Anticipatory Spending: Pre-Payment MPCs and Consumer Types*

Scott R. Baker[1] Michael Gelman[2] Lorenz Kueng[3] Seung Hyeong Lee[4]

Preliminary; click here for the most recent version.

March 26, 2026

Abstract

We quantify consumption heterogeneity by jointly estimating pre-payment MPCs – spending increases before payment arrival – and traditional post-payment MPCs, leveraging multiple pre-announced income changes per individual from various fiscal policies. We find that these MPCs exhibit higher within-individual persistence than one would expect from models that feature ex-ante identical individuals. Building on this idea of persistent characteristics, we use k-means clustering to assign individuals into one of four distinct consumer types: *Consumption smoothers* with positive but moderate pre- and post-MPCs (about 45% of the sample), the *unresponsive* with near-zero pre- and post-MPCs (21%), *late spenders* with high post- but near-zero pre-MPCs (24%), and *always spenders* with high pre- and post-MPCs (9%). Finally, we use additional moments untargeted by the algorithm (liquidity, event salience, transfer size) to map each consumer type to established consumption models, including canonical buffer-stock behavior, inattention, mental accounting, and present bias.

JEL classification: D12, D14, E21, G51

Keywords: Anticipation effects, intertemporal MPCs, excess sensitivity, excess smoothness, consumption smoothing, inattention, mental accounting, present bias, near rationality

*We thank Matthew Rognlie, Pascal Noel, Jonathan Parker (discussant), seminar participants at the ASSA, the Universities of St. Gallen and Vienna, and our colleagues at Wisconsin, Claremont, Lugano, and Northwestern for their valuable feedback.

[1]University of Wisconsin and National Bureau of Economic Research (NBER); Department of Finance, Insurance and Banking 975 University Avenue, Madison, WI 53706, US. E-mail: scott.baker@wisc.edu.

[2]Claremont McKenna College; The Robert Day School of Economics and Finance, 500 E. 9th Street, Claremont, CA 91711, US. E-mail: mgelman@cmc.edu.

[3]USI Lugano, Swiss Finance Institute (SFI), and Centre for Economic Policy Research (CEPR); Department of Economics, Via Buffi 13, 6900 Lugano TI, Switzerland. E-mail: lorenz.kueng@gmail.com.

[4]Northwestern University; Department of Economics, 2211 Campus Drive, Evanston, IL 60208, US. E-mail: seunghyeong.lee@u.northwestern.edu.

1 Introduction

Heterogeneity in marginal propensities to consume (MPCs) features prominently in recent macroeconomic models. It governs the transmission of fiscal transfers to aggregate demand, shapes the potency of monetary policy in heterogeneous-agent New Keynesian (HANK) models, and can serve as a sufficient statistic for the welfare effects of stabilization policies (Kaplan and Violante, 2014; Auclert, 2019). A large empirical literature has therefore focused on estimating MPCs and documented that they vary substantially across households (see Jappelli and Pistaferri 2010; Havranek and Sokolova 2020; Sokolova 2023 for surveys). While some of this variation correlates with observable financial characteristics, particularly with income and liquid and illiquid wealth, most is left unexplained (Kaplan and Violante, 2022).

A fundamental question remains unresolved: are households' consumption responses to income changes primarily driven by their *current economic circumstances* – such as temporarily low income or binding liquidity constraints (e.g., models building on Zeldes 1989; Deaton 1991; Carroll 1997) – or by *persistent behavioral traits* (e.g., spender-saver models building on Campbell and Mankiw 1989)? This distinction has first-order consequences for both theory and policy. If high MPCs reflect transitory financial circumstances, then the population of high-MPC households shifts with the business cycle, fiscal multipliers are state-dependent, and models calibrated to one macroeconomic regime may fail in another. If, instead, high MPCs primarily reflect persistent household types, then consumption responses to policy are more stable and predictable across environments, but single-dimensional models that generate MPC heterogeneity mainly through liquidity channels and ex-post heterogeneity in realized income shocks are misspecified. Distinguishing between these explanations requires richer data and new sources of variation beyond standard observables that the existing literature has relied upon, such as income, liquidity, or credit access.

This paper introduces such a new source of variation: households' anticipatory spending responses, *pre-payment MPCs* (or pre-MPCs for short), measured in the one week prior to the arrival of predictable income changes. It complements the post-payment MPCs typically estimated in the literature (post-MPCs), which measure spending within a given time period after receiving the income payment, like a week, month or quarter. Canonical macroeconomic models such as the permanent income hypothesis and buffer stock models impose strong restrictions on the joint behavior of pre- and post-MPCs. In these models, unconstrained consumers smooth consumption by responding when they learn about future income, not when the money arrives; constrained consumers do the opposite. The set of consumers who respond in anticipation should therefore be *distinct* from those who respond at arrival. Estimating pre-MPCs alongside the more familiar post-MPCs thus allows us to quantify the extent to which standard models can account for observed heterogeneity in individuals' consumption dynamics, and – where they cannot – offer a basis for better classifying households into consumer types.

Measuring anticipatory spending is empirically demanding. Announcements and awareness of fiscal transfers typically precede payments by only weeks or even days, requiring high-frequency transaction-level data with minimal temporal measurement error. Pre-payment re-

sponses are predicted to be small even under canonical models, since the annuity value of most transfer payments is tiny at typical spending frequencies of a week, month or quarter, necessitating large samples for statistical power. And to separate persistent traits from transitory circumstances, the data must feature an extended panel dimension so that both pre- and post-MPCs can be estimated across multiple predictable income changes within the same individual.

We turn to high-frequency financial transaction data covering more than 50,000 U.S. individuals from 2012 to 2023 and leverage a range of exogenous income increases stemming from federal and state transfer policies: three rounds of COVID-19 stimulus payments, the expanded Child Tax Credit, annual Alaska Permanent Fund Dividends, and federal and state tax refunds. These transfers are announced well ahead of the payment transaction date, allowing us to estimate pre- and post-MPCs. The payment events vary in size, frequency, salience, and timing, and most individuals in our data experience multiple events over the sample period. We validate the external validity of our data by closely replicating the main post-MPC estimates of five prior studies which have analyzed the same events but used different data sources. We develop measures of event salience using newspaper mentions and Google searches to quantify differences in awareness across events and over time.[1]

Utilizing this high-frequency data across multiple fiscal transfers, we document four sets of findings. First, we find anticipatory spending responses across all income events we study. These pre-MPCs in the week prior to income arrival typically range from 0.5 to 1.5% of the payment. They are an order of magnitude smaller than the corresponding post-MPCs.[2] While these estimates may represent a lower bound due to difficulty of measuring the exact announcement date of each event and varying household awareness, their small magnitude suggests they contribute negligibly to aggregate demand. Consequently, neglecting the pre-announced nature of most fiscal stimulus programs likely introduces only minimal downward bias in estimates of their total impact on the economy.

We find that events with higher salience as measured by news coverage and Google search intensity feature modestly larger anticipatory responses and somewhat smaller post-payment responses, consistent with more salient events producing smoother consumption paths over pre-announced income changes. While the robustness of these small average anticipatory effects across income events is notable, our main finding is what the *heterogeneity* in pre-MPCs across individuals reveals about the nature of individuals' consumption dynamics.

