

Consumption Behavior Across the Distribution of Liquid Assets

James Graham*
University of Sydney
Centre for Applied Macroeconomic Analysis, ANU

Robert A. McDowall
L.E.K. Consulting

June 26, 2023

Abstract

We study consumption responses to predictable income using transaction-level data from a U.S. financial institution. We find that consumption responses are front-loaded to income receipt, decline moderately with liquidity, but are significant even for households with substantial liquid assets. To rationalize these facts, we develop a model of mental accounts in which consumption choices are partitioned across current income and current assets. Our calibration suggests households are moderately averse to dissaving out of assets, which generates the timing and distribution of consumption responses observed in the data. Finally, we explore implications of the model for the design of fiscal policy.

Keywords: Consumption, excess sensitivity, fiscal stimulus

JEL: E21, E62, E70

*James Graham: james.a.graham@sydney.edu.au. Robert A. McDowall: rm3564@nyu.edu. We are grateful to Alberto Bisin and Diego Perez for their invaluable advice and support, and to Jess Benhabib, Andrew Caplin, and Basil Williams for their continued discussions and suggestions. We would also like to thank Tim Cogley, Peter Ganong, William Gamber, Victoria Gregory, Fiona Greig, Amar Hamoudi, Laurent Mathevet, Pascal Noel, Pablo Ottonello, Jonathan Parker, Thomas Philippon, Ryan Stevens, and seminar participants at NYU Stern and NERA for helpful comments and discussions. The underlying data used in this paper was made available by a financial institution. The data is de-identified and selectively made available for academic research. All opinions are the authors' own and do not reflect those of the financial institution. While working on this paper, Robert A. MacDowall was compensated for providing research assistance on public reports published by the financial institution. Any errors are our own.

1. Introduction

Household consumption responses to income fluctuations are fundamental for understanding the dynamics of individual behavior and are at the core of any model in macroeconomics. Informed by an extensive empirical literature (Hall, 1978; Zeldes, 1989; Johnson, Parker, and Souleles, 2006; Parker et al., 2013) documenting deviations from permanent income theory, structural models of household consumption have largely focused on the role of liquidity constraints (Deaton, 1991; Carroll, 1997; Kaplan and Violante, 2014; Mankiw, 2000) to rationalize the excess sensitivity of household consumption to predictable forms of income. The HANK model (Kaplan, Moll, and Violante, 2018), a workhorse currently employed to study the impacts of macro-policy, considers a two asset framework with large liquidation costs, obtaining average consumption responses approaching those observed in the data by endogenously inducing a significant proportion of wealthy agents to be cash-constrained. In this paper we document empirically that liquidity constraints and transaction costs are insufficient to explain household consumption responses to predictable income. The goal of this paper is to document the relationship between liquidity and excess sensitivity, propose a structural model rationalizing salient features of household consumption responses, and assess its positive implications for the design of fiscal stimulus policies.

In this paper we utilize a novel administrative dataset of de-identified transaction data from a panel of 17.2 million households obtained from a large U.S. financial institution. At the daily frequency, we study expenditure and balance sheet responses to predictable forms of income (federal and state tax refunds, regular paychecks, and bonus paychecks) as well as likely innovations in information at tax filing, for households with varying levels of liquidity and income. In line with previous findings (Johnson, Parker, and Souleles, 2006; Kaplan and Violante, 2014; Kueng, 2018), we find that, on average, households expend 25 percent of their tax refunds and 20 percent of their bonus checks towards non-durables in the quarter following receipt. However, we find that these consumption responses decline moderately in levels of liquidity and are significant even for households with high levels of income and substantial liquid assets. Further, we document that 70 percent of the five month non-durable consumption response generated by tax refunds occurs within thirty days of receipt and anticipatory expenditure is insignificant.

This pattern of consumption responses - insignificant spending in anticipation, excess sensitivity at receipt, and a response that is immediate and short-lived -

is pervasive. We show that these results are robust to the form and magnitude of the income studied, the category of expenditure considered, holds when one large inflow shortly follows another, and is restricted to receipt, rather than likely information innovation, such as at tax filing. Additionally, this pattern persists across age, levels of income, and liquid assets. In accordance with the predictions of standard models, we observe a tight negative correlation between liquid balances relative to income and the estimated consumption responses. However, contrary to the predictions of standard models, we document that excess sensitivity is significant for households with substantial liquidity.

We show that a standard buffer-stock life-cycle model, estimated to match the magnitude of the cross-sectional one month consumption responses that we observe in the data, is unable to jointly obtain the degree of consumption front-loading or the liquidity available to households observed empirically. Additionally, the empirical results are inconsistent with the predictions of several notable classes of behavioral models. At odds with models of rational inattention,(Reis, 2006; Gabaix, 2014) household’s consumption responses in the data are immediate and short-lived, and there is no systematic innovation in consumption at the date of tax filing.¹ In contrast with models of temptation(Laibson, 1997; Gul and Pesendorfer, 2001; Laibson, Repetto, and Tobacman, 2007) and forward-looking reference-dependent utility (Kőszegi and Rabin, 2006) consumption responses are delayed until receipt, even for households with substantial liquid wealth and those with access to credit cards.²

In order to interpret the empirical results, we develop a dynamic and tractable model of mental accounts (Shefrin and Thaler, 1988; Farhi and Gabaix, 2020). Motivated by recent empirical work documenting cases in which the fungibility of otherwise interchangeable resources fails (Hastings and Shapiro, 2012; Hastings and Shapiro, 2018), we consider a model in which agents are averse to dissaving and partition their consumption choice sets between current income and current assets. Agents in the model face a decline in marginal utility when spending out of savings and therefore take advantage of high income realizations by over-consuming relative to an agent whose income and assets are perfectly fungible.

¹Filing 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 likely be indicated by innovations in consumption on this day. Baugh et al. (2021) document similar spending inactivity at filing.

²Further, we find that consumption is delayed until receipt even as the predictable future income increases in magnitude.

The degree of mental accounting frictions in the model is governed by a single dissavings aversion preference parameter, and the model nests canonical buffer-stock and hand-to-mouth households as limiting cases.

We structurally estimate the model's preference parameters to jointly match the high frequency non-durable consumption responses we observe in the liquid wealth-to-income cross section as well as low frequency life-cycle wealth accumulation profiles. The nested extremes generate stark predictions for the cross-section of consumption responses to transitory fluctuations in predictable income. Whereas standard buffer-stock agents display large consumption responses only when liquidity constrained, hand-to-mouth agents consume the entirety of their income each period. The moderate decline in consumption responses across levels of liquidity that we observe in the data suggests an intermediate case. We document that the dissavings aversion preference parameter is crucial for rationalizing the magnitude and timing of consumptions responses amongst households with high levels of liquidity. These results suggest a re-evaluation of traditional aggregate business cycle models through the lens of mental-accounting frictions.

We employ the estimated model to consider a number of counter-factual experiments. We decompose the extent to which household savings decisions are due to precautionary, life-cycle, and mental accounting motives, respectively. Consistent with standard buffer-stock theory, we find that household savings decisions are driven by precautionary motives until they approach retirement. Captive to the structure of mental accounts, however, consumption in the model tracks closer to income than in the standard setup. These deviations add up - during working life the median mental accounting agent's liquid savings buffer is roughly 40 percent less than it would be if this friction was relaxed, and he enters retirement with around 20 percent less liquid assets.

We next assess the model's positive implications for the design of fiscal stimulus payments. In comparison to the standard buffer-stock model, a pre-announced redistributive stimulus to liquidity-constrained households generates an aggregate consumption response that is 53 percent smaller. This is because all households in the economy are sensitive to temporary fluctuations in income, and households with substantial liquidity decrease consumption to finance the lump-sum taxes they face. In this way, lump-sum transfers have a distortionary effect on the path of household consumption.

We compare three distinct budget-equivalent policies: 1.) an un-targeted \$100 stimulus to all households in an economy (akin to the Bush tax rebates of 2001

and 2008³), 2.) a \$500 stimulus targeted to households experiencing the bottom 20 percent of income realizations at announcement (akin to unemployment insurance or workers' compensation payments), and 3.) a \$500 stimulus targeted to households in the bottom quintile of liquid asset holdings (akin to means-tested programs, such as SNAP⁴ or TANF.⁵). We show that, 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 - income-targeted and asset-targeted stimulus policies are, respectively, 47 percent and 33 percent more effective than an un-targeted stimulus.

