c}_i = \alpha_i + \delta \bar{I}_i \tag{E.6}$$

Subtracting (E.6) from (E.5) extracts household fixed effects from the underlying data

$$c_{i,t} - \bar{c}_i = \delta(I_{i,t} - \bar{I}_i) + \epsilon_{i,t} \tag{E.7}$$

Note the residual term has not changed. Define $\tilde{x}_{i,t} \equiv x_{i,t} - \bar{x}_t$ so that (8) collapses to

$$\tilde{c}_{i,t} = \delta \tilde{I}_{i,t} + \epsilon_{i,t} \tag{E.8}$$

Note the centered values $\tilde{x}_{i,t}$ are all mean zero. When this residualized data is processed by a regression package such as `biglm` the n_{α_i} household-dependent intercept terms are not accounted for. The standard error of the estimate output by the `biglm` regression package is

$$s_{\hat{\delta}}^{biglm} = \sqrt{\frac{1}{n-1} \frac{\sum_{i,t} \epsilon_{i,t}^2}{\sum_{i,t} (\tilde{I}_{i,t} - \bar{\tilde{I}}_{i,t})^2}}$$

Where $n-1$ denotes the sample size minus number of regressors (δ). Because $\tilde{I}_{i,t}$ is mean zero then

$$\begin{aligned} s_{\hat{\delta}}^{biglm} &= \sqrt{\frac{1}{n-1} \frac{\sum_{i,t} \epsilon_{i,t}^2}{\sum_{i,t} (\tilde{I}_{i,t})^2}} \\ &= \sqrt{\frac{1}{n-1} \frac{\sum_{i,t} \epsilon_{i,t}^2}{\sum_{i,t} (I_{i,t} - \bar{I}_{i,t})^2}} \end{aligned}$$

and so the only difference between $s_{\hat{\delta}}^{biglm}$ and $s_{\hat{\delta}}$ is that the former does not correctly account for the degrees of freedom soaked up by the n_{α_i} household fixed effects. To correct $s_{\hat{\delta}}^{biglm}$ perform the following:

$$s_{\hat{\delta}} = \sqrt{(s_{\hat{\delta}}^{biglm})^2 \cdot \frac{n-1}{n - n_{\alpha_i} - 1}}$$

Appendix E.5. Under the hood of `biglm`

It is important to understand what `biglm` is doing as it updates the regression object on each loop. Here I present an overview due to Portugués (2019). Decompose the least squares estimator $\hat{\beta} = (\mathbf{X'X})^{-1}\mathbf{X'Y}$, and the estimator sans the $i - th$ observation $\hat{\beta}_{-i} = (\mathbf{X}'_{-i}\mathbf{X}_{-i})^{-1}\mathbf{X}'_{-i}\mathbf{Y}_{-i}$. The estimator is updated via

$$\hat{\beta} = \hat{\beta}_{-i} + (\mathbf{X'X})^{-1}\mathbf{x}_i(Y_i - \mathbf{x}'_i\hat{\beta}_{-i}) \tag{E.9}$$

Note the only great computational difficulty, once $\hat{\beta}_{-i}$ is obtained in chunk zero, is in the inversion of $\mathbf{X'X}$, which employs the entire dataset. The trick to `biglm`'s efficiency is updating from $(\mathbf{X}'_{-i}\mathbf{X}_{-i})^{-1}$ to $(\mathbf{X}'\mathbf{X})^{-1}$ on each iteration without having to re-invert this object. Making use of the following equalities

$$\begin{aligned}
\mathbf{X'X} =& \mathbf{X}'_{-i}\mathbf{X}_{-i} + x_i x'_i \\
\mathbf{X'Y} =& \mathbf{X}'_{-i}\mathbf{Y}_{-i} + x_i Y'_i \\
(\mathbf{A} + \mathbf{bb}')^{-1} =& \mathbf{A}^{-1} - \frac{\mathbf{A}^{-1}\mathbf{bb}'\mathbf{A}^{-1}}{1 + \mathbf{b}'\mathbf{A}^{-1}\mathbf{b}}
\end{aligned}$$

Where the first two equalities are tautological, and the third equality is the Sherman-Morrison formula. From the first and third equations it follows

$$(\mathbf{X'X})^{-1} = (\mathbf{X}'_{-i}\mathbf{X}_{-i} + x_i x'_i)^{-1} \tag{E.10}$$

$$= (\mathbf{X}'_{-i}\mathbf{X}_{-i})^{-1} - \frac{(\mathbf{X}'_{-i}\mathbf{X}_{-i})^{-1}\mathbf{x}_i\mathbf{x}'_i(\mathbf{X}'_{-i}\mathbf{X}_{-i})^{-1}}{1 + \mathbf{x}'_i(\mathbf{X}'_{-i}\mathbf{X}_{-i})^{-1}\mathbf{x}_i} \tag{E.11}$$

and so only the inverted object $(\mathbf{X}'_{-i}\mathbf{X}_{-i})^{-1}$ need be kept in memory. It is straightforward to see how the `biglm` algorithm works. On chunk zero $\hat{\beta}_{chunk0}$ is obtained along with $(\mathbf{X}'_{chunk0}\mathbf{X}_{chunk0})^{-1}$. Adding chunk one's data, $(\mathbf{X}'_{chunk1}\mathbf{X}_{chunk1})^{-1}$ is obtained via (10) and $\hat{\beta}_{chunk1}$ is then obtained via (8). This iteration proceeds until all chunks are processed. See https://bookdown.org/egarpor/PM-UC3M/lm-iii-bigdata.html for a more thorough overview.

Appendix F. Tax Refund Season, Aggregates

This section documents estimates of the aggregate impact of tax refunds on quarterly consumption. Figure B.6 plots the total amount of refunds processes by week from 2014 to 2016. Weekly variation in 2016 relative to prior years caused by tax day falling on Monday, April 18th of that year (week 16). On average, roughly $225 billion in refunds are processed within the first quarter of each year.

A back of the envelope calculation, utilizing a combination of BEA non-seasonally adjusted aggregate Personal Consumption Expenditures data, the IRS reports of

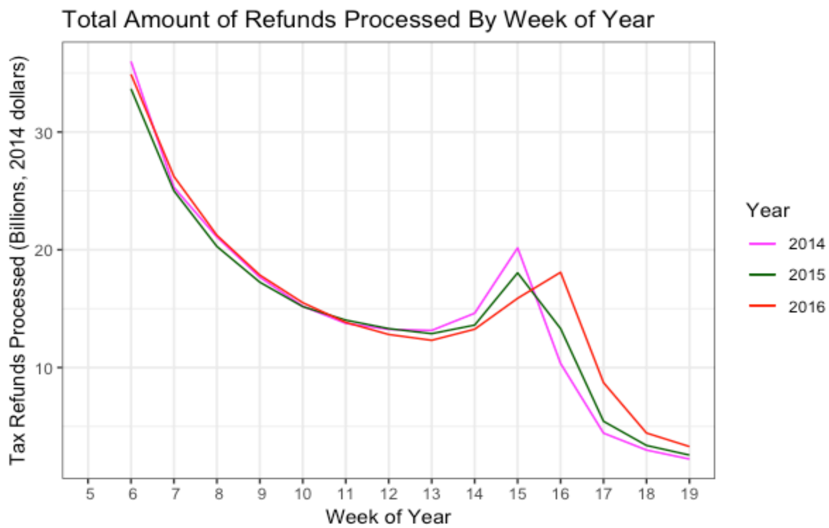

Figure F.3: Refunds By Week, Source: IRS

refunds processed by week, each adjusted for inflation; and the dynamic paths of expenditure estimated in the transaction data, imply consumption out of tax refunds on average accounts for as much as 3.3% of aggregate expenditure in the first quarter of each year. This, compared to an upper bound of 7.4% of aggregate expenditure if the 1 week MPC were equal to one.

Appendix G. Tax Refund Responses

Appendix G.1. Event Identification

Tax refunds and 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

[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 G.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:

I report the cross-sectional non-durable consumption responses obtained before the imputation procedure (Figure F.20).

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

[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.

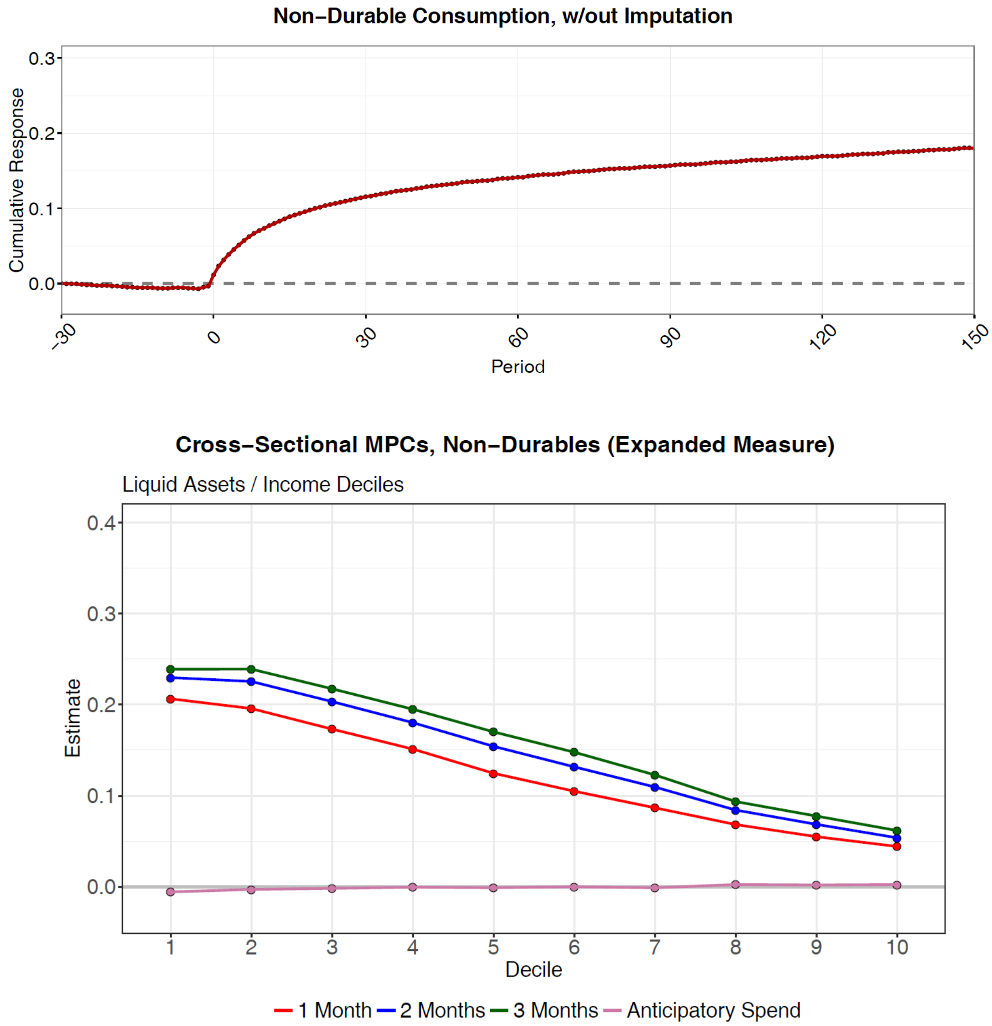

Figure G.4: Non-Durable Responses in the Cross-Section, Expanded Measure Before Imputation

for cross-sectional quantile q at lag j are then obtained via $\delta_{t-l}^{\mathcal{ND}_{\mathcal{I}},q} = \delta_{j}^{\mathcal{ND}_{\mathcal{E}},q} + \xi^{q} \cdot \delta_{j}^{\mathcal{C},q}$. 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 F.21 plots the total expenditure results.

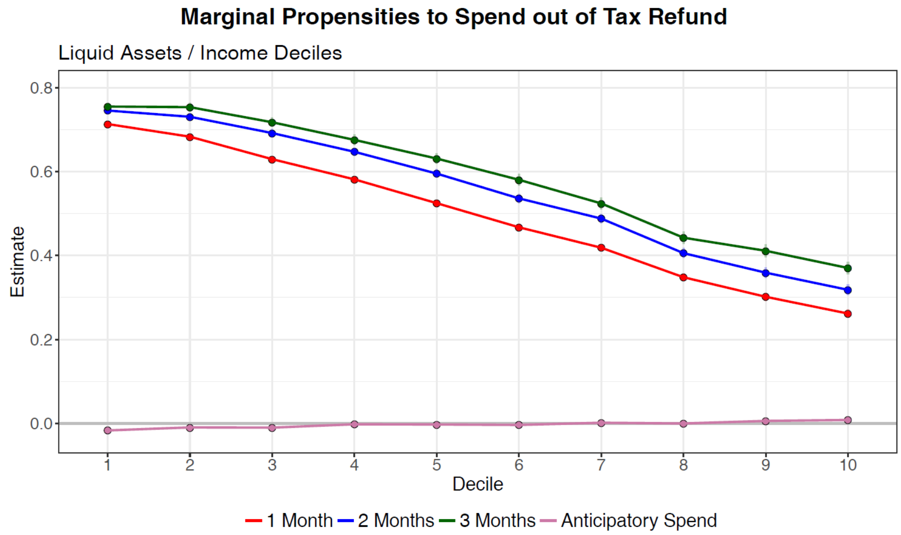

Figure G.5: Total Expenditure Responses in the Cross-Section

Appendix G.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).

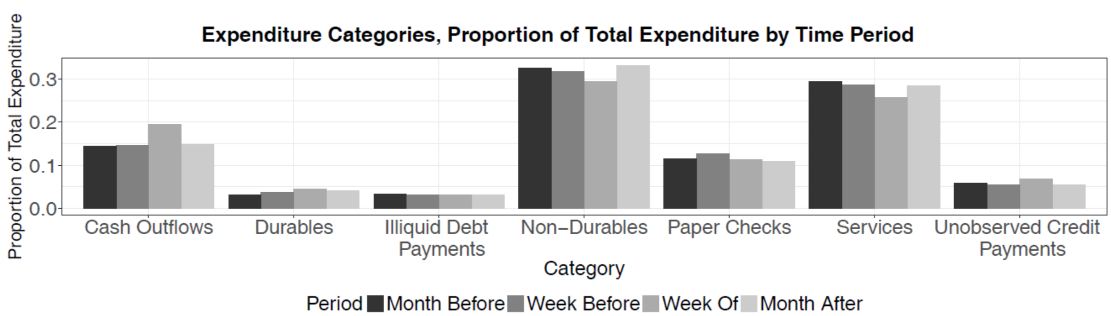

Figure G.6: Weekly Consumption Basket Around Tax Refund

However it is notable that the expenditure response is not simply story of timing large durable purchases to refund receipt. Average expenditure towards 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. Figures 8 and 9 plot proportions of consumption basket around refund receipt and average weekly expenditure by consumption category respectively. This large jump are likely some combination of households supplementing cash balances and financing consumption that is roughly contemporaneous with withdrawal.

