Transitory Shocks and Consumption Dynamics: Pent-Up Demand or Hand-to-Mouth?

Martin Brown, Mohamed Hamoud, Jan Toczynski

Working Paper 26.01

This discussion paper series represents research work-in-progress and is distributed with the intention to foster discussion. The views herein solely represent those of the authors. No research paper in this series implies agreement by the Study Center Gerzensee and the Swiss National Bank, nor does it imply the policy views, nor potential policy of those institutions.

Transitory Shocks and Consumption Dynamics: Pent-Up Demand or Hand-to-Mouth? *

Martin Brown Mohamed Hamoud Jan Toczynski

This Version: May 2026

Abstract

We examine the recovery of U.S. consumer spending following the COVID-19 recession using monthly transaction data for over 34,000 individuals. We find limited evidence that cutbacks in spending on in-person services or durables during the recession generated pent-up demand. Instead, spending dynamics are largely consistent with hand-to-mouth behavior: expenditures during and after the recession are correlated with income changes, even more so among households with low liquidity buffers and lower income levels. These results underscore the central role of income and liquidity, rather than sentiment or forced savings, in driving the post-recession dynamics of consumer spending.

Keywords: Consumer spending, Transitory income shocks, Hand-to-mouth behavior, Pent-up demand, COVID recession

JEL Codes: D12, E21, E32.

*Martin Brown; Study Center Gerzensee and University of St.Gallen martin.brown@szgerzensee.ch. Mohamed Hamoud; University of Zürich and Swiss Finance Institute, mohamed.hamoud@df.uzh.ch. Jan Toczynski; Queen Mary University London, j.toczynski@qmul.ac.uk. We thank Michael Barczay, Diana Bonfim, Edouard Challe, Philip Coyle, Dirk Krueger, Felix Kübler, Winfried Königer, and Fabio Trojani for helpful comments.

1. Introduction

What drives the response of consumer spending to transitory shocks? This question has gained renewed relevance in light of the COVID-19 recession, during which households faced multiple significant, but short-lived shocks, including income disruptions, public health restrictions, and heightened uncertainty. The resulting variation in shocks across households provides a unique opportunity to revisit prominent hypotheses about the dynamics of consumption and the underlying mechanisms that govern household spending.

Two leading narratives have been advanced to explain consumption dynamics in the face of temporary shocks. The first centers on the idea of *pent-up demand* (Beraja and Wolf, 2021). In this view, consumers delay purchases, particularly of durable goods or luxuries, when confronted with uncertainty or direct restrictions, such as lockdowns. They accumulate excess savings during these periods and subsequently "catch up" on consumption once confidence is restored or constraints are lifted.

The second narrative emphasizes *hand-to-mouth behavior*. Here, consumers hold persistently low liquidity buffers and exhibit a high marginal propensity to consume (MPC) out of transitory income (Aguiar et al., 2024). In such a framework, spending falls with income and only recovers as income rebounds, especially for those at the lower end of the income distribution or with weak liquidity buffers. This view is the empirical backbone of Heterogeneous-Agent New Keynesian (HANK) models, in which aggregate consumption responses to transitory shocks are governed primarily by the spending of hand-to-mouth households rather than by intertemporal substitution among the unconstrained (Kaplan et al., 2018, Auclert, 2019). A key insight of this literature is that hand-to-mouth households include not only the asset-poor but also a sizable share of "wealthy hand-to-mouth" households who hold substantial illiquid wealth but little liquidity, and that both types display large marginal propensities to consume out of transitory income (Kaplan and Violante, 2014).

The COVID-19 recession provides a unique empirical setting to assess the relevance of these competing mechanisms underlying consumption dynamics. The pandemic generated multiple transitory shocks, including stringent public health interventions, sharp declines in employment and income, and abrupt shifts in consumer confidence (see Figure 1). The intensity of these shocks varied considerably across households, creating rich heterogeneity that can be exploited to identify the drivers of consumption responses. Large-scale fiscal transfers offer additional variation to evaluate how income dynamics shaped spending dynamics.

Figure 1: The COVID Recession: Employment, Spending, Saving and Sentiment

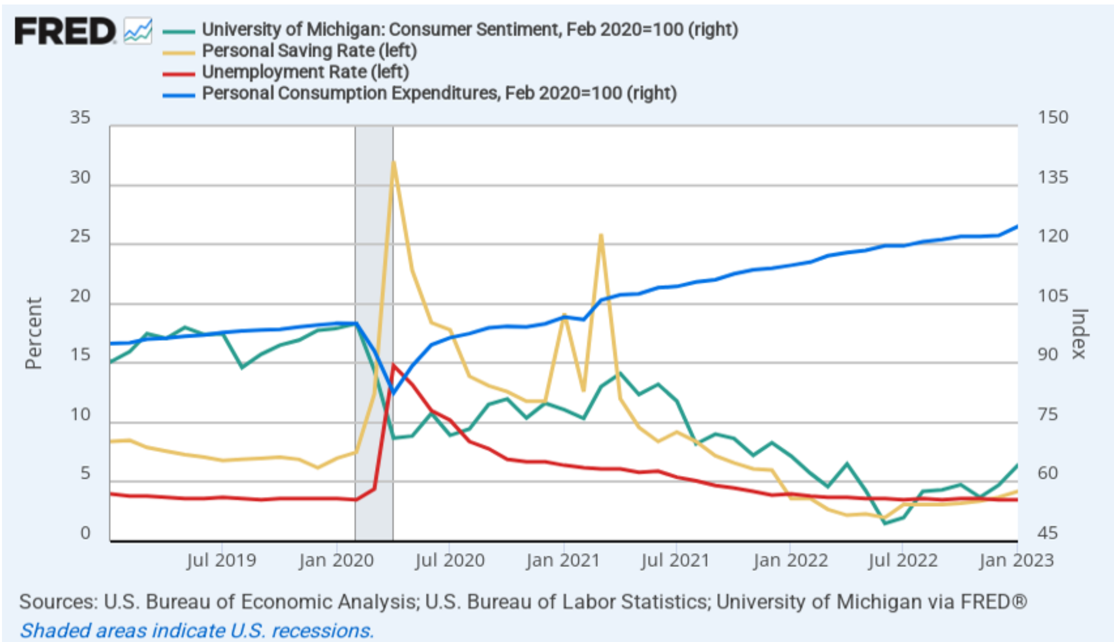

In this paper, we use transaction-level data on 34,453 U.S. consumers from April 2019 to October 2021 to study the dynamics of consumption during and after the COVID recession. We aggregate spending at the consumer × category × month level and relate this to individual indicators of income shocks, liquidity positions, local public health measures and sociodemographic information. This allows us to relate both the magnitude of the initial spending shock and its decomposition, as well as income dynamics to the post-recession recovery in total spending at the consumer level.

We report seven cross-sectional observations on spending dynamics through the COVID recession. Taken together, these observations build a narrative that income dynamics, rather than pent-up demand govern, consumer spending dynamics. First, a majority of consumers experience a sharp spending decline at the onset of the COVID pandemic. Indeed, in April 2020 59% of our sample experienced a decline in total spending of at least 10% compared to

April 2019. Second, for a majority of consumers, memorable services account for a significant share of the spending shock. Third, consumers with larger spending shocks do not build up larger liquidity reserves. Fourth, spending recovers faster for consumers with weaker spending shocks. Fifth, consumers with larger spending shocks also suffer larger income shocks. Sixth, income recovers faster for consumers with weaker income shocks. Seventh, the recovery of spending is correlated with the recovery of income.

We substantiate the findings of our descriptive analysis with a cross-sectional local-projection analysis in which we examine the recovery of total consumer spending at 3-18 months following the onset of the COVID pandemic. Our estimates strongly favor the hand-to-mouth interpretation of post-recession spending dynamics over the pent-up demand narrative. The magnitude and composition of the initial spending shock—the variables most relevant to pent-up demand—predict slower, not faster, recovery. By contrast, the cumulative recovery of disposable income is the dominant predictor of spending recovery, its importance grows over the recovery horizon, and its inclusion improves the model's ability to account for cross-sectional variation. Heterogeneity analysis reveals that the role of income recovery is particularly strong for ex-ante lower income and more liquidity constrained consumers. This heterogeneity is consistent with (Weidner et al., 2014), who emphasize that liquidity, rather than net worth alone, identifies high-MPC households.

These findings are consistent with a narrative that a large share of consumers are liquidity-constrained and their spending paths are paced by the dynamics of their income, including fiscal transfers, rather than by the release of stored liquidity from pent-up demand for services and durable goods.

Our contribution is twofold. First, we provide new evidence on the dynamics of consumption responses to transitory shocks using a high-frequency dataset of category-level spending that spans a large and heterogeneous population of consumers. This advances the literature on the cyclical behavior of consumer spending (Beraja and Wolf, 2021, Browning and Crossley, 2009, Kaplan and Violante, 2014) and contributes to growing evidence examining business cycle and spending dynamics during the pandemic (Baker et al., 2020) (Goolsbee and Syverson, 2021, Chetty et al., 2023, Stock and Watson, 2025).

Second, we contribute to the literature on intertemporal spillovers in consumption preferences. Our findings on memorable in-person services inform models of habit formation or memorable consumption (Meghir and Weber, 1996, Gilboa et al., 2016), where certain categories of consumption carry utility that exhibits spillovers between periods. Our analysis of durables and memorables sheds new light on the extent to which goods and services are subject to delay and catch-up dynamics (Dynan, 2000, Hai et al., 2020).

Our results confirm the central role of household-level liquidity constraints in shaping private consumption, even in a pandemic-driven recession. This finding is consistent with Stock and Watson (2025), who conclude that conventional macro dynamics persisted in the COVID recession but were masked by a massive pandemic-specific shock. Our findings are relevant to the design and targeting of fiscal policy. The dynamics of spending we document at the consumer level underscore the importance of accounting for liquidity constraints when designing stimulus policies. Our findings also complement studies that estimate MPCs out of specific COVID-era stimulus payments (Karger and Rajan, 2020, Coibion et al., 2020, Parker et al., 2022, Toczynski, 2023, Ganong et al., 2024). This literature documents large average MPCs out of pandemic transfers and pronounced heterogeneity in spending responses, with liquidity-constrained households exhibiting the highest propensities to consume. Rather than estimating the response to a specific transfer, we trace consumer-level spending dynamics over the full recession-recovery cycle, allowing us to weigh the relative importance of income flows and accumulated liquidity in shaping the recovery. We show that a similar logic operates over a full recession-recovery cycle with large income disruptions and stimulus transfers.

The remainder of the paper is organized as follows. Section 2 develops a stylized theoretical framework that contrasts various consumer types and derives testable predictions. Section 3 describes the transaction-level data and the construction of our key variables for spending, income, and liquidity. Section 4 documents seven cross-sectional observations on spending dynamics through the COVID recession, organized around the predictions of the theoretical framework. Section 5 presents the regression methodology and results. Section 6 concludes.

2. Theoretical Framework

In this section, we present a stylized model that serves as a conceptual framework for organizing our empirical hypotheses. We build on models of intertemporal choice featuring multiple categories of goods. On the one hand, we relate to models featuring non-durable vs. durable goods in which the consumption benefits of durables are spread out over future periods after the purchase (Browning and Crossley, 2009).[1] On the other hand, we relate to models that emphasize the memorable nature of some non-durable services, e.g., eating out, entertainment, travel. (Gilboa et al., 2016, Hai et al., 2020). Common to both model types are the intertemporal spillovers of utility following the purchase of one type of good (durables, memorables) which impacts its spending dynamics compared to standard non-durable goods and services with time-separable utility.

Note that many goods which we label as "memorables" are also typically categorized as luxury goods (travel, eating out, entertainment). Compared to necessities, luxury goods feature more cyclical spending due to the higher intertemporal elasticity of substitution for such goods (Browning and Crossley, 2000). In our model we abstract from differences in income elasticity across goods and focus on differences in intertemporal separability. For this reason we adopt the terminology of "durables" vs. "non-durables", whereby we subsume durable goods and memorable services under "durables", while we subsume non-durable goods and non-memorable services under "non-durables". In our empirical analysis we leverage our transaction-level spending data to compare the dynamics of durable goods and memorable services with non-durables. Moreover, we also control for the level of individual income to account for differences in the income elasticity of demand across goods categories.

2.1 Model

We consider an intertemporal consumption model with discrete time and a long, finite time horizon T. In each period t, the consumer chooses consumption of nondurables n_t and

[1] For simplicity we abstract from lumpy purchases and adjustment costs in the spirit of Caballero (1993).

durables d_t to maximize lifetime utility:

$$U = \sum_{t=0}^{T} \delta^t \left(\ln(n_t) + \alpha \ln(d_t + \gamma d_{t-1}) \right), \tag{1}$$

where δ is the discount factor, $\alpha > 0$ reflects the relative preference for durables, and $0 \leq \gamma \leq 1$ captures intertemporal nonseparability in durable consumption. The utility function is separable across goods within each period. For $\gamma = 0$, preferences for durables are time-separable; for $\gamma > 0$, consumption of durables exhibits a "durable" feature, where past consumption reduces current marginal utility.[2]

The nondurable good serves as the numeraire, and the relative price of durables is constant at $p_t = p$. The consumer can save or borrow at an interest rate r, implying the intertemporal budget constraint:

$$\sum_{t=0}^{T} \frac{n_t + p_t d_t}{(1+r)^t} \leq \sum_{t=0}^{T} \frac{y_t}{(1+r)^t}, \tag{2}$$

where y_t denotes income in period t. For tractability, we assume constant income $y_t = y$ and $\delta(1+r) = 1$.