Second, and more importantly, we estimate pre- and post-MPCs at the individual level for each income event, and leverage the fact that we observe multiple events per person to charac-

[1]We also find qualitatively similar but quantitatively larger pre- and post-MPCs in response to regular paychecks, consistent with previous research (e.g., Olafsson and Pagel 2018; Gelman 2022). However, as this literature has documented, those pre-paycheck MPCs are often the result of individuals timing their regular bill payments (e.g., Gelman et al. 2014; Baker and Yannelis 2017; Gelman et al. 2020, 2022). The income payments we study in this paper are not as regular or frequent and therefore do not lend themselves to this type of liquidity management.

[2]Post-payment responses are much easier to detect: the t-statistic of the typical post-MPC estimate is more than four times larger than that of the pre-MPC for the same event, helping to explain why prior studies – including our own – have mostly found insignificant anticipatory effects (e.g., Broda and Parker 2014; Baugh et al. 2021; Graham and McDowall 2025).

terize persistent household-level differences in consumption responses. Averaging across events within individual substantially reduces noise while preserving the large dispersion in both pre- and post-MPCs across households.

Our key empirical finding concerns the joint distribution of pre- and post-MPCs within households. Canonical macroeconomic models – buffer stock models, the permanent income hypothesis with occasionally binding liquidity constraints – predict a negative correlation: a household with a small but positive pre-MPC signals that it is unconstrained and should therefore smooth consumption by spending about the same post payment (similar pre- and post-MPCs), while a household with a large post-MPC should be constrained and thus unable to spend in anticipation (near-zero pre-MPC). Instead, on average we find a robust and significant *positive* correlation between pre- and post-MPCs within individual, both at the event level and when averaging across events. Many consumers who increase spending in advance of income also increase spending substantially upon its arrival. This pattern is inconsistent with models where liquidity constraints are the sole source of MPC heterogeneity. Some form of behavioral heterogeneity – differences in household "type" – seems needed to account for this joint distribution.

In line with prior literature studying heterogeneity in post-MPCs, we find that observable financial characteristics, such as income or liquidity, predict which households tend to have high post-MPCs. However, this strong cross-sectional relationship between post-MPCs and financial characteristics is substantially attenuated when we include person fixed effects and only use within-person variation to absorb persistent individual characteristics. Together, these findings point to persistent behavioral traits as major drivers of heterogeneity in individual consumers' spending dynamics, consistent with Parker (2017) and Gelman (2022).

Third, motivated by this evidence of persistence consumption dynamics within person, we use the new concept of pre-MPCs to group individuals into unique consumer types. Pioneering work by Campbell and Mankiw (1989) and Mankiw (2000) uses two types – savers and spenders – to describe the one-dimensional distribution of individuals' post-MPCs. To describe the two-dimensional *joint* distribution of within-person average pre- and post-MPCs, we follow a similar approach by using k-means clustering to group individuals into four distinct consumer types. While the algorithm is agnostic about economics, we assign the clusters descriptive labels that reflect the consumption dynamics they represent: *Consumption smoothers* (about 45% of the sample) with positive but moderate and similar pre- and post-MPCs, who show no discernible consumption jump after the payment and thus represent the savers in Mankiw's terminology; the *unresponsive* (21%) with a flat consumption profile reflected by near-zero pre- and post-MPCs, with no discernible consumption increase also in the pre-period relative to their baseline spending; the *late spenders* with high post-MPCs but near-zero pre-MPCs (24%), showing a strong consumption response at payment arrival; and the *always spenders* with both high pre- and post-MPCs (9%), showing both payment sensitivity but also anticipatory spending behavior.

The resulting classifications prove stable within person: at our preferred specification of four clusters, *normalized purity* – the share of events for which an individual is assigned to their

modal cluster, scaled to equal zero under random assignment and one under perfect persistence – averages 0.55, well above what temporary circumstances alone would produce. This stability reinforces the second finding that persistent behavioral traits, rather than transitory financial conditions, are the primary drivers of MPC heterogeneity, and lends credibility to the four-type classification as a meaningful empirical decomposition of our sample.

Fourth, rather than reverse-engineer a new theory, we map these consumer types as closely as we can to consumption models that are already well established in the literature. We do this by examining how the four clusters differ along auxiliary dimensions other than the pre- and post-MPCs used by the algorithm to form them, including financial and event characteristics such as income, liquidity, average propensity to consume (APC), event salience, or transfer size.

Always spenders and late spenders both have relatively low levels of income and liquidity and exhibit high APCs with frequent overdraft fees – characteristics associated with time-inconsistency and impulsiveness typical of present-biased consumers (Kuchler and Pagel, 2021; Gelman, 2022; Aguiar et al., 2025). However, late spenders are distinctively unresponsive to liquidity or to the salience of upcoming payments: they do not anticipate spending even when they have temporarily more financial resources or when news coverage of an event is high. Consumption behavior of always spenders is therefore well represented by present-biased consumers who are impatient and time-inconsistent, wanting to pull spending forward but then being tempted to also splurge when the cash hits their bank account. Late spenders, on the other hand, could be constrained hand-to-mouth consumers, but because their relatively low liquidity is not a temporary circumstance but rather a persistent state they find themselves in, their behavior is better described as that of mental accountants: they choose not to pull spending forward even when they have the necessary resources, showing a behavioral reluctance to spend before income is "in hand" (Thaler, 1985; Shefrin and Thaler, 1988).

Consumption smoothers and the unresponsive in turn both have high liquidity, few overdrafts, and low APCs. Consumption smoothers are the most responsive to news and salience and are therefore best described as the rational consumers of canonical macroeconomic models, i.e., the savers in Mankiw (2000). The unresponsive resemble those rational consumers on financial observables but respond only to large payments relative to income, consistent with models of inattention and near-rationality (Akerlof and Yellen 1985; Cochrane 1989; Kueng 2018).

This analysis highlights the key role of the pre-MPC dimension in the classification of consumer types. With only two clusters, the algorithm recovers the standard spender–saver split central to much of the existing literature, with spenders having low and savers high liquidity. But with four clusters, the pre-MPC separates households *within* these two groups that are otherwise indistinguishable on observables, including liquidity. Among high-post-MPC households, those with elevated pre-MPCs can by classified as the traditional spenders in the framework of Mankiw (2000) – hence labeled *always* spenders and best modeled as present-biased –, while those with near-zero pre-MPCs are a different class we call *late* spenders and best described as mental accountants. Similarly, low-post-MPC households representing traditional savers in the framework of Mankiw (2000) are split into those with small positive pre-MPCs, which we classify

as rational consumption smoothers, while those with near-zero responses on both dimensions are classified as the unresponsive and best described by models of inattention or near rationality.

This methodology of utilizing pre-MPCs can in principle be applied to any setting with multiple income shocks and high-frequency spending data, providing a portable approach to identifying behavioral types. While this is not a feature of the k-means clustering algorithm, we find that in our context, varying k reveals a nesting structure, although not perfectly. Going from two clusters – savers and spenders – to three splits savers into consumption smoothers (Mankiw's savers) and the unresponsive, while going to four further splits Mankiw's spenders into always spenders and late spenders. Going above four cluster mostly splits the large group of consumption smoothers into ever finer subgroups without obvious mapping to additional consumption models.