Below I report the expenditure composition around refund receipt for the first and fifth quintiles of liquidity:

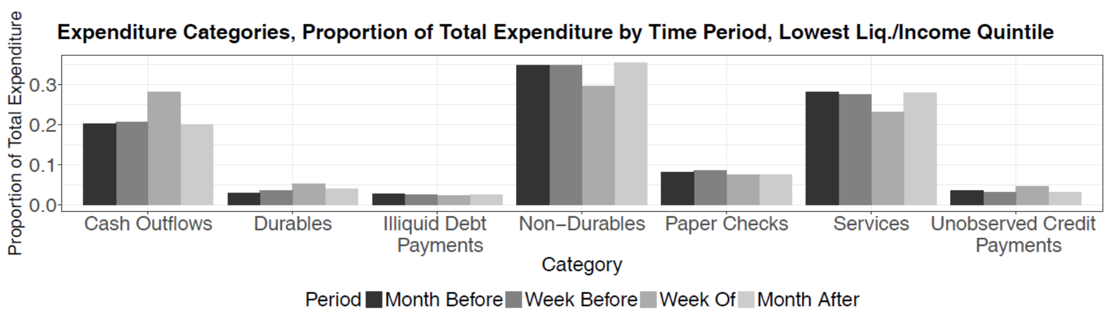

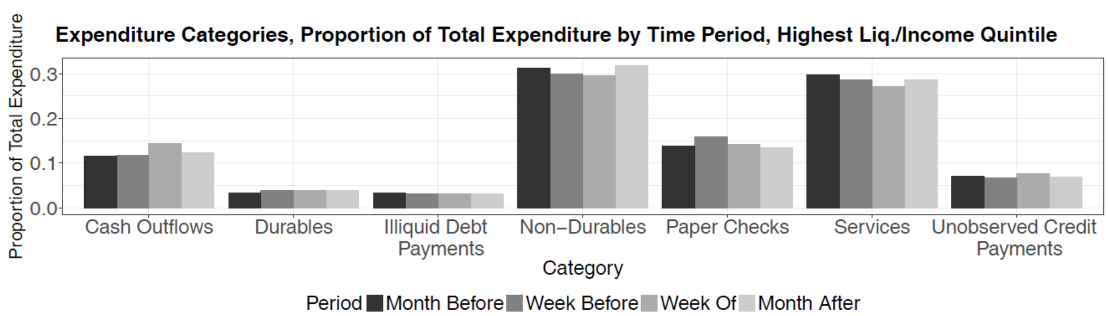

Appendix G.4. Tax Refund Expenditure Responses

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.

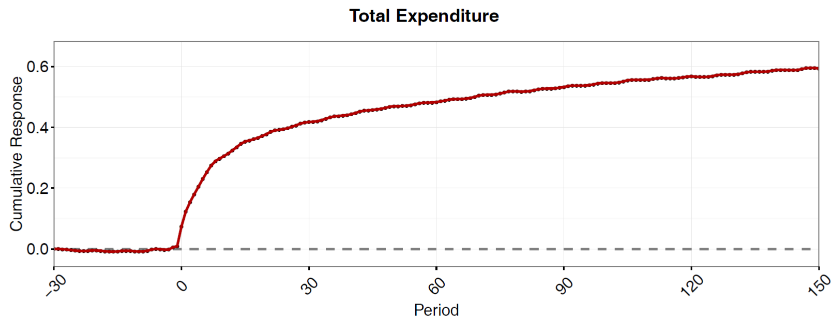

The residual total expenditure not spent towards non-durables or debt payments are classified as durables. I report this category below:

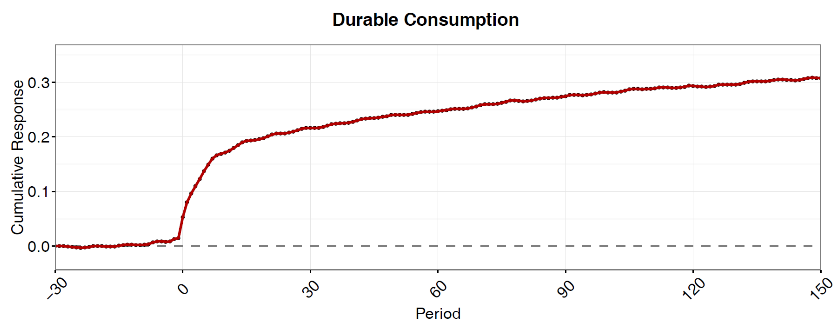

The findings are qualitatively robust to the category of expenditure considered. Crucially, they hold for food services expenditures (restaurants and bars), which can be largely attributed to contemporaneous consumption.

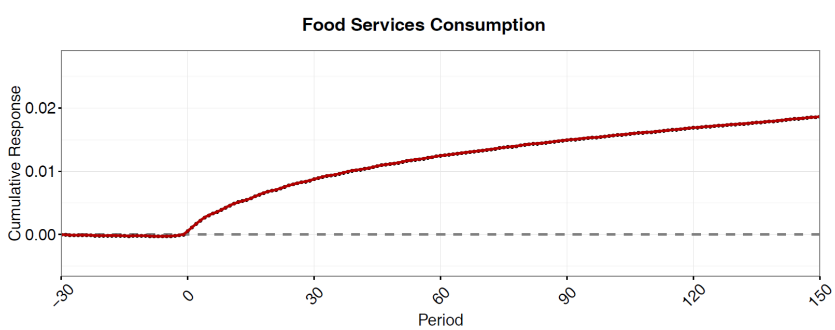

Below I show that credit card balance accumulation in advance of receipt is negligible with a less than one to two cents of every refund dollar spent towards debt repayment.

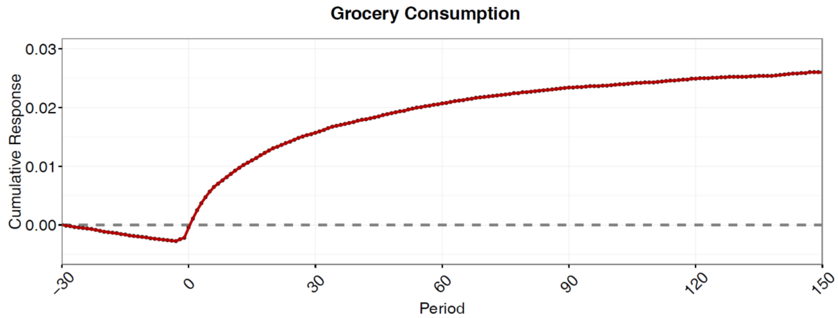

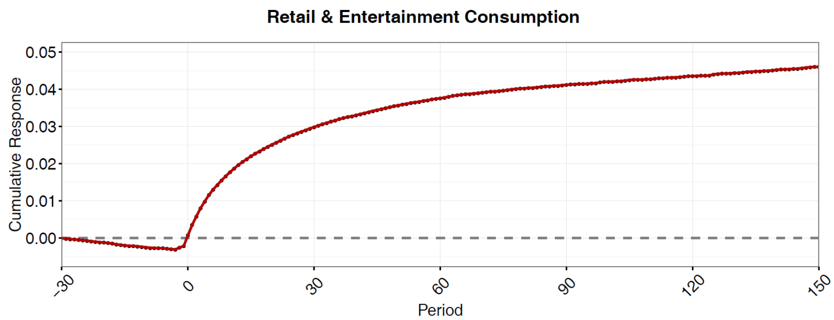

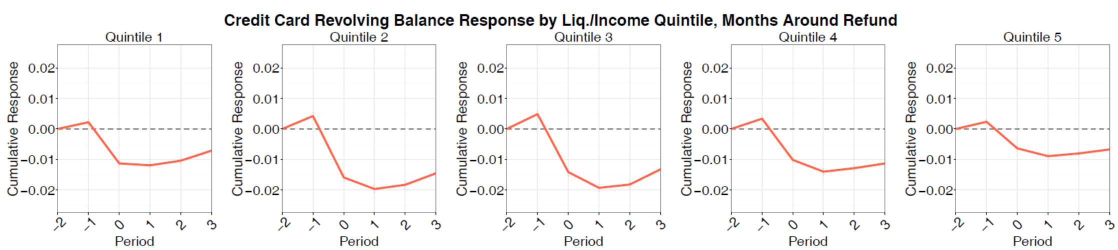

Figure G.7: Credit Card Balances around Refund Receipt

Appendix G.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 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 9 plots the estimated cumulative total expenditure responses proportional to each refund.

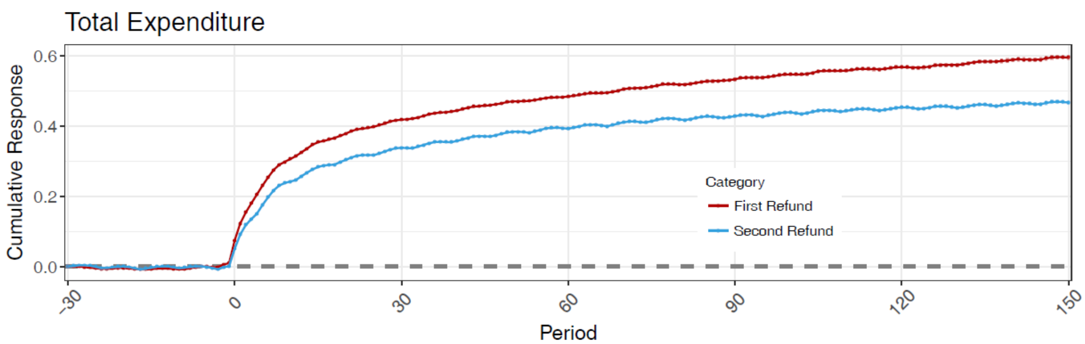

Figure G.8: Tax Refund Responses, Multiple Refunds

With respect to 10 days prior to receipt, the cumulative expenditure out of the second refund is just 0.29 cents of every refund dollar, while the response on the day of receipt is 5.2 cents. There is thus no significant evidence that the first refund receipt is used to smooth through the second - ie. both generate a large degree of excess sensitivity. It is, however, the case that receipt of the first refund mutes the overall expenditure response to the second. Both of these findings are consistent across categories of expenditure. While the first result could not be generated by liquidity constraints alone, the latter is consistent with any model in which the consumption policy function is concave in current assets.

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

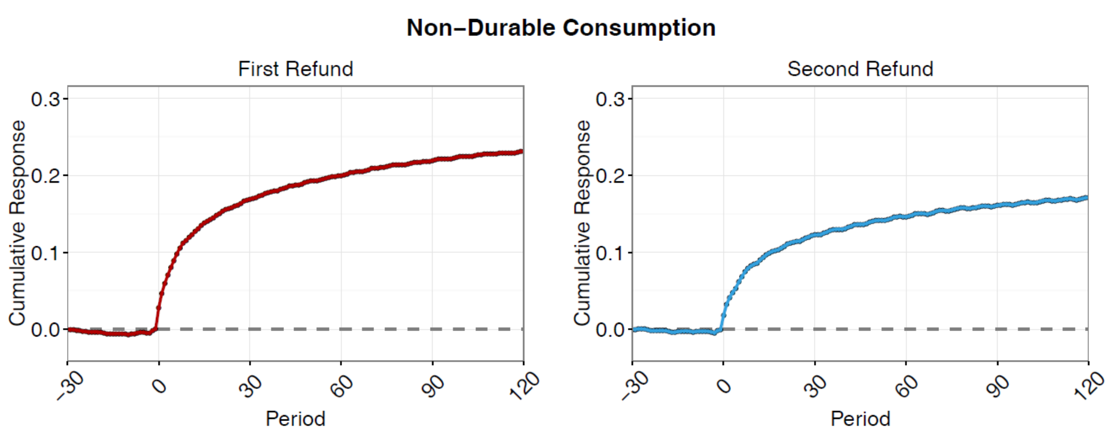

Figure G.9: Tax Refund Non-Durable Consumption Responses, Multiple Refunds

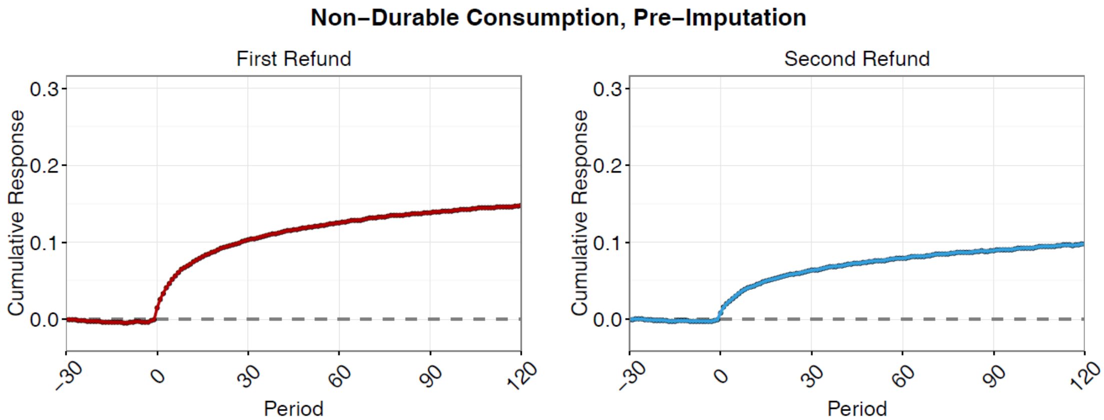

Figure G.10: Tax Refund Non-Durable Consumption Responses, Multiple Refunds

Appendix G.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.

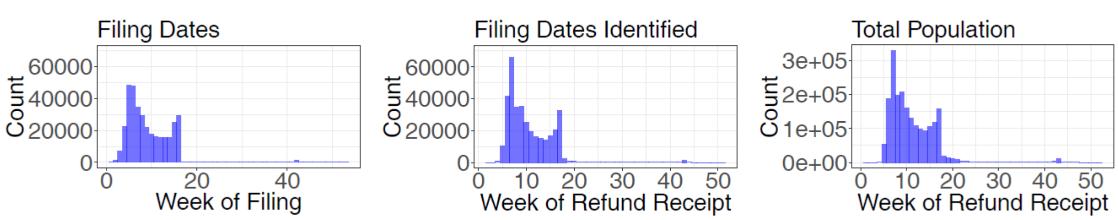

Figure G.11: Filing Date Sub-Population Comparison

I estimate specification (1), where I_i^1 is the first tax refund received in the

Table G.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

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 10 and 11 report the cumulative total expenditure (less payments to tax service providers at the filing date) and non-durable consumption response 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.

These consumption responses are consistent with a model in which households face externally imposed liquidity constraints. However, this result holds for households with substantial liquid wealth (Figure F.29) 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

[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.

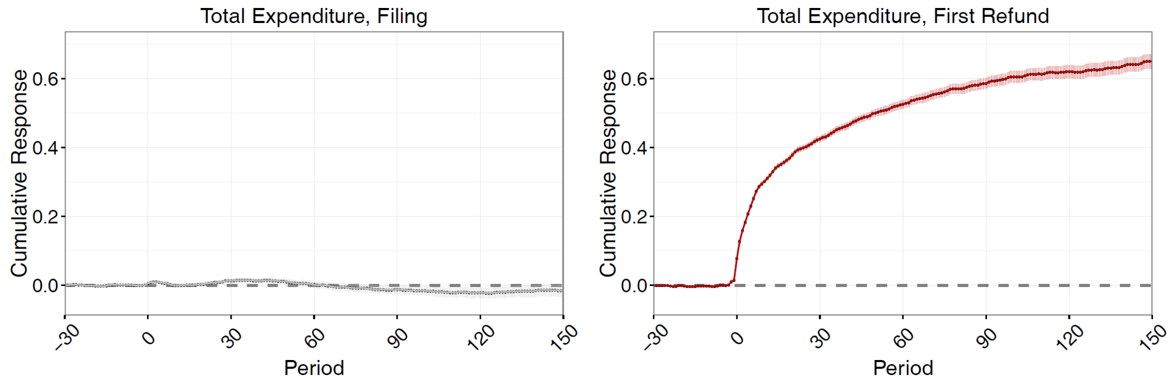

Figure G.12: Expenditure Response at Tax Filing

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.