2.2 Pre-shock Consumption

The first order conditions for optimal consumption in our model imply the following consumption structure in any period t:

$$d_t = \left(\frac{d_t}{d_t + \gamma d_{t-1}} + \frac{d_t \gamma \delta}{d_{t+1} + \gamma d_t} \right) \frac{\alpha_t}{p_t} n_t, \tag{3}$$

where the term in parentheses simplifies to 1 in the case of time-separable preferences for durables $\gamma = 0$.

Our assumptions imply constant spending and a constant ratio of durable to nondurable

[2]With $\gamma < 0$, the model would instead exhibit habit formation, where past consumption increases current marginal utility.

spending absent shocks:

$$n^\star = n_t = n_{t+1}$$
$$d^\star = d_t = d_{t+1}$$
$$\frac{d^\star}{n^\star} = \frac{(1+\gamma\delta)}{(1+\gamma)}\frac{\alpha}{p}$$

2.3 Transitory Shocks and Consumption Dynamics

Our stylized model provides intuition for the response to three types of transitory shocks which are relevant to the COVID recession:

First, a temporary shock to the preference for durables at time $\tau : \alpha_\tau < \alpha$ will lead to a temporary change in the level and composition of spending which will depend on the intertemporal separability of preferences for durables. Following Beraja and Wolf (2021), such a preference shock could be triggered by a decline in consumer confidence which temporarily reduces the relative preference for durables vs. nondurables.

Second, a temporary shock to the relative price for durables at time $\tau : p_\tau > p$ will lead to a temporary change in the level and composition of spending which likewise depends on the intertemporal separability of preferences for durables.[3]

Third, a temporary income shock $y_\tau < y$ at time τ will lead to intertemporal consumption smoothing in the absence of credit constraints. Credit-constrained consumers cut spending by the full amount of the income loss.

Figure 2a presents a simulation for the reaction of durable spending to a temporary preference or price shock. The figure shows that stronger intertemporal spillovers of durable consumption, γ, intensify the contraction of durable consumption associated with the transitory shock and the subsequent overshoot. Consider first the consumer with time-separable preferences $\gamma = 0$. For this consumer, a temporary decrease (increase) in the preference (price) for durables implies that durable consumption drops temporarily in period $t = \tau$, but reverts to the pre-shock level immediately after the shock. By contrast, consider the

[3] In the context of the pandemic, temporary public health measures (e.g. a lockdown) can be interpreted as a substantial price shock to goods and services subject to those restrictions. In our framework "durables" subsumes both durable goods and memorable services such as restaurant dining, travel, and entertainment.

consumer with non-separable utility of durable consumption $\gamma > 0$. Following the same shock to durables, durable consumption goes down in the shock period and overshoots after the shock, only then adjusting back to pre-shock level.

2.4 Predictions

2.4.1 Standard Consumers

We define standard consumers as those with time-separable preferences for durables and no credit constraints. For a transitory shock in period τ, the model yields the following predictions:

> A transitory decrease in the preference or increase in the price for durables lowers durable spending in period τ only. Consumption reverts to its pre-shock composition immediately thereafter. The reduction of spending in period τ leads to an increase in liquidity which is gradually spent, slightly raising the consumption level. See Figure 2a ($\gamma = 0$).
>
> A temporary income decline leads to a small, persistent decline in nondurable and durable spending, as consumers smooth the shock over time. See Figure 2b ($\gamma = 0$).

2.4.2 Pent-up Demand Consumers

"Pent-up demand" consumers have time-inseparable preferences for durables but are not credit constrained. For a transitory shock in period τ, the model predicts:

> A transitory decrease in the preference or increase in the price for durables causes a sharper drop in durable spending in period τ than in the standard case. This is followed by an overshooting of spending on durables in the subsequent period. Liquidity rises more strongly than in the standard case, part of which is spent as consumption overshoots, with the remainder smoothed over time. See Figure 2a ($\gamma > 0$).
>
> A temporary income decline leads to a small, persistent decline in nondurable and durable spending as in the standard case. See Figure 2b ($\gamma > 0$).

Figure 2: Consumption response to a transitory shock

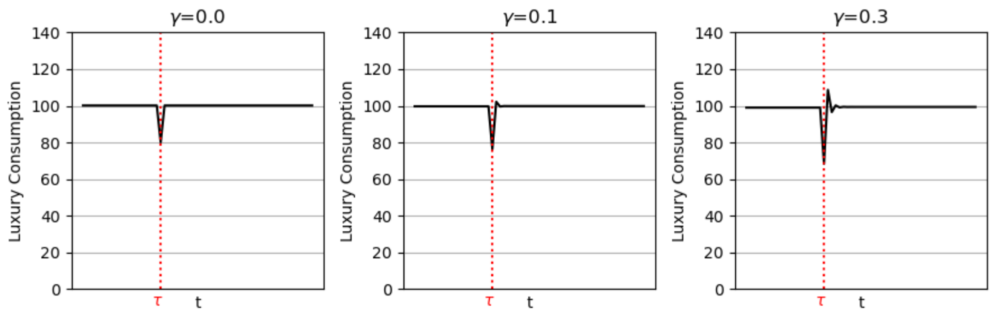

(a) Price shock

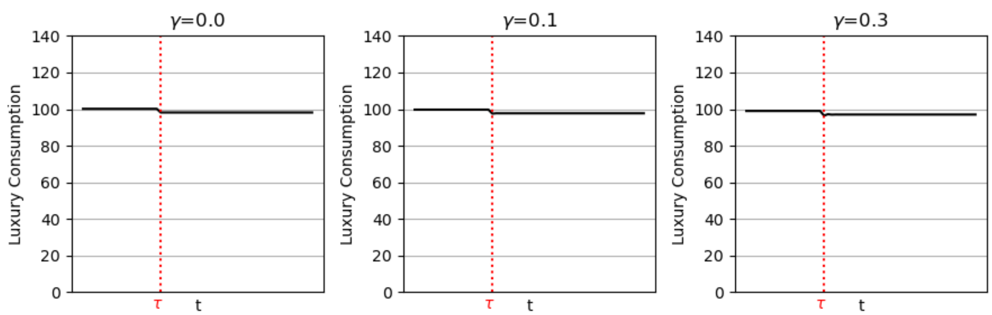

(b) Income shock

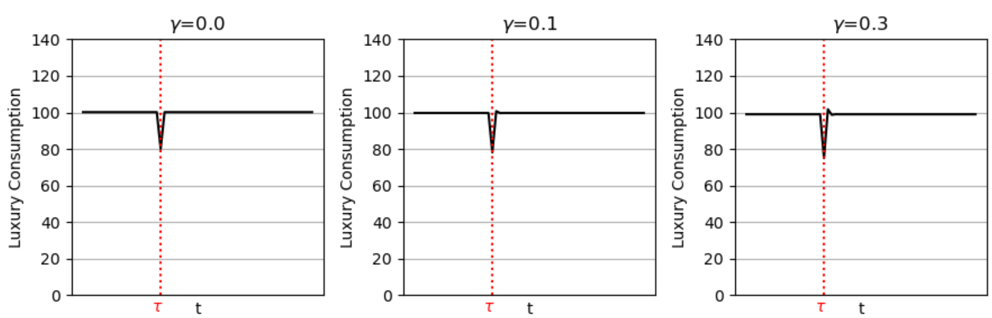

(c) Income shock (Hand-to-mouth)

2.4.3 Hand-to-mouth Consumers

In our baseline model, we assume that consumers can borrow and save at interest rate r, facing no credit constraints. Hand-to-mouth consumers on the other hand are liquidity constrained. For simplicity, we assume that they cannot store wealth at all; instead of the intertemporal budget constraint in equation (2), they face a period-by-period constraint and thus always consume $n_t + p_t d_t = y_t$[4] (see Figure 2c). For transitory shocks to these consumers the model predicts:

> A transitory decrease in the preference or increase in the price of durables reallocates spending towards nondurable goods in period τ without affecting total spending.
>
> A temporary income decline leads to an equivalent drop in total spending.

2.5 Testable Implications

The two consumer types yield contrasting predictions along four observable dimensions that we can bring to the data. First, regarding **liquidity accumulation**: pent-up demand consumers who cut spending on durables or memorables should see a corresponding rise in liquid balances, the "forced savings" that later fund the catch-up. By contrast, hand-to-mouth consumers who cut spending because of income losses should display no systematic increase in liquidity, since lower spending is matched by lower income. Second, regarding **recovery speed and shock magnitude**: the pent-up demand hypothesis implies that consumers with the largest initial spending cuts, and hence the largest stock of deferred consumption, should display the strongest rebound. For hand-to-mouth consumers, by contrast, spending recovery is mechanically tied to income recovery. If, as is plausible, larger income shocks are also more persistent, then consumers with the largest initial spending cuts should display the *slowest* recovery—the opposite of the pent-up demand prediction. Third, regarding **the relationship between shock composition and recovery**: the pent-up demand model predicts that, conditional on the total magnitude of the spending shock, consumers whose contraction is concentrated in durables or memorables (high γ goods) should recover faster,

[4] Alternatively, we could assume that consumers can save, but not borrow at r and have strongly increasing income profiles over the life-cycle.

as the intertemporal non-separability of utility generates an overshooting rebound once the shock dissipates. The hand-to-mouth model makes no such prediction. Fourth, regarding **the comovement of spending and income**: for hand-to-mouth consumers, the spending shock should be correlated with a concurrent income shock, and the subsequent spending recovery should track income recovery. For pent-up demand consumers, by contrast, spending recovery should eventually decouple from income as accumulated savings are drawn down to finance catch-up consumption.

3. Data

We study the COVID-19 recession in order to examine how consumer spending responded to the magnitude and composition of initial spending shocks as well as to the dynamics of household income. To this end, we use transaction-level data providing comprehensive information on both outgoing and incoming payments for a large sample of U.S. consumers during 2019-2021.

As outlined in Section 2, our theoretical framework contrasts two broad categories—durables (encompassing both durable goods and memorable services) and non-durables (encompassing non-durable goods and routine services)—unified by the key distinction between goods whose utility exhibits intertemporal spillovers ($\gamma > 0$) and those with time-separable utility ($\gamma = 0$). In our empirical analysis, however, we exploit the granularity of our transaction-level data to disaggregate the "durable" category of the model into its two constituent components: durable goods and memorable services.

There are several reasons why this finer decomposition is informative. First, during the COVID-19 pandemic the two sub-categories were differentially exposed to public health restrictions. Memorable services—restaurant dining, travel, live entertainment—were directly curtailed by lockdowns and social-distancing mandates, whereas many durable goods could still be purchased online, and some categories (e.g., home office equipment, electronics) even saw increased demand while others (e.g., apparel) declined. Thus, the effective "price shock" implied by public health measures was far larger and more uniform for memorables than for durables. Second, even in normal times, the degree of intertemporal non-separability may differ between the two: the lingering utility from a memorable vacation potentially has a different persistence parameter γ than the service flow from a durable good such as furniture or a vehicle.[5] By separating the two empirically, we can test whether their post-shock dynamics differ in ways that a single composite category would obscure. Meanwhile, we group non-durable goods and other (non-memorable) services together as "non-durables," since both categories exhibit time-separable utility in the model.

3.1 Transaction-level Data

We use anonymized transaction-level data gathered by a U.S. fintech application named "Albert." The application's primary function is account aggregation and financial management. Albert's users connect their checking and savings accounts and credit cards to the service, which aggregates financial information from these multiple sources in one place. Notably, we can observe outgoing and incoming transactions to the linked accounts, providing insight into users' spending behavior, income flows, and changes in their liquidity.

We obtain an unbalanced panel of approximately 56,000 active users from April 2019 to January 2022. To be selected for the panel, users must (i) have been on the app since at least early 2020, (ii) have linked their main checking account, and (iii) have logged into the application in the last month of our observation (October 2021). For our panel of users, we observe all inflow and outflow transactions from their linked accounts, including their amount, date, category, and a transaction description. In addition, the data includes a rich set of user-level characteristics such as age, self-reported income, or a 5-digit zip code. For more details about the dataset, see Toczynski (2023) and Benetton et al. (2025). Table A1 presents the summary statistics of available user-level variables.

To construct the main variables used in the empirical analysis, we first combine information from the category variable and transaction text descriptions to classify financial flows into different types of spending and income. For example, to define the variable *Food Delivery*, we create a dictionary of keywords associated with food delivery services, such as "Uber Eats" or "DoorDash," and then perform a keyword search on the transaction database. Using this approach, we classify outflow transactions into 25 types of consumption. In a similar

[5] See e.g. Hai et al. (2020), Harmenberg and Öberg (2021), Cao et al. (2025).

way, among the inflows, we identify regular salary payments, social benefit payments, tax refund payments, and stimulus payments, which provide measures of income shocks. The residual unclassified transactions are primarily financial (fees, investments, savings) and transfers between accounts.