Our findings have direct implications for macroeconomic modeling and policy design. HANK models typically generate MPC heterogeneity through a single channel: the endogenous distribution of liquid wealth (Kaplan et al., 2018). Our results suggest that while this approach captures an important dimension of heterogeneity, the positive within-household correlation between pre- and post-MPCs and the existence of distinct behavioral types within liquidity groups indicates that a richer specification of household heterogeneity is needed – one that incorporates variation in discount factors, attention, or mental accounting rules alongside variation in asset positions. The shares of types we quantify empirically can help discipline the calibration of such models.

Our findings also speak to the forward guidance puzzle (Carlstrom et al. 2015; Kiley 2016; McKay et al. 2016; Del Negro et al. 2023) – the well-known disconnect between the large effects of forward guidance as a policy tool predicted by representative-agent and many heterogeneous-agent models and the more modest effects observed empirically. Pre-payment MPCs directly measure households' willingness to adjust consumption in response to anticipated future income changes. Our finding that roughly 24% of households are well described as mental accountants who do not pull spending forward even when having the financial capacity to do so, and that an additional 21% appear largely inattentive to these income events, provides micro-level evidence for why the aggregate anticipatory consumption response is attenuated relative to full-information rational-expectations benchmarks.

Literature.— This paper contributes to three main strands of literature on dynamic consumer behavior. First, it provides evidence on spending responses prior to future income increases. While there is a smaller and older literature that tests for anticipation effects, it mostly relies on macroeconomic time series data (e.g., Blinder 1981; Flavin 1981; Deaton 1986; Poterba 1988; West 1988; Campbell and Deaton 1989). This literature has coined the term *excess smoothness* to describe the pattern that (aggregate) consumption responds too little to new information. Based on the logic of the permanent income hypothesis, people should respond to future income changes when they learn about them, and recent survey-based studies find that individuals react to hypothetical future income events (e.g., Fuster et al. 2021; Colarieti et al. 2024).[3] However,

[3]Other work has also shown that households may respond in advance to anticipated non-income events such as future tax changes. For instance, households pull spending forward prior to anticipated

empirical findings on anticipatory spending behaviors in response to future income changes using observational microdata are mixed (e.g., Agarwal and Qian 2014; Broda and Parker 2014; Baugh et al. 2021; Caldwell et al. 2023; Thakral and Tô 2024; Graham and McDowall 2025). Using detailed transaction-level data, we estimate households' anticipatory spending behaviors for several income events that differ in amounts, timing, and frequency, and develop measures of event salience using newspaper mentions and Google searches to quantify differences in salience across events and over time.

Second, this paper relates to the large modern literature estimating MPCs out of exogenous and predictable income changes and tests for excess sensitivity relative to canonical models' null hypothesis of consumption smoothing. While previous studies, including some of our own, mostly focused on single events, this study examines numerous events with varying characteristics and analyzes how households' responses vary across different events using high-frequency data. It thereby contributes to the large literature explaining household heterogeneity in MPCs from anticipated income changes (see the survey articles mentioned above as well as more recent studies such as Andreolli and Surico 2021; Crawley and Kuchler 2023; Commault 2024; Indarte et al. 2024; Koşar et al. 2024). Our key advance over this literature is the use of pre-MPCs as an additional dimension that, combined with post-MPCs, can distinguish behavioral mechanisms that are observationally equivalent when only post-MPCs are available.

Finally, this paper contributes to a fast-growing literature that incorporates behavioral frictions into dynamic models of individual consumption, such as near-rationality, inattention, mental accounting, present bias, and hyperbolic discounting (e.g., Parker 2017; Kuchler and Pagel 2021; Maxted forthcoming; Lee and Maxted 2023; Mijakovic 2023; Gelman and Roussanov 2024; Hamilton et al. 2024; Indarte et al. 2024). We advance this literature by providing the first empirical quantification of the population shares that fall into each of four canonical consumer types, based on revealed consumption behavior rather than hypothetical survey responses or model-imposed assumptions. This classification, grounded in the observable joint distribution of pre- and post-MPCs, offers a new empirical benchmark for calibrating and testing structural models of consumption with behavioral heterogeneity.

The rest of the paper proceeds as follows. Section 2 describes the data and measurements and provides institutional background for the different fiscal interventions. Section 3 discusses the empirical strategy and presents estimates of average pre- and post-MPCs and their joint distribution. Section 4 uses individual-level pre- and post-MPCs to classify consumers into types and maps these types to consumption models using variables not targeted by the clustering algorithm. Section 5 concludes.

sales tax increases (Baker et al. 2018, 2021; Angelucci et al. 2024) and prior to federal income tax changes (Kueng, 2014).

2 Data and Events

2.1 Transaction Data from Linked Financial Accounts

We mainly use de-identified proprietary account-level financial transaction data from a major U.S. financial aggregation and analytics firm to construct our dataset. Primarily contracting with financial institutions, our data provider aggregates financial information across users' accounts and assists these institutions in offering personal financial management services. We use the terms user, individual and household interchangeably. As part of 'open banking,' the platform's data access is based on agreements with bank and non-bank partners rather than consumers, ensuring comprehensive coverage and mitigating selection biases when consumers must opt-in to data sharing. Consequently, no additional selection bias occurs among the population of individuals with financial accounts once users' financial institutions form an agreement.

The data enable us to track bank, credit card, and debit card account transactions. For each transaction, we observe the precise date and amount, the transaction category (from one of 43 categories), and a transaction description. Additionally, we capture the merchant's name and its location for most transactions. Although we do not observe the user's demographic information (age, gender, and race), we can infer the user's residence at the zip code level, along with detailed financial characteristics (e.g., balances and overdrafts).

The data cover the complete financial flows across millions of Americans. Due to the platform's rapid increase in clients during the first two years, we focus on the period from January 2012 to September 2023. The dataset covers numerous financial institutions, allowing us to observe income and spending for a given user, potentially across several separate financial institutions. However, it is still possible that these accounts do not represent the totality of users' total income and spending, as some customer accounts may be excluded if held at institutions not contained within our database. Therefore, we follow Aiello et al. (2023a) to construct a representative sample of high-quality users based on the quality measure of the data provider and the tenure of the account. For computational convenience, we work with a 2% random sub-sample of households from January 2012 to September 2023.[4]

Table 1, panel A, reports the summary statistics for the spending households in the final dataset, which is made up of the daily spending and income transactions of over 50,000 users. We define consumption based on the categories of consumption, as described in Appendix Table A.1, and aim to follow the definition of other studies (Kueng, 2018; Di Maggio et al., 2022; Graham and McDowall, 2025). Our main consumption variable is spending in non-durable and service categories, while total spending combines spending on non-durable goods and services, durable goods, and other expenditures such as charitable giving and gifts as well as ATM withdrawals and check payments.