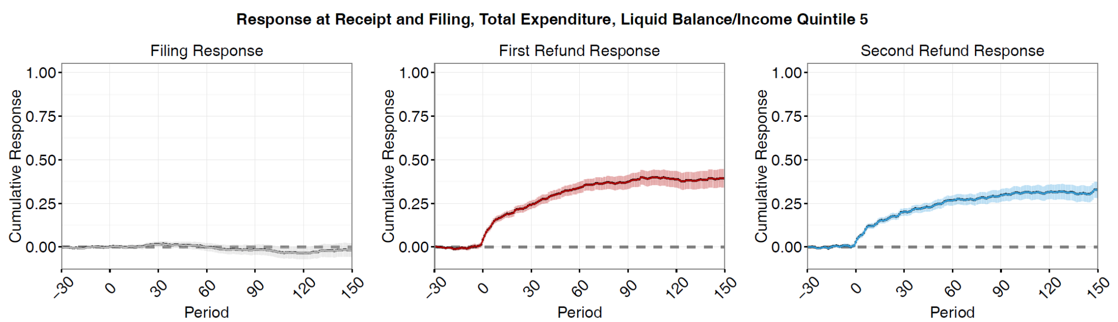

Figure G.13: Expenditure Response, Multiple Refunds and Filing, Highest $\frac{Liquid\ Asset}{Income}$ Quintile

Appendix G.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

[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.

Table G.8: Summary Statistics by Income Level

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

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 18). Responses in levels predictably 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 (Table F.16) 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 - by obtaining residuals from a first-stage regression of log income on a polynomial in age and then estimating an $AR(1)$ process.

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

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

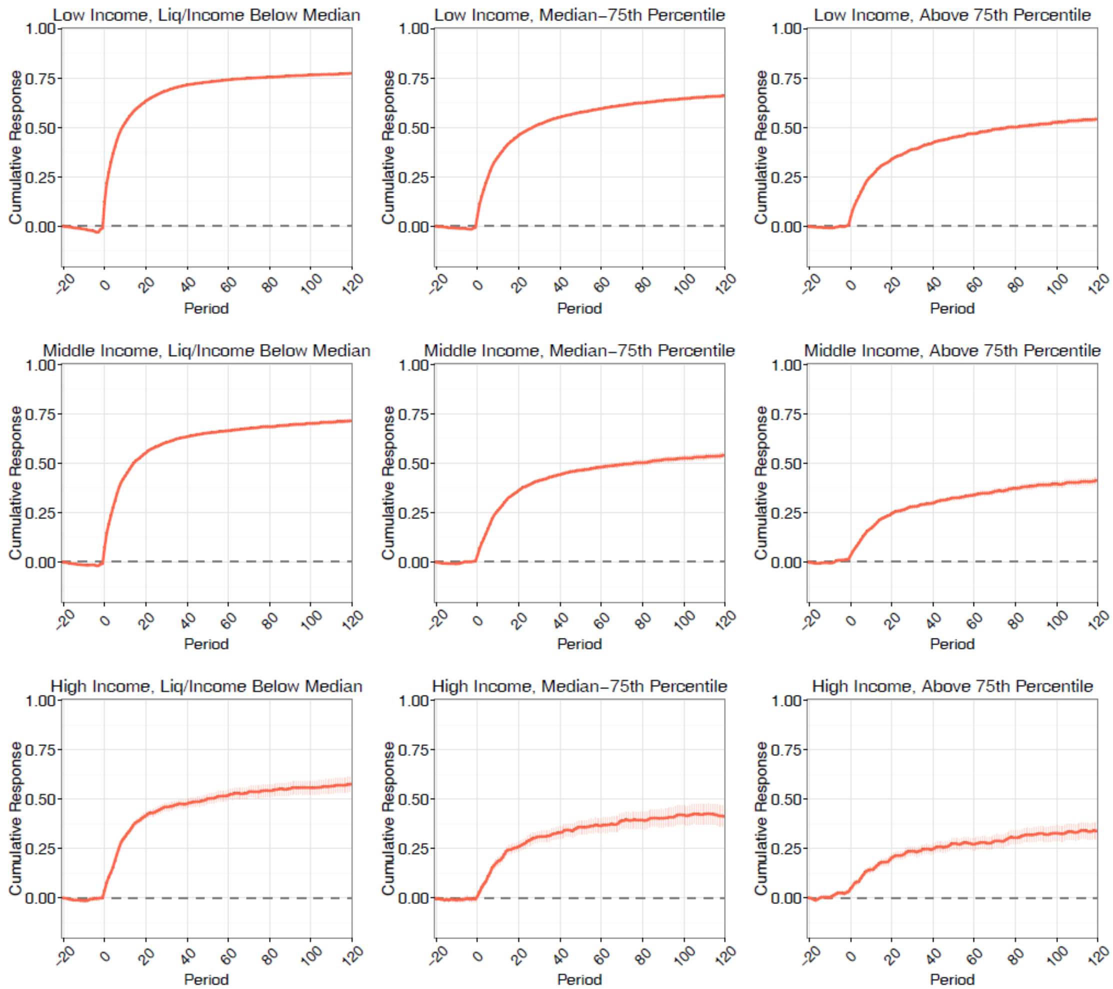

Figure G.14: Expenditure Responses, Income and Asset Distributions

(ie. stock market participation).

Appendix G.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 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

Table G.9: Monthly Income Estimates by Bracket

		Categorized Income		**Total Income**	
	$\frac{Income}{LiquidBalance}$	ρ	σ^2_ϵ	ρ	σ^2_ϵ
Low Income	Low	0.823	0.037	0.783	0.052
($< 40k$)	Middle	0.835	0.035	0.790	0.050
	High	0.857	0.035	0.818	0.048
Middle Income	Low	0.859	0.04	0.790	0.064
($40k - 120k$)	Middle	0.861	0.043	0.791	0.064
	High	0.859	0.043	0.795	0.062
High Income	Low	0.786	0.074	0.672	0.098
($> 120k$)	Middle	0.774	0.079	0.663	0.107
	High	0.810	0.079	0.699	0.133

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 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 G.10: 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.

I begin by estimating specification (1) for the entire non-hand-to-mouth population. These results are reported in the accompanying Appendix and are consistent

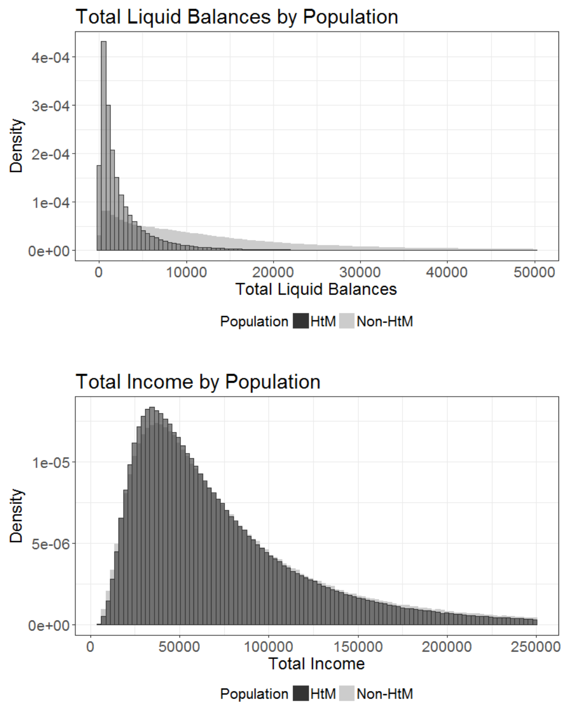

Figure G.15: Hand-to-Mouth and Non-Hand-to-Mouth Comparisons

with those of previous sections. Notably, these non-hand-to-mouth households display average 150 day responses 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. Likewise, for non-durable goods, cumulative responses average of 6.43 cents of every refund dollar, with 56% of the cumulative response occurring within the first 30 days.

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 12 reports these quantiles. Above the 80th percentile, the average refund totls over \$6500 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 later address this 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 test non-durable consumption responses as well.

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

Table G.11: 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

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.

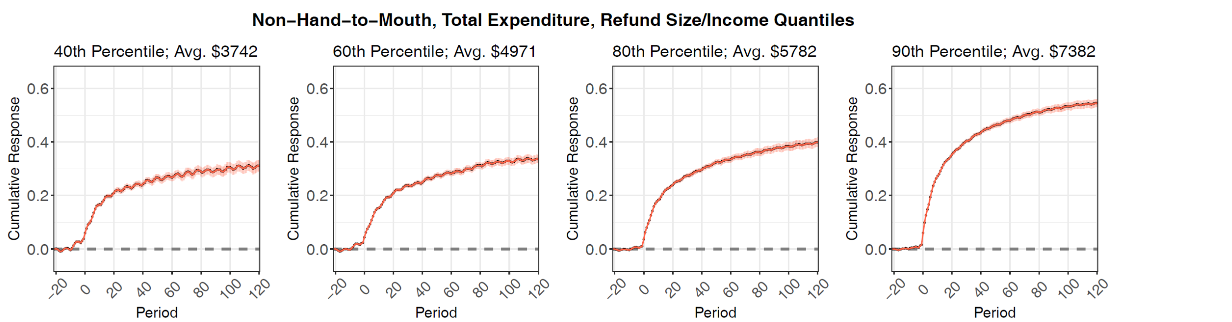

Figure G.16: Expenditure Responses of the Non-Hand-to-Mouth Across Refund Size Relative to Income

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 H. 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

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

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 H.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 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

[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.

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 22) and across quintiles of the liquid balance-to-income ratio (Figure 23). 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 H.12: Cross-Section of Liquid Balance-to-Income, Regular Paychecks (Refund Population)

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

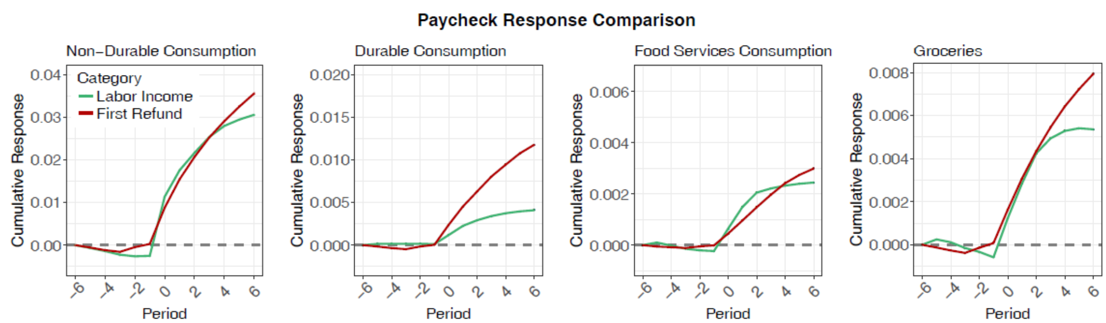

Figure H.17: Consumption Responses Around Regular Paychecks

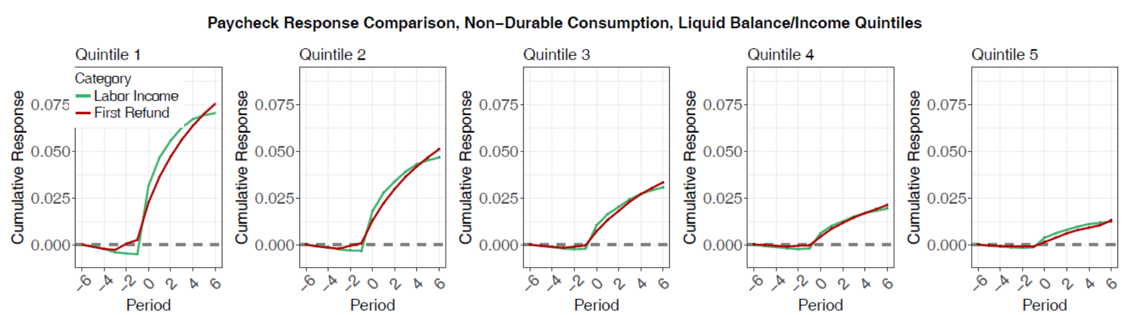

Figure H.18: Non-Durable Goods Responses Around Regular Paychecks, Cross-Section

In dollar terms, the average household amongst the first quintile spends an 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 24[18]).

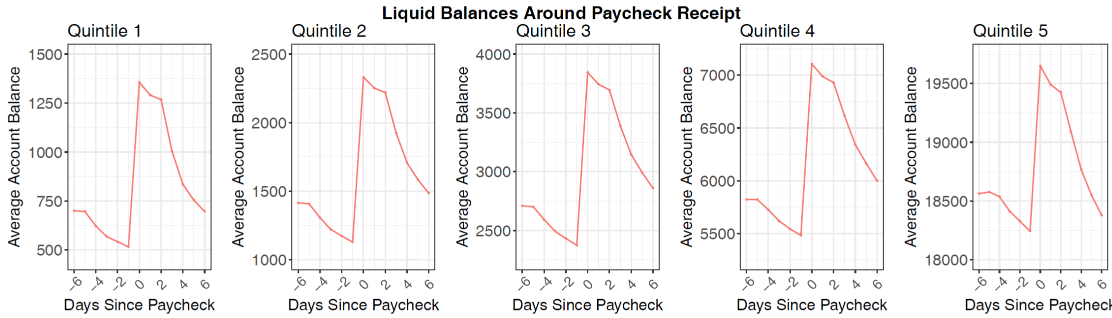

Figure H.19: Path of Liquid Balances Around Regular Paychecks

The consumption responses at payday are somewhat surprising, given the liquid 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 G.29) is further evidence of short-term budgeting via internally imposed constraints on the part of households.

Appendix H.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.

[18]Figure 24 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).

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

Figure G.30 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).

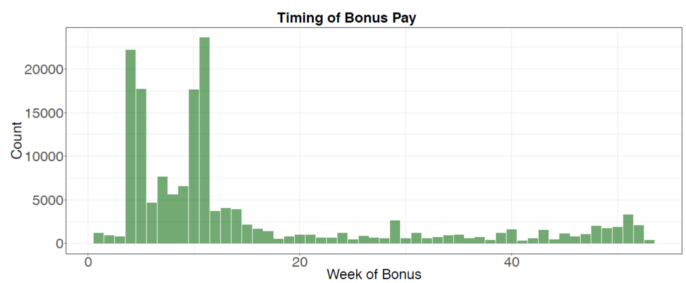

Figure H.20: Timing of Bonus Paychecks

Bonus checks are significant, 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 H.13: 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 in-

[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.

cludes insignificant anticipatory spend and a large degree of excess sensitivity at receipt, with a majority of the response coming in the first thirty days.