While we are confident that our classification is largely accurate, we acknowledge that keyword-based methods can occasionally lead to misclassification. In some cases, it is also impossible to unambiguously assign a transaction to a single category. In the next subsection, we describe how we aggregate the classified transactions into the main variables of interest for our analysis. We also verify our spending data against representative spending data from the Survey of Consumer Expenditures.

3.2 Spending

We want to measure the magnitude of the total spending contraction in the COVID recession as well as its composition across key categories. To this end, we aggregate the observed spending transactions at the consumer $\times$ month level into the following three categories:

- *Non-durables*: Spending on non-durable goods and services (e.g., groceries, gas and fuel, public transportation).
- *Durables*: Spending on (semi-) durable goods (e.g., apparel, furniture, or car payments).
- *Memorables*: Services that provide memorable experiences (e.g., restaurant spending, travel, or recreation).

For each consumer i, we thus measure spending in three aggregate categories $S_{i,j}$ at the monthly level. We refer to the sum of spending on the three categories as *Total Spending* (S_i). We focus our analysis on the 34,453 consumers with at least $1,000 of total spending in April 2019 and with positive spending in every subsequent month, thereby restricting the sample to those with sufficiently comprehensive coverage. Table A1 shows that our selected sample of users is similar in terms of sociodemographics to the underlying panel. Table A2 reports

Figure 3: Evolution of Spending and Income Over Time

(a) Total Spending and Total Income

summary statistics for the three focal categories and their components in the final sample. On average, consumers spend about \$1,998 per month on non-durable goods and services, followed by memorable services (\$919) and durable goods (\$541). Average total monthly spending in our sample is 3,459 USD. Note that we exclude spending on housing due to lack of accurate data on mortgage interest payments (as opposed to principal repayments). By comparison, the U.S. Consumer Expenditure Survey reports monthly average non-housing expenditures of 3,692 USD for 2021.[6]

[6]See https://www.bls.gov/cex/tables.htm#topline

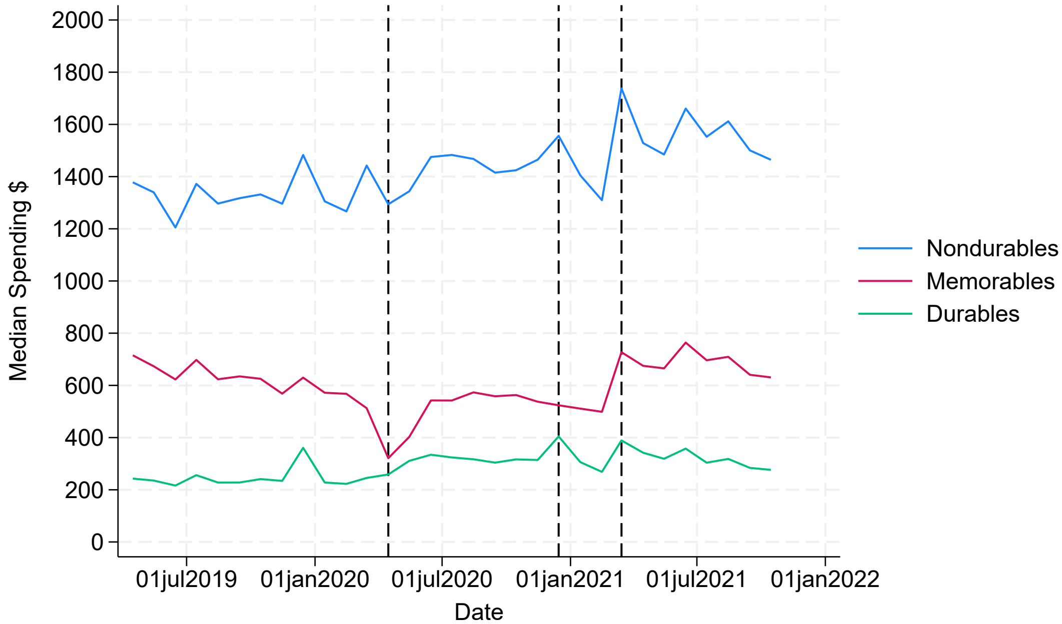

(b) Spending By Category

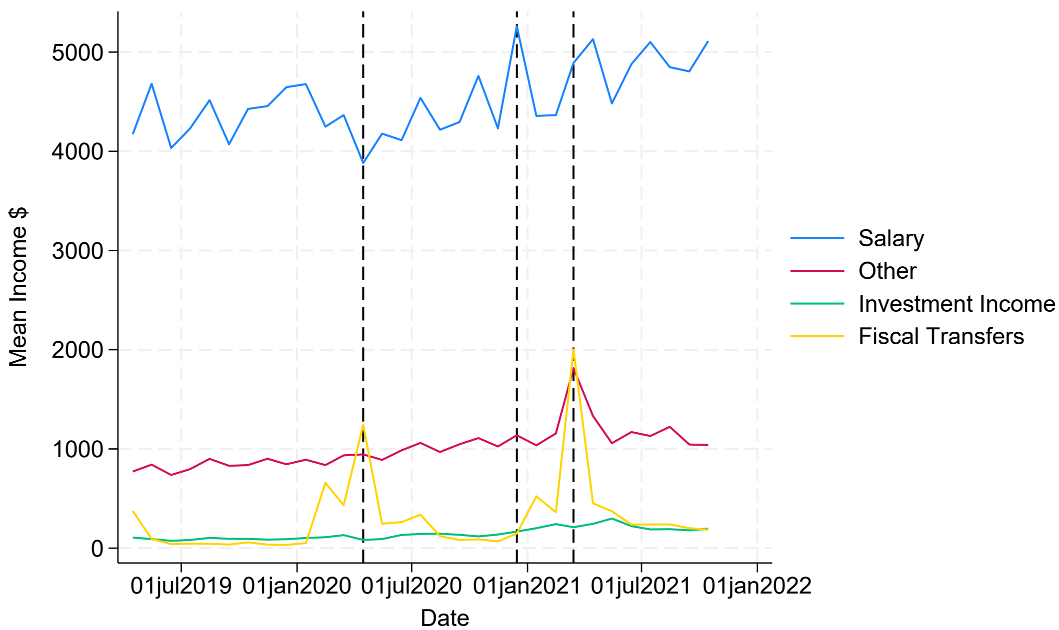

(c) Income by Source

Note: Panel A plots the mean *Total Spending* and *Total Income* over time. Panel B plots the evolution of median spending in three focal categories: (1) Nondurables, (2) Memorable services, and (3) Durables. The vertical lines mark the three waves of fiscal stimulus payments. Panel C plots the evolution of income split by source: (1) Salary, (2) Other, (3) Investment Income, (4) Tax Refunds (including stimulus payments).

Figure 3 plots the evolution of median total spending and spending by category over the sample period. There is a noticeable drop in total spending in April 2020 following the initial pandemic shock and first-wave lockdown. As documented by previous studies this drop is concentrated in memorables where median spending falls by about 50%.

3.3 Income and Liquidity

To construct our measure of *Total Income (I_i)* at the consumer × month level, we combine (i) identified payroll and benefits payments, (ii) tax refunds (including fiscal stimulus payments), and (iii) other regular income transactions. The latter include payroll flows we were unable to classify based on keywords alone as well as recurring payments from other sources, such as transfers from family members. As such, *Total Income* captures both regular income and windfall transfers. Figure 3, Panel A shows that - in contrast to spending - there is no significant fall in average income at the onset of the pandemic. This is related to the immediate rollout of fiscal transfers in April 2020. Moreover, over our observation period we observe sharp increases in average income that coincide with subsequent disbursements of stimulus payments to households (Toczynski, 2023). Panel C of Figure 3 shows that the fiscal payments mask a significant decline in payroll income at the onset of the pandemic.

Throughout our analysis we employ measures of consumer liquidity to study pent-up demand and hand-to-mouth consumption. Unfortunately, our data does not include exact balances of each account in each month. However, we can measure the *changes* in liquid balances between any two points in time by taking the difference between net cumulative inflows and outflows between these periods. To this end, we define variable *Cumulative Net Liquidity Change* – the net cumulative flow to and from linked bank accounts in specific months during our observation period. Mean liquid assets sharply increase at the onset of the pandemic. Panel A in Figure 4 compares the distribution of monthly changes in liquid balances in April 2020 to average monthly changes in 2019. Further, Panel B of 4 plots the distribution of *Cumulative Net Liquidity Change* between April 2019 and December 2019. In our analysis, we use this measure to separate *Hand-to-Mouth* vs. *Other* users. We classify a user as hand-to-mouth if their cumulative net liquidity change between April 2019 and December 2019 is no larger than $300, i.e. less than 10% of average monthly total spending

Figure 4: Liquidity Changes

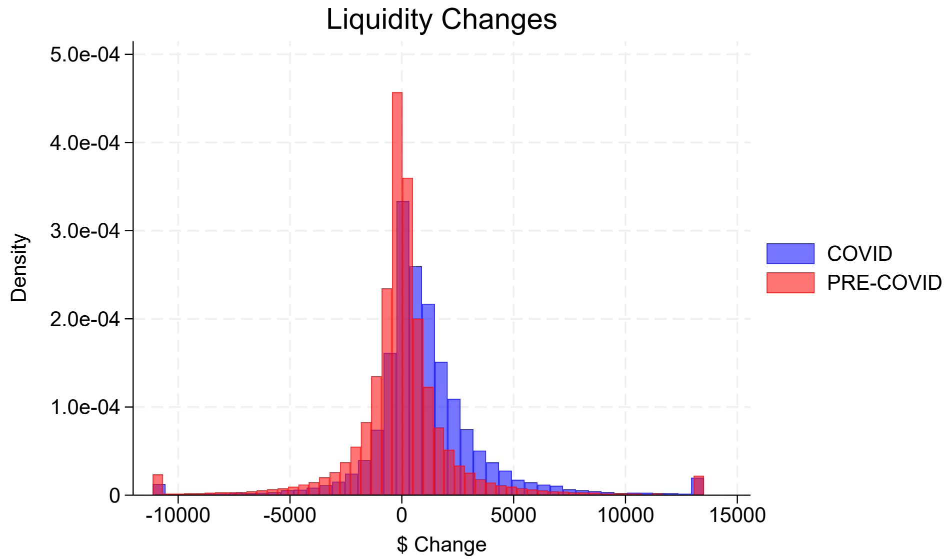

(a) Liquidity Changes in April 2020 vs Average Monthly Change in 2019

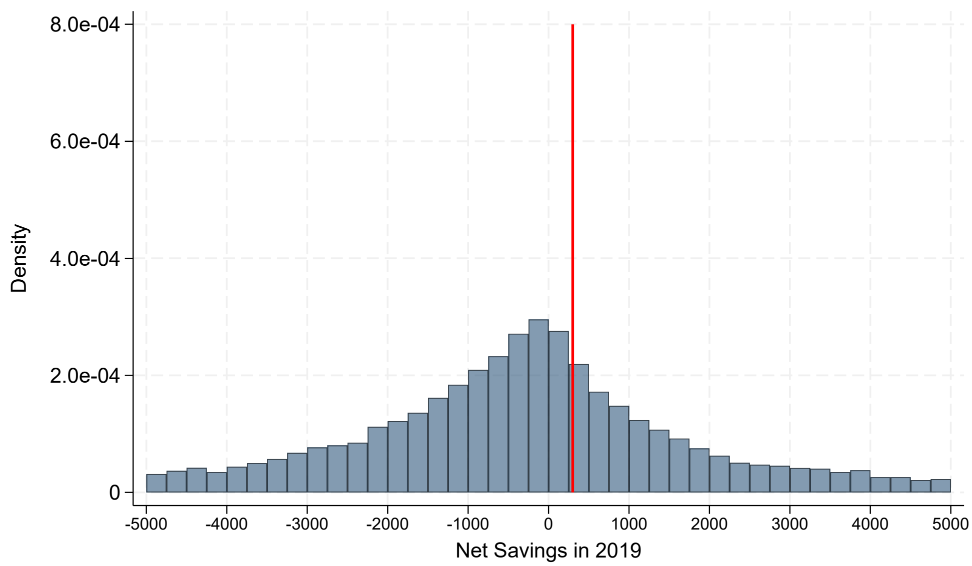

(b) Net Savings Distribution, April 2019 - December 2019

Note: Panel A presents the distribution of the month-to-month changes in liquid assets in linked bank accounts in April 2020 (COVID, in blue) vs. April 2019 (Pre-COVID, in red). Panel B presents the distribution of cumulative net liquidity changes in April - December 2019. The vertical line indicates the threshold we use to define Hand-to-Mouth users – those with cumulative liquidity increases below \$300.

in our sample over this nine-month period.

4. Cross-Sectional Observations on Spending Dynamics

In this section, we document seven cross-sectional observations on the dynamics of consumer spending through the COVID-19 recession. These observations are organized around the four empirical dimensions derived from our theoretical framework in Section 2: the composition of the initial spending shock and the associated liquidity dynamics (Observations 1-3), the relationship between shock magnitude and recovery speed (Observation 4), the comovement of spending and income shocks (Observation 5), and the link between spending recovery and income recovery (Observations 6-7). Together, the patterns that emerge challenge pent-up demand as the primary driver of spending recovery and instead point toward income dynamics as the dominant force shaping consumption paths. The observations documented here motivate the formal regression analysis in the following section.