1.1. Related Literature

This paper contributes to a large empirical literature investigating the excess sensitivity of household consumption to predictable income. This includes papers documenting the tracking of consumption to income in aggregate data (Hall, 1978; Campbell and Mankiw, 1990), as well as a more recent literature documenting deviations from permanent income theory in micro-data. This literature includes Zeldes (1989), who studies the interaction of liquidity constraints and consumption in the PSID, and the works of Johnson, Parker, and Souleles (2006) and Parker et al. (2013), studying consumption responses to the government stimulus programs of 2001 and 2008, respectively. These papers document a large degree of excess sensitivity to receipt of stimulus that is consistent with models of liquidity constraints and has meaningful implications for macroeconomic aggregates.

However, a recent literature has documented contexts in which consumption responses cannot be explained by liquidity constraints alone. Kueng (2018) studies the excess sensitivity of high income consumers to payments from the Alaska Permanent Fund and finds responses that are inconsistent with buffer-stock behavior and rational inattention. Olafsson and Pagel (2018) document consumption patterns in the week before and after paydays for a population of Icelandic households. They show that households with the credit space to smooth expenditure in the days immediately in advance of income display excess sensitivity to receipt. Ganong and Noel (2019) document a drop in household consumption at the ex-

³In actuality these payments were \$300 for single filers and \$600 for married couples.

⁴The Supplemental Nutrition Assistance Program program jointly requires household income to be below 130 percent of the poverty line (\$2,252 per month for a family of three in 2019) and assets of \$2,250 or less. These figures vary from state to state. See <https://www.cbpp.org/>.

⁵Temporary Assistance for Needy Families

haustion of unemployment benefits and show that this behavior is inconsistent with liquidity constraints, but can be rationalized by present-bias or myopia on the part of a population of households.

This contributes to this literature by tightly characterizing how the magnitude and timing of consumption responses vary in the cross-section of liquidity. We show that this characterization extends to a variety of contexts - even when likely information innovation (as proxied by the date of tax filing) and manner of disbursement (as in the case of households receiving federal and state refunds separately) are accounted for. Additionally, we make progress towards reconciling the findings of these papers within a structural framework.

This paper also contributes to a literature concerned with the implications of excess sensitivity for aggregate fluctuations and stabilization policies. In large part, structural models since Friedman (2018) have focused on the role of liquidity constraints (Deaton, 1991; Carroll, 1997; Mankiw, 2000). Recently, Kaplan and Violante (2014) rationalize the large average consumption responses observed by Johnson, Parker, and Souleles (2006) in a two asset framework with costly liquidation. This mechanism endogenously increases the proportion of households that are liquidity-constrained, as a significant population of ‘wealthy hand-to-mouth’ agents hold their wealth in illiquid form. Informed by the empirical findings that we establish, we develop a structural model in which households with substantial liquid assets display consumption responses to predictable income in line with those of the data. We show that this higher-order cross-sectional dispersion in consumption responses has important implications for policy design.

A literature studying mental accounts and the resulting positive predictions for consumption behavior dates to Thaler (1985) and Shefrin and Thaler (1988). More recently, papers have documented behavior consistent with mental accounting frictions in the context of gasoline prices (Hastings and Shapiro, 2012), food stamps (Hastings and Shapiro, 2018), household budgeting (Koszegi and Matejka, 2018), and optimal taxation (Farhi and Gabaix, 2020). In each case, mental accounting frictions break the fungibility of otherwise interchangeable resources.

In this paper we consider transfer payments in a dynamic environment where mental accounting frictions affect household consumption/savings decisions. Shefrin and Thaler (1988) posit a life-cycle model incorporating a dual preference (planner/doer) setup and mental accounts that break the fungibility of resources across current assets and present and future income. With a similar structure of mental accounts, this paper develops a setup that directly nests the traditional buffer-stock and hand-to-mouth agents. We further discipline the model by

performing a structural estimation procedure to identify household’s preference parameters. We show that this class of models is able to jointly match life-cycle liquid asset accumulation, as well as the timing and magnitude of the household consumption responses that we observe in the data.

2. Empirical Evidence

We utilize an administrative dataset of de-identified bank transaction data obtained from a large American financial institution. We study household expenditure and balance sheet responses to predictable forms of income including state and federal tax refunds, bonus checks, and regular paychecks. We begin by describing the data and then present the main empirical results.

2.1. Data & Empirical Strategy

The dataset includes a panel of 17.2 million U.S. households, with active checking accounts from 2012 to 2019. In addition to individual transactions, we observe checking, savings, and credit card balances, as well as non-transaction account vehicles such as money market accounts, brokerage accounts, and certificates of deposit held at the bank. Throughout the paper we aggregate all accounts to the primary account holder level, and restrict my analysis to households where the primary account holder is of working age (24 to 64).

The data allows me to track deposit, debit card, and credit card inflows and outflows at the daily frequency, providing for clean identification of household expenditure responses to income. We categorize these transactions according to IRS Merchant Category Codes (MCCs) in close accordance with the NIPA Handbook.

Crucially, we take steps to ensure that we observe the primary checking account of a household by restricting to those with at least five deposit account outflows in each month of a given calendar year. Below we provide evidence that this filtering procedure largely serves to exclude households whose primary checking accounts are held at other financial institutions. This is necessary to ensure a reliable view of household’s day-to-day financial activity and to address concerns that these households break up large portions of their expenditure across multiple banks. Table 1 reports summary statistics for the broad population of households meeting the exclusion criteria.

We aggregate checking, savings, and non-transaction account balances to generate a measure of the total liquid assets available to households. The median

Table 1: Summary Statistics

		Mean	25th	Median	75th
Demographics	Age	42.2	32	41	52
	Account Users	1.4	1	1	2
Income	Total Income	5935	2273	3923	6782
	Labor Income	4022	1835	2957	4754
Balances	Total Liquid	8673	473	1835	6442
	Checking	4955	341	1255	3691
	Savings	2302	0	0	263
	Revolving Credit	920	0	0	0




Figure 1: Survey of Consumer Finances Comparison, 2016

household in the data holds around 3 weeks of monthly income in its liquid accounts, two thirds of which resides in its checking accounts. In order to assess the external validity of the results and directions of possible bias in the results that follow, we benchmark key data moments and distributions to representative national surveys of household assets, income, and expenditure. Notably, the transaction data excludes two tails of U.S. households in terms of wealth - the unbanked portion (6.5% of U.S. households in 2017⁶) as well as those employing forms of private wealth management. As such, median levels of 2016 annual post-tax income observed in the Survey of Consumer Finances (SCF) and transaction data are roughly equal (Table 2). Median checking account balances are also comparable with the survey data, however, the transaction data likely understates the total liquid assets available to households. Additional moments of these distributions, including the procedure for post-tax adjustment of the SCF, are reported in the Online Appendix.

⁶Source: FDIC

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

⁸SCPC respondents report in which of eighteen intervals their combined household income falls. Here we report the median of the interval housing each percentile.

Table 2: Benchmark Comparisons, Quantiles, 2016

Source	Annual Income			Checking			Liquid Balances ⁷		
	25th	Median	75th	25th	Median	75th	25th	Median	75th
SCF	29863	50569	85632	300	1700	5600	800	3800	16000
SCPC	26499 ⁸	54999	86499	200	1000	3500	-	-	-
BANK	31754	50568	82484	338	1251	3687	459	1796	6182

The transaction data tracks closely to the Consumer Expenditure Survey (CEX) micro-data for non-durable goods and food services (Table 3). The transaction data’s coverage of durable goods purchases is markedly lower than the external benchmark. This is due to a combination of factors making these purchases difficult to identify: a large extensive margin and installment payment structures. An additional complication, affecting all categories of expenditure, is transactions made via paper check, cash, or unobserved credit cards.

Table 3: Expenditure Comparison, Monthly Averages, 2016

Source	Total	Non-Dur.	Durables	Services	Food Svcs.	Groceries
CEX	4776	981 ⁹	634 ¹⁰	2387 ¹¹	337 ¹²	303
BANK	5348	1059	168	1252	306	220

These unclassified expenditures represent roughly a third of average monthly expenditure (Figure 2)¹³. In order to assess models of consumption and derive direct comparisons to the literature, we develop an imputation procedure to assign these unclassified to transactions to non-durables and durables. We describe this procedure below. We compare expenditure figures obtained in the transaction data to additional external benchmarks in the Online Appendix. We measure expenditure and consumption responses to predictable income at the daily frequency

⁹CEX measure includes: Food at home, laundry and cleaning, postage/stationery, apparel, motor oil/gasoline, entertainment, smoking supplies, and drugs.