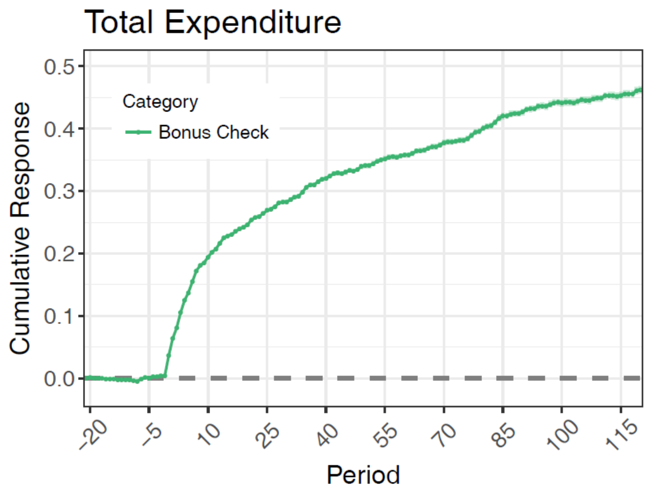

Figure H.21: Total Expenditure Response to Bonus Checks

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.

Table H.14: 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

I estimate the total non-durable responses to bonus receipt by quintile and report these a similar manner to the refund population (Figure G.32). The tight correlation between $\frac{Liquid\ Asset}{Income}$ and consumption responses, along with significant responses amongst the highly liquid, is evident here.

Appendix I. Further Cross-Sections and Variance Decomposition

In this section I report consumption responses to tax refunds across deciles of observables at the monthly frequency and perform a variance decomposition to identify which correlates account for variation in marginal propensities to consume. Table H.22 reports averages across observables by decile. I estimate

$$y_{i,t} = \alpha_i + \lambda_t + \sum_{j=t-l}^{t+L} \delta_j I_{i,j} + \epsilon_{i,t} \tag{I.1}$$

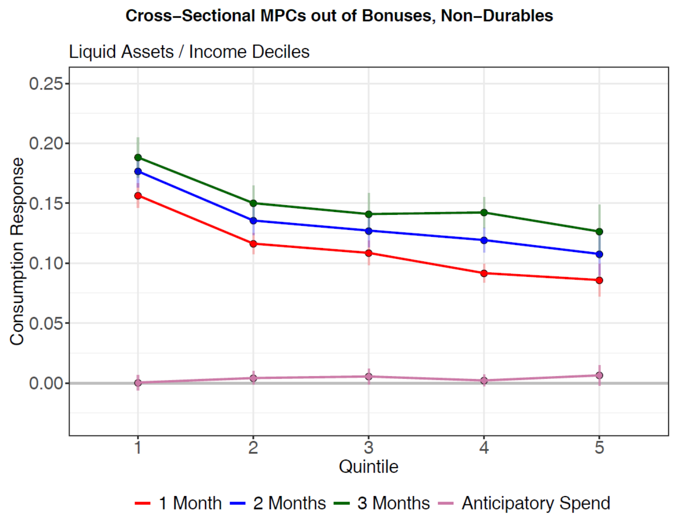

Figure H.22: Non-Durable Response to Bonus Checks

at the monthly frequency, where $y_{i,t}$ denotes non-durable goods or total expenditure responses. Intercepts α_i and λ_t are household and month fixed effects, respectively, and $I_{i,j}$ represents the amount of the tax refund at lead/lag j days received by household i.

I estimate specification $H.1$ for a series of observables by decile. Variables in figure $H.35$ denoted *average* measure means across the nine months prior to refund receipt. Notably, there are right negative correlations between consumption responses and age, income, and liquid balances. Consumption responses are relatively flat across the distribution of CVs of income, and dips only for those households above the 80th percentile in income volatility. Consumption responses are upward sloping in account logins (mobile and online), a proxy for account engagement and attention.

Table I.15: MPC Correlates, Averages within deciles

Decile	1	2	3	4	5	6	7	8	9	10
Age	25	29	33	37	41	45	49	53	57	62
Income	2534	2675	3026	3450	3949	4543	5307	6356	8020	13725
Liquid Assets	133	378	635	968	1434	2134	3261	5264	9579	33176
C.V. of Income	0.08	0.15	0.19	0.22	0.25	0.31	0.39	0.50	0.70	1.25
Daily Logins	0.07	0.23	0.42	0.60	0.77	0.95	1.17	1.45	1.90	3.35

I perform a variance decomposition across observables in order to ascertain which correlates account for explainable variation in consumption responses. I

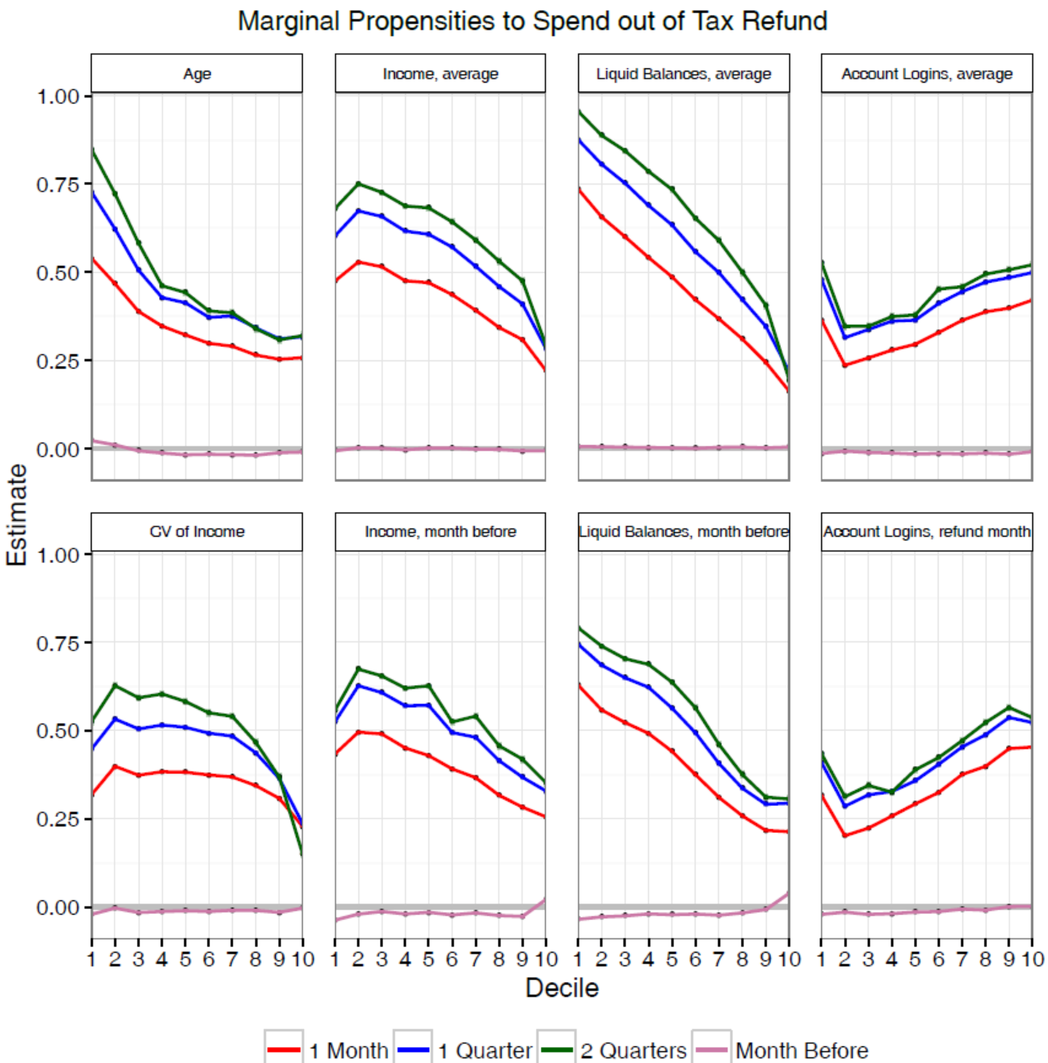

Figure I.23: Observable MPC Correlates

calculate monthly and quarterly MPCs non-parametrically by calculating

$$MPC_{i,t} = \frac{\Delta c_{i,t}}{\Delta y_{i,t}} = \frac{c_{i,t} - c_{i,t-1}}{Refund_{i,t}}$$

Where $Refund_{i,t}$ denotes the tax refund amount and $c_{i,t}$ denotes consumption in the quarter or month t of refund receipt. For any variable x_{it} denote the baseline average level by

$$\bar{x}_{it} = \frac{\sum_{j=t-T}^{t} x_{ij}}{T}$$

and deviation from baseline

$$\widetilde{x_{it}} = x_{it} - \bar{x}_{it}$$

For this variance decomposition exercise I use the 9 months prior to refund receipt to calculate a baseline measure. Households are identified as homeowners if they contribute at least $1000 to mortgage payments during the baseline period. Differencing derives an orthogonal component. I estimate the following specification

$$\begin{aligned} MPC_{it} = \alpha + \gamma_1 age_{i,t} + \gamma_{2a}\widetilde{liq\_bal}_{i,t} + \gamma_{2b}\overline{liq\_bal}_{i,t} + \gamma_{3a}\widetilde{credit\_bal}_{i,t} + \gamma_{3b}\overline{credit\_bal}_{i,t} \\ + \gamma_{4a}\widetilde{income}_{i,t} + \gamma_{4b}\overline{income}_{i,t} + \gamma_{5a}\widetilde{logins}_{i,t} + \gamma_{5b}\overline{logins}_{i,t} \\ + \gamma_6 \mathbb{I}_{home,i} + \epsilon_{i,t} \end{aligned}$$

and perform a variance decomposition of the resulting estimates. Due to correlation in RHS variables, a variance decomposition is sensitive to the ordering of regressors. When regressors are uncorrelated, one can remove regressors one by one and record the difference in R^2 as the variance of the response explained by each subsequent regressor.

In order to address this issue I use the method of Lindeman, Merenda, and Gold (1980) which derives the variance decomposition from sequential sums of squares averaged over all permutations of the RHS variables. Table E.18 presents the results of this procedure.

Table I.16: Variance Decomposition, Tax Refund Consumption Responses

	MPC Total Expenditure		**MPC Non-Durables**	
	Proportion of Variation	Correlation	Proportion of Variation	Correlation
Age	0.0053	+	0.0168	+
Liquid Balance	0.5841	−	0.5006	−
Liquid Balance, Deviation	0.0008	+	0.0013	+
Credit Card Balance	0.0317	−	0.0007	+
Credit Card Bal., Deviation	0.0036	−	0.0012	−
Total Income	0.1667	−	0.1333	−
Total Income, Deviation	0.0571	−	0.0515	−
Account Logins	0.0538	+	0.1399	+
Account Logins, Deviation	0.0480	+	0.038	+
Home Owner	0.0487	−	0.1147	−

Notably, the baseline measure of liquid balances accounts for over half of the explainable variation in consumption responses. Once controlling for liquid balances and income, age has very little explanatory power for consumption responses.

Appendix I.1. Non-Parametric Approach & Self-Selection

In this section I further explore the panel dimension of tax refund responses by exploiting within household variation in liquidity across years. As above, I construct monthly non-durable MPCs non-parametrically, ie. $\frac{\Delta e_{i,t}^{\mathcal{ND}}}{\Delta Income_{i,t}} = \frac{\Delta e_{i,t}^{\mathcal{ND}}}{Refund_{i,t}}$. I decile households by $\frac{Liquid\ Assets}{Income}$ as in the event study design explore in the main text. I regress the non-parametric MPCs on this grouping as follows

$$\frac{\Delta e_{i,t}^{\mathcal{ND}}}{Refund_{i,t}} = \alpha_i + \gamma_t + \sum_{q=2}^{10} \beta_q \frac{Liquid\ Assets}{Income}^q_{i,t} + \epsilon_{i,t} \tag{I.2}$$

Where α_i represents a household fixed effect and γ_t a time fixed effect. The coefficients of interest, $\{\beta_q\}_{q=2}^{10}$, measure the average MPC of each liquidity decile with respect to the omitted group. Figure $H.41$ plots the results against those obtained in the event study design.

Notably, the negative correlation between liquidity and consumption responses remains. Additionally, the non-durable responses of the highly liquid remain significant. The results indicate that the large responses of the liquidity constrained are more likely to be driven by self-selection, indicating a role for preference heterogeneity in explaining household proximity to liquidity constraints.

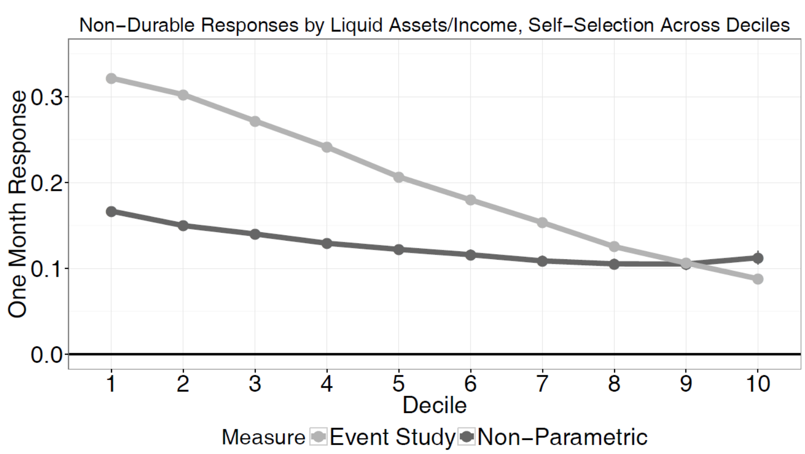

Figure I.24: Non-Parametric and Event Study Approaches

I also consider an imputed version in which I run the reduced form regression $H.2$ with $\frac{\Delta e_{i,t}^{\mathcal{NC}}}{Refund_{i,t}}$ on the RHS, obtaining coefficients $\{\beta_q^C\}_{q=2}^{10}$ and then obtaining imputed coefficients via $\beta_q + \xi^q \cdot \beta_q^C$ for each decile, where , as above ξ^q represents the identified proportion expended towards non-durables for households in population q in the month prior to refund receipt.

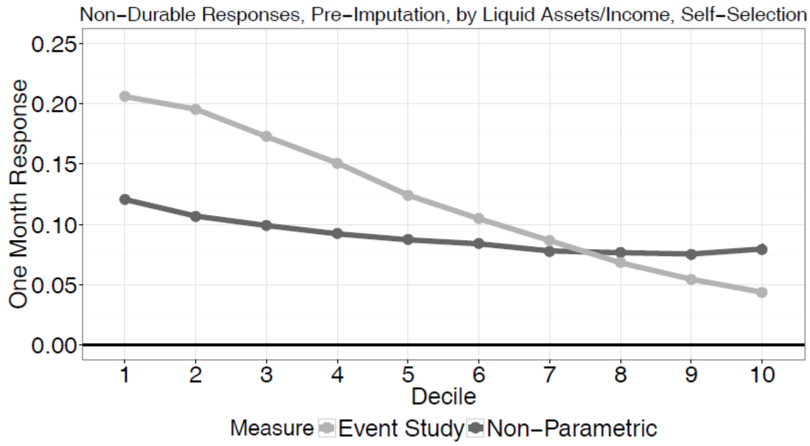

Figure I.25: Non-Parametric and Event Study Approaches, Pre-Imputation

Appendix J. 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 Appendix

Below all appendices regarding the modeling sections of the paper.