Observation 1: A majority of consumers experience a sharp spending decline

We begin by characterizing the severity of the initial spending shock across consumers. We define $ShockMagnitude_i$ as the percentage change in total spending in April 2020 relative to April 2019:

$$ShockMagnitude_i = \frac{S_{i,t=2020:04} - S_{i,t=2019:04}}{S_{i,t=2019:04}}. \tag{4}$$

Figure 5 presents the distribution of $ShockMagnitude_i$. The distribution is heavily left-skewed: approximately two-thirds of consumers experienced a negative spending shock at the onset of the pandemic. A substantial majority—roughly 59% of the sample, or 20,486 consumers—experienced a decline of at least 10%. These consumers form the basis of our subsequent analysis of spending recovery dynamics. At the consumer level, the magnitude of the spending shock is strongly correlated with the local severity of the COVID-19 pandemic, as measured by the Social Distancing Index (see Appendix 1), consistent with the region-level evidence documented by Chetty et al. (2023). The pervasiveness of the spending contraction underscores the breadth of the pandemic's impact on household consumption and motivates the question of what mechanisms governed the subsequent recovery.

Figure 5: Distribution of Spending Shocks in the COVID Recession

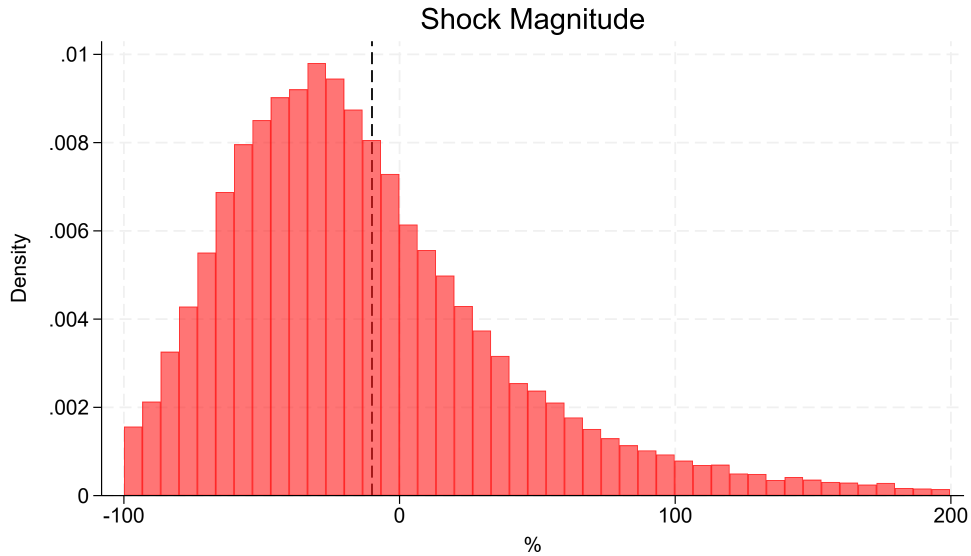

Note: The figure presents the distribution of $ShockMagnitude_i$, the percentage change in total spending in April 2020 compared to April 2019. The vertical line denotes the cutoff of -10% for inclusion in the main analysis sample. The sample comprises 34,453 consumers with at least $1,000 in spending in April 2019 and positive spending in all months from April 2019 through October 2021.

Observation 2: For a majority of consumers, memorables account for a large share of the spending shock

Given the breadth of the spending contraction documented in Figure 5, a natural question is what categories of spending drove these declines. To characterize the composition of each consumer's spending shock, we define the share attributable to category j as:

$$y_{i,j} = \frac{S_{i,j,\,2020:04} - S_{i,j,\,2019:04}}{S_{i,\,2020:04} - S_{i,\,2019:04}}, \tag{5}$$

where $y_{i,j} \in \{MemorablesShare_i,\, DurablesShare_i\}$. A higher value indicates that cutbacks in category j account for a larger fraction of the consumer's total spending decline.

Figure 6 presents a stacked bar chart showing, for each shock-magnitude group, the composition of consumers by the share of their spending shock attributable to memorable services. The x-axis groups consumers by shock magnitude. Segments —ranging from "No contribution" (memorables spending did not fall) to "80%+" (memorables account for more than 80% of the total spending drop)— show the within-bin distribution.

The figure reveals that for the vast majority of consumers, a cutback in memorable services such as restaurants, travel, and recreation accounts for a substantial portion of the initial spending decline. The 0–40% and 40–80% segments dominate every bar, indicating that most consumers saw between a moderate and large share of their spending shock driven by reduced spending on memorable services. The "No contribution" segment is small in every bin, confirming that almost all consumers cut back on memorables to some extent. This pattern is consistent across all shock magnitude bins, though consumers with milder shocks (-30 to -10%) tend to have slightly more dispersed memorables shares. The concentration of the spending shock in memorable categories is precisely the pattern that would, under a pent-up demand narrative, predict a strong subsequent rebound: consumers who were forced to forgo memorable experiences should accumulate excess savings and "catch up" once restrictions are lifted. The following observations, however, reveal that the cross-sectional data do not support this prediction.

Figure 6: Memorables Share of the Spending Shock vs. Shock Magnitude

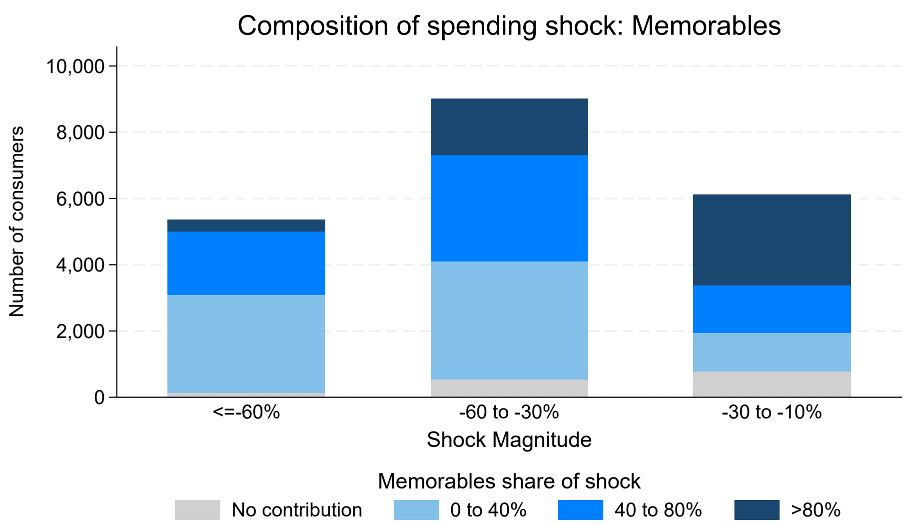

Note: The figure shows the composition of consumers by the share of their spending shock attributable to memorables, separately for three shock-magnitude bins ($\leq -60\%$, -60 to -30%, -30 to -10%). Bar height gives the number of consumers in each shock-magnitude bin; segments correspond to memorables-share categories (No contribution, 0–40%, 40–80%, 80%+). The sample covers consumers with a spending shock of at least 10%.

Figure 7 presents an analogous decomposition for durable goods. The pattern for durables differs markedly from that of memorables. The "No contribution" and "0–40%" segments

Figure 7: Durables Share of the Spending Shock vs. Shock Magnitude

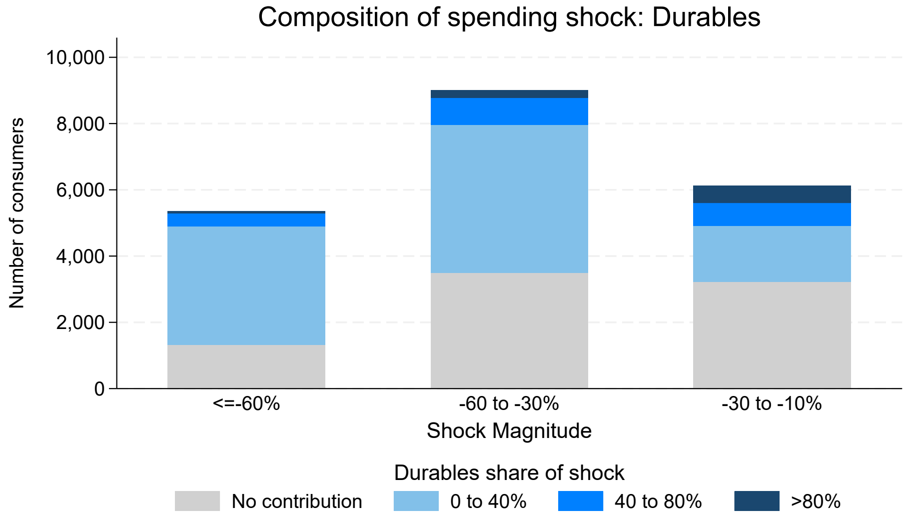

Note: The figure shows the composition of consumers by the share of their spending shock attributable to durables, separately for three shock-magnitude bins ($\leq -60\%$, -60 to -30%, -30 to -10%). Bar height gives the number of consumers in each shock-magnitude bin; segments correspond to durables-share categories (No contribution, 0–40%, 40–80%, 80%+). The sample covers consumers with a spending shock of at least 10%.

dominate each bar, indicating that for a majority of consumers, cutbacks in durable spending either did not contribute to the overall shock or accounted for only a modest share. The 40–80% and 80%+ segents are small across all shock-magnitude bins. This asymmetry between memorables and durables is consistent with the differential exposure of these categories to pandemic restrictions: memorable services were directly curtailed by lockdowns and social-distancing mandates, whereas many durable goods could still be purchased online. Taken together, Figures 6 and 7 establish that the COVID spending shock was predominantly driven by foregone memorable services rather than deferred durable purchases.

Observation 3: Consumers with larger spending shocks do not build up larger liquidity reserves

A central mechanism underlying pent-up demand is the accumulation of "forced savings": consumers who cut spending more should see correspondingly larger increases in their liquid balances, which then fund the subsequent spending rebound. Figure 8 plots the change in consumer liquidity in April 2020 against the magnitude of the spending shock, with liquidity changes normalized by average 2019 spending.

Contrary to the pent-up demand prediction, consumers with larger spending shocks do not display larger liquidity increases. By contrast, the relationship is slightly negative—those who cut spending the most accumulated somewhat less, not more, liquidity. When read alongside the spending shock composition in Figures 6, this observation implies that while memorables dominated the spending decline, the foregone spending did not translate into a liquidity buffer increase that could fuel a subsequent rebound.

Figure 8: Change in Liquidity vs. Spending Shock Magnitude

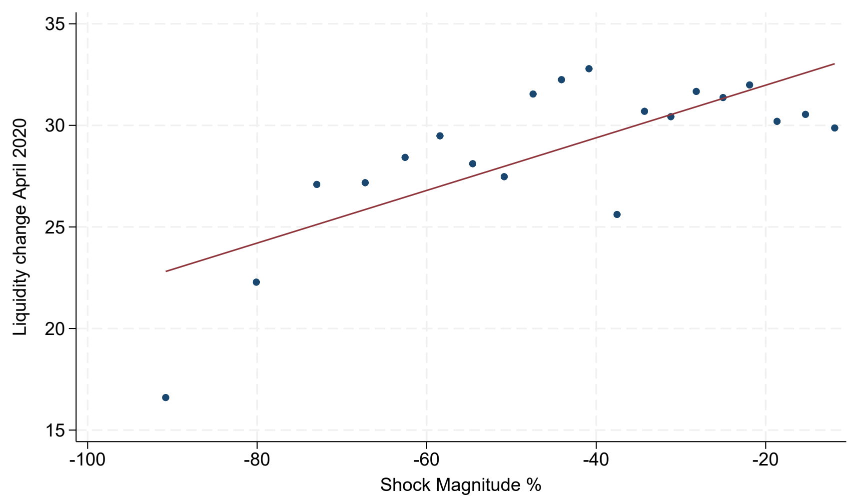

Note: The figure plots the mean change in liquid assets in April 2020 against the spending shock magnitude (April 2020 vs. April 2019). Liquidity changes are normalized by average total monthly spending in 2019 per consumer.

Observation 4: Spending recovers more slowly for consumers with larger spending shocks

To track the dynamics of the post-recession recovery, we define a cumulative measure of spending recovery at the consumer level. At horizon T, $SpendingRecovery_{i,T}$ is the cumulative deviation of total spending from its pre-recession reference level, scaled by average monthly spending in 2019:

$$SpendingRecovery_{i,T} = \frac{\sum_{t=2020:04}^{T} \Delta S_{i,t}}{S_{2019}^{AVG}}, \tag{6}$$

where $\Delta S_{i,t} = S_{i,t} - S_{i,t}^{ref}$ and the reference period for each post-recession month is the same calendar month in the pre-COVID period (April 2019–February 2020). This measure has a straightforward interpretation: negative values indicate that cumulative lost consumption has not yet been recovered, zero corresponds to full recovery, and positive values imply that lost consumption has been more than made up.