Panel B shows monthly income of households. We define income as including not only salary

[4]To obtain an adequate sample size for our estimates of the spending response to the Alaska Permanent Fund Dividend, we oversample residents of the state of Alaska. Since the population of Alaska is less than 0.25% of the total U.S. population, this does not affect the estimation of the responses to other income transfers.

but also other types of income, such as pensions, bonuses, and interest income. In addition, it presents several household financial characteristics (at the household-by-event level) related to liquidity and credit constraints: bank account balances, the amount of credit available across credit cards, and the number of bank overdrafts. Balance and credit data are obtained via snapshots of accounts for a subset of users. For each account, a maximum of four snapshots are available: May 2022, November 2022, May 2023, and September 2023. Bank balances and available credit levels before and after these dates are imputed on the basis of total transaction flows within the accounts. Bank overdrafts are identified from transaction description strings within bank accounts that describe an overdraft fee being charged.

Given that the announcement and occurrence of income changes from fiscal policy events are close in time, it is essential to use transaction-level data to measure anticipatory spending effects. However, one concern is the representativeness of the transaction data. To address this, previous studies that use the same dataset for different settings compare it to the Census Retail Sales Surveys and demonstrate its representativeness (e.g., Balyuk and Williams 2021; Di Maggio et al. 2022; Aiello et al. 2023a,b).[5] Except for households without bank accounts, our sample broadly represents U.S. customers in terms of income, spending, and geography. Additionally, we compare our dataset to those used in other studies employing different transaction-level and survey data to demonstrate its representativeness (see Table 3). We discuss the results of these replications of prior MPC estimates in Section 3.1.

2.2 Salience and Attention: Google Search and Newspaper Data

In order for households to respond in anticipation of these events, they need to first be aware of the impending arrival of income. We use two additional sources of data to develop proxies for the salience of the various income shocks. The first utilizes data from Google Trends to build a metric that tracks the frequency of Google search queries regarding particular terms. The second uses data from newspapers to measure the fraction of newspaper articles that are written about a given event. These measures are designed to track both active household searching for information about a specific event and the prominence of news about a specific event.

For the first metric, after providing a search term and a date window, Google Trends returns a daily metric called relative search volume, which ranges from 0 to 100. It is not an absolute search count, but rather a normalized metric that represents the proportion of searchers for a term relative to the total search volume within a specified time frame.[6]

Prior work, such as Da et al. (2011), Choi and Varian (2012), Baker and Fradkin (2017), and Baker et al. (2021) demonstrate the utility of using Google Trends to measure contemporaneous attention across households for a wide range of economic and financial topics.

For our newspaper based metric, we turn to the Access World News Newsbank database. We use the same search term and date window as the Google search data to find relevant articles for each event that we study. We then normalize the daily number of articles by the total number of

[5] See Baker and Kueng (2022) for a review of recent studies that use transaction-level data.

[6] The search terms are stimulus, child tax credit, ctc, permanent fund, pfd, and tax refund.

articles published in all states for national level events like the COVID stimulus programs and the Child Tax Credit. For the Alaska Permanent Fund Dividend events we only normalize the measure by total articles published in Alaska.

2.3 Policy-Induced Income Changes

Table 2 lists several types of fiscal income transfers we recognize to estimate anticipatory spending effects. We select events discussed in prior literature that are observable in our data during the sample period 2012–2023. The various transactions are identified within the data based on the description field, and if necessary also merchant and payment size.

The first COVID-19 stimulus has a description denoting a generic tax refund, so additional restrictions are employed based on the precise dollar amounts of the stimulus check. For the second and third COVID-19 stimulus checks, transaction descriptions are more precise, identifying the relevant transactions as 'TAXEIP2' or 'TAXEIP3', and dispensing with the need for utilizing other filters. The Child Tax Credit has a description denoting each transaction as a 'CHILDCTC' payment. Federal tax refunds and payments are identified based on merchant and description, allowing for transactions of any size (e.g., 'USATAXPYMT' or 'IRS TREAS 310 TAX REF'). The Alaska Permanent Fund Dividend disbursements typically are denoted as a 'PFD' payment, but in some years this description is absent, necessitating us to identify the transactions based on the precise dollar amount of the dividend in a given year (e.g., multiples of $1,884.00 for 2014).

Finally, regular salary transactions are identified as those categorized as "Salary/ Regular Income." For replicating existing MPC estimates, we utilize all transactions. However, for our main specifications, we focus on periods during which households receive income as monthly payments instead of bi-weekly or weekly payments, ensuring that the pre- and post-periods of regular income payments do not overlap and have sufficient gaps between them.

Panel C of Table 1 shows the distribution of each income transfer. The transfers differ in their magnitude and also in their frequency and periodicity. This heterogeneity allows us to examine how individuals respond to different income magnitudes and to the salience of each event. Additionally, by observing multiple events for each individual, we can analyze responses both across and within individuals.

2.3.1 Announcement and Payment Dates

Table 2 shows the payment date indicating the date at which the payments are distributed to most households, the announcement date based on occurrences such as the passing of a legislative bill or the official announcement of the policy, and the predicted announcement date according to two different salience measures together with the two salience measures. We define the predicted anticipation date as the date prior to the payment date that has the highest value of our coverage metric.

The table shows that of the 15 official announcement dates, the earliest of the two predicted announcement dates based on Google searches and newspaper mentions identifies nine events exactly and three events within two weeks of the official announcement. There are two events

that are off by over a month.[7] The average absolute difference between official and predicted announcement date is three days once we exclude the two large outliers. This alignment provides credibility to our salience measure.

As described above, our newspaper-based salience measure consists of the number of articles that mention the event, normalized by either a national or state measure based on the event. For example, newspaper articles related to COVID stimulus payments are normalized by the total number of articles in all states while newspaper articles related to the Alaska Permanent Fund Dividend are normalized only by the number of articles in Alaska.

In terms of news salience, the COVID stimulus payments have the largest measures reflecting the very publicized nature of the events. The child tax credit has the smallest measure, suggesting it is not as widely reported as the other events. Looking at the Alaska Permanent Fund Dividend, the largest values for search occur in 2018 and 2019 while the the largest values for news occur in 2021 and 2022. These were all years in which the dividend amount was higher relative to the surrounding years (see Appendix Figure A.1).

The table also shows that, excluding the two outliers, there are on average 13 days between announcement and payment dates. The relatively short time between announcement and payment dates highlights why high frequency data is essential for estimating anticipation effects. Furthermore, this relatively short timeline means that we have to limit the estimation of the anticipation effects to one week or two in most cases.

2.3.2 Salience of Events

One precondition to any anticipatory spending response is the extent to which consumers are aware of the upcoming payments. We measure salience of these events using three different methods. The first uses Google search data regarding a particular event and is a more active measure of what information individuals are actively seeking out. One drawback of this measure is that it is a relative measure across time within a search term. This feature makes it difficult to compare salience across events.

To address this concern, we also use the number of newspaper articles that mention a particular event. This is a more passive measure of what information is available to individuals without knowing if this information is processed by consumers. However, it allows us to better compare salience across events. Finally, we are able to observe whether a given consumer has had experience with a particular event type previously. For instance, whether a consumer has previously received an Alaska Permanent Fund Dividend. We use an indicator for being a repeat event as a final measure of salience at an consumer-specific level.