¹⁰Housekeeping and other household supplies, furnishings, and equipment; reading; medical supplies; auto repairs; and vehicle purchases.

¹¹Food away from home, alcoholic beverages, transportation, insurance, education, housing services, personal services, telecommunications, and other bills.

¹²Food away from home and alcoholic beverages.

¹³We explore a subgroup for which the sum of payments to unobserved credit cards and transfers to unobserved deposit accounts is less than \$100 each calendar year. This subsample represents roughly 40% over the overall population. The results that follow are robust to this choice of subsample.




Figure 2: Monthly Household Expenditures

via the following distributed lags specification

$$c_{i,t} = \alpha_i + \lambda_t + \sum_{n=0}^N \sum_{j=t-l}^{t+L} \delta_j^n I_{i,j}^n + \epsilon_{i,t} \quad (1)$$

Where $c_{i,t}$ denotes the outcome variable of interest (ie. total expenditure, food services consumption). Intercepts α_i and λ_t are household and time specific respectively, and $I_{i,j}^n$ represents the amount of the n th inflow at lead/lag j days for household i (ie. the receipt of a second tax refund in the same calendar year).

Estimated parameters δ_j^n measures the proportional change in $c_{i,t}$ associated with an increase in inflow I^n at lag j . Identification relies on variation in calendar time t and time between each event n . We adopt the convention of the literature in referring to the estimated parameters, $\{\{\delta_j^n\}_{t-l}^{t+L}\}_{n=1}^N$, as marginal propensities to consume or expend. To control for extreme outliers we trim the top $1 \cdot 10^{-5}$ expenditure days. The results are robust to more stringent Winsorization.

2.2. Consumption Responses to Income

We consider consumption responses to various forms of predictable income including tax refunds, regular paychecks, and bonus paychecks. Notably, we find that excess sensitivity pervades the liquid wealth and income distributions. This excess sensitivity is restricted to receipt of income (rather than potential news at tax filing), is evident when a large inflow is received shortly after another (as in the case with households receiving state and federal refunds), and is evident across all sources of income. Further, consumption responses are highly front-loaded to receipt, meaning that, absent the large consumption response in the month or so following receipt, expenditure is largely smoothed.

We begin by documenting household consumption responses to tax refund receipt. According to the IRS, roughly 80% of tax filers each year receive federal refunds, and 20% make federal tax payments. The magnitude of a household's tax

refund is largely determined jointly by the IRS income tax withholding tables and worker’s withholding allowances. Workers can reduce withholdings at any time by claiming allowances. These adjustments can be claimed for any number of reasons, including changes in marital status, status of dependents, etc. Given claimed allowances employers calculate withholdings using worker’s pay frequency and wages. This, in combination with household’s other income (capital gains/losses, tax credits, etc.) determines the magnitude of their tax refund. In aggregate, the IRS processes roughly \$200 billion in tax refunds by the end of March each year, with an additional \$75 billion processed by the end of tax season in May¹⁴.

The size of an individual’s tax refunds is a source of uncertainty that is largely resolved given their prior year’s income and previous returns, and is all but certain at the date of filing (but for complications or mistakes in the individual’s return). Additionally, the exact date of refund arrival is unknown, though 94% of refunds arrive within 30 days of filing, with 97% arriving within 60 days. Due to the size of the data we restrict the analysis to households receiving tax refunds in 2014 or 2015. The resulting sample is 1.7 million distinct households observed across the two years.

Table 4: Summary Statistics, Monthly, Tax Refund Recipients

	Mean	25th	Median	75th
Liquid Balances	7279	581	1828	5699
Income	5259	2425	3868	6245
Expenditure	4949	2323	3677	5882
Tax Refund	2072	360	1120	2993

Table 7 reports summary statistics for the subpopulation observed receiving tax refunds. We report liquid balances and income as measured from the average across the 9 months to 1 month prior to the initial refund receipt in order mitigate both contamination from response to the future refund and tax activity in the previous year. This population earns median monthly income that is 1.4% lower and liquid balances that are just \$7 less than the broader population (see Table 1). Roughly 30% of those households observed receiving one refund also receive a second.¹⁵ In what follows we estimate consumption responses to the first tax refund a household receives in a calendar year.¹⁶ Together, the combination of state and federal refunds represent over 3 weeks of income on average. The refunds

¹⁴As much as 3.3% of Q1 aggregate consumption.

¹⁵The actual proportion is likely much higher. Households might receive a second refund that is not identified as such. In any case the identifications arguments still apply.

¹⁶In the Online Appendix we consider those experiencing multiple refund events

observed in the data are of similar magnitude to those reported by the IRS. For federal refunds, the IRS reports that the average refund is \$2860, versus \$2844 in transaction data. For state and local refunds these figures are \$1622 and \$1218, respectively.

Using specification (1) we estimate household balance sheet responses to tax refund receipt and plot the resulting vectors $\{\delta_j^n\}_{t-l}^{t+L}$ for each regression along with their 95 percent confidence intervals. Each regression has a sample size of 882.3 million observations. We separate the flow of tax refund dollars towards expenditures (roughly 60% over 150 days post-receipt),¹⁷ non-transaction accounts transfers (20%),¹⁸ and those that remain in household’s transaction (checking and savings) accounts (12.5%). A small proportion goes towards paying down credit card balances both at the bank and elsewhere (6%),¹⁹ with the rest going towards illiquid debt payments (student and auto loans) or being transferred to demand deposit accounts at other financial institutions.



Figure 3: Balance Sheet Response, Tax Refunds

In what follows we restrict my focus to the non-durable component of total expenditure (denote this measure $e_{\mathcal{ND}}^i$, where i indexes a particular household). To construct this measure we aggregate all household expenditures at the daily frequency that represent purchases of non-durable goods²⁰ and ser-

¹⁷Defined as all account (credit and deposit) outflows, excluding credit card balance payments for which card purchases are observable and account transfers.

¹⁸Defined as transfers to brokerage, money market, retirement and certificates of deposit

¹⁹This includes some double counting, as observable credit card expenditures assigned to the time of purchase are included in the total expenditure measure, while this panel measures excess payments towards credit card balances. We report the path of revolving balances around refund receipt in the Online Appendix

²⁰Including groceries, entertainment, fuel, discount and drug stores, direct market catalogs.

vices.²¹ Additionally, we perform an imputation procedure to assign cash, unclassified checks, and payments to unobserved credit card accounts for individual i (denoted e_c^i) to contemporaneous non-durable expenditure. Define the mean observable proportion assigned to non-durables for individuals in population q , $\xi^q \equiv \frac{1}{N} \sum_{i=1}^N \frac{e_{\mathcal{N}\mathcal{D}}^{i,q}}{e^{i,q} - e_c^{i,q}}$, where $e^{i,q}$ denotes total expenditure. The imputed non-durable consumption responses ($\mathcal{N}\mathcal{D}_{\mathcal{I}}$) for population q at lag j are then obtained via $\delta_{t-l}^{\mathcal{N}\mathcal{D}_{\mathcal{I}},q} = \delta_j^{\mathcal{N}\mathcal{D},q} + \xi^q \cdot \delta_j^{c,q}$. Expenditures used to compute the expenditure share ξ are taken from the month prior to tax refund receipt.

This procedure relies on two assumptions. First, 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, second, that the excess response of these categories at refund receipt scales proportionally. For the first assumption, 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. For the second assumption, we show in the Online Appendix that shifts in the composition of expenditure around refund receipt are small. For specifications employing the imputation procedure, we also report pre-imputation values in the Online Appendix.

We estimate household non-durable expenditure responses to tax refund receipt. Figure 3 plots the resulting parameters for the thirty days before and 150 days after income is received. This figure is illustrative of the consumption responses we observe across all categories of expenditure and forms of income. Notably, on average, anticipatory expenditure is insignificant, responses are immediate (in the present case, 4.1 cents of every dollar received are expended within a day of receipt, 19.3 cents are expended within 30 days, and 28 within 150 days), and highly front-loaded to income receipt (70% of the 150-day non-durable response to tax refunds occurs within 30 days of receipt). The response of non-durables represents 47% of the average total expenditure response over the five months post-refund receipt. These estimates are in line with average non-durable consumption responses to forms of predictable income previously reported in the literature.²²

This characterization of consumption responses is not special to tax refunds,

²¹Including utilities, telecommunications, insurance, health expenses, other bills, food services, travel services and other personal and professional services.

²²See the Online Appendix for comparisons to the literature.