Appendix .1. Mental Accounts Life-Cycle Model

In this section I explore the implications of mental accounting frictions in a life-cycle context and assess the quantitative performance of the model. I structurally estimate the model's preference parameters, calibrating these parameters to match two sets of moments: low frequency life-cycle liquid asset accumulation and high frequency consumption responses to a pre-announced payment. In order to assess the counter-factual implications of the model in comparison to a standard buffer-stock case, I also estimate the model with mental accounting frictions turned off. I allow the standard buffer-stock model to attempt to match both sets of moments jointly, and to match only the observed consumption responses.

I consider the canonical buffer-stock life-cycle model studied by Carroll (1996), Parker and Gourchinas (1999), and Cagetti (2003), amongst others, calibrated to the monthly frequency, and augmented by the mental accounting utility function described above. Households in the model retire at $t = 480$ (40 years of working life) and perish with certainty at $T = 660$ months (80 years of age). They solve

$$\begin{aligned}
\max \mathbb{E}\Bigg[&\sum_{t=0}^{T} \beta^t \nu(\mathbf{c}_t) + \beta^{T+1} \upsilon_{T+1}(\mathbf{a}_{T+1})\Bigg] \\
&st. \\
&\mathbf{c}_t + \mathbf{a}_{t+1} \leq \mathbf{y}_t + \mathbf{a}_t(1+r) \\
&\mathbf{a}_{t+1} \geq \underline{a}
\end{aligned}$$

Utility, $\nu(c)$, is as define above and $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$. The function $\upsilon_{T+1}(\cdot)$ represents a bequest function of the form $\kappa \frac{(\mathbf{a}_{T+1})^{1-\gamma}}{1-\gamma}$. The income process during working life is determined by $\mathbf{y}_t = \mathbf{p}_t y_t$. Where $\mathbf{p}_{t+1} = \Gamma_{t+1}\mathbf{p}_t$ and the process $\{\Gamma_t\}_{t=0}^{T}$ is a deterministic life-cycle income growth profile. During working life household face income uncertainty. When a household is employed their income, y_t, follows an AR(1) process with persistence ρ and volatility $\epsilon_t \sim (0, \sigma_\epsilon^2)$. With exogenous probability p_u households become unemployed and receive unemployment insurance u_i. Households regain employment with probability p_e.

I begin by assuming no borrowing ($\underline{a} = 0$). Writing the model recursively and normalizing by $\{\Gamma_t\}_{t=0}^T$, as in Carroll (2012), yields

$$\begin{aligned} v_t(y_t, a_t) = &\max_{c_t} \left\{ \nu(c_t) + \beta \cdot \hat{\beta}_t \cdot \mathbb{E}_t[\Gamma_{t+1}^{1-\gamma} v_{t+1}(y_{t+1}, a_{t+1})] \right\} \\ &st. \\ &c_t + a_{t+1} \leq y_t + a_t \frac{(1+r)}{\Gamma_{t+1}} \\ &a_{t+1} \geq \underline{a} \end{aligned}$$

The model in levels is recovered via $\mathbf{x}_t = \mathbf{p}_t \cdot x_t$. Here I introduce the discount factor correction term, $\{\hat{\beta}_t\}_{t=0}^T$, studied by Attanasio et al. (1999), which deterministically adjusts the period discount factor for the number of adults and children in the household. The income process $(\{\rho, \sigma_\epsilon^2, \{\Gamma_t\}_{t=0}^T\})$ is determined via a first-stage estimation. Initial assets, a_0, are chosen to match quintiles of the liquid asset holdings amongst 24 year old SCF respondents. I proceed by estimating the four preference parameters $\{\beta, \gamma, \psi, \kappa\}$ via the method of simulated moments to match life-cycle wealth accumulation observed in the SCF (8 moments, one for each five year interval of working life) and high frequency consumption responses observed in the transaction data (10 moments - one for each decile of liquidity). Table 6 summarizes the model's parameters.

Table .17: Model Parameters, Monthly Calibration, Liquid Assets

	Parameter	Symbol	Value	Source
Primitives	Rate of Return	r	0.78%	60-Month CD[20]
	Initial Endowment	a_0	·	SCF
	Demographic Correction	$\{\hat{\beta}_t\}_{t=0}^T$	·	Cagetti (2003)
Income	UI Replacement	u_i	0.463	OUI
	Job-Finding Probability	p_e	0.48	CPS
	Separation Probability	p_u	0.011	CPS
	Stimulus Payment	T	0.6966	BANK
	Deterministic Income	$\{\Gamma_t\}_{t=0}^T$	·	First Stage
	Income Persistence	ρ	0.8962	First Stage
	Income Volatility	σ_ϵ	0.073	First Stage
Preferences	Time Preference	β	0.9344	SMM
	Risk Aversion	γ	2.48	SMM
	Dissaving Aversion	ψ	0.346	SMM
	Bequest Motive	κ	239	SMM

Standard arguments for parameter identification apply for determination of the time preference and risk aversion parameters. Agent's motives for saving are both precautionary and to smooth the income drop at retirement. As agents approach retirement and income uncertainty is resolved, the importance of the former motive (driven by the degree of risk aversion, γ) gives way to the latter (driven by the degree of impatience, β). Identification of the dissaving aversion parameter, ψ, relies on the model-generated cross-section of consumption responses to predictable income. As shown above, this parameter largely governs the consumption response of households with high levels of liquidity and the degree of consumption front-loading to the period of income receipt.

Appendix .1.1. First Stage 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, 2005, 2010, and 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). 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 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. In the Model Appendix E. I consider an estimation procedure in which I match total household net worth instead. Figure 9 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).

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

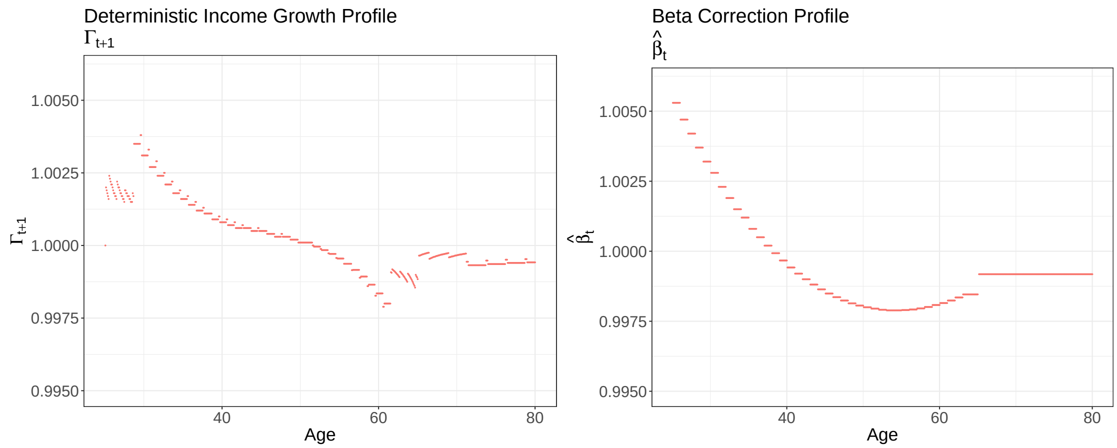

Figure .26: Deterministic Profiles

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 following $AR(1)$ process

$$y_{i,a} = \rho y_{i,a-1} + \epsilon_{i,a}$$

I obtain values for ρ and σ_ϵ^2 of 0.8962 and 0.073, respectively. 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.

Appendix .1.2. Preference Parameter Estimation

In order 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. Employing the SCF sample weights, Figure 12 plots median wealth for individuals within eight age groups constructed from five year increments, along with the 25th and 75th percentiles.

Due to the large skewness in wealth profiles, I follow the literature and perform the estimation to match the median wealth within each age group. For high

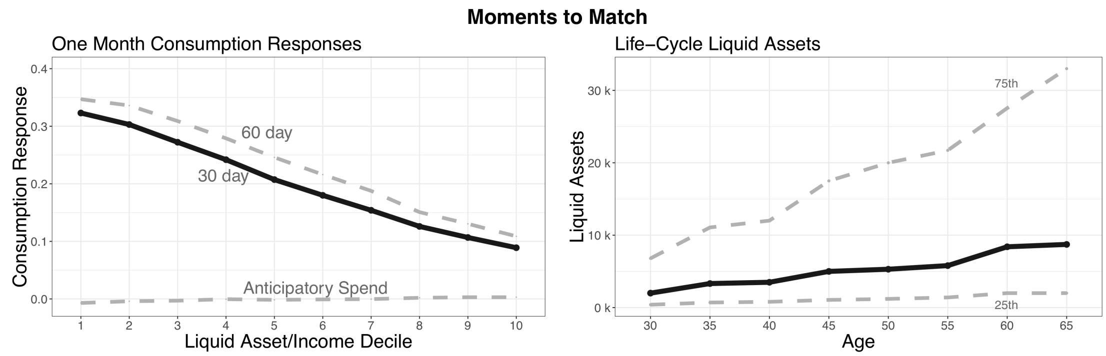

Figure .27: Data moments

frequency consumption responses, I utilize the estimates obtained from the transaction data for tax refund recipients. I plot the set of eighteen moments to be matched in Figure 12. In the model, agents are alerted l periods in advance to an inflow, T, to be deposited at time t. The size of the inflow is calibrated to match the average tax return relative to income observed in the transaction data. Parameters are calibrated to satisfy[21] :

$$\min_{\beta,\gamma,\psi,\kappa} \Theta \sum_{a}^{8} \mid d_a^{liq} - m_a^{liq}(\beta,\gamma,\psi,\kappa) \mid + (1-\Theta) \sum_{j}^{10} \mid d_j^{mpc} - m_j^{mpc}(\beta,\gamma,\psi,\kappa) \mid . \tag{.2}$$

This objective function includes a life-cycle liquid assets component (liq) and a cross-sectional consumption responses component (mpc). The relative importance of these two components is governed by the parameter $\Theta \in [0,1]$, which is adjusted to account for the relative magnitudes of each term. The first component measures the median absolute distance between d^a, the median value of liquid assets for SCF respondents amongst each each group, and m_a^{liq} denotes the model-constructed liquid wealth holdings for each age group at the parameter values $\{\beta,\gamma,\psi,\kappa\}$.

In the case of the second component, j indexes deciles of the liquid wealth

[21] As suggested by Carroll (2012), I also consider the following objective, which allows for SCF measurement error:

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

Where d_i^a denote the empirical value of liquid assets for SCF respondent i and ω_i^a denotes the weight assigned to each observation, obtained from SCF sample weights. The results are robust to this alternate objective.

distribution in advance of income receipt. This component measures the median absolute deviation of model implied average consumption responses by decile ($\{m_j^{mpc}\}_{j=1}^{10}$) from their empirical counterparts $\{d_j^{mpc}\}_{j=1}^{10}$. In order to obtain the former, I simulate the model and decile households by liquid wealth in advance of stimulus announcement. Consumption responses are then obtained non-parametrically from the simulated data (ie. for a stimulus announced l periods in advance of time t, $mpc_{j,i} = \frac{c_{t,j,i} - c_{t-l-1,j,i}}{R_{t,j,i}}$). This amounts to running the same reduced form regressions on the model generated data as the transaction data. I search across the parameter space via a four-dimensional Sobol sequence. I solve the model for each set of generated parameters, simulate the lifetimes of forty thousand agents and obtain the model-generated moments. I choose the set of parameters that minimize the objective (4).

Appendix .2. Estimation Results & Life-Cycle Savings Decomposition

In this section I report the results for the structural estimation procedure and perform a decomposition of life-cycle savings. Additionally, I compare these results to those obtained from a structural estimation of the standard buffer-stock model. This establishes the unique role for mental accounts frictions in rationalizing the data.

The results suggest moderate levels of impatience and risk aversion (Table 7). Dissaving aversion is found to be significant (0.346). On the interval between buffer-stock ($\psi = 0$) and hand-to-mouth ($\psi = 1$), the results suggest households tend towards the former. Figure 11 plots the model generated moments against the data. Average one month consumption responses observed in the data and those generated by the model are 0.2 and 0.17, respectively. The average responses amongst the tenth decile of liquid asset holders is 0.09 in the data, compared to 0.12 in the model. Along the dimension of liquid assets, the model matches both the level and profile of median life-cycle liquid asset accumulation. In addition to generating a correlation between assets-to-income and consumption responses in line with the data, the mental accounts model is able to capture the degree of consumption front-loading to receipt.

Table .18: Estimation Results

Model	β (Annual)	γ	κ	ψ
Mental Accounts	0.9344	2.481	238.7	0.346
Buffer-Stock	0.8994	2.330	278.0	·
Buffer-Stock, $\Theta = 0$	0.7480	1.043	287.8	·

To assess parameter identification, I report contours for each set of moments, plotting the median absolute distance against key model preference parameters (Figure 14). For each plot I fix the absent parameters to a small window around

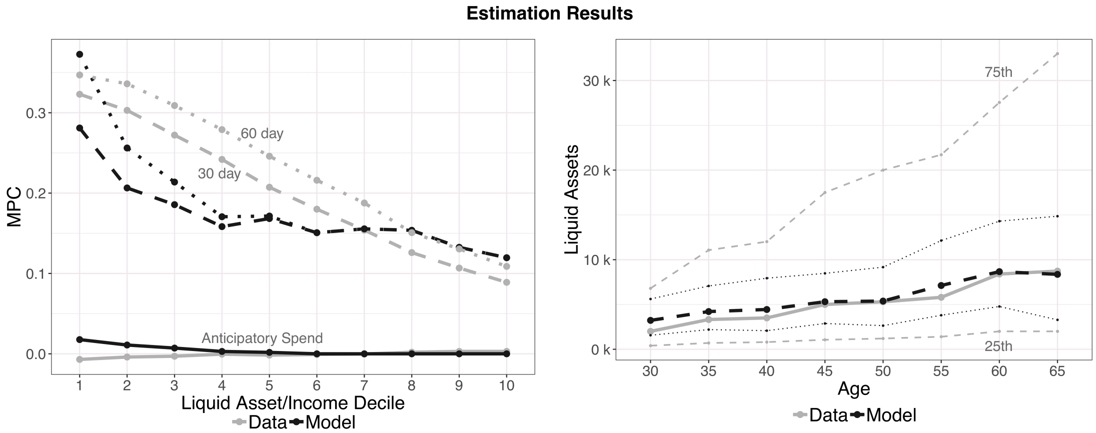

Figure .28: Estimation Results

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.