The pent-up demand hypothesis makes a clear prediction about recovery dynamics: consumers who experienced the deepest spending cuts should display the strongest and fastest recovery as they draw down accumulated savings to "catch up" on foregone consumption. Figure 9 plots the evolution of $SpendingRecovery_{i,T}$ over 18 months, with consumers grouped by the magnitude of their initial spending shock.

The data reveal the opposite pattern. Consumers with the weakest spending shocks (10–30%) recovered their pre-recession spending levels by approximately October 2020. Those with moderate shocks (30–60%) recovered by mid-2021. In contrast, consumers who experienced the most severe initial contraction (shock > 60%) had not fully recovered by the end of our observation period. This negative relationship between shock magnitude and recovery speed directly contradicts the pent-up demand prediction. Combined with the failure of the liquidity accumulation channel documented in Figure 8, these two observations jointly undermine the core mechanisms of the pent-up demand narrative: consumers who cut spending the most neither built up larger savings buffers nor recovered faster.

Observation 5: Consumers with larger spending shocks also suffer larger income shocks

To complement our measures of spending dynamics, we define the income shock analogously to $ShockMagnitude_i$:

$$IncomeShock_i = \frac{I_{i,t=2020:04} - I_{i,t=2019:04}}{I_{i,t=2019:04}}. \tag{7}$$

This variable captures the percentage change in total income in April 2020 relative to April 2019 at the consumer level.

The failure of the liquidity accumulation test (Observation 3) suggests that spending

Figure 9: Spending Recovery by Initial Shock Magnitude

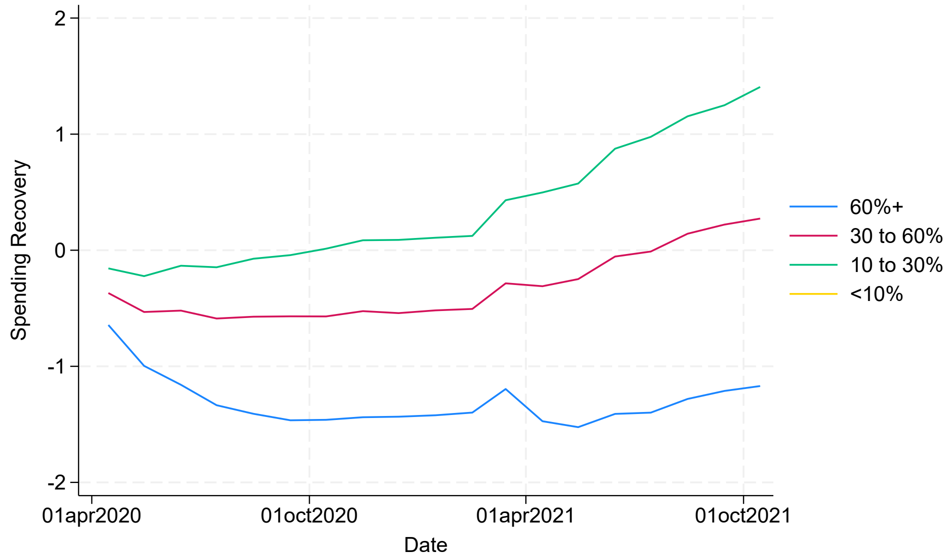

Note: The figure plots the evolution of median $SpendingRecovery_{i,T}$ by groups of consumers split by the magnitude of their April 2020 spending shock. $SpendingRecovery_{i,T} = \frac{\sum_{t=2020:04}^{T} \Delta S_{i,t}}{S_{2019}^{AVG}}$. The sample covers 20,486 consumers who experienced a spending shock of at least 10%.

declines during the pandemic were not purely "voluntary" postponements but were at least partly driven by concurrent income losses. Before examining the relationship between spending and income shocks, it is useful to characterize the distribution of income shocks in our sample. Figure 10 presents the distribution of $IncomeShock_i$. The distribution is wide, with substantial mass in both tails. This heterogeneity in income shocks provides the cross-sectional variation needed to assess whether income dynamics can account for the spending patterns documented above.

Figure 11 examines this channel directly by plotting the spending shock against the income shock experienced in April 2020 at the consumer level. The figure reveals a clear positive relationship: consumers who experienced larger drops in spending also tended to suffer larger declines in income. This correlation, while by no means perfect, is economically meaningful and suggests that for many households the ability to spend during the pandemic was constrained by reductions in income rather than solely by restricted access to services. This pattern is consistent with hand-to-mouth behavior, where consumption tracks disposable income because households lack sufficient liquid savings to smooth through income shocks.

When viewed alongside the memorables-dominated spending composition in Figure 6 and the absent liquidity buildup in Figure 8, a coherent alternative narrative emerges: the pandemic simultaneously reduced both income and spending opportunities, and the consumers hit hardest on the spending side were also those most exposed to income losses.

Figure 10: Distribution of Income Shocks April 2020 vs. April 2019

Note: The figure presents the distribution of the *Income Shock*. The sample covers 20,486 users who experienced a spending shock of at least 10% relative to the pre-recession period.

Observation 6: Income recovers faster for consumers with weaker income shocks

To parallel our measure of spending recovery, we define cumulative income recovery at the consumer level as:

$$IncomeRecovery_{i,T} = \frac{\sum_{t=2020:04}^{T}(Income_{i,t} - Income_{i,t}^{ref})}{Income_{2019}^{AVG}}, \tag{8}$$

where, as with spending, the reference period for each month is the same calendar month in the pre-COVID period. This measure captures the cumulative shortfall (or surplus) in income relative to pre-recession levels, expressed in units of average pre-recession monthly

Figure 11: Spending Shock vs. Income Shock in April 2020

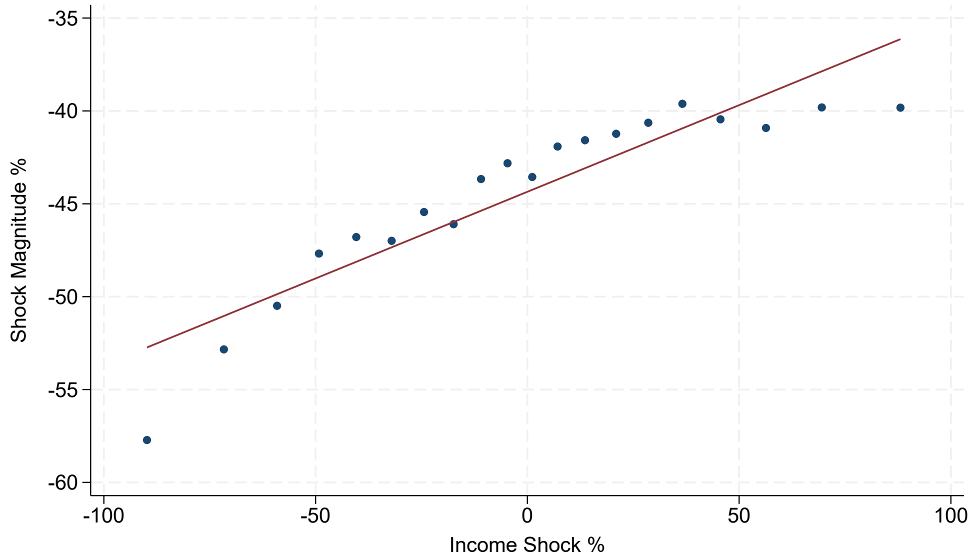

Note: The figure plots the spending shock against the income shock at the consumer level. Both shocks are measured as percentage changes in April 2020 relative to April 2019.

income.

Having established that spending and income shocks are correlated in the cross-section, we now examine the dynamics of income recovery. Figure 12 plots the evolution of $IncomeRecovery_{i,T}$ over 18 months, with consumers grouped by the magnitude of their initial income shock. The pattern mirrors the spending recovery documented in Figure 9: consumers with the mildest initial income shocks recovered their pre-recession income levels relatively quickly, while those who experienced the deepest income declines saw a further deterioration before any recovery, and had not fully recovered cumulative lost income by October 2021.

The parallel between the spending recovery dynamics in Figure 9 and the income recovery dynamics in Figure 12 is striking. In both cases, the most adversely affected consumers recover slowest. This symmetry is exactly what one would expect if consumption is substantially driven by current income: the persistence of income losses translates directly into persistent spending shortfalls, and the speed at which spending recovers is paced by the speed at which income recovers. The observation that large income shocks are associated with slow income recovery further helps explain why consumers with the largest spending shocks (who also suffer the largest income shocks, per Observation 5) display the slowest spending recovery (Observation 4).

Figure 12: Income Recovery by Initial Income Shock Magnitude

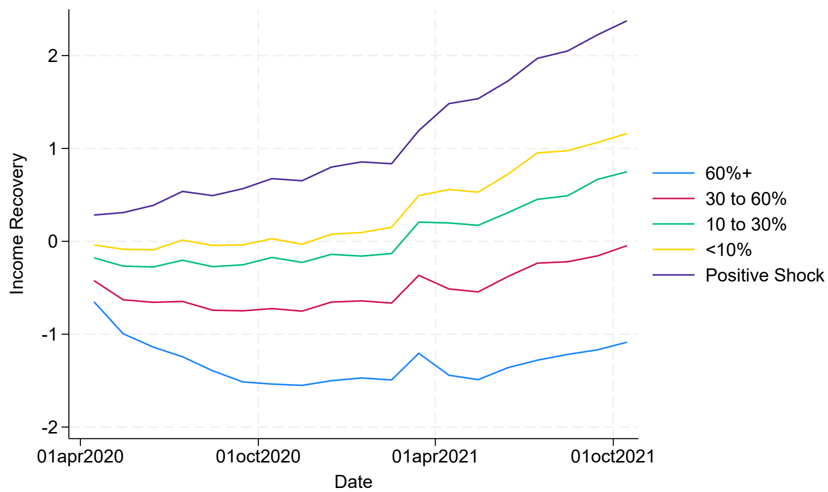

Note: The figure plots the evolution of median $IncomeRecovery_{i,T}$ by groups of consumers split by the magnitude of their April 2020 income shock relative to April 2019. $IncomeRecovery_{i,T} = \frac{\sum_{t=2020:04}^{T}(Income_{i,t} - Income_{i,t}^{ref})}{Income_{2019}^{AVG}}$. The sample covers 20,486 consumers who experienced a spending shock of at least 10%.

Observation 7: The recovery of spending is correlated with the recovery of income

The preceding observations have documented parallel patterns in the cross-sectional dynamics of spending and income. We now directly examine the relationship between spending recovery and income recovery at the consumer level. Figures 13 and 14 plot $SpendingRecovery_{i,T}$ against $IncomeRecovery_{i,T}$ at horizons of 3 and 18 months, respectively.

Both figures reveal a remarkably tight and positive relationship between the two recovery measures. At the 3-month horizon (Figure 13), the strong correlation indicates that the early trajectory of spending recovery is closely aligned with income dynamics. At the 18-month horizon (Figure 14), the correlation persists, demonstrating that the link between income and spending recovery is not merely a short-run phenomenon but a durable feature of the data. The persistence of this relationship over 18 months is particularly notable: if pent-up demand were the dominant force, one would expect spending recovery to eventually decouple from income recovery as consumers draw down accumulated savings to catch up on deferred consumption. Instead, spending continues to track income closely, consistent with hand-to-mouth behavior in which consumption is persistently governed by the flow of disposable income.

Taken together, the seven observations documented in this section paint a coherent picture. While the COVID recession produced widespread and memorables-dominated spending declines (Observations 1–2), the cross-sectional patterns in liquidity accumulation and recovery dynamics are inconsistent with pent-up demand as the primary recovery mechanism (Observations 3–4). Instead, spending shocks are correlated with income shocks (Observation 5), income and spending recovery dynamics display strikingly parallel patterns (Observation 6), and spending recovery is tightly linked to income recovery at both short and medium horizons (Observation 7). These descriptive findings motivate the formal regression analysis in the next section, where we test whether the composition of the initial spending shock (pent-up demand) or the dynamics of income (hand-to-mouth consumption) better explains the cross-sectional variation in spending recovery, controlling for the magnitude of the initial shock and a rich set of consumer characteristics.

Figure 13: Spending Recovery vs. Income Recovery at 3 Months

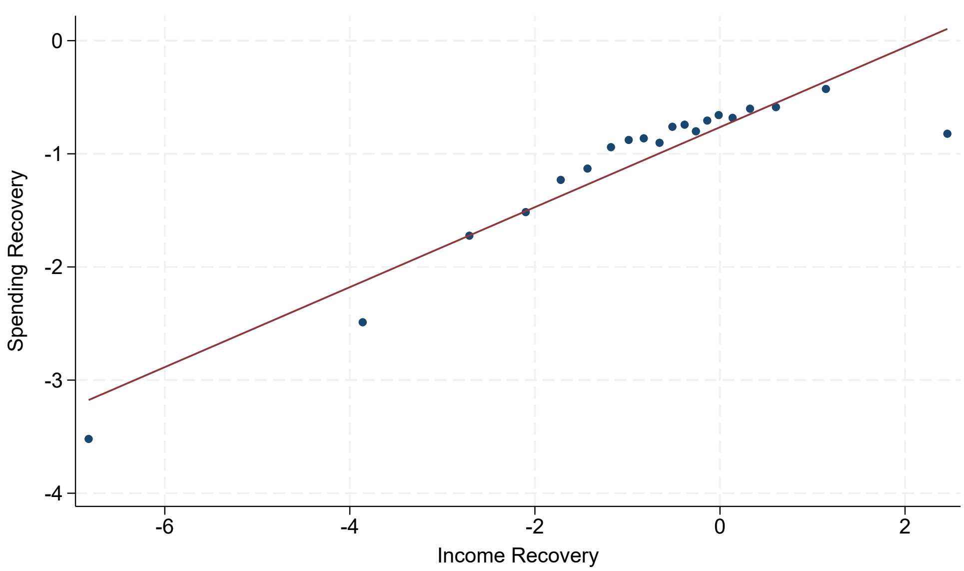

Note: The figure plots $SpendingRecovery_{i,T}$ against $IncomeRecovery_{i,T}$ at $T = T_0 + 3$ months (July 2020). Each point represents a bin of consumers.