[7] The two events are the child tax credit and the 2018 Alaska Permanent Fund Dividend. For the Child Tax Credit, the mismatch is likely related to the fact that the official announcement was four months ahead of the disbursement. Furthermore, the Child Tax Credit was jointly announced with COVID Stimulus 3 as part of the American Rescue Plan. It is likely that most of the consumer and media attention focused on the stimulus payments rather than the child tax credit. For the Alaska Permanent Fund Dividend, 2018 was an anomaly as the dividend amount was announced in May rather than the typical September timing although the disbursement was consistent with previous years. The early announcement was tied to special legislation that was passed regarding the use of the Alaska Permanent Fund.

We consider these methods complementary. Google search data captures the demand for information and newspaper article data the supply of information about an event. Moreover, Google search data can be thought of as measuring salience within events and newspaper articles measuring salience across events. For example, a high Google search data value leading up to the payment date of an event shows that individuals are actively seeking information about the payment before it occurs. On the other hand, newspaper article data can provide insight into how well-covered different events are relative to each other.

The final two columns of Table 2 show the two salience measures for each payment event. As expected, our newspaper salience values are higher for the various COVID stimulus payments compared to the child tax credit as the stimulus was widely reported on in the news media. Similarly, we find that the Alaska Permanent Fund receives more newspaper coverage in years where its dividend payment is larger.

3 Pre- and Post-MPCs

Because our primary goal is to study consumer behavior and discriminate between different theories or models, we estimate marginal propensities to consume (MPCs) by focusing on spending on nondurables and services. In other specifications, we use total consumer expenditures by adding in spending on durable goods to estimate marginal propensities to spend (MPXs), because they are typically the object of interest for policymakers managing aggregate demand (Laibson et al., 2022).

3.1 Replication of Prior Post-MPC Estimates

While the focus in this paper is on understanding the anticipatory behavior of consumers in response to forecastable income payments, the consumption response *following* the particular income payments (i.e., following the 'cash flow date') used in this paper – what we call the post MPC in the next section – has been previously estimated by researchers in a range of studies. Partly to demonstrate the external validity of our estimates and representativeness of our sample, we first show that, when following the relevant methodologies of the prior studies, our data can recover these 'post' MPCs documented by these researchers.

Table 3 shows that we can replicate the results of these previous papers when mirroring the specifications those authors laid out. For each fiscal transfer, we chose a leading published article that also uses transaction-level data, except for the Child Tax Credit, where we use an unpublished study using the Consumer Expenditure Survey (CEX) as our benchmark (Schild et al. 2023), because we are not aware of published articles that study this policy using transaction-level data. For consistency, if available, we replicate the monthly MPCs or MPXs and the papers' main specification otherwise. We match the sample periods of each study when it is available. We are able to replicate previous estimates closely, despite the fact that the data source, sample, and even time span are often quite different between our paper and these other studies. Moreover, while we closely followed the data cleaning and sample selection processes of these papers, the consumption measures could differ slightly due to variations in how the datasets categorize

transactions.

For example, previous studies used various private transaction-level and public survey datasets, as well as different sample periods (e.g., periods earlier than January 2012). Graham and McDowall (2025), for instance, utilize transaction-level data from a single U.S. financial institution to estimate the post-payment of tax refund arrival on household spending across the liquidity distribution. We mirror their approach, calculating the MPC across all expenditures in the three-months after refund arrival. While they find an MPC of approximately 0.42 across 2014 and 2015, we recover an estimate of 0.37 in our own data.

Schild et al. (2023), in contrast, utilize data from the CEX to investigate the post-payment of the Child Tax Credit on both overall and child-related spending. While we cannot precisely mirror their specification given the differences in measurement inherent to the CEX as compared to our transaction data, we recover MPXs ranging from 0.21–0.40, depending on the categories included in our measure of consumption, as compared to their estimate of 0.44.

Finally, we replicate three of our own studies which estimate the MPC out of regular paycheck income (Gelman et al., 2014), out of the annual Alaska Permanent Fund Dividend (Kueng, 2018), and out of the first COVID stimulus payments in 2020 (Baker et al., 2023).[8] Despite using a different source for the transaction-level income and spending data, we obtain very similar estimates, 0.06 vs. 0.07 for the MPC during the first day after receiving a regular paycheck, 0.10 vs. 0.11 for the monthly MPC out of the Alaska dividend, and we match the monthly MPX of 0.22 out of the first COVID stimulus payments exactly despite using a different source of transaction-level data.

Overall, we take these results as evidence for the broad utility of this data source in providing reliable and externally valid measurements of individual consumption behavior.

3.2 Pre- and Post-Payment MPCs

Moving beyond these replication exercises, we next provide standardized MPCs for each income transfer, holding constant our specification rather than mirroring the disparate approaches for

[8]We are aware of only two working papers that study the spending response to Stimulus 2 or Stimulus 3, and their estimates diverge somewhat. Karger and Rajan (2021) use transaction-level data and report an MPC out of Stimulus 2 of 39% over a two-week period. Their sample over-represents lower-income households, which potentially leads to a higher estimate than in the full population. Parker et al. (2022) use the nationally-representative CEX to study the response to all three COVID stimuli and report relatively low MPCs over a three-month period, although the estimates' standard errors are very large because of the small sample of CEX respondents and the considerable measurement error. They find MPCs of 0.102 and 0.083 out of Stimulus 1 and 2, receptively, and only 0.009 for Stimulus 3. However, the authors document substantial under-reporting of payment receipt by survey respondent, which has a large effect on the point estimates. After imputing payment receipt and size using the government's distribution rule, these three MPCs increase to 0.250, 0.262, and 0.126, respectively. These corrected estimates closely align with our estimates of three-months MPCs reported in Table 4 of 0.241 and 0.159 for Stimulus 1 and 3. Because Stimulus 2 occurred only 11 weeks before Stimulus 3, we do not estimate the three-month MPC. However, the one-month MPCs out of all three COVID stimuli are similar, suggesting that a three-month MPC of 0.262 for Stimulus 2 is plausible.

each prior study. We estimate post-payment MPCs using the following regression specification:

$$c_{i,t} = \beta_{post} \frac{Post_{i,t} \times Amount_i}{Days_{i,t}} + \alpha_{i,dom(t)} + \alpha_t + u_{i,t}, \tag{1}$$

where $c_{i,t}$ is daily non-durable (respectively total) spending of individual i in date t, $Post_{i,t}$ is an indicator that equals one if date t is after individual i receives the income transfer and zero otherwise, $Amount_i$ is the amount of the income transfer, $Days_{i,t}$ is the total time period, in days, over which we estimate the MPC (respectively MPX), e.g., 7 days for the one-week and 30 days for the one-month MPCs. $\alpha_{i,dom(t)}$ are individual-by-day of month fixed effects and α_t are date fixed effects. We use $\alpha_{i,d(t)}$ to account for regular spending patterns that occur on a specific day of the month (e.g., utility bills or rent and mortgage payments). Standard errors are clustered at the level of individual accounts.