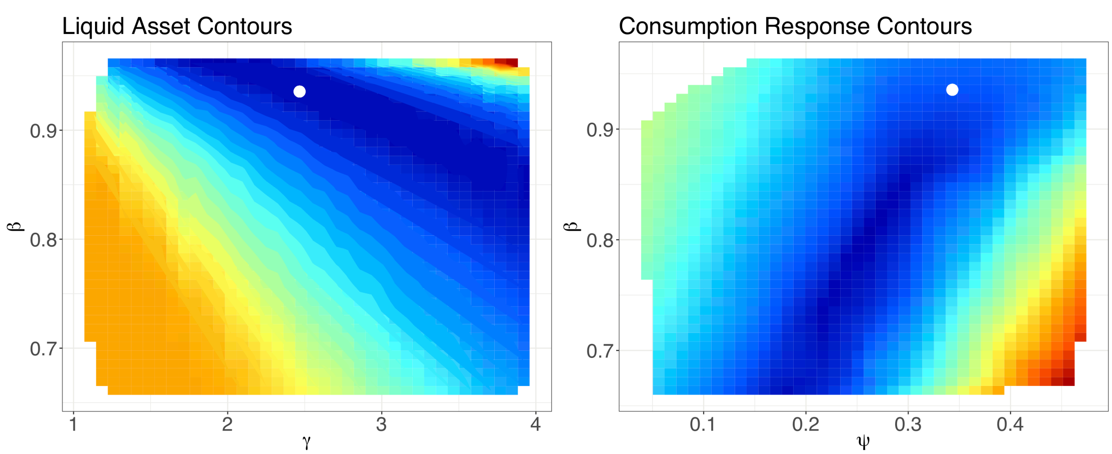

Figure .29: Preference Parameter Contour Plots

I re-estimate the model with the restriction $\psi = 0$. This is the standard buffer-stock case. I estimate the restricted model, first to match the same objective as the mental accounts model, and, second, to match the profile of consumption responses only ($\Theta = 0$). The latter serves to give the standard model the best odds of obtaining the dispersion in consumption responses observed in the data. The estimation results are reported in Table 7. Figure 12 plots the cross-sectional consumption and median liquid asset life-cycle profiles against those of the data for each of the three estimations. Additionally, the distribution of liquid assets in advance of stimulus announcement is reported, along with the median (solid line) and the 20th and 80th percentiles (dashed lines) of liquid asset holdings.

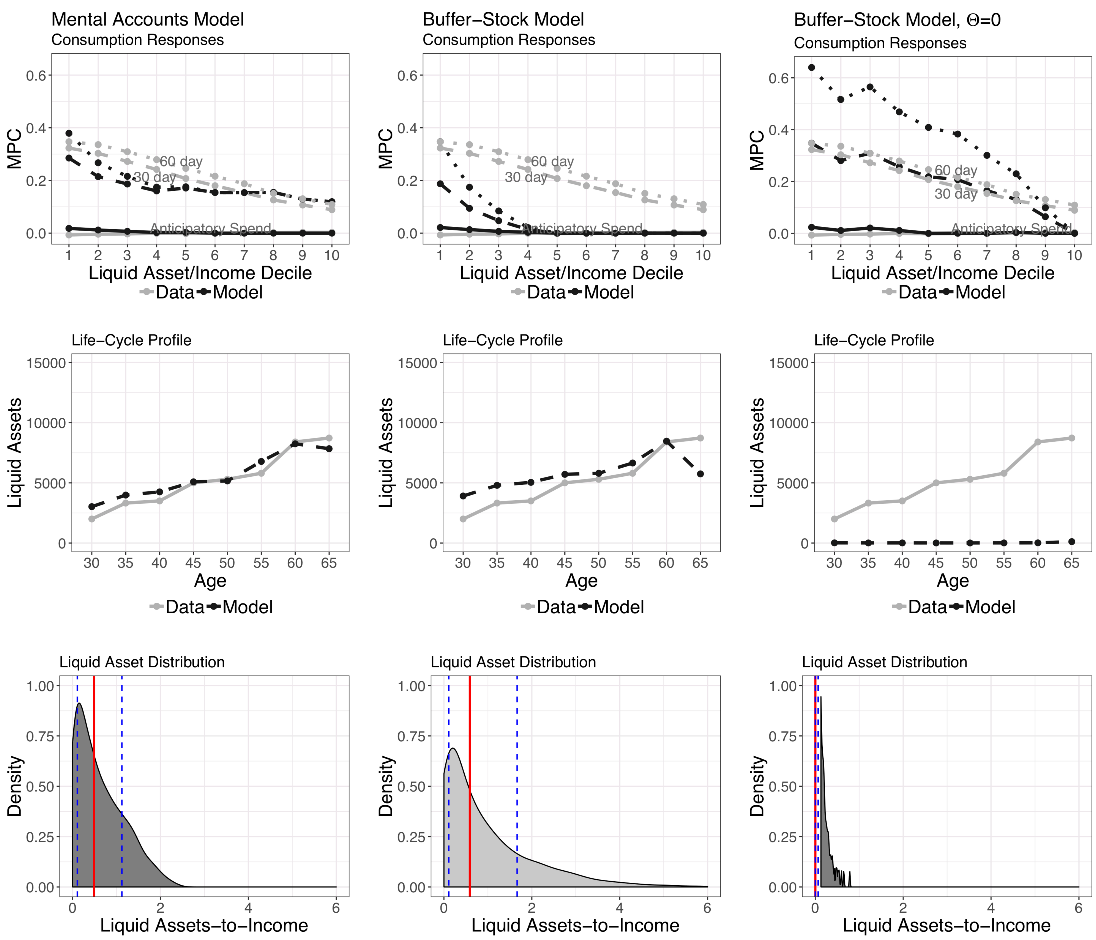

Figure .30: Model Comparison

The standard buffer-stock model is unable to obtain a significant degree of excess sensitivity amongst highly liquid households. The standard model is able to generate a dispersion in thirty day consumption responses in line with the data, but at the cost of counter-factually collapsing the liquid wealth distribution. Under this parameterization households at the 80th percentile of liquidity hold less than 7% of monthly income in liquid assets. Additionally, the standard model is unable to capture the degree of consumption front-loading evident in the data. In Model Appendix A. I further illustrate the role of each preference parameter by varying each in turn. Additionally, I report contours, plotting each parameter pair against the objective distance they generate.

All else equal, the introduction of mental accounts push estimates for both the patience and risk aversion parameters upwards. The latter implies a decreased sensitivity of household consumption growth to changes in the rate of return on savings. It is instructive to disentangle the role of mental accounting frictions from other factors that affect life-cycle consumption/savings decisions. By shutting down income uncertainty, relaxing the ad-hoc constraint on borrowing, and turning off the mental accounting friction in turn, I decompose the extent to which household savings decisions are due to precautionary, life-cycle, and mental accounting motives, respectively.

I first shut down the mental accounting friction and obtain the counter-factual life-cycle savings profile (Figure 14). Captive to the structure of mental accounts, households in the model construct their budgets on a monthly basis and their consumption tracks closer to income than a traditional buffer-stock agent. These deviations add up - during working life the median mental accounting agent's savings buffer is roughly forty percent less than a traditional agent, and he enters retirement with about twenty percent less in liquid assets. I next uncover the proportion of household savings held for precautionary reasons by shutting down income uncertainty and solving the model at the estimated parameters. Consistent with previous studies, I find that household savings decisions are driven by precautionary motives until households approach retirement age. This result is robust to whether or not mental accounts are present.

Implementing the 'no mental accounts' allocation in an economy with mental accounts frictions is straightforward - a consumption subsidy of $\frac{1}{1-\psi}$ when $c > y$ and 0 otherwise. However, such a policy is likely to be infeasible in practice. An alternative approach is for resources to be manually reassigned across mental accounts. Here I consider a linear savings rule by which a proportion of current period income is immediately earmarked for savings:

$$a^d = a + \delta y$$

where δ parameterizes the proportion of current period income set aside as savings. In practice, this might represent an automatic transfer, δy, of a household's income to a savings or money market account. Note the distinction between a savings rule of this nature and contributions to a retirement account or pension fund - the former are fully liquid.

A savings rule of this form pushes a household's month-to-month consumption reference point from $c^d = y$ to $c^{d,\delta} = y(1-\delta)$. Due to the cognitive cost of dissaving,

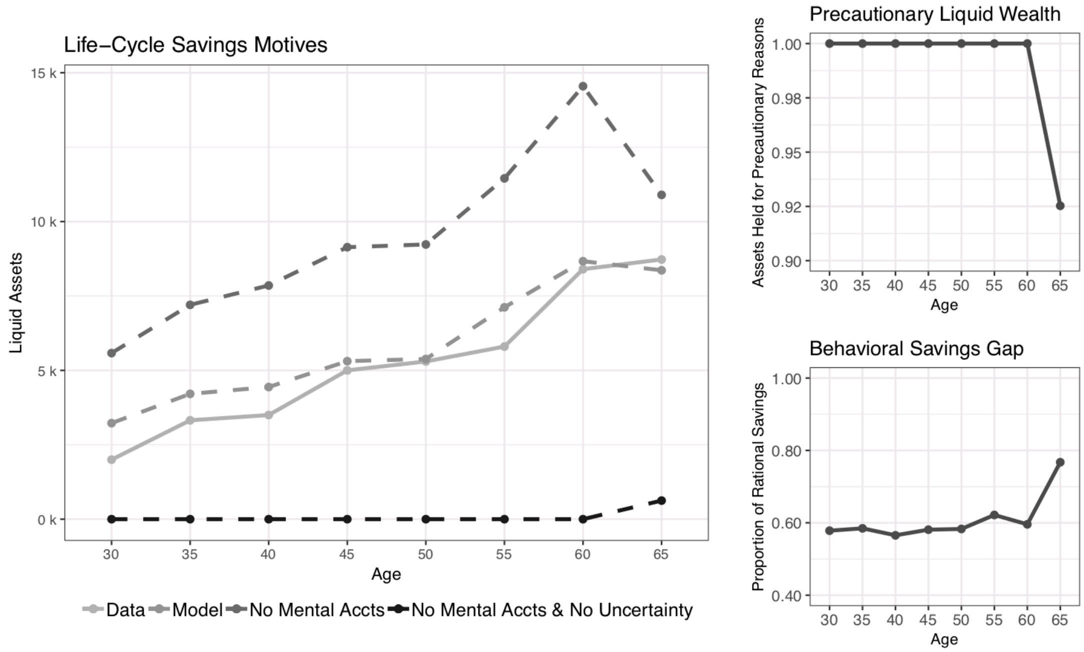

Figure .31: Life-Cycle Savings Decomposition

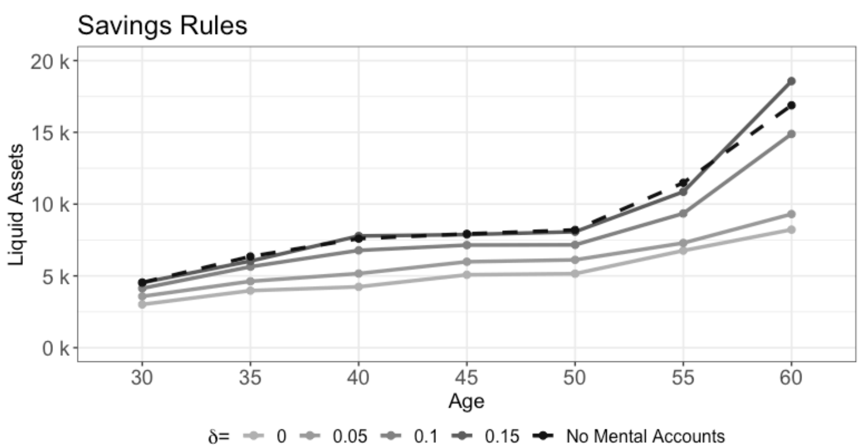

Figure .32: Linear Savings Rules

agents subject to a savings rule ($\delta > 0$) are similarly excessively sensitive to present income, taking advantage of large transitory income realizations by over-consuming relative to a traditional buffer-stock agent. At the extreme ($\psi = 1$) an agent subject to a savings rule consumes $c = (1 - \delta)y$, accumulating savings each period $\Delta a = \delta y$. It is straightforward to see that that a linear savings rule fails to dampen the excess sensitivity of consumption to present income (where the consumption subsidy described above succeeds) but acts to promote long-term savings.

Table .19: Savings Rules and Partial Insurance

Parameters	ψ	0.346	0.346	0.346	0.346	0.0
	δ	0.0	0.05	0.1	0.15	0.0
Partial Insurance	$\frac{\bar{\sigma_c}}{\sigma_y}$	0.768	0.735	0.606	0.559	0.557
MPC	$\frac{\partial c}{\partial T}$	0.259	0.273	0.281	0.283	0.014

I compare the consumption responses to a pre-announced stimulus and the partial insurance provided by liquid assets[22] for buffer-stock households and mental accounts households with various parameterizations of the linear savings rule.

A restrictive savings rule serves to decrease the frequency with which households find themselves close to the ad-hoc borrowing constraint ($a \geq 0$). Over the course of the life-cycle, the average households benefits from these autonomous savings. A reassignment of 15% of each month's income to the asset account recovers the partial insurance levels of the standard buffer-stock model (Table 9) and recovers a similar right-skewed distribution of liquid assets (Figure 18). An absence of such savings rules, when households manage their cashflow as if they are subject to the restrictions of mental accounts, might explain the prevalence of liquidity-constrained households that is observed empirically.

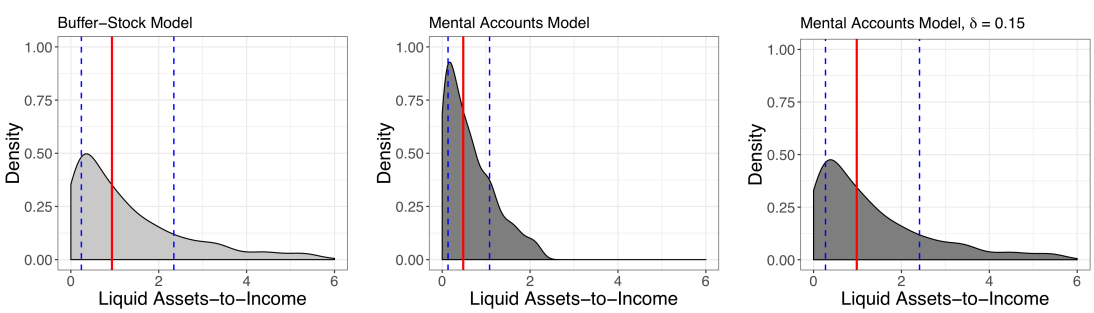

Figure .33: Comparison of Liquid Asset Distributions

[22]Partial insurance is measured as the average of the standard deviation of households' simulated monthly consumption paths relative to their income paths from ages 25 to 60, ie. $\frac{\bar{\sigma_c}}{\sigma_y} = \frac{1}{n}\sum_i^n \frac{\sigma_{c,i}}{\sigma_{y,i}}$. A value $\frac{\bar{\sigma_c}}{\sigma_y} = 1$ represents full pass-through of income fluctuations to consumption, whereas $\frac{\bar{\sigma_c}}{\sigma_y} = 0$ represents full insurance.

Appendix A. Implications for Fiscal Stimulus Policies

In this section I assess the model's positive implications for the design of fiscal stimulus payments by performing two experiments. In the first experiment I consider a redistributive policy in which the fiscal authority taxes households with substantial liquid wealth, and uses the revenue to provide a targeted stimulus to the most constrained households in the economy. In the second experiment I endow fiscal authority with a fixed amount of resources and consider three budget-equivalent policies: an un-targeted stimulus, an asset-targeted stimulus, and an income-targeted stimulus. In both cases I compare the aggregate consumption responses from the estimated mental accounts model to the counter-factual economy in which mental accounts are turned off, and a baseline buffer-stock economy estimated to match life-cycle liquid asset accumulation.