Figure 14: Spending Recovery vs. Income Recovery at 18 Months

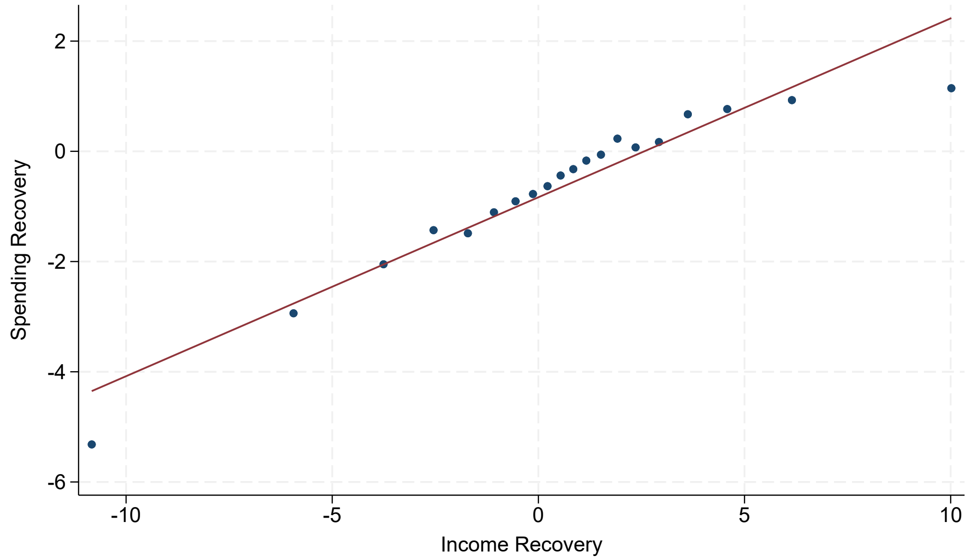

Note: The figure plots $SpendingRecovery_{i,T}$ against $IncomeRecovery_{i,T}$ at $T = T_0 + 18$ months (October 2021). Each point represents a bin of consumers.

5. Regression Methodology & Results

5.1 Specification

The cross-sectional patterns documented in the previous section suggest that income dynamics, rather than the composition of foregone spending, govern the speed at which consumers recover from the COVID-19 recession. To move beyond descriptive correlations, we now test these competing explanations in a unified regression framework. We estimate cross-sectional local projection models of the form:

$$\begin{aligned} SpendingRecovery_{i,T} = \alpha &+ \beta_1\, IncomeRecovery_{i,T} + \beta_2\, ShockMagnitude_i + \beta_3\, IncomeShock_i \\ &+ \beta_4\, MemorablesShare_i + \beta_5\, DurablesShare_i + \gamma\, Controls_i + \epsilon_{i,T}, \end{aligned} \tag{9}$$

where the horizon T ranges from 3 to 18 months after April 2020 (T_0). To test the pent-up demand channel, we include $MemorablesShare_i$ and $DurablesShare_i$, defined as the ratio of the change in memorables (durables) spending to the change in total spending between April 2019 and April 2020, so that higher values indicate a larger contribution of memorables (durables) cutbacks to the overall spending shock. To test the hand-to-mouth channel, we include $IncomeRecovery_{i,T}$, which measures the cumulative recovery of disposable income through horizon T. All variables enter the same specification alongside $ShockMagnitude_i$ and $IncomeShock_i$ as continuous variables, together with self-reported income, age, gender, and three lags of pre-recession spending growth. We estimate the model at six horizons—$T = T_0 + 3, 6, 9, 12, 15, 18$ months—to trace how the predictive power of each mechanism evolves over the recovery period.

This specification is designed to distinguish between the competing mechanisms identified by our theoretical framework. The coefficient β_2 captures the relationship between the magnitude of the initial spending shock and the speed of recovery: under the pent-up demand hypothesis, β_2 should be positive, as consumers with deeper initial cuts accumulate larger stocks of deferred demand and therefore rebound more strongly. For hand-to-mouth consumers, by contrast, spending recovery is mechanically tied to income recovery; if larger income shocks are also more persistent, then β_2 should be negative—the opposite of the pent-up demand prediction. The coefficients β_4 and β_5 test whether the composition of the shock matters for recovery conditional on its magnitude: the pent-up demand model predicts positive coefficients, as a higher concentration of the spending decline in high-γ categories (memorables and durables) should generate stronger overshooting once the shock dissipates. Under hand-to-mouth behavior, shock composition should be irrelevant once income dynamics are controlled for, yielding insignificant or zero coefficients. Finally, β_1 captures the hand-to-mouth channel directly: a positive and growing coefficient on $IncomeRecovery_{i,T}$ would indicate that spending recovery is paced by the flow of disposable income rather than by the release of stored liquidity. The joint inclusion of both sets of variables allows us to assess their relative explanatory power within a single framework.

5.2 Results

Table 1 reports the estimates. Five sets of results stand out, corresponding to the key regressors.

Table 1: Spending Recovery: The Role of the Initial Shock and Disposable Income

	(T_0+3)	(T_0+6)	(T_0+9)	(T_0+12)	(T_0+15)	(T_0+18)
$IncomeRecovery$	0.168***	0.232***	0.267***	0.281***	0.292***	0.299***
	(0.013)	(0.011)	(0.010)	(0.009)	(0.008)	(0.008)
$ShockMagnitude$	-3.127***	-3.798***	-3.946***	-4.734***	-5.488***	-5.692***
	(0.073)	(0.100)	(0.123)	(0.147)	(0.175)	(0.200)
$IncomeShock$	-0.046***	-0.071***	-0.076***	-0.095***	-0.094***	-0.100***
	(0.011)	(0.015)	(0.019)	(0.022)	(0.026)	(0.029)
$MemorablesShare$	-0.065***	-0.137***	-0.193***	-0.202***	-0.179***	-0.195***
	(0.022)	(0.034)	(0.041)	(0.050)	(0.061)	(0.073)
$DurablesShare$	-0.060**	-0.061	-0.081*	-0.119**	-0.076	-0.079
	(0.026)	(0.038)	(0.049)	(0.058)	(0.072)	(0.086)
N. observations	19,751	19,751	19,751	19,751	19,751	19,751
R^2	0.260	0.243	0.227	0.217	0.204	0.192
Controls	Y	Y	Y	Y	Y	Y

Note: Each column corresponds to a cross-sectional regression estimated at a different horizon T after April 2020 (T_0). $ShockMagnitude$, $MemorablesShare$, and $DurablesShare$ enter as continuous variables. Controls include three monthly lags of total spending growth, self-reported income, age, and gender. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Shock magnitude. The coefficient on $ShockMagnitude_i$ is large, negative, and highly significant at every horizon, growing in absolute value from -3.13 at 3 months to -5.69 at 18 months. Because $ShockMagnitude$ is measured as a negative percentage change, the negative coefficient means that consumers who experienced deeper initial spending cuts recovered more slowly, not faster. This result formalizes the pattern visible in Figure 9, where consumers with shocks exceeding 60% had not recovered cumulative lost spending by October 2021, while those with shocks of 10–30% had recovered by late 2020. Combined with the evidence from Figure 8 that consumers with larger spending cuts did not accumulate larger liquidity buffers, the regression confirms that the two core mechanisms of the pent-up demand narrative—forced savings accumulation and subsequent catch-up spending—fail in the cross-section.

Memorables share. The coefficient on $MemorablesShare_i$ is negative and statistically significant across all six horizons, indicating that—even after controlling for income dynamics—consumers whose initial spending shock was more concentrated in memorable services experienced a *slower* recovery of total spending. This finding directly contradicts the pent-up demand prediction: if consumers with the largest cutbacks in memorables were accumulating deferred demand, they should exhibit stronger recoveries, yielding a positive coefficient. Instead, the negative sign aligns with the descriptive evidence in Figure 6, which showed that virtually all consumers cut back on memorables, and with the observation that the group with the fastest recovery was the one for whom memorables did *not* contribute to the initial shock. The persistently negative coefficient confirms that the composition of the initial shock does not generate the catch-up dynamics predicted by pent-up demand; if anything, consumers whose spending decline was more heavily concentrated in memorable services faced a persistent spending disadvantage.

Durables share. The coefficient on $DurablesShare_i$ is negative at every horizon, though it is not statistically significant only at selected horizons. The pattern is qualitatively similar to that of $MemorablesShare$: a larger concentration of the initial spending shock in durable goods is associated with a slower, not faster, recovery. This result is again inconsistent with the pent-up demand prediction that deferred purchases of high-γ goods should generate catch-up dynamics. The weaker statistical significance relative to $MemorablesShare$ likely reflects the fact that durables played a smaller role in the overall spending contraction during the pandemic, as documented in Figure 7, where the majority of consumers saw little or no contribution of durables to their spending shock. Together, the negative coefficients on both composition variables indicate that neither memorables- nor durables-led pent-up demand can account for the cross-sectional variation in spending recovery.

Income recovery. The coefficient on $IncomeRecovery_{i,T}$ is positive, large, and significant at the 1% level across all horizons, rising steadily from 0.168 at 3 months to 0.299 at 18 months. At the 18-month horizon, a one-unit increase in cumulative income recovery (equivalent to one month of average pre-recession income) is associated with roughly a 0.30-unit increase in cumulative spending recovery. The growing magnitude of the coefficient indicates that the link between income and spending dynamics strengthens, rather than weakens, as the recovery unfolds—precisely the pattern one would expect if consumption is persistently governed by the flow of disposable income. This result echoes the tight positive relationship visible in Figures 13 and 14, which showed that spending and income recovery are closely aligned at both short and medium horizons, and formalizes the parallel recovery dynamics documented in Figures 9 and 12, where the most adversely affected consumers recovered slowest on both dimensions simultaneously.

Model fit. The R^2 of the specification in Table 1 remains stable between 0.260 at the 3-month horizon and 0.192 at 18 months. This stands in stark contrast to specifications that exclude income recovery (see Table A4 in the Appendix) and rely solely on the initial shock characteristics to explain spending dynamics, where R^2 decays from 0.205 to 0.072 over the same horizons. The fact that income dynamics account for a substantial and growing share of the cross-sectional variation in spending recovery—variation that the magnitude and composition of the initial shock alone cannot explain—is itself informative: it is consistent with a model in which the recovery is paced by the evolving flow of income rather than by a one-time stock of deferred demand.

Sample splits by liquidity status. To further probe the hand-to-mouth mechanism, we split the sample by pre-recession liquidity status. Table 2 re-estimates equation (9) separately for hand-to-mouth consumers—those who accumulated less than \$300 in net liquidity over April–December 2019—and for others. The results sharpen the full-sample findings. The coefficient on $IncomeRecovery$ is larger for hand-to-mouth consumers at both the short and long horizon (0.196 vs. 0.145 at $T_0 + 3$; 0.363 vs. 0.268 at $T_0 + 18$). The difference in estimated coefficients between the two groups is statistically significant at both horizons, confirming that the reliance on current income to fund consumption is most pronounced among the most liquidity-constrained households. These patterns are consistent with the theoretical prediction that hand-to-mouth consumers, who cannot smooth through income shocks, display the tightest coupling between income and spending paths.