The coefficient of interest, β_{post}, estimates cumulative MPCs after receiving an anticipated income transfer by including indicators for the days after the income transfer, where each day's indicator is equal to 1 times the amount of the transfer divided by the number of days in the relevant period. It flexibly captures the excess spending for a given period following the arrival of the income. We scale by the size of the income transferred to each individual to allow responses to vary with the amount of money received. We can examine the cumulative MPC over different time periods by including more or fewer daily indicators before or after the event.

Similarly, to estimate pre-payment MPCs, β_{pre}, we use the following regression specification:

$$c_{i,t} = \beta_{pre} \frac{Pre_{i,t} \times Amount_i}{Days_{i,t}} + \alpha_{i,dom(t)} + \alpha_t + u_{i,t}, \tag{2}$$

where $Pre_{i,t}$ is an indicator that equals one if date t is before individual i receives transfer payment $Amount_i$ and zero otherwise.

Specifications (1) and (2) employ several sources of identifying variation in both the receipt, timing, and amount of income transfers conditional on receipt. Each of the events are only received by a subset of individuals, many of the event types have substantial variation in the amount of income received, and several events also have substantial variation in when precisely the income was received by an individual (see for example Figure 2 in Parker et al. (2022) for the variation generated by the three COVID Stimulus programs, reproduced in Appendix Figure A.9 for convenience).[9]

3.2.1 Excess Sensitivity: Post-Payment MPCs

Table 4 lays out the nondurables MPCs out of the listed income transfers for three different post-payment periods: (i) the week, (ii) the month, and (iii) the three months following the income transfer. For each of these post-payment MPCs, we find broadly consistent consumption

[9] As an alternative estimation approach, following the recent literature on two-way fixed effects models with staggered treatment (e.g., Borusyak et al. 2024), we also estimate a fully dynamic distributed lag specification at daily frequency surrounding each event date. We then cumulate the daily responses to a one-week (see Appendix Figure A.2), one-month, and three-month MPCs, finding results consistent with our more aggregated approach outlined above.

responses across payment events. MPCs generally range from 0.03 to 0.11 in the week following the income's arrival to approximately 0.15 to 0.39 during the first three months after the transfer.[10] These ranges of post-payment MPCs are in line with those of a substantial prior literature related to excess sensitivity tests, which test the prediction of the canonical models such as the permanent income hypothesis that consumption of individuals with sufficient liquidity should not systematically respond to predictable income changes (e.g., Jappelli and Pistaferri 2010; Fuchs-Schündeln and Hassan 2016; Carroll et al. 2017; Havranek and Sokolova 2020; Sokolova 2023).

3.2.2 Anticipatory Spending: Pre-Payment MPCs

Table 4 also examines the pre-payment MPCs for one week prior to income arrival. We do find that households increase spending in the week prior to the arrival of income, but that magnitudes for most events that are are 70%–90% smaller than their spending responses following arrival.

To the best of our knowledge, we are the first to document small but statistically significant anticipatory spending effects on average, and robustly across different income transfers. The majority of studies report insignificant anticipatory spending effects (Broda and Parker, 2014; Baugh et al., 2021; Graham and McDowall, 2025), and only very few document significant effects (Agarwal and Qian, 2014; Kueng, 2014; Caldwell et al., 2023).[11] A likely reason is that even if individuals' spending behavior follows that of canonical consumption models, it is empirically challenging to estimate the anticipatory spending results, as the theoretically predicted response size is small, even for responding households, because the annuity value of income changes is typically very small, leading to spending responses that would be even smaller due to consumption smoothing (Kaplan and Violante, 2014). Furthermore, there is heterogeneity in the salience of events and it is likely that individuals differ in their ability to acquire and process information. They also differ in the subjective probability that they assign to the event occurring. Measuring this individual heterogeneity is very difficult, which makes it empirically challenging to estimate these anticipation effects, requiring data with a large sample size and with observations recorded at high frequency and with minimal measurement error.

Appendix Table A.2 reports the corresponding MPXs by looking at total spending rather than nondurable spending responses, finding similar patterns of smaller, but non-zero, anticipatory spending responses (pre-MPXs) and sizable responses following income arrival (post-MPXs). Magnitudes of MPXs are mechanically larger given the larger absolute dollars that we track across total spending relative to non-durable spending.

Appendix Table A.3 shows the response of spending in the week before and after the public announcement of a future income transfer. Overall, we find small and often insignificant spending responses in anticipation of these announcements, though the the post-announcement period often features a spending response on the order of 0.01–0.03.

[10] Differences between MPCs reported in Table 3 and post-payment MPCs in Table 4 stem from following each previous paper's main specification in Table 3 (including sample selection, data cleaning, variable definition (e.g., MPC vs. MPX), etc.) while using a common specification across all events in Table 4.

[11] Thakral and Tô (2024) note a relationship between a post-payment MPCs and the duration of time between announcement and income receipt, but does not find evidence for sizable anticipatory spending.

3.3 Event-Specific Individual MPCs

To understand how consumption responses vary across events and individuals, and to identify the determinants of these heterogeneous responses, we estimate MPCs at the individual-event level by leveraging the dataset's structure, in particular its long panel dimension and the fact that all individuals experience multiple events. Specifically, we estimate individual event-level pre- and post-payment MPC using the following regression specification:

$$c_{i,t} = \beta_{post,i} \frac{Post_{i,t} \times Amount_i}{Days_{i,t}} + \alpha_{i,dom(t)} + \alpha_t + u_{i,t}, \tag{3}$$

$$c_{i,t} = \beta_{pre,i} \frac{Pre_{i,t} \times Amount_i}{Days_{i,t}} + \alpha_{i,dom(t)} + \alpha_t + u_{i,t}, \tag{4}$$

where $\beta_{pre,i}$ and $\beta_{post,i}$ are allowed to vary across individuals and payment events and thus estimate heterogeneous cumulative MPCs for each individual and each event separately.

3.3.1 Unconditional and Within-Person Distributions of MPCs

Figure 1 plots several distributions of one month post-payment nondurables MPCs across our sample. (Appendix Figure A.3 shows the same when using MPXs.) First, the red solid line denotes the kernel density of the marginal propensity to consume out of pre-announced one-time income changes in the month after the arrival of income, allowing for multiple observations per individual. Some fiscal policies have multiple payments across years, such as tax refunds or the Alaska dividends. We treat each annual payment as a separate event. In line with previous studies, we find that the individual consumption response coefficients exhibit a relatively wide dispersion around their mean. The distribution peaks near zero but has a pronounced positive skew, consistent with the positive post-MPC seen in Table 4.

These thick tails, similar to the post-payment MPCs out of the 2001 and 2008 stimulus checks computed from CEX survey data in Misra and Surico (2014) and Lewis et al. (2024), are partly a result of noise in consumption measured at relatively high-frequency levels. In fact, the thickness of tails using actual spending data shows that measurement error in survey data is unlikely the main driver of this dispersion. Similarly, the distribution of individual post-payment MPCs estimated in Boehm et al. (2025) based on a randomized controlled trial (RCT) with customers of a French bank – and also using transaction data as in this paper – shows similar dispersion even after constraining the distribution to have only positive support. Jappelli and Pistaferri (2014) also document considerable heterogeneity in MPCs based on responses to hypothetical-scenario questions.