Figure 4: Non-Durable Expenditure Response, Tax Refunds

where the magnitude and timing of receipt are, to a degree, driven by self-selection. We consider a population of 163,300 households receiving large bonus checks. The average bonus recipient holds over six weeks of income in their liquid accounts. Compared to the average household receiving a tax refund, the average bonus recipient pay holds roughly 74% more liquid assets and earns 56% more in monthly income.

Table 5: Summary Statistics, Monthly, Bonus Check Recipients

	Mean	25th	Median	75th
Liquid Balances	12690	2024	4930	13609
Income	8246	4207	6182	9396
Expenditure	8306	4042	6237	9626
Bonus Check	11445	3290	5733	10802

The features of the non-durable consumption response to bonus checks is similar to that of tax refunds (Figure 5). Namely, insignificant anticipatory spending and a large degree of front-loading to receipt, with 64% of the 150 day response coming in the first thirty days.




Figure 5: Non-Durable Expenditure Response, Bonus Checks

In the Online Appendix we show that households also display excess sensitivity

to their regular paychecks, as shown by Olafsson and Pagel (2018) and Gelman et al. (2014). Additionally, we show that the results above are robust to the category of expenditure considered. In combination, the observed responses to predictable forms of income (tax refunds, bonuses, regular paychecks) indicate that households excess sensitivity is neither a direct bi-product of the population receiving each inflow; nor driven by self-selection in the timing or magnitude of cash flow events, nor unpredictability in the timing of receipt. However, these responses are consistent with the predictions of models in which a large subset of each population is liquidity-constrained. We further explore this dimension in the section that follows.

2.3. Consumption Responses & Liquidity

In order to assess the impact of liquidity constraints on consumption responses we classify households according to available liquid balances relative to monthly income. Standard models predict a tight correlation between this measure of liquidity and consumption responses to income. This ratio is calculated according to average of the ratio of month end total liquid balances (checking, savings, money market, brokerage, and certificates of deposit) to total monthly income in the calendar year prior to tax refund receipt. This averaging serves to avoid short-run endogenous responses in anticipation of refund receipt significantly contaminating the ratio.²³ In order to prevent households with low incomes distorting the measure, we restrict to those with average monthly income of \$500 or more in the baseline period (roughly 98% of households in the raw sample of refund recipients).

As documented above, this measure likely understates the liquidity available to households, especially amongst higher income individuals. Additionally, the measure is in terms of gross liquid assets and not liquid net worth. Un-securitized debts are excluded for two reasons: 1.) the data covers credit card accounts held at the financial institution, and households might carry revolving debts on unobserved credit cards;²⁴ and 2.) as documented above, observable credit card debts are held by a minority of households at any given time and, while substantial, are relatively short lived.²⁵ In the Online Appendix we consider additional measures

²³In a variance decomposition of refund expenditure responses we find that once the long-run average is controlled for, the marginal explanatory power provided by the orthogonal component of income and cash balances immediately in advance of receipt is negligible.

²⁴The results are quantitatively similar after restricting to the subset of households with no payments to unobserved credit card accounts describe above.

²⁵The half life of revolved balances, on average, is less than 4 months. Additionally, a variance decomposition across observables suggests credit card revolving balances account for less than




Figure 6: Cross-Sectional Comparisons

of liquidity, deciling liquid assets in levels, deciling liquid assets relative to monthly expenditure, and by terciling households according to levels of annual income in levels before partitioning according by liquid assets to income. The results are robust to these alternative measures.

Household’s balance sheets display significant variation in $\frac{\text{Liquid Assets}}{\text{Income}}$, ranging from around a week of income available in liquid accounts at the first decile, to over 20 weeks of income amongst households in the tenth decile (Table 6). The level of average monthly income increases across the lowest deciles and flattens out above the median, while average liquid balances are roughly increasing across deciles (Figure 6). The results are robust to deciling by transaction account balances relative to income.²⁶

Table 6: Cross-Section of Liquid Balances-to-Income, Refund Recipients, Means

Decile	1	2	3	4	5	6	7	8	9	10
$\frac{\text{Trans. Assets}}{\text{Income}}$	0.11	0.20	0.28	0.38	0.50	0.65	0.96	1.45	2.55	4.91
$\frac{\text{Liquid Assets}}{\text{Income}}$	0.21	0.36	0.49	0.64	0.80	1.0	1.35	1.91	3.09	5.11
Total Income	4348	4860	5163	5420	5680	5909	6026	6088	6017	5840
Tax Refund	2049	1999	2020	2029	2096	2132	2161	2154	2110	2097

We estimate specification (1) by deciles of $\frac{\text{Liquid Assets}}{\text{Income}}$. Figure 7 plots the results. Notably, the non-durable expenditure response are significant across the liquid wealth-to-income distribution and decline moderately in levels of liquid wealth. One month MPCs ranging from 0.32 at the first decile to 0.09 at the tenth decile. Further, responses display significant front-loading to receipt, from 89% of the 120 day response coming in the first thirty days at the first decile to 66% at the tenth decile. In dollar terms, excess total expenditures across deciles

4% of explainable variation in consumption responses.

²⁶The correlation of deciles across the two measures is 0.92 and the results are virtually identical. Likewise, the results are robust to including alternate measures of income in the numerator (ie. only labor income or categorized income)

range from an additional \$1,347 to \$580 in the quarter of receipt.




Figure 7: Non-Durable Responses in the Cross-Section

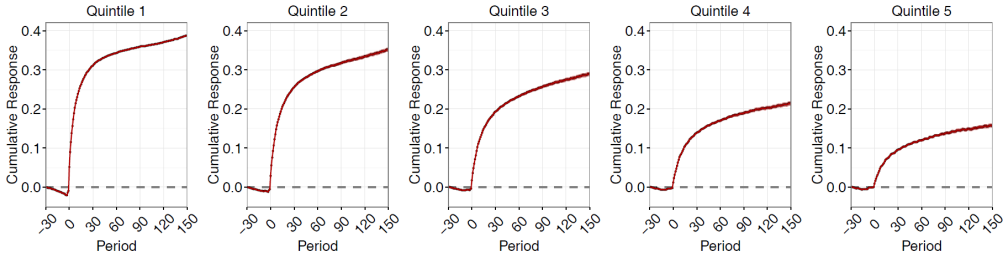


Figure 8: Dynamic Non-Durable Responses, Cross-Section

The cross-sectional responses (Figures 7 and 8) nest several salient facts that are helpful in distinguishing across models of consumption. Firstly, the large degree of excess sensitivity at receipt combined with no significant anticipatory spending that is observed amongst households with large amounts of liquidity is inconsistent with a model of externally imposed borrowing constraints. In contrast with the predictions of models of rational inattention (Reis, 2006; Gabaix, 2014), household's consumption responses are immediate and short-lived. That the bulk of the consumption response comes in the month of receipt for the unconstrained, as opposed to in advance, is in stark contrast with the predictions of models of temptation (Laibson, 1997; Gul and Pesendorfer, 2001), and forward-looking reference-dependent utility (Kőszegi and Rabin, 2006).

The results characterizing household consumption responses:

- Are robust to the form (tax refunds, bonus checks, and regular paychecks) and magnitude of the predictable income (as measured relative to baseline monthly income and monthly expenditure).
- Are not driven by the non-durable measure. The results hold for purchases of non-durable and durable goods as well as categories more aligned with contemporaneous consumption, such as grocery purchases and food services.
- Are robust to the measure of liquidity employed (We consider levels of liquid assets, $\frac{\text{Liquid Assets}}{\text{Expenditure}}$, income brackets and deciles of $\frac{\text{Liquid Assets}}{\text{Income}}$).

We highlight the robustness of the results to the form of income because of two notable features of tax refunds. One is possible endogeneity in refund size (as households can self-select these magnitudes to a degree), and another is that the exact date of refund arrival is the source of some uncertainty (however, as documented above, over 97% of refunds arrive within thirty days). In contrast, the dates of arrival for bonuses and paychecks are known or ascertainable in advance of receipt. The results are also robust across populations and consumption responses are restricted to income receipt, that is they:

- Hold across age groups, levels of income, and the cross-section of income volatility, and are robust to different measures of liquidity.
- Persist even when one large inflow comes shortly after another (in the case of households receiving state and federal tax refunds in the same year).
- Is restricted to receipt as opposed to the date of tax filing, even for households with substantial liquidity. The date of tax filing jointly represents a resolution of uncertainty and an allocation of household attention.