I assume a small open economy. Agents take the rate of return, r, as given and the liquid savings technology, a, is external to the economy. I parameterize economies using the preference parameter estimates from above (Table 8). Each agent solves the problem described in Section 3.3, subject to the parameterization described in Table 7.

Index each generation of agents by the period of birth j. Each generation begins working life with initial assets distributed to match the $10th$, $30th$, median, $70th$, and $90th$ percentiles of liquid assets amongst 24 year olds as reported in the SCF. Agents live from $t = j$ to $j + 660$ periods and retire at $t = j + 480$. I simulate the economy for 480 months with 250 agents born each month. The government is endowed with M resources and seeks to redistribute them to households. It faces the budget constraint

$$M = \sum_t \int_i T_t^i \cdot \Gamma_t^i di$$

In all cases, the requisite stimulus policy is announced to agents in the economy $t - 1$ periods in advance of implementation. Households fully internalize the announcement before making their time $t - 1$ consumption and savings decisions.

Appendix A.1. A Redistributive Stimulus

In this section I consider a redistributive stimulus policy, in which the government issues a targeted stimulus to households holding low levels of liquid assets by taxing those with high levels of liquidity. As before, the government announces the time t policy at time $t - 1$. The government has no resources ($M = 0$) and must implement a balanced budget policy ($\sum_t \int_i T_t^i \cdot \Gamma_t^i di = 0$).

Here I consider an example in which the government implements a lump sum transfer system, redistributing \$2500 households amongst the fifth quintile ($q = 5$) of liquid balances at the time of announcement to households in the bottom quintile

$(q = 1)$. That is, $\sum_{q=1}^{5} \int_i T_t^{i,q} \cdot \Gamma_t^i di = 0$, with $\int_i T_t^{i,1} di > 0$, $\int_i T_t^{i,5} di < 0$ and $\sum_{q=2}^{4} \int_i T_t^{i,q} di = 0$. Ex-post, the \$2500 transfer amounts to roughly a third of monthly income for agents in quintile 1.

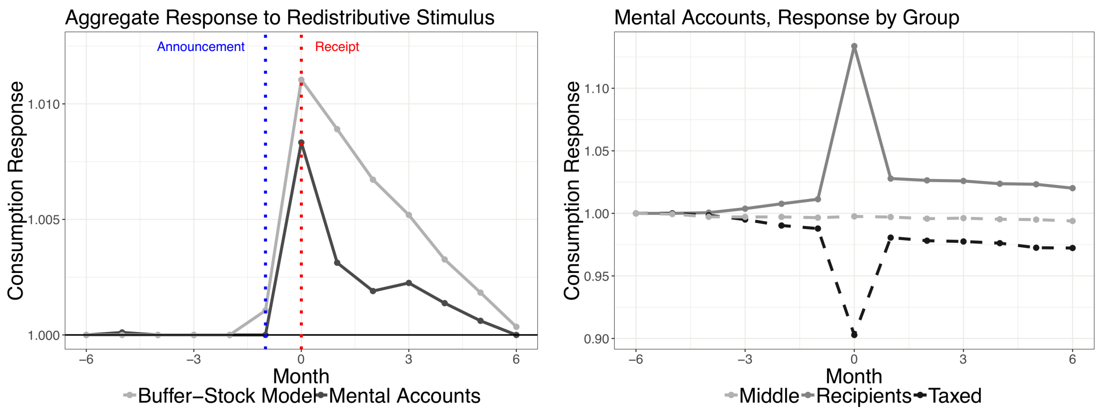

Figure A.34: A Redistributive Stimulus

Compared to the standard buffer-stock case, the redistributive stimulus policy is 53% less effective over two quarters under mental accounts (Figure 17, left panel). In the standard model agents with high levels of liquidity remain at their target consumption levels even in the face of a substantial lump-sum tax on their liquid wealth. This is in symmetry with their muted responses to predictable increases in income (see Figure 14). However, in the presence of mental accounts even lump-sum taxes are distortionary. Due to their reluctance to consume out of their current asset accounts, the agents facing a wealth tax adjust their consumption downwards during the month of payment (Figure 17). This downward adjustment offsets the high degree of excess sensitivity among constrained households receiving a transfer.

The results stand in stark contrast with the predictions of models in which heterogeneity in marginal propensities to consume are driven by households' proximity to an externally-imposed liquidity constraint. Whereas in the standard buffer-stock model a redistribution from households with high levels of liquidity to those with low levels of liquidity is essentially '*for free*', under mental accounts the decline in consumption amongst highly liquid households facing the lump-sum tax has a first order effect, depressing the government spending multiplier.

Appendix A.2. A Targeted Stimulus

Stimulus policy design is largely a matter of evaluating which population segments will have the largest propensity to expend the funds they receive. Another component is evaluating the feasibility of such targeting and the gains from doing so. In this section I explore these dimensions through the lens of the mental accounts model. I compare three distinct budget-equivalent policies[23]:

1. Un-targeted Stimulus: A \$100 blanket stimulus to all households in the economy
2. Income-targeted Stimulus: A \$500 stimulus targeted to households experiencing the bottom 20 percent of income realizations at announcement
3. Asset-targeted Stimulus: A \$500 stimulus targeted to households in the bottom quintile of liquid asset holdings

The un-targeted stimulus policy is comparable to the Bush tax rebates of 2001 and 2008. In the case of the former, passage of the Bush tax cuts triggered a rebate of \$300 − \$600 for all taxpayers that filed a return in the previous year. An income-targeted stimulus policy is comparable to unemployment insurance or workers' compensation policies. Unemployment programs replace roughly half of a workers pre-unemployment income, regardless of their current asset position. An asset-targeted stimulus is akin to means-tested programs, such as SNAP or TANF. The former jointly requires household income to be below 130 percent of the poverty line and assets of \$2,250 or less.

In comparison to a standard buffer-stock economy in which the income-targeted and asset-targeted policies are 6 and 8 times more effective than a blanket stimulus, respectively; under mental accounts the gains to targeting are significantly reduced and the relative effectiveness of these targeted policies is reversed. Consistent with the data, the decline in consumption responses across levels of liquidity is moderate, meaning the increased "bang for the buck" coming from targeting is stifled.

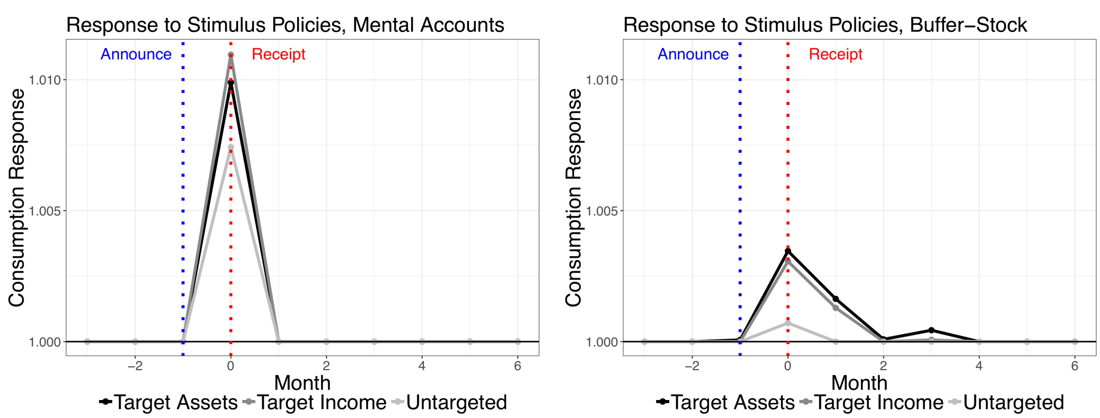

Figure A.35: Aggregate Consumption Responses Across Policies by Model

[23] Assuming 100 million U.S. households, each policy requires $M = \$10$ billion

Income-targeted and asset-targeted stimulus policies are, respectively, 47 percent and 33 percent more effective than an un-targeted stimulus in economies where mental accounting frictions are present. This reversal is the direct result of households' sensitivity to fluctuations in income, regardless of their current asset positions.

These results have direct implications for policy. Namely, the relative gains to targeting under mental accounts are roughly an eighteenth of those in the standard buffer-stock case. This indicates the feasibility, both financially and politically, of implementing such policies likely carries greater weight than previously thought. Additionally, the mental accounts model generates a key positive prediction: policies targeted towards households with temporarily low income, regardless of their current liquid asset positions, are particularly potent. This indicates a role for unemployment insurance programs with increased income replacement rates during economic downturns. Figure 19 plots the resulting cross-sectional consumption responses from asset-targeted and income-targeted policies.

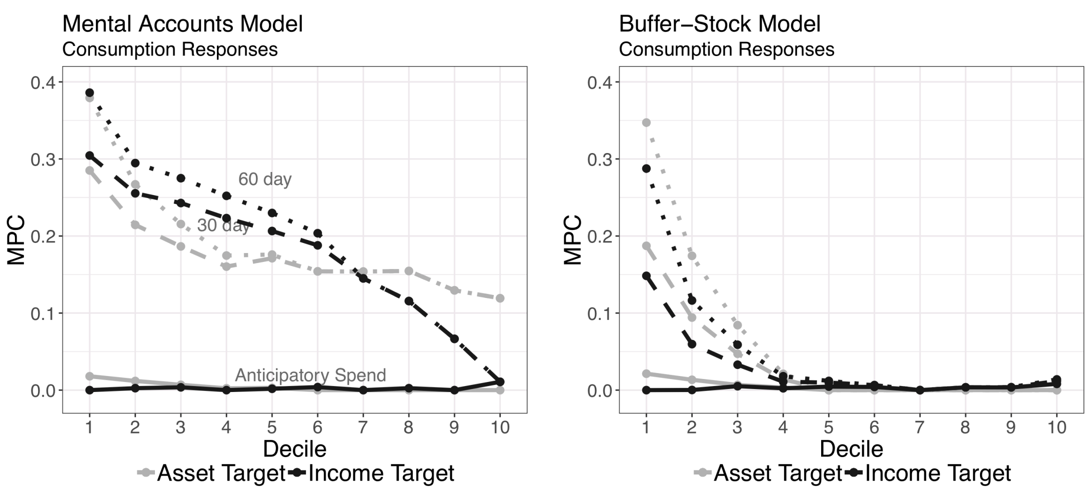

Figure A.36: Targeted Stimulus Policies

Appendix B. Structural Estimation

Appendix C. Parameter Identification

Estimate standard buffer-stock model to the objective satisfy

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

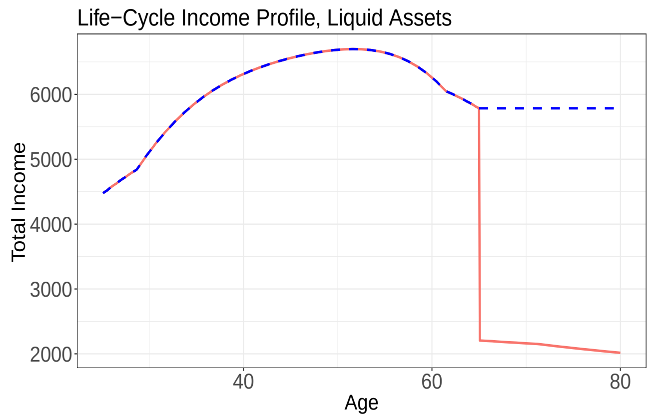

Figure B.37: Deterministic Component of Income, Monthly Frequency

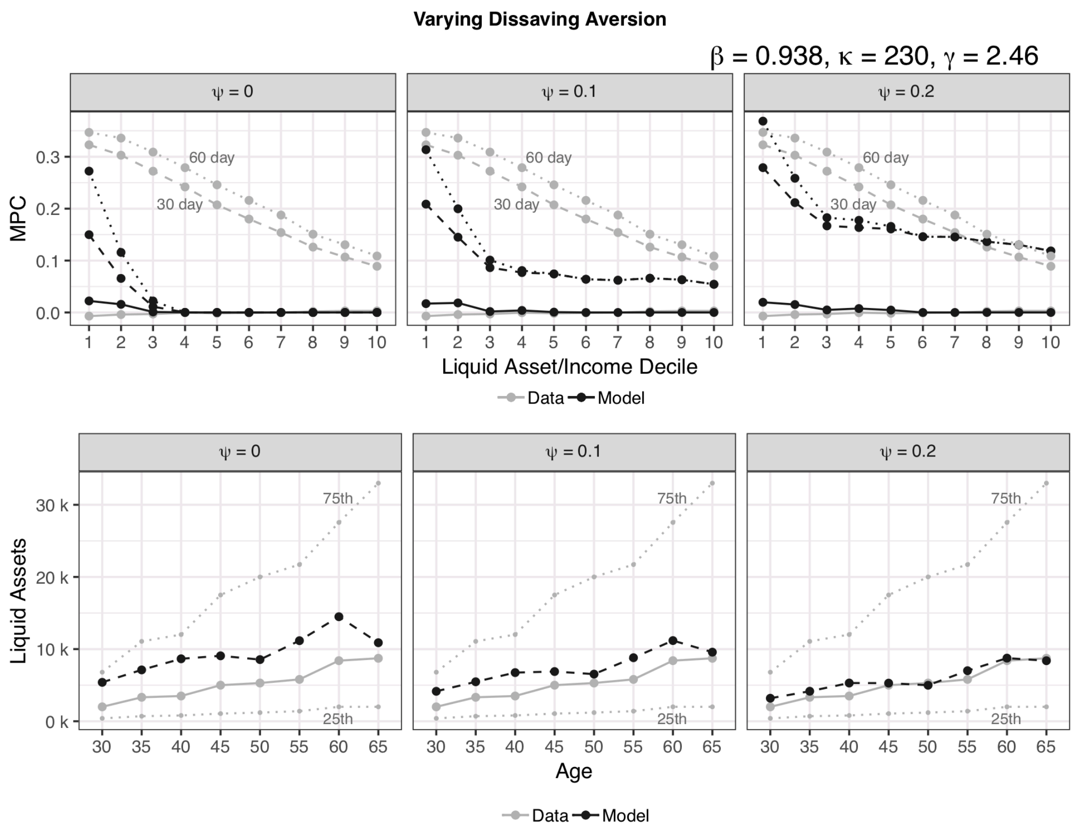

Figure B.38: Varying Dissaving Aversion, ψ

Where third argument, ψ, is set to zero.