We then extend this analysis further in Figure 1. The blue dashed line plots the same distribution after *averaging* these user-events across all events within a user. Averaging coefficients across events within individual removes much of the noise contained in the unconditional distribution of single-event MPCs, yielding a more compressed, but still positively skewed distribution with thinner tails.

Finally, the black dotted line represents the distribution of averaged user-level MPCs for

only those users where we observe 10 or more income events. We note a further narrowing of the distribution. Despite this narrowing after averaging over many events within a user, the distribution of within-person average post-MPCs shows sizable dispersion, consistent with persistent personal traits playing an important role in determining consumption responses.[12]

3.3.2 The Joint Distribution of Pre- and Post-Payment MPCs

The large average post-payment MPCs and the small but statistically significant average pre-payment MPCs shown in Table 4 across several fiscal transfers events is consistent with canonical macroeconomic models of rational consumption choice under uncertainty, such as one- and two-asset buffer stock models (e.g., Carroll 2001; Kaplan and Violante 2014).

If excess sensitivity of consumption to predictable income payments was mainly a product of credit and liquidity constraints as in those models, we would expect to see a *negative* relationship between pre- and post-payment MPCs, with constrained individuals unable to increase spending prior to the actual arrival of income. That is, the canonical models impose strong testable restrictions on the joint distribution of pre- and post-MPCs: For rational forward-looking consumers $\beta_{pre,i} > 0$ and $\beta_{post,i} \approx \beta_{post,i}$ due to the desire to smooth consumption over time, while for rational but credit-constrained consumers, $\beta_{pre,i} = 0$ due to the constraint and $\beta_{post,i} \gg 0$.

Figure 2 plots the relationship between these two objects across individuals. (Appendix Figure A.4 shows the same when using MPXs instead.) There is a strong *positive* relationship between the average pre-payment MPC for an individual and their average post-payment MPC. We see that post-payment MPCs are highest for those who also have high levels of spending in the week before income arrival, contrary to the predictions of the canonical consumption models.

Appendix Table A.4 displays results of regressions of one-week pre-payment MPCs on one-week and one-month post-payment MPCs, confirming this graphical relationship while accounting for event and household fixed effects, as well as other time-varying individual financial characteristics such as bank balances, available credit, income, recent overdrafts, and after controlling for the person's average propensity to consume (APC) and average salary MPCs.

3.3.3 Predictive Power of Pre- and Post-MPCs Across Events

Table 5 examines the persistence of these pre- and post-payment MPCs within household across events. We regress pre- and post-payment MPCs on lagged values from prior events, including event specific fixed effects and household-level financial characteristics.

In columns 1 and 3, We find that there is a significant persistence in the size of a households' consumption response, both prior to income arrival and following arrival. That is, households that increase spending in advance of (or following) income arrival for one event tend to do so in the context of the following event, as well.

Columns 2 and 4 add interactions of the lagged MPC value with two measures of the salience of these payments. 'High News' indicates that both a given event and the preceding event were in the top quartile of news coverage across all events, as measured by the fraction of newspaper

[12]Figure A.3 displays the same distributions for MPX's rather than MPCs.

articles in the household's state that discussed the payment. 'Repeat Event' denotes an event type with which the household had prior experience and thus was more familiar with. For both of these interactions, we find that higher levels of salience or familiarity tended to increase the persistence of the observed Pre- and Post-MPCs approximately doubling the magnitude of the relationship. Intuitively, households need to be aware of the income shock in order for persistent characteristics to drive similarity in spending responses.

3.4 Income, Liquidity, and MPC Heterogeneity

We now turn to a more in depth examination of factors that can explain variation in MPCs across and within individuals. Estimating MPC heterogeneity is crucial for both policy and macroeconomic modeling. For policy, it enables targeting households with high MPCs to maximize aggregate demand responses. For macroeconomic models, matching the observed heterogeneity in MPCs is key to validate model predictions, potentially falsify certain frameworks (e.g., Kaplan and Violante 2022; Auclert et al. 2023, 2024). The distribution of MPC can even serve as a sufficient statistic for evaluating the welfare effects of policy changes in a manner robust to different specifications in a class of supporting structural models (Auclert, 2019).

Table 6 examines the relationship between both pre- and post-payment MPC and the two main drivers of MPC heterogeneity in the canonical models, liquidity and income, which many previous papers have studied. This wide ranging literature includes both empirical contributions (Baker 2018; Gelman 2021; Baker et al. 2023; Crawley and Kuchler 2023) and theoretical contributions (Deaton 1991; Carroll 2001; Kaplan and Violante 2014; Gelman et al. 2022; Kaplan and Violante 2022) examining the interactions of financial resources and consumption behavior. Overall, a major finding is that a strong relationship exists between financial resources and MPC.

We advance this literature by leveraging the fact that we observe the same individual over several distinct payment events. This allows us to decompose the explanatory power of liquidity into the within-person variation that reflects changes in liquidity due to temporary circumstances, such as transitory low income or large expenditure shocks, and permanently low liquidity that reflects persistent household traits, such as high subjective discount rates, impatience, or present bias from hyperbolic discounting or similar psychological behavior. Furthermore, having access to financial transaction data we can expand on the typical measure of liquidity – bank balances, typically the sum of checking and savings accounts – used in previous studies based on survey data and also study the role of unused credit or credit card utilization, measured as the available credit before reaching the credit card's limit, and the number of bank overdrafts, which we can identify from the description string of the transaction.

Consistent with prior literature, columns (1) and (2) show strong effects of income and all three metrics of liquidity on the MPC following the receipt of income in the cross section and when controlling for event-specific fixed effects. Post-payment MPCs are negatively correlated with bank balances and available credit on the card, while positively related to the number of overdrafts across the four quarters prior to the arrival of the payment. Similarly, consumers with higher incomes have lower post-MPCs and hence their consumption dynamics exhibit less excess

sensitivity to the predictable payments as predicted by canonical macroeconomic models.

However, this finding is not robust to controlling for household fixed effects in column (3). While adding event effects does not materially affect the other coefficients, column (3) shows that when focusing on within-individual variation by adding household fixed effects, the coefficients for the role of liquidity and income either become substantially smaller (bank balances and overdrafts) or even flip sign (available credit and income).

There are relatively few studies that investigate the within-individual relationship between post-MPC and liquidity or income largely due to limited opportunities to survey the same individuals repeatedly (Jappelli and Pistaferri, 2020) or observe multiple events per individual (Gelman, 2021; Toczynski, 2022). The vast number of articles that investigate the relationship between post-MPCs and liquidity or income study the cross-sectional relationship. This is because financial transaction data with high-frequency observations around the payment event and with a sufficiently long panel dimension has become available only recently (Baker and Kueng, 2022). Previously, studies had to rely mostly on expenditure survey data, which typically consists either of repeated cross-sections or has only a very limited panel dimension such as the British Expenditure and Food Survey (EFS, formerly called the Family Expenditure Survey, FES) or the CEX, respectively. The result that the relationship between the post-MPC and liquidity and income is weaker when isolating the within-individual variation is consistent with Gelman (2021).