These results are consistent with a model in which households follow rules-of-thumb for consumption/savings and make budgeting decisions on a short-term basis. Households likely employ such rules in order to internally impose constraints on present consumption and prevent over-spending from month-to-month. Household responses to paycheck receipt illustrate this point particularly saliently. Households among the fourth and fifth quintiles of $\frac{\text{Liquid Assets}}{\text{Income}}$ lump an extra day's worth of non-durable goods expenditure into the week following paycheck receipt relative to the week before, even though they have the requisite liquidity to smooth these discretionary expenses across paychecks. Similarly, household responses to surplus cash coming in the form of tax refunds or bonus checks are immediate and

short-lived, indicating that fungibility between these inflows and a household's liquid balances only fails for a short period of time.

3. A Model of Mental Accounts

We develop a dynamic and tractable model of mental accounts that rationalizes both the timing and magnitude of consumption responses in the cross-section of liquidity documented above. Fundamentally, we break the fungibility between an agent's current income and current assets that is assumed in standard consumption models. In the mental accounts model, agents follow simplified rules of thumb for savings guided by one salient reference point: current assets. Overriding these internal processes entails utilitarian costs in the form of declines in marginal utility. This friction serves to partition the consumption choice set across 'mental accounts' pertaining to current assets and current income, in the spirit of Shefrin and Thaler (1988). We begin by presenting the mental accounts utility function, followed by a motivating exercise. We proceed by performing a structural estimation of the model's preference parameters in a life-cycle context, and conclude by exploring the resulting positive implications for stimulus policy design.

3.1. Mental Accounts Utility

The mental accounts utility function is of the form

$$\nu(c) \equiv u(c) + \psi \cdot d(a', a^d) \quad (2)$$

Where $u(\cdot)$ denotes the usual consumption utility ($u'(c) > 0$, $u''(c) < 0$) and $d(\cdot)$ a savings deviation function, accepting as arguments an agent's current savings decision (a') and some default rule-of-thumb for savings (a^d). Agents face a utility cost when deviating below their default, the intensity of which is parameterized by ψ . The following form for the savings deviation function is assumed:

$$d(a', a^d) \begin{cases} = 0 & \text{if } a' \geq a^d \\ < 0 & \text{if } a' < a^d \end{cases}$$

and $\frac{\partial d(a', a^d)}{\partial a'} \Big|_{a' < a^d} > 0$, so that in the loss domain (where $a' < a^d$) deviation losses decrease as the agent approaches the default. Agents face the usual constraint $c + a' \leq y + a(1 + r)$. It is useful to define the consumption allocation at the

default, $c^d = y + a(1 + r) - a^d$.

We impose the following structure. First, we assume agents' rule-of-thumb for savings is guided by their current asset positions, $a^d = a(1 + r)$. Second, we assume the following functional form for the deviation function

$$d(a', a^d) = \begin{cases} 0 & \text{if } a' \geq a^d \\ -(u(c) - u(c^d)) & \text{if } a' < a^d \end{cases} \quad (3)$$

This functional form (Figure 9) is continuous and induces a natural bound for $\psi \in [0, 1]$. When $\psi = 0$ then $v(c) = u(c)$ and the standard permanent income agent is recovered. At the other extreme, when $\psi = 1$, agents face marginal utility $u'(c)(1 - \psi) = 0$ when $c > y$, and are hand-to-mouth in each period ($c = y$).²⁷ For intermediate values of $\psi \in (0, 1)$ agents exhibit *dissavings aversion*.




Figure 9: Utility Representation

The savings deviation function induces a kink in the standard utility function, partitioning an agent's consumption choice set between mental accounts pertaining to current income and current assets. Agents consider the income they receive each period separately from the assets they have accumulated over time. Unconsumed resources obtained as income in period t are re-labeled as current assets in future periods. Agents face the standard utility function when consuming out of current income, and a decline in marginal utility when consuming out of their current asset balances.

²⁷In environments where income is allowed to be 0, reference points may be determined by $c^d = \max\{y, \epsilon\}$ for ϵ arbitrarily small. This prevents indeterminacy without altering the economics

3.2. Consumption Responses Under Perfect Foresight

We first study the influence of mental accounts on consumption behavior in a simple life-cycle model of consumption and savings. Households live for T months, they face no uncertainty, and have perfect foresight over the stream of income they will receive. The household is endowed with initial assets a_0 , and income $\{y_t\}_{t=0}^T$. At time $t = 0$, the household learns that they will receive a one-time increase in income in period h : $y_h + \Delta$. The household decision problem is:

$$\begin{aligned} \max_{c_t, a_{t+1}} \sum_{t=0}^{T-1} \beta^t & \left[\frac{c_t^{1-1/\gamma}}{1-1/\gamma} + \psi d(a_{t+1}, a_t(1+r)) \right] \\ c_t + a_{t+1} &= y_t + a_t(1+r) \\ a_T &= 0 \end{aligned}$$

where c_t are consumption choices, a_{t+1} are asset choices, r is the net interest rate, β is the discount factor, the deviation function $d(\cdot, \cdot)$ is from Equation (3), and flow utility is due to the CRRA function with intertemporal elasticity of substitution given by γ .

The first order conditions yield the following Euler equations:

$$(1 - \psi \mathbb{1}_{c_t > y_t}) c_t^{-1/\gamma} = \beta(1+r)(1 - \psi \mathbb{1}_{c_{t+1} > y_{t+1}}) c_{t+1}^{-1/\gamma} \quad (4)$$

where $\mathbb{1}_{c_t > y_t}$ is an indicator function describing whether the household is dissaving in period t and therefore whether the savings deviation function is triggered. Dissaving in period t reduces the marginal utility of consumption, which tends to increase saving and shift consumption from period t to period $t + 1$.

In Online Appendix B.1 we show that combining the Euler equations (4) with the intertemporal budget constraint yields the consumption function:

$$c_t = \frac{(\beta(1+r))^{\gamma t} (1 - \psi \mathbb{1}_{c_t > y_t})^\gamma}{\sum_{s=0}^{T-1} \beta^{\gamma s} (1+r)^{(\gamma-1)s} (1 - \psi \mathbb{1}_{c_s > y_s})^\gamma} \left(a_0(1+r) + \sum_{s=0}^{T-1} \frac{y_s}{(1+r)^s} \right) \quad (5)$$

Note that the consumption function is not everywhere continuously differentiable with respect income. This is because changes in income may lead households to consume out of current assets thereby triggering dissaving aversion.

Model parameters and the path of income determine whether the household is dissaving at any particular date t , and thus whether the mental accounts mechanism is triggered. So to compute consumption responses to income receipts, we

need to make assumptions about these model primitives. In Online Appendix B.1 we consider two special cases. First, we consider the simple case where $c_t = y_t$ for all t in the absence of the announced income inflow. Second, we consider the case where the initial path of income follows a hump-shaped life-cycle profile. To ease notation and intuition, we consider the first special case here.

Assume that absent the income inflow, model parameters and the path of income $\{y_t\}_{t=0}^T$ are such that $c_t = y_t$ for all t . When the inflow is announced at date 0 to be delivered at date h , the consumption smoothing motive induces households to save at h so that $c_h < y_h + \Delta$. In addition, households increase their consumption in every other period so that $c_t > y_t$ for all $t \neq h$. Thus, dissavings aversion is triggered in all periods $t \neq h$. Along the consumption path with the income inflow, Equation (5) becomes:

$$c_t = \frac{(1 - \theta)(\beta(1 + r))^{\gamma t}(1 - \psi \mathbb{1}_{t \neq h})^\gamma}{(1 - (1 - \psi)^\gamma)(1 - \theta)\theta^h + (1 - \psi)^\gamma(1 - \theta^T)} \times W$$

where $\theta = \beta^\gamma(1 + r)^{\gamma-1}$ and $W = \left(a_0(1 + r) + \sum_{s=0}^{T-1} \frac{y_s}{(1+r)^s} \right)$ is the present discounted value of life-time wealth.

We then compute the marginal propensity to consume (MPC) at t out of income inflow Δ at h as:

$$\begin{aligned} MPC_t^h &\equiv \lim_{\Delta \rightarrow 0^+} \frac{c_t(y_h + \Delta) - c_t(y_h)}{\Delta} \\ &= \frac{(1 - \theta)\beta^{\gamma t}(1 + r)^{\gamma t - h}(1 - \psi \mathbb{1}_{t \neq h})^\gamma}{(1 - (1 - \psi)^\gamma)(1 - \theta)\theta^h + (1 - \psi)^\gamma(1 - \theta^T)} \end{aligned} \quad (6)$$

We construct the MPC using the consumption function as the inflow Δ approaches zero from above. This reflects our study of anticipated income tax refunds, but it is also consistent with our argument that the inflow causes households to consume less than current income in period h and more than current income in all periods $t \neq h$.