Figure D.42 plots the resulting contours for each term of the objective and indicates a the fundamental trade-off in the standard model between matching

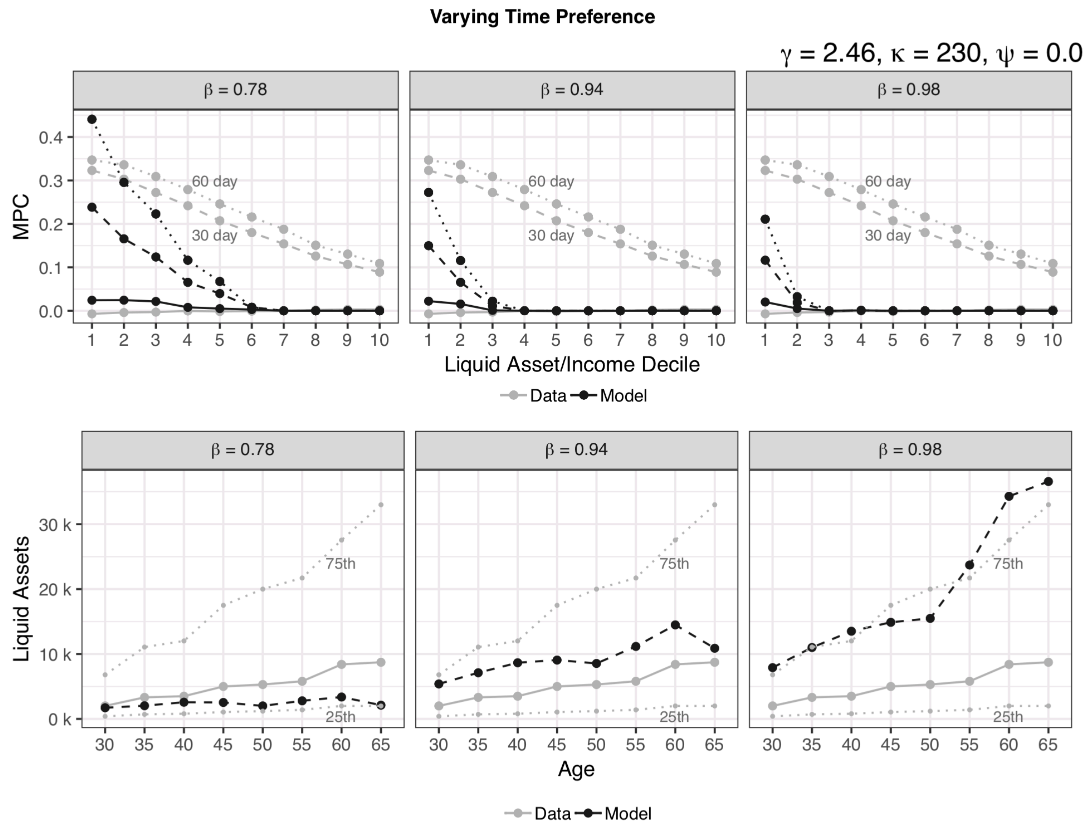

Figure B.39: Varying Time Preference, β

liquid asset balances and consumption responses.

Appendix D. Fiscal Stimulus

Here I compare the aggregate responses generated by a lump-transfer to all households in the mental accounts economy versus the standard buffer-stock case. All households receive two weeks of income at time t ($T_t^i = 0.5\ \forall i$). Below I plot the aggregate responses for the six months before and after stimulus receipt. As documented above, the models have distinct predictions for both the timing and magnitude of household consumption responses.

In comparison to the standard buffer-stock case, the mental accounts economy generates a response that is 4.35 times larger on impact, 2.55 times larger over one quarter, and 1.82 times larger once the stimulative effects die out after seven months. The magnitude of these aggregate responses has

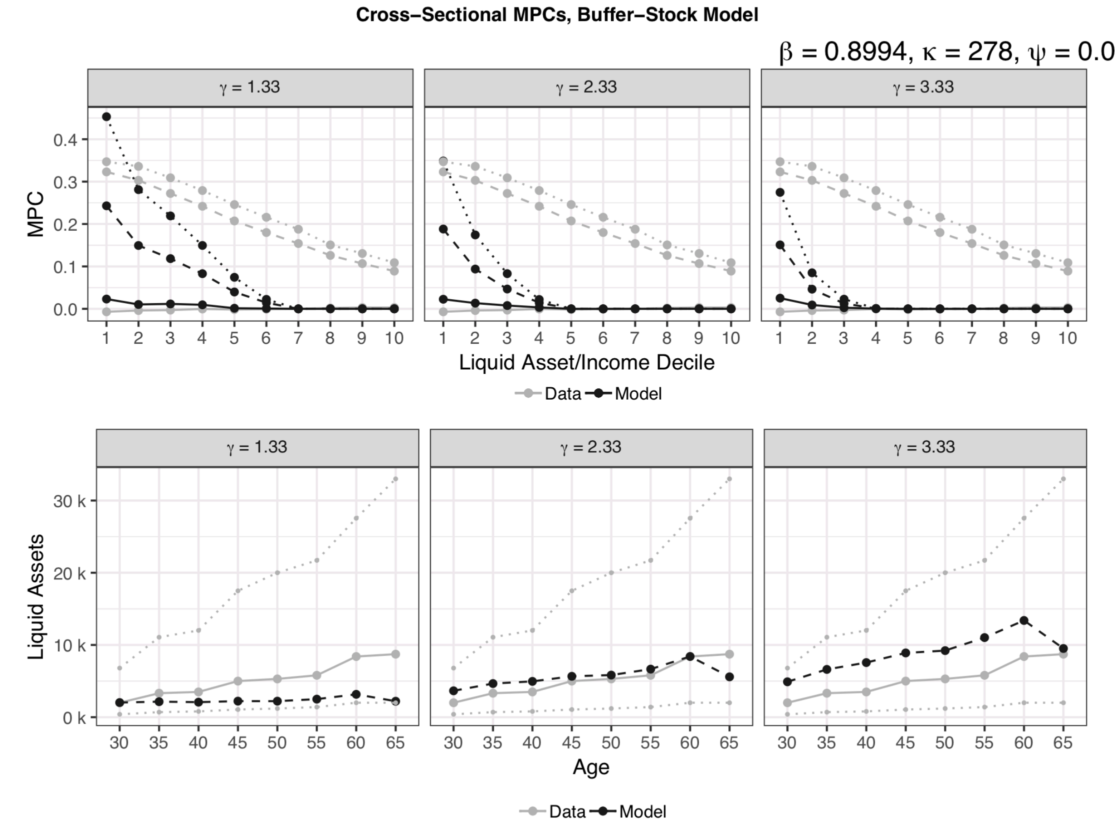

Figure B.40: Varying Risk Aversion, γ

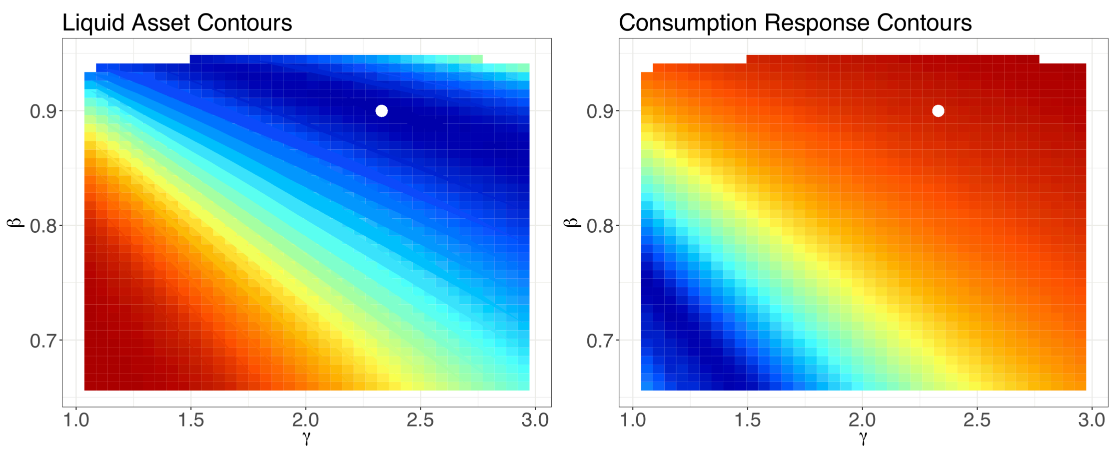

Figure C.41: Parameter Identification Contours, Standard Buffer-Stock Model

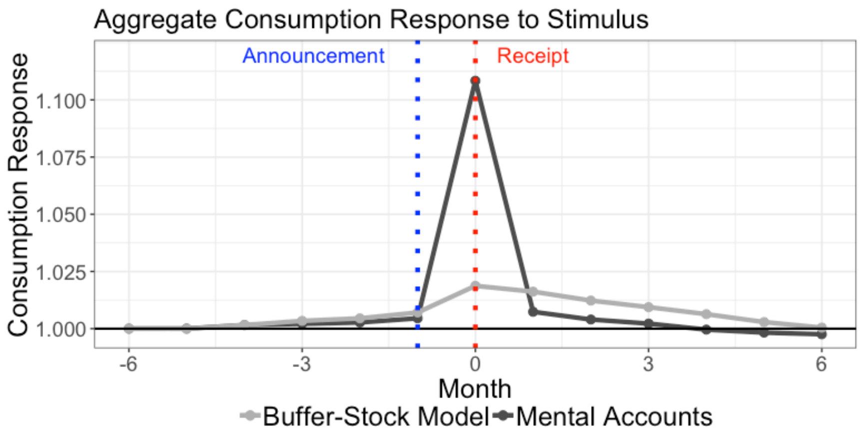

Figure D.42: Pre-Financed Stimulus Responses

Table D.20: Pre-Financed Stimulus Response Comparison

Cumulative	Mental Accts.	Buffer-Stock
Announcement	0.009	0.014
Receipt	0.226	0.052
One Quarter	0.24	0.094

Appendix E. Alternate Economies

Appendix E.1. Infinite Horizon

Here I solve a consumption/savings problem with mental accounts at the infinite horizon. Agents solve

$$\max_{\{c_t\}_{t=0}^{\infty}} \mathbb{E}_0[\sum_{t=0}^{\infty} \beta^t \nu(c_t)]$$
$$st.$$
$$c_t + a_{t+1} \leq y_t + (1+r)a_t$$

Writing the problem recursively

$$V(a,y) = v(c) + \beta \mathbb{E}[V(a',y')]$$
$$st.$$
$$c + a' \leq y + (1+r)a$$

Let $y \in \{y_{low}, y_{high}\}$. Here I consider a savings default such that the agent attempts to save a proportion $1-\delta$ of each periods income, $a^d = a(1+r) + +(1-\delta)y$,

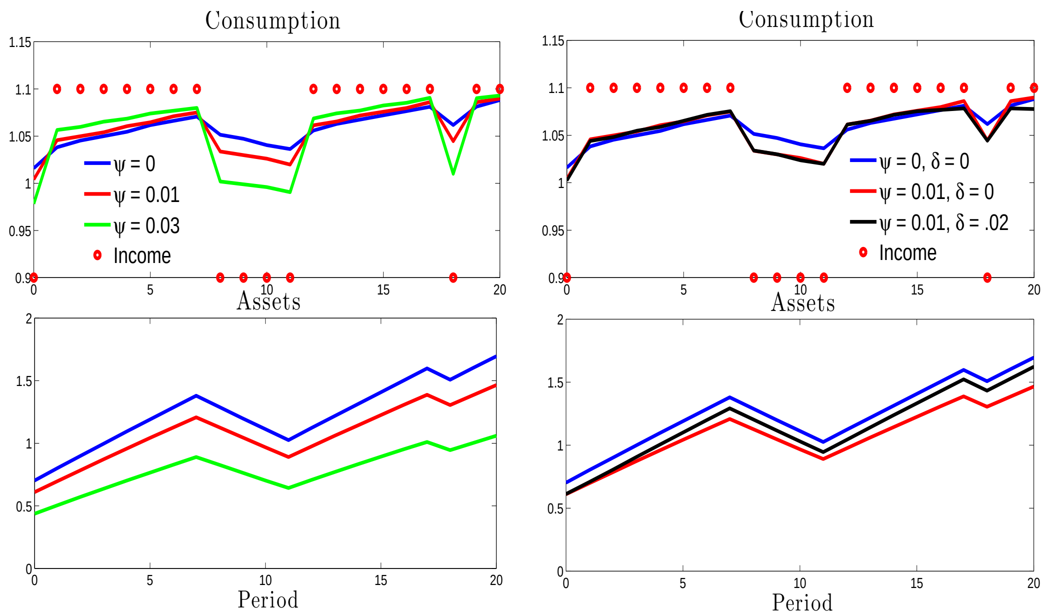

Figure E.43: Simulated Paths Under Uncertainty

pinning down default allocation $c^d = (1-\delta)y$. In the computational example I use log preferences. Income process $y_{low} = 0.9$, $y_{high} = 1.1$, and transition probabilities $p_{ii} = .8$, $p_{ij} = .2$. I solve the model for a variety of parameter values $\psi \in [0, .25]$ and $\delta \in [0, .075]$.

Figure 3 plots a typical simulated time series of consumption and assets in this economy. In the left panels I restrict to $\delta = 0$, pinning down a hand-to-mouth rule of thumb $c^d = y$. Under this parameterization the agent is dissaving averse for $\psi > 0$. As the psychological cost of deviation increases the agent more closely adheres to the hand-to-mouth rule, and consumption more closely tracks to income. Additionally, during periods of persistently low income the agent dissaves at faster rate for lower levels of ψ. In the right panels I vary δ, allowing the agent to following a savings accumulation rule-of-thumb. For a given ψ, levels of assets are increasing in δ.

Figure 4 plots the consumption and savings policy functions in this economy for $\psi = 0$ and $\psi = .05$. Mental accounting frictions dominate for low levels of assets. Agents dislike dissaving, and thus consume less than otherwise in the low state, and more than otherwise in the high state. For sufficient levels of savings, agents in the high income state consume their present income. For large levels of assets the policies of the mental accounting agents coincide with those of the non-mental accounting agents.

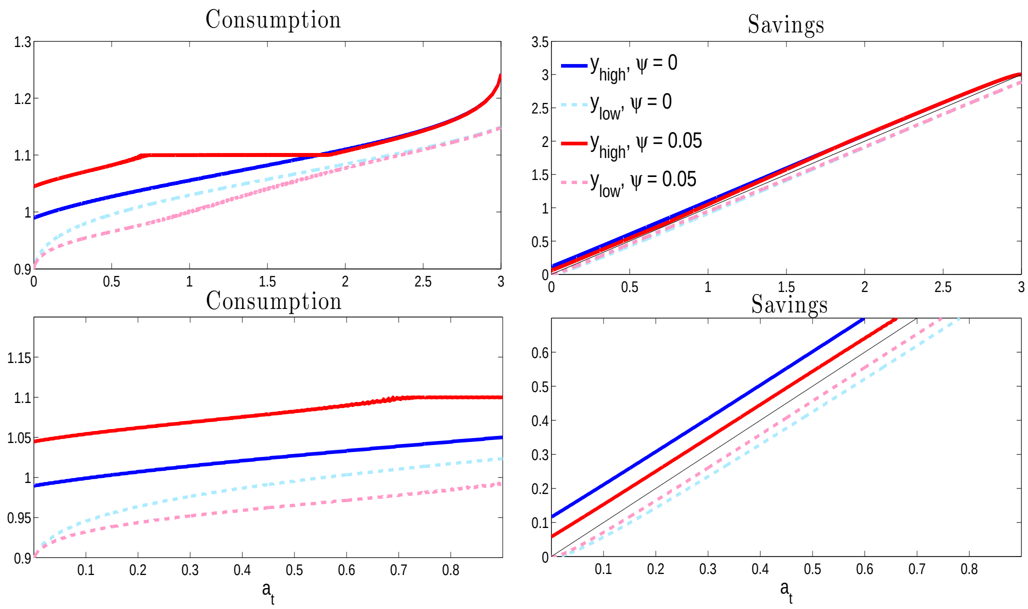

Figure E.44: Policy Functions (Top: Entire Grid, Bottom: Zoomed In)