Sample splits by income level. Table 3 re-estimates the specification separately for consumers with self-reported income at or below \$50,000 and those above. The coefficient on $IncomeRecovery$ is positive and significant for both groups, confirming that hand-to-mouth behavior is not confined to low-income households. However, the estimated coefficient at 18

Table 2: Spending Recovery by Hand-to-Mouth Status

	HtM = 1		HtM = 0		Difference	
	$(T_0 + 3)$	$(T_0 + 18)$	$(T_0 + 3)$	$(T_0 + 18)$	$(T_0 + 3)$	$(T_0 + 18)$
$IncomeRecovery$	0.196***	0.363***	0.145***	0.268***	0.050**	0.095***
	(0.017)	(0.011)	(0.018)	(0.012)	(0.025)	(0.016)
$ShockMagnitude$	-2.414***	-4.268***	-3.225***	-5.794***	0.811***	1.526***
	(0.068)	(0.229)	(0.119)	(0.328)	(0.137)	(0.400)
$IncomeShock$	-0.029**	-0.192***	-0.048***	-0.038	0.019	-0.154**
	(0.012)	(0.035)	(0.018)	(0.049)	(0.022)	(0.060)
$MemorablesShare$	-0.055**	-0.130	-0.071*	-0.112	0.016	-0.018
	(0.023)	(0.082)	(0.039)	(0.126)	(0.045)	(0.150)
$DurablesShare$	-0.040	-0.198*	-0.032	0.187	-0.008	-0.385**
	(0.026)	(0.102)	(0.041)	(0.136)	(0.049)	(0.170)
N. observations	10,280	10,280	7,234	7,234		
R^2	0.256	0.211	0.247	0.181		
Controls	Y	Y	Y	Y		

Note: This table re-estimates the specification of Table 1 separately for hand-to-mouth (HtM = 1) and non-hand-to-mouth (HtM = 0) consumers. A consumer is classified as hand-to-mouth if their cumulative net liquidity change over April–December 2019 is no larger than \$300. Panel *Difference* reports the difference between the HtM and non-HtM group coefficients, with standard errors from a Wald test of equality across the two subsamples. Controls are the same as in Table 1. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

month horizon is significantly larger for the low-income group, mirroring the pattern from the sample split by liquidity status. One might suspect that the negative full-sample coefficient on $ShockMagnitude$ is driven by liquidity-constrained low-income consumers and that pent-up demand could still operate among higher-income households with greater discretionary spending. The income split rules this out: the coefficient on $ShockMagnitude$ is not only negative for the higher-income group but substantially *more* negative (-4.034 vs. -2.086 at 3 months; -7.481 vs. -3.383 at 18 months), so that the failure of pent-up demand is most pronounced precisely where one would most expect it to operate. The $MemorablesShare$ coefficient is insignificant for lower-income consumers and negative for higher-income consumers—the opposite sign of the pent-up demand prediction—while $DurablesShare$ is uninformative for both groups. Across the income distribution, the dominant predictor of spending recovery remains the cumulative recovery of disposable income.

Table 3: Spending Recovery by Income Level

	Income $\leq$ 50k		Income > 50k		Difference	
	(T_0+3)	(T_0+18)	(T_0+3)	(T_0+18)	(T_0+3)	(T_0+18)
$IncomeRecovery$	0.154***	0.350***	0.171***	0.278***	-0.017	0.072***
	(0.023)	(0.013)	(0.015)	(0.010)	(0.028)	(0.016)
$ShockMagnitude$	-2.086***	-3.383***	-4.034***	-7.481***	1.948***	4.098***
	(0.076)	(0.228)	(0.118)	(0.314)	(0.140)	(0.388)
$IncomeShock$	-0.021**	-0.119***	-0.071***	-0.106**	0.050**	-0.013
	(0.010)	(0.031)	(0.023)	(0.053)	(0.025)	(0.061)
$MemorablesShare$	-0.007	0.068	-0.110***	-0.338***	0.102**	0.406***
	(0.022)	(0.080)	(0.037)	(0.119)	(0.043)	(0.143)
$DurablesShare$	-0.049	-0.188*	-0.055	0.054	0.006	-0.241
	(0.030)	(0.097)	(0.039)	(0.132)	(0.049)	(0.163)
N. observations	9,079	9,079	10,672	10,672		
R^2	0.192	0.202	0.260	0.185		
Controls	Y	Y	Y	Y		

Note: This table re-estimates the specification of Table 1 separately for consumers with self-reported income at or below \$50,000 and those above \$50,000. Panel *Difference* reports the difference between the low income and high income group coefficients, with standard errors from a Wald test of equality across the two subsamples. Controls are the same as in Table 1. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Sample splits by credit score. Table 4 re-estimates the baseline specification separately for users with self-reported "poor" or "average" (Bad) and those with "good" or "excellent" (Good) credit score. This exercise is meant to capture differences in borrowing capacity across individuals: a user with limited liquid wealth might still be able to smooth a transitory shock if they have access to credit. Thus, the sample split is an additional test of whether the income-recovery channel is truly driven by those most financially constrained. Mirroring the results for the liquidity split, the coefficient on $IncomeRecovery$ is positive and significant for both groups, but substantially larger for the Bad credit score group with the difference being statistically significant at both 3- and 18-month horizons. This result indicates that spending recovery tracks income recovery most closely for the consumers who cannot smooth through shocks by borrowing.

Taken together, the regression results strongly favor the hand-to-mouth interpretation of post-recession spending dynamics over the pent-up demand narrative. The magnitude and composition of the initial spending shock—the variables most relevant to pent-up demand—

Table 4: Spending Recovery by Self-Reported Credit Score Category

	Bad Credit Score		Good Credit Score		Difference	
	(T_0+3)	(T_0+18)	(T_0+3)	(T_0+18)	(T_0+3)	(T_0+18)
$IncomeRecovery$	0.243***	0.396***	0.141***	0.240***	0.102***	0.156***
	(0.030)	(0.013)	(0.018)	(0.012)	(0.035)	(0.018)
$ShockMagnitude$	-2.754***	-4.346***	-3.288***	-5.942***	0.534***	1.596***
	(0.113)	(0.318)	(0.114)	(0.320)	(0.160)	(0.451)
$IncomeShock$	-0.051***	-0.162***	-0.033**	-0.030	-0.018	-0.132**
	(0.018)	(0.045)	(0.017)	(0.045)	(0.024)	(0.063)
$MemorablesShare$	-0.010	0.023	-0.092***	-0.224*	0.081*	0.247
	(0.034)	(0.108)	(0.035)	(0.121)	(0.049)	(0.162)
$DurablesShare$	-0.032	0.057	-0.071*	-0.150	0.039	0.208
	(0.040)	(0.136)	(0.039)	(0.137)	(0.056)	(0.193)
N. observations	6,511	6,511	8,573	8,573		
R^2	0.294	0.261	0.247	0.160		
Controls	Y	Y	Y	Y		

Note: This table re-estimates the specification of Table 1 separately for consumers with self-reported credit score which is either "poor" or "average" (Bad) or "good" or "excellent" (Good). Panel *Difference* reports the difference between the Bad and Good credit group coefficients, with standard errors from a Wald test of equality across the two subsamples. Controls are the same as in Table 1. Robust standard errors in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

predict slower, not faster, recovery. By contrast, the cumulative recovery of disposable income is the dominant predictor of spending recovery, its importance grows over the recovery horizon, and its inclusion improves the model's ability to account for cross-sectional variation. The sample splits by liquidity status and income level further reinforce this conclusion: the link between income and spending recovery is strongest for hand-to-mouth and lower-income consumers, precisely the groups for whom liquidity constraints bind most tightly. These findings are consistent with a world in which a large share of consumers are liquidity-constrained and their spending paths are paced by the dynamics of their income, including fiscal transfers, rather than by the release of deferred demand for memorable services and durable goods.

6. Conclusion

This paper investigates household consumption dynamics during and after the COVID-19 recession using high-frequency transaction-level data on more than 34,000 U.S. consumers.

The unique coincidence of transitory shocks during the pandemic—income disruptions, public health measures, and fiscal transfers—offers an ideal setting to revisit two competing hypotheses of consumer behavior, namely *pent-up demand* and *hand-to-mouth* responses.

Our findings provide little support for pent-up demand. While memorable in-person services such as hospitality, travel, and entertainment experienced sharp declines in the spring of 2020, our data reveals that the corresponding cutback in spending did not spur a faster or stronger subsequent recovery of spending at the consumer level. Nor do we find systematic evidence of pent-up demand led by a cut-back in purchases for (semi-)durables during the recession. These results stand in contrast to models in which consumers postpone spending due to a drop in confidence or spending restrictions and later catch-up.

By contrast, our consumer-level evidence is consistent with hand-to-mouth consumption. Spending declines correlate with income losses in the COVID recession. The subsequent recovery of spending is strongly associated with the rebound of income. This finding aligns with the literature emphasizing high marginal propensities to consume out of transitory income shocks among financially constrained households.

Taken together, our analysis advances two main conclusions. First, standard liquidity-based mechanisms remain central to understanding consumption dynamics even in the presence of extraordinary - in our case pandemic-related - shocks. Second, policies aimed at stabilizing household income appear to be more effective in sustaining aggregate demand than those predicated on stimulating deferred purchases. Finally, our results invite further research on the specific conditions under which contractions in spending on durable goods and "memorable" services during recessions may spur pent-up demand during subsequent recoveries.

References

Aguiar, M., Bils, M., and Boar, C. (2024). Who are the hand-to-mouth? *The Review of Economic Studies*, 92(3):1293–1340.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–2367.

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., and Yannelis, C. (2020). How does household spending respond to an epidemic? consumption during the 2020 covid-19 pandemic. *The Review of Asset Pricing Studies*, 10(4):834–862.

Benetton, M., Mullins, W., Niessner, M., and Toczynski, J. (2025). Celebrity persuasion. *Available at SSRN*.

Beraja, M. and Wolf, C. K. (2021). Demand composition and the strength of recoveries. *NBER working paper*, 29304.

Browning, M. and Crossley, T. F. (2000). Luxuries are easier to postpone: A proof. *Journal of Political Economy*, 108(5):1022–1026.

Browning, M. and Crossley, T. F. (2009). Shocks, stocks, and socks: Smoothing consumption over a temporary income loss. *Journal of the European Economic Association*, 7(6):1169–1192.

Caballero, R. J. (1993). Durable goods: An explanation for their slow adjustment. *Journal of Political Economy*, 101(2):351–384.

Cao, J., Cui, C., Dinger, V., Holm, M. B., and Kang, S. (2025). Identifying the depreciation rate of durables from marginal spending responses. *Journal of Money, Credit and Banking*, 57(1):223–241.

Chetty, R., Friedman, J. N., Stepner, M., and Team, O. I. (2023). The economic impacts of covid-19: Evidence from a new public database built using private sector data*. *The Quarterly Journal of Economics*, 139(2):829–889.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). How did us consumers use their stimulus payments? Technical report, National Bureau of Economic Research.

Dynan, K. E. (2000). Habit formation in consumer preferences: Evidence from panel data. *American Economic Review*, 90(3):391–406.

Ganong, P., Greig, F., Noel, P., Sullivan, D. M., and Vavra, J. (2024). Spending and job-finding impacts of expanded unemployment benefits: Evidence from administrative micro data. *American Economic Review*, 114(9):2898–2939.

Gilboa, I., Postlewaite, A., and Samuelson, L. (2016). Memorable consumption. *Journal of Economic Theory*, 165:414–455.

Goolsbee, A. and Syverson, C. (2021). Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. *Journal of Public Economics*, 193:104311.

Hai, R., Krueger, D., and Postlewaite, A. (2020). On the welfare cost of consumption fluctuations in the presence of memorable goods. *Quantitative Economics*, 11(4):1177–1214.

Harmenberg, K. and Öberg, E. (2021). Consumption dynamics under time-varying unemployment risk. *Journal of Monetary Economics*, 118:350–365.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Karger, E. and Rajan, A. (2020). Heterogeneity in the marginal propensity to consume: evidence from covid-19 stimulus payments.

Meghir, C. and Weber, G. (1996). Intertemporal nonseparability or borrowing restrictions? a disaggregate analysis using a u.s. consumption panel. *Econometrica*, 64(5):1151–1181.

Parker, J. A., Schild, J., Erhard, L., and Johnson, D. S. (2022). Economic impact payments and household spending during the pandemic. *Brookings papers on economic activity*, 2022(2):81–156.

SpatialDataLab (2020). Social Distancing Index from University of Maryland (US).

Stock, J. H. and Watson, M. W. (2025). Recovering from covid. *NBER working paper series*, 33857.

Toczynski, J. (2023). Spend or invest? analyzing mpc heterogeneity across three stimulus programs. *SFI working paper*, 23-57.

Weidner, J., Kaplan, G., and Violante, G. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, 1:77–153.

Online Appendix

Transitory Shocks and Consumption Dynamics: Hand-to-Mouth vs. Pent-Up Demand

Martin Brown, Mohamed Hamoud, and Jan Toczynski

Appendix 1. Covid Pandemic Severity: Social Distancing Index

The transaction-level dataset includes the five-digit ZIP code that users reported when they first subscribed to the app. We use this information to assign county-level measures of Covid shock intensity to each individual in the sample. To do so, we employ the Social Distancing Index prepared by the University of Maryland (SpatialDataLab 2020). This county-level index measures the extent of social distancing in a given area—a value of 0 indicates no social distancing, while a value of 100 represents maximum social distancing. Using these data, we define county-level Covid Pandemic Severity for county j as the ratio between the mean index value during April and May 2020 and that in January and February 2020:

$$Covid\ Pandemic\ Severity_j = \frac{SDI_{j,04-05.2020}}{SDI_{j,01-02.2020}} \tag{10}$$

Figure A2 plots Covid Pandemic Severity at the county level.

Figure A1 visualizes the distribution of the Social Distancing Index for the pre-covid period (January and February 2020) and the covid period (April and May 2020). Before the pandemic, the index is tightly concentrated around a low value, indicating consistently minimal social distancing across counties. In contrast, the distribution for the covid period is flatter and shifted to the right, with a much wider spread. This illustrates a significant overall increase in social distancing and, crucially, a greater variation in the intensity of this change across different counties. This heterogeneity in the response to the pandemic is what our Covid Pandemic Severity measure is designed to capture.