Under our assumptions, the simple model highlights the role that the mental accounts mechanism plays in MPCs out of anticipated income receipts. In Equation (6), the dissaving aversion parameter ψ dictates the degree of consumption front-loading relative to the date of income receipt h . In the extreme case where $\psi = 1$, the MPC is:

$$MPC_t^h = \begin{cases} 1 & \text{if } t = h \\ 0 & \text{if } t \neq h \end{cases}$$

The household acts like a hand-to-mouth consumer at the date of income receipt ($t = h$), but is unresponsive to anticipated income receipts ($t < h$) and spends nothing out of past income receipts ($t > h$).

At the other extreme, $\psi = 0$ and the household faces no aversion to dissaving. The MPC becomes:

$$MPC_t^h = \frac{(1 - \theta)\beta^{\gamma t}(1 + r)^{\gamma t - h}}{(1 - \theta^T)}$$

Under log-utility, the consumption response at $t = 0$ is

$$MPC_0^h = \frac{(1 - \beta)}{(1 - \theta^T)}(1 + r)^{-h}$$

which is the consumption response under the permanent income model, and is a function of the familiar annuity factor.

Thus the mental accounts model nests both standard consumption smoothing behavior as well as hand-to-mouth behavior as edge cases with respect to the dissaving aversion parameter ψ .

Note the independent roles of β and ψ in this environment. At their respective extremes ($\psi = 1, \beta = 0$) these parameters induce the same path of life-time consumption ($\frac{\partial c_0}{\partial y_0} = 1$). On the intermediate interval, however, the paths they induce are distinct. Mental accounts have an asymmetric effect on the path of consumption (as unconsumed resources move from the current income account in the initial period to the current asset account in all future periods). In contrast, the time preference parameter induces a path of consumption that decays geometrically (ie. $\frac{\partial c_{t+1}}{\partial y_0} = \frac{\partial c_t}{\partial y_0} [\beta(1 + r)]^{\frac{1}{\gamma}}$).

Given a time horizon T , a value for the risk aversion parameter, γ , and the empirically observed relative consumption responses ($\frac{\partial c_0}{\partial y_0}, \frac{\partial c_{t+1}}{\partial y_0} / \frac{\partial c_t}{\partial y_0}$ for all $t > 0$) it is straightforward to back out the remaining preference parameters, β and ψ . For this back-of-the-envelope calculation we use the average consumption responses of the highest cash-on-hand quintile receiving tax refunds, since this population's consumption decisions are the least likely to be distorted by their proximity to an externally imposed liquidity constraint. Amongst this group $\frac{\partial c_4}{\partial y_0} / \frac{\partial c_3}{\partial y_0} \approx 0.986$. Choosing $(1 + r) \approx 1$ (a reasonable assumption at the monthly frequency) and $\gamma = 1$ (log utility), it follows that $\beta = 0.986$. For the same group of households $\frac{\partial c_0}{\partial y_0} \approx 0.095$. Plugging the recovered value of β into this expression and solving for ψ yields a dissaving aversion factor of 0.867.

This stylized economy is illustrative of how household's time preference and




Figure 10: Mental Accounts Consumption Response

dissaving aversion parameters can be disentangled. Figure 10 depicts the consumption responses generated by the mental accounts model alongside those generated by the data.

4. A Quantitative Life-Cycle Model With Mental Accounts

We now embed the mental accounts framework in an otherwise standard quantitative life-cycle model. The goal of this exercise is to assess ability of the model to match our empirical results in Section 2. Specifically, we task the model with explaining the strong response of consumption to anticipated income receipts across the household distribution of liquid asset holdings.

It is possible for standard life-cycle models to generate large MPCs if most households are close to their borrowing constraints. However, this implies a counterfactually low rate of savings, especially among the wealthiest households. For this reason, we also require that our model matches the life-cycle profile of liquid asset holdings. In doing so, our model produces both low- and high-wealth households as well as the high rates of consumption spending we observe across the distribution of wealth.

We compare our framework to the benchmark life-cycle model in the absence of mental accounts frictions. We first study a model calibrated to match both the life-cycle savings profile and our estimated cross-sectional consumption responses. We then calibrate the same model to match consumption responses only. We show that this model produces a counterfactually low rate of life-cycle savings.

4.1. Setup

The idiosyncratic state variables for a household are $\mathbf{s} = [t, a, y, e, i]$, where t is age, a is liquid assets, y is an idiosyncratic productivity shock, e is employment status, and i is the date of an anticipated inflow of funds.

Households live for T periods, where a model period corresponds to one month. They earn uncertain labor income during working life, and enter retirement after age \tilde{T} . Households save in the liquid asset a and earn returns given by the risk-free interest rate r . They cannot borrow, and face a constraint on the choice of savings: $a' \geq 0$.

Income is a function of age, productivity, and employment status. At each age t , income is proportional to a deterministic life-cycle component of income given by Γ_t . During working life, idiosyncratic income productivity evolves according to a log-AR(1) with persistence ρ_y and standard deviation of shocks given by σ_y . Also during working life, employment status follows a two-state Markov chain, where the probability of becoming unemployed (i.e. the separation rate) is π_u , and the probability of finding a job is π_e . Unemployed households receive unemployment insurance given by a simple replacement rate ω_u relative to their regular income. Upon retirement at age $\tilde{T}+1$, households continue to earn the life-cycle component of income but no longer face uncertainty over employment or productivity. From retirement until the end of life, idiosyncratic productivity is held constant at its value in the last period of working life $y_{\tilde{T}}$. Thus, the income function is given by

$$m_t(y_t, e_t) = \begin{cases} (\mathbb{1}_{\{e=1\}} + \mathbb{1}_{\{e=0\}}\omega_u) y_t \Gamma_t & \text{if } t \leq \tilde{T} \\ y_{\tilde{T}} \Gamma_t & \text{if } t > \tilde{T} \end{cases}$$

Households may be alerted h periods in advance to a future inflow of funds. The state variable i keeps track of the number of months until inflow receipt. For $i < 0$, the household anticipates an inflow at date $t-i$. When $i = 0$, the household receives the inflow in the current period. And for $i = 1$, the household no longer anticipates any future inflows. The inflow itself is given by $\Delta m_t(y, e)$, which is proportional to their usual income. Thus the inflow mimics the receipt of tax refunds, as discussed in Section 2.

The household decision problem is described by the value function:

$$\begin{aligned}
V_t(a, y, e, i) = \max_{c, a'} & \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \psi d(a', a(1+r)) + \beta \hat{\beta}_t \mathbb{E} [V_{t+1}(a', y', e', i')] \right\} \\
\text{s.t. } & c + a' = m_t(y, e) + a(1+r) \\
& i' = \max\{i + 1, 1\} \\
& a' \geq 0
\end{aligned}$$

where γ is the risk aversion parameter, $d(\cdot, \cdot)$ is the deviation function described in Equation (??), β is the common discount factor, and $\hat{\beta}_t$ is a discount factor correction term that adjusts for the number of adults and children in a household at each age t (see Attanasio et al., 1999; Cagetti, 2003b). Finally, households enjoy a warm-glow bequest motive that governs the desirability of wealth holdings late in life (see De Nardi, 2004). This bequest function pins down the terminal value function, $V_{T+1}(a) = \kappa \frac{a^{1-\gamma}}{1-\gamma}$.

4.2. Calibration

The income process ($\{\rho, \sigma_\epsilon^2, \{\Gamma_t\}_{t=0}^T\}$) is determined via a first-stage estimation (described in the Online Appendix). Initial assets, a_0 , are chosen to match quintiles of the liquid asset holdings amongst 24 year old SCF respondents. We 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). These moments are plotted in Figure 11.




Figure 11: Data moments

Table 6 summarizes the model's parameters. Standard arguments for parameter identification apply for determination of the time preference and risk aversion

Table 7: Model Parameters

Parameter	Symbol	Value	Source
<i>Panel (a): Primitives</i>			
Real return on liquid assets (%)	r	0.0075	OECD, 2000-2016
Fraction receiving endowment	π_a	0.9045	SCF, 1998–2019
Endowment as fraction of income	λ_a	0.7163	SCF, 1998–2019
Demographic correction	$\{\beta_j\}_{j=1}^J$.	Cagetti (2003a)
<i>Panel (b): Income</i>			
Deterministic income	$\{\Gamma_j\}_{j=1}^J$.	SCF, 1998–2019
AR(1) income, persistence	ρ_y	0.8962	Bank data
AR(1) income, std. dev. shocks	σ_y	0.0730	Bank data
UI replacement rate	ω_u	0.5000	DOL, 2016
Job finding rate	π_e	0.2444	CPS, 1999-2019
Job separation rate	π_u	0.0130	CPS, 1999-2019
<i>Panel (c): Preferences</i>			
Discount factor	β	0.9400	Internal calibration
Risk aversion	γ	2.0000	Internal calibration
Dissaving aversion	ψ	0.0000	Internal calibration
Weight on retirement utility	κ	5	Internal calibration

Notes: Return on liquid assets reported at annualized percentage rate. Discount factor reported at annualized rate. Job finding and separation rates reported as monthly probabilities. AR(1) income persistence and shocks estimated on monthly Bank data and reported at monthly frequency.

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 governs the consumption response of households with high levels of liquidity and the degree of consumption front-loading to the period of income receipt.

4.3. Preference Parameter Estimation

In order to structurally estimate the model’s four preference parameters $\{\beta, \gamma, \psi, \kappa\}$ we utilize a simulated method of moments approach and estimate the model’s preference parameters to fit the moments described in Figure 11. 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:²⁸

$$\min_{\beta, \gamma, \psi, \kappa} \Theta \sum_a^8 |d_a^{liq} - m_a^{liq}(\beta, \gamma, \psi, \kappa)| + (1 - \Theta) \sum_j^{10} |d_j^{mpc} - m_j^{mpc}(\beta, \gamma, \psi, \kappa)|. \quad (8)$$

This objective function includes a life-cycle liquid assets component (liq) and a cross-sectional consumption responses component (mpc), and the median absolute distance between the data (d) and model-generated (m) moments of each. 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.

In the first component a denotes 5-year age groups. In the case of the second component, j indexes deciles of the liquid wealth distribution in advance of income receipt. Model-generated consumption responses are obtained non-parametrically from 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. We search across the four-dimensional parameter space via a Sobol sequence, solving the model for each set of generated parameters, and simulating the lifetimes of forty thousand agents to obtain the model-generated moments. We choose the set of parameters that minimize the objective (6).

5. Estimation Results & Life-Cycle Savings Decomposition

The results suggest moderate levels of impatience and risk aversion (Table 8). Dissaving aversion is found to be significant (0.346). Average one month consumption responses observed in the data and those generated by the model are 0.2 and

²⁸As suggested by Carroll (2011), we also consider the following objective, which allows for SCF measurement error:

$$\min_{\beta, \gamma, \psi, \kappa} \Theta \sum_i^N \omega_i^a |d_{i,a}^{liq} - m_a^{liq}(\beta, \gamma, \psi, \kappa)| + (1 - \Theta) \sum_j^{10} |d_j^{mpc} - m_j^{mpc}(\beta, \gamma, \psi, \kappa)|. \quad (7)$$

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.

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 captures the degree of consumption front-loading to receipt.

Table 8: 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	.

We re-estimate the model for the standard buffer-stock case (restricting $\psi = 0$), 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.

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.

The standard buffer-stock model is unable to obtain a significant degree of excess sensitivity amongst highly liquid households. It 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.

All else equal, the introduction of mental accounts push estimates for both the patience and risk aversion parameters upwards relative to the baseline buffer-stock case. By shutting down income uncertainty, relaxing the ad-hoc constraint on borrowing, and turning off the mental accounting friction in turn, we decompose the extent to which household savings decisions are due to precautionary, life-cycle, and mental accounting motives, respectively. 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-

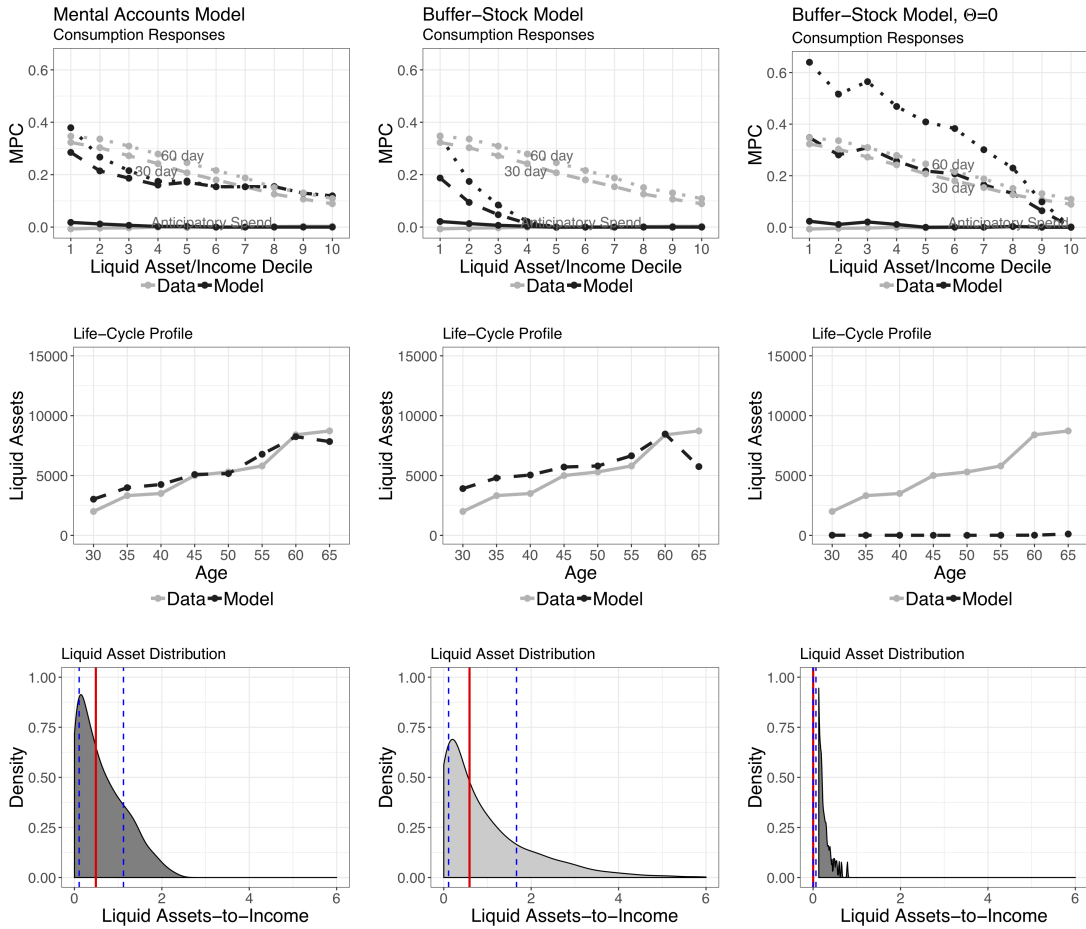


Figure 12: Model Comparison

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. Consistent with previous studies, we find that household savings decisions are driven by precautionary motives until households approach retirement age. This cash-flow management behavior -households budgeting month-to-month *as if* they are subject to the restrictions of mental accounts, might explain the prevalence of liquidity-constrained households that is observed empirically.




Figure 13: Life-Cycle Savings Decomposition

6. Implications for Fiscal Stimulus Policies

In this section we assess the model’s positive implications for the design of fiscal stimulus payments via two experiments. In the first we consider a redistributive policy in which a fiscal authority taxes households with substantial liquid wealth, and uses the revenue to provide a targeted stimulus to constrained households in the economy. In the second, we endow the 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.

We assume a partial equilibrium, overlapping generations setup. Agents take the rate of return, r , as given. We parameterize economies using the preference parameter estimates from above (Table 8). Index each generation of agents by the period of birth j . Each generation begins working life with initial assets distributed to match the quintiles of liquid asset holdings amongst 24 year olds in the SCF. Agents live from $t = j$ to $j + 660$ periods and retire at $t = j + 480$. In each case, 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.

6.1. A Redistributive Stimulus

In this section we 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. 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 we 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.