Figure A1: Social Distancing Index Distribution

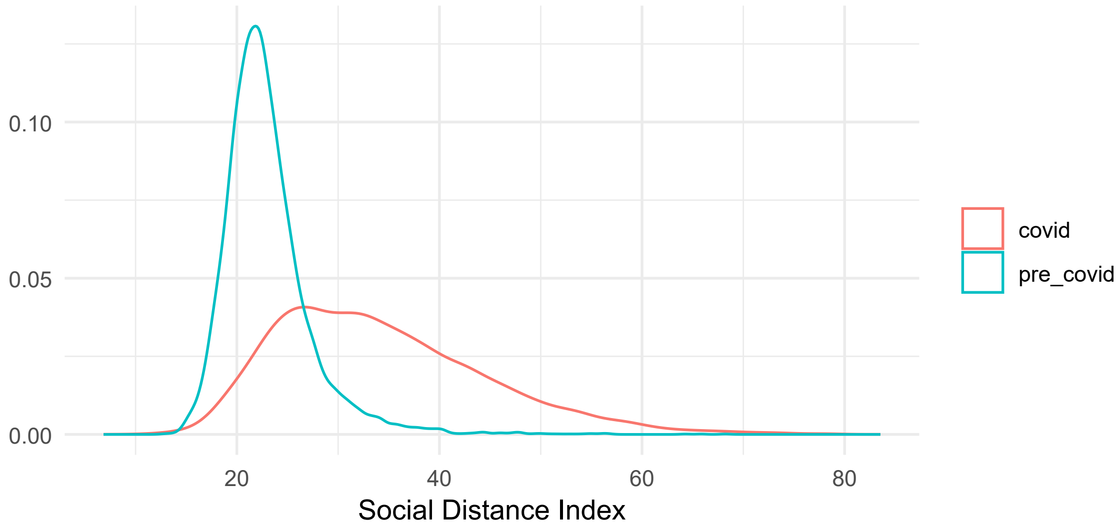

Figure A2: Covid Pandemic Severity per County

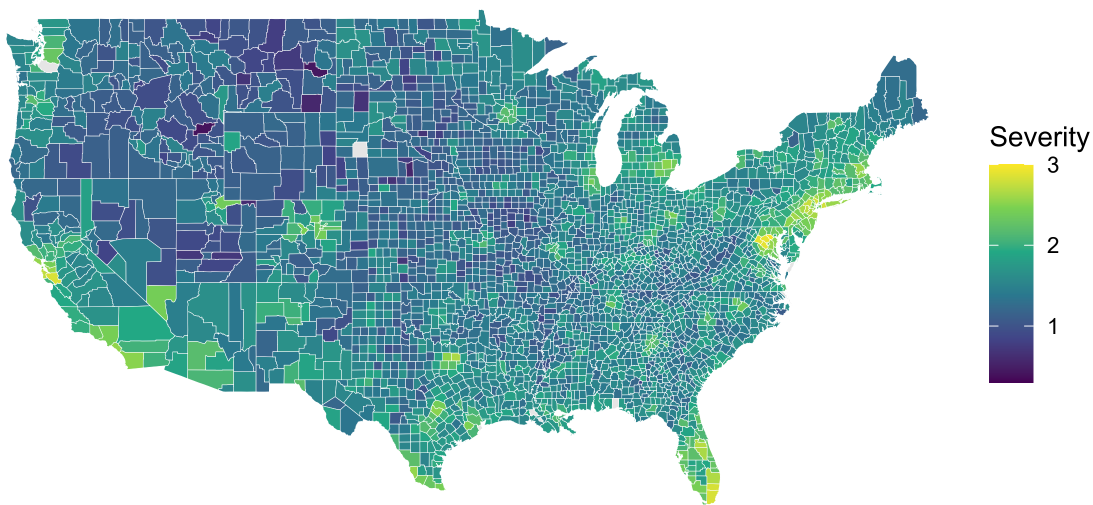

Figure A3 shows a strong correlation between the local severity of the COVID pandemic and the magnitude of the spending shock at the consumer level. Our transaction-level dataset includes the five-digit ZIP code that users reported when they first subscribed to the app. We use this information to assign county-level measures of pandemic intensity to each individual in the sample. In counties with a stronger increase in social distancing, consumers display a larger drop in their total spending.

Figure A3: COVID Pandemic Severity vs Spending Shock

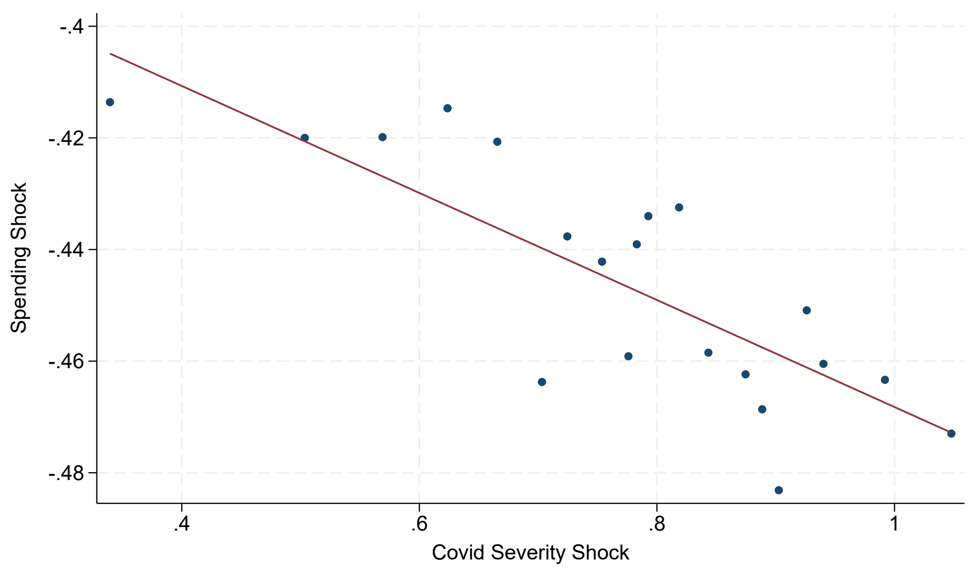

Note: the figure plots the spending shock (April 2020 vs. April 2019) against county-level changes in social distancing (April / May 2020 vs. January / February 2020).

Appendix 2. Tables

Table A1: User summary statistics - aggregator app

	Count	Mean	Std	p10	p50	p90
Demographic Variables						
Final Sample						
Gender: Male	34,453	0.31	0.46	0.00	0.00	1.00
Gender: Female	34,453	0.48	0.50	0.00	0.00	1.00
Gender: Other/not-reported	34,453	0.21	0.41	0.00	0.00	1.00
Age	34,434	31.01	7.87	23.00	29.00	41.00
Income ($1,000s)	34,401	54.71	50.32	12.00	44.00	102.00
Kids	34,346	0.54	1.02	0.00	0.00	2.00
Credit score: missing	34,453	0.29	0.46	0.00	0.00	1.00
Credit score: poor	34,453	0.18	0.38	0.00	0.00	1.00
Credit score: average	34,453	0.19	0.39	0.00	0.00	1.00
Credit score: good	34,453	0.14	0.35	0.00	0.00	1.00
Credit score: excellent	34,453	0.20	0.40	0.00	0.00	1.00
Married	34,453	0.24	0.43	0.00	0.00	1.00
End balance main account ($1,000s)	34,447	5.20	21.51	0.17	2.01	9.33
Full Sample						
Gender: Male	56,215	0.32	0.46	0.00	0.00	1.00
Gender: Female	56,215	0.48	0.50	0.00	0.00	1.00
Gender: Other/not-reported	56,215	0.21	0.41	0.00	0.00	1.00
Age	56,185	30.96	7.91	23.00	29.00	41.00
Income ($1,000s)	56,130	54.38	50.02	12.30	43.04	100.95
Kids	56,041	0.54	1.03	0.00	0.00	2.00
Credit score: missing	56,215	0.29	0.46	0.00	0.00	1.00
Credit score: poor	56,215	0.18	0.39	0.00	0.00	1.00
Credit score: average	56,215	0.19	0.39	0.00	0.00	1.00
Credit score: good	56,215	0.14	0.35	0.00	0.00	1.00
Credit score: excellent	56,215	0.20	0.40	0.00	0.00	1.00
Married	56,215	0.24	0.43	0.00	0.00	1.00
End balance main account ($1,000s)	56,207	5.01	20.25	0.16	1.98	8.97

Note: the table presents the summary statistics for the user-level variables. The first panel reports them for the final sample used in the analysis. The second panel covers the entire dataset.

Table A2: Monthly Spending – Summary Statistics

	Count	Mean	Std	p10	p50	p90
Durables						
Apparel	34,453	240.85	267.87	45.57	173.57	501.19
Car Payments	34,453	117.74	251.95	0.00	0.00	385.76
Household	34,453	182.65	542.67	12.73	77.10	411.32
Durable Goods	34,453	541.24	761.83	105.59	370.27	1097.04
Non-Durables:						
Gas and Fuel	34,453	161.09	260.84	30.56	117.38	327.18
Groceries	34,453	322.76	377.06	70.13	236.93	663.88
Retailer	34,453	188.12	216.20	20.91	120.01	432.09
Shopping Other	34,453	220.00	619.47	47.21	139.98	409.01
Bills & Utilities Other	34,453	107.24	456.62	0.00	20.72	236.18
Business Services Other	34,453	219.77	1507.71	16.34	84.30	401.99
Communication	34,453	241.26	2538.39	35.15	169.36	463.86
Education	34,453	46.40	816.85	0.00	3.81	83.40
Medical care	34,453	255.66	426.39	41.35	185.03	499.38
Private Transportation	34,453	146.41	301.96	23.52	96.30	308.14
Public Transportation	34,453	8.00	41.44	0.00	0.32	19.00
Energy and Utilities	34,453	61.55	182.73	0.00	24.95	169.29
Goods: Total	34,453	1998.40	4170.47	799.80	1531.23	3377.25
Memorables:						
Coffee Shops	34,453	25.53	37.40	1.40	14.44	62.59
Fast Food	34,453	91.15	4381.20	8.06	47.22	149.69
Food Delivery	34,453	40.58	72.93	0.00	13.74	110.60
Hotels	34,453	59.39	131.65	0.18	22.77	140.95
Pets	34,453	20.15	46.49	0.00	3.34	58.11
Recreation	34,453	144.55	2304.48	30.66	94.14	257.28
Restaurants & Bars	34,453	442.98	685.63	124.07	325.00	826.15
Sports	34,453	59.63	186.88	3.22	31.35	133.19
Travel	34,453	55.99	123.31	0.00	22.63	138.11
Memorable Services: Total	34,453	919.79	6806.96	316.49	690.54	1573.38
Consumption Spending	34,453	3459.44	9136.49	1455.54	2676.28	5763.94

Note: the table presents the breakdown of monthly user-level spending by category.

Table A3: Summary statistics

Panel A

	$\Delta Consumption$			
$MemorablesShare$	**≤-60%**	**-60 to -30%**	**-30 to -10%**	**Total**
no contribution	129	534	780	1,443
0 to 40%	2,967	3,574	1,165	7,706
40% to 80%	1,904	3,204	1,435	6,543
>80%	358	1,697	2,744	4,799
Total	5,358	9,009	6,124	20,491

Panel B

	$\Delta Consumption$			
$DurablesShare$	**≤-60%**	**-60 to -30%**	**-30 to -10%**	**Total**
No contribution	1,322	3,493	3,227	8,042
0 to 40%	3,571	4,465	1,694	9,730
40% to 80%	403	823	680	1,906
>80%	62	228	523	813
Total	5,358	9,009	6,124	20,491

This table tabulates the distribution of *MemorablesShare* and *DurablesShare* against the size of the spending shock. The shares are defined by Equation 5.

Table A4: Spending Recovery: The Role of the Initial Shock and Disposable Income

	(T_0+3)	(T_0+6)	(T_0+9)	(T_0+12)	(T_0+15)	(T_0+18)
$ShockMagnitude$	-3.301***	-4.111***	-4.354***	-5.235***	-6.041***	-6.256***
	(0.076)	(0.105)	(0.131)	(0.157)	(0.187)	(0.214)
$IncomeShock$	0.047***	0.081***	0.112***	0.142***	0.170***	0.189***
	(0.010)	(0.015)	(0.019)	(0.023)	(0.026)	(0.030)
$MemorablesShare$	-0.078***	-0.159***	-0.226***	-0.245***	-0.219***	-0.224***
	(0.023)	(0.036)	(0.045)	(0.054)	(0.065)	(0.077)
$DurablesShare$	-0.050*	-0.044	-0.066	-0.097	-0.061	-0.061
	(0.027)	(0.041)	(0.053)	(0.063)	(0.077)	(0.091)
N. observations	19,751	19,751	19,751	19,751	19,751	19,751
R^2	0.205	0.150	0.114	0.100	0.086	0.072

Note: This table re-estimates the specification of Table 1 without including *IncomeRecovery* among the independent variables. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.