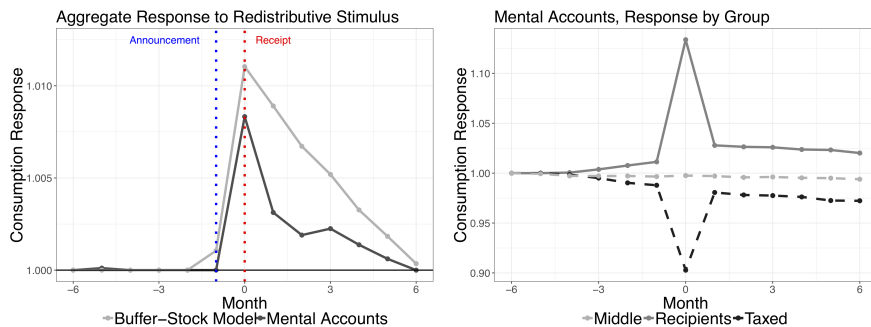


Figure 14: A Redistributive Stimulus

Compared to the standard buffer-stock case, the redistributive stimulus policy is 53% less effective over two quarters under mental accounts (Figure 15, 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 12). 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 15, right panel). This downward adjustment offsets the high degree of excess sensitivity among constrained households receiving a transfer. Ganong and Noel (2019) provide some suggestive evidence to this effect - they show that, regardless of their asset level, households cut expenditure due to a predictable shock to income as with the exhaustion of unemployment benefits.

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.

6.2. 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. Consider a government endowed with M resources. It seeks to distribute these resources to households via lump-sum transfer, facing budget constraint $M = \sum_t \int_i T_t^i \cdot \Gamma_t^i di$. Here we compare three distinct budget-equivalent policies²⁹:

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 - just 47 percent and 33 percent more effective under income-targeting and asset-targeting policies, respectively. 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.

²⁹Assuming 100 million U.S. households, each policy requires $M = \$10$ billion

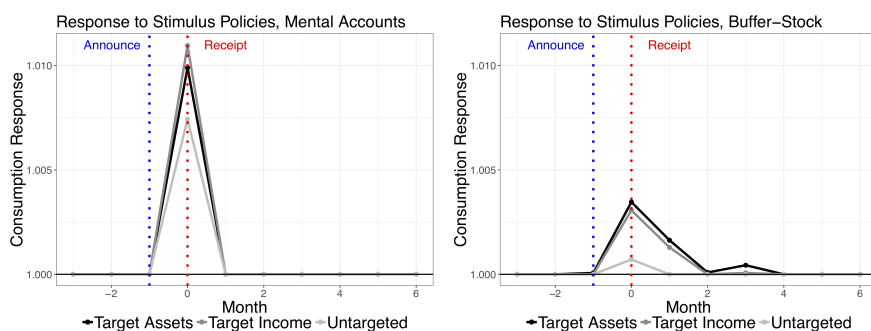


Figure 15: Aggregate Consumption Responses Across Policies by Model

These results have positive implications for policy. 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 should be weighted accordingly. Additionally, 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.

7. Conclusion

In this paper we document the relationship in micro-data between household liquidity and consumption responses to predictable forms of income. We document significant responses amongst households with high levels of liquidity that are highly front-loaded to receipt. In order to rationalize these dimensions of consumption responses we propose a model of mental accounts in which households are averse to dissaving. We show that the model nests standard buffer-stock and hand-to-mouth agents as limiting cases and that the data is consistent with an intermediate case, in which households are moderately averse to dissaving.

The model represents a gentle, parsimonious, and tractable departure from full rationality, and is able to generate consumption responses in line with the data. Through the lens of this model we re-evaluate fiscal stimulus policy design. Notably, we show that a redistributive stimulus to liquidity constrained households is approximately 50% less effective relative to a standard buffer-stock economy. Additionally, we show that the gains to moving from a blanket stimulus policy to a targeted one are significantly less than indicated by the canonical one asset setup, and policies targeted toward households with temporarily low income are

particularly stimulative. The results suggest a re-assessment of business cycle theory through the lens of mental accounting frictions.

References

- Attanasio, Orazio P, James Banks, Costas Meghir, and Guglielmo Weber, “Humps and bumps in lifetime consumption”, *Journal of Business & Economic Statistics* 17 (1999), 22–35.
- Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A Parker, “Asymmetric consumption smoothing”, *American Economic Review* 111 (2021), 192–230.
- Cagetti, Marco, “Wealth accumulation over the life cycle and precautionary savings”, *Journal of Business & Economic Statistics* 21 (2003), 339–353.
- *Wealth accumulation over the life cycle and precautionary savings [dataset]*, 2023-06-21, Data retrieved from <http://www.cagetti.com/Research/Precsavjbes.zip>. Archived at <http://web.archive.org/>, 2003.
- Campbell, John Y and N Gregory Mankiw, “Permanent income, current income, and consumption”, *Journal of Business & Economic Statistics* 8 (1990), 265–279.
- Carroll, Christopher D, “Buffer-stock saving and the life cycle/permanent income hypothesis”, *The Quarterly journal of economics* 112 (1997), 1–55.
- “Solution methods for microeconomic dynamic stochastic optimization problems”, *World Wide Web*: <http://www.econ.jhu.edu/people/ccarroll/solvingmicrodsops.pdf>. (accessed August 30, 2011). Cited on (2011), 28.
- De Nardi, Mariacristina, “Wealth inequality and intergenerational links”, *The Review of Economic Studies* 71 (2004), 743–768.
- Deaton, Angus, “Saving and liquidity constraints”, *Econometrica* 59 (1991), 1221–1248.
- Farhi, Emmanuel and Xavier Gabaix, “Optimal taxation with behavioral agents”, *American Economic Review* 110 (2020), 298–336.
- Friedman, Milton, *Theory of the consumption function*, Princeton university press, 2018.
- Gabaix, Xavier, “A sparsity-based model of bounded rationality”, *The Quarterly Journal of Economics* 129 (2014), 1661–1710.
- Ganong, Peter and Pascal Noel, “Consumer spending during unemployment: Positive and normative implications”, *American economic review* 109 (2019), 2383–2424.
- Gelman, Michael, Shachar Kariv, Matthew D Shapiro, Dan Silverman, and Steven Tadelis, “Harnessing naturally occurring data to measure the response of spending to income”, *Science* 345 (2014), 212–215.

- Gul, Faruk and Wolfgang Pesendorfer, “Temptation and self-control”, *Econometrica* 69 (2001), 1403–1435.
- Hall, Robert E, “Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence”, *Journal of political economy* 86 (1978), 971–987.
- Hastings, Justine and Jesse M Shapiro, “How are SNAP benefits spent? Evidence from a retail panel”, *American Economic Review* 108 (2018), 3493–3540.
- *Mental accounting and consumer choice: Evidence from commodity price shocks*, tech. rep., National Bureau of Economic Research, 2012.
- Johnson, David S, Jonathan A Parker, and Nicholas S Souleles, “Household expenditure and the income tax rebates of 2001”, *American Economic Review* 96 (2006), 1589–1610.
- Kaplan, Greg, Benjamin Moll, and Giovanni L Violante, “Monetary policy according to HANK”, *American Economic Review* 108 (2018), 697–743.
- Kaplan, Greg and Giovanni L Violante, “A model of the consumption response to fiscal stimulus payments”, *Econometrica* 82 (2014), 1199–1239.
- Koszegi, Botond and Filip Matejka, “An attention-based theory of mental accounting”, *CERGE-EI mimeo* (2018).
- Kőszegi, Botond and Matthew Rabin, “A model of reference-dependent preferences”, *The Quarterly Journal of Economics* 121 (2006), 1133–1165.
- Kueng, Lorenz, “Excess sensitivity of high-income consumers”, *The Quarterly Journal of Economics* 133 (2018), 1693–1751.
- Laibson, David, “Golden eggs and hyperbolic discounting”, *The Quarterly Journal of Economics* 112 (1997), 443–478.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman, *Estimating discount functions with consumption choices over the lifecycle*, 2007.
- Mankiw, N Gregory, “The savers-spenders theory of fiscal policy”, *American economic review* 90 (2000), 120–125.
- Olafsson, Arna and Michaela Pagel, “The liquid hand-to-mouth: Evidence from personal finance management software”, *The Review of Financial Studies* 31 (2018), 4398–4446.
- Parker, Jonathan A, Nicholas S Souleles, David S Johnson, and Robert McClelland, “Consumer spending and the economic stimulus payments of 2008”, *American Economic Review* 103 (2013), 2530–53.
- Reis, Ricardo, “Inattentive consumers”, *Journal of monetary Economics* 53 (2006), 1761–1800.

- Shefrin, Hersh M and Richard H Thaler, “The behavioral life-cycle hypothesis”, *Economic inquiry* 26 (1988), 609–643.
- Thaler, Richard, “Mental accounting and consumer choice”, *Marketing science* 4 (1985), 199–214.
- Zeldes, Stephen P, “Consumption and liquidity constraints: an empirical investigation”, *Journal of political economy* 97 (1989), 305–346.