Columns (4)–(6) repeat the analysis for pre-payment MPCs. The cross-sectional relationships in column (4) are weaker when compared to those of the post-payment MPC, with the income consistently positively related to pre-MPCs and smaller relationships with bank balances and overdrafts. That is, both higher-income individuals (columns 4 and 5) and especially times when individuals have high income (column 6), are associated with high pre-payment MPCs. Adding household fixed effects in column (6) shows a clearer relationship between the three liquidity measures and the pre-MPCs. When isolating within-individual variation in liquidity, we find that times when households have higher liquidity are times when they are able to pull spending forward to the pre-event period. This suggests that liquidity constraints may occasionally bind to prevent individuals from increasing consumption before income is received.

The result that the relationship between the post-MPC and liquidity and income is different in the cross section compared to within individual – and the new finding that this also applies to the pre-MPC – is consistent with the idea that persistent individual characteristics or 'types' are important in driving cross sectional MPC heterogeneity (Parker, 2017). We pursue this idea further in the next section by attempting to categorize individuals into different groups based on their pre- and post-payment MPCs.

3.5 News, Salience, and Responses to 'News Shocks'

We next turn to investigate how the difference between pre- and post-payment MPCs vary with the amount of news and attention or with whether it was a repeated event. In general, we find that repeated events or events with more news coverage tend to prompt larger consumption responses surrounding the transfer itself.

As noted above, Table 5 examines how the salience or knowledge about an upcoming income transfer affects the degree of persistence of the consumption response from one event to the next within an individual. While average consumption responses within individual are strongly persistent over time, this persistence is magnified for events which we believe the individual is more aware of. In this sense, non-salience of some events induces some noise in our estimates. Moreover, given a likely jump in awareness or salience upon income arrival, low levels of salience of an upcoming event may contribute to the large discontinuities in consumption behavior at arrival, as well.

Figure 3 displays the full distribution of the difference between pre-MPCs and post-MPCs across individual-event observations for two separate sets of events. We split events by whether they are in the top or bottom half of news coverage relative to the entire sample. We find that events with low levels of news coverage – salience – tend to have both a higher mean and wider distribution of spending changes at income arrival (e.g., post-MPC − pre-MPC). Events with higher levels of salience tend to have more stable consumption paths in the weeks surrounding the arrival of income, suggesting that individuals are increasing consumption in anticipation of income arrival.

Table 7 reinforces this tendency in a regression format. We note that higher levels of salience, as measured by either Google Search activity or news coverage, is associated especially with increased levels of pre-MPCs at an individual-event level. For events with higher levels of news coverage, we also observe lower levels of spending in the 1-week following income arrival. For both measures of salience, we see negative relationships between the jumps in spending around income arrival (columns 5–6), though this difference is insignificant in the case of Google Search.

4 Grouping Consumers into Types

When exploring MPC heterogeneity, it is common to split the sample based on observable characteristics like income and liquidity or age as in Section 3.4. In this section, we advance the study of MPC heterogeneity by taking advantage of the granular nature and the long panel dimension of our data combined with the analysis of multiple events per person, which allows us to calculate pre- and post-payment MPCs at the individual level. Our study is the first to use these novel measures to split the sample into different groups.

Instead of splitting the sample based on equally spaced percentiles, we use k-means clustering to remain agnostic about the cutoff values and sizes for the different groups of consumers.[13]

[13] A small but growing literature has begun to characterize the distribution of individual-level post-payment MPCs, motivated by the recognition that MPC heterogeneity is a central object in modern macroeconomic models of monetary and fiscal transmission (e.g., Kaplan and Violante 2014; Auclert 2019). Recent contributions include Misra and Surico (2014), Lewis et al. (2024), Boehm et al. (2025), and Colarieti et al. (2024). An early and robust finding of this literature—which our analysis strongly corroborates—is that MPCs exhibit substantial yet systematic heterogeneity that is difficult to reconcile with any single model of consumer behavior. Misra and Surico (2014) were among the first to estimate the distribution of individual-level MPCs, using quantile regressions applied to a single exogenous income event: the 2008 fiscal stimulus. Lewis et al. (2024) extend this approach by employing a parametric Gaussian mixture model to address the substantial measurement error in CEX spending data. However, their reliance on

k-means clustering is a computationally efficient unsupervised machine learning algorithm to group individuals into clusters of similar *consumer types* based on individual-level characteristics (Forgy 1965; Bonhomme et al. 2022). The algorithm chooses membership into one of k clusters by minimizing the total within-cluster variation. In our context, we limit the algorithm's inputs to the individual's pre- and post-payment MPCs to form these groups. We then assign each resulting group to a well-known consumer type based on models previously developed in the literature. Finally, we use additional observables not used by the algorithm, such as socio-economic demographics, to assess "out-of-sample" how well these types represent the characteristics one would expect from the corresponding theory.

Figure 4 graphs the density contour plot of the joint distribution of pre- and post-payment MPCs (i.e., the underlying distribution of the bin scatter plot in Figure 2) while also indicating the cluster each point belongs to for our preferred specification with $k = 4$. We arrive at our preferred value of four clusters because it creates groups that map best to consumption models that are widely used in the literature, such as rational consumers in canonical macroeconomic models, or consumers that are subject to behavioral frictions, such as present bias, inattention or mental accounting. Appendix A.1 discusses in more detail how we choose our k and shows the effect of using different numbers of clusters, with k ranging from 2 to 6 (see Appendix Figure A.8).

Table 8 splits individuals by both cluster (columns 1–4) and by liquidity (columns 5–6) to better see these differences in responses as well as socio-economic characteristics across types.[14] Our main results are the first two rows of Panel A which represent the 1-week post- and pre-MPC for consumers assigned to the four groups in columns 1–4 by the k-means clustering algorithm. For each cluster, we separately re-estimate the different MPCs (the post-MPC, pre-MPC, and post-announcement MPC) within these groups in order to calculate correct standard errors within cluster. This re-estimation within group leads to small difference when comparing with the average of the pre- and post-MPCs used to initially calculate the four clusters as shown in Figure 4. We discuss these results for each cluster in more detail below.

A natural question is whether these cluster assignments reflect stable individual characteristics or are primarily driven by temporary circumstances. To assess this, we examine within-person purity — the share of events for which an individual is assigned to their modal cluster — using the normalized measure defined in Appendix A.2, which equals zero under random assignment and one under perfect persistence. Table A.6 reports mean normalized purity by number of clusters k. At our preferred specification of $k = 4$, mean normalized purity is approximately 0.55, indicating that cluster assignments are far more stable than random reshuffling would predict. Purity peaks at $k = 4$ and remains above 0.4 through $k = 7$ before declining gradually to around 0.37 at $k = 10$. This declining pattern is expected: with more clusters, idiosyncratic

cross-sectional variation and a single income event precludes the use of nonparametric clustering methods such as k-means. Boehm et al. (2025) instead exploit a randomized controlled trial conducted with a French bank and apply deconvolution techniques to recover the distribution of individual MPCs from experimental variation. Finally, while they do not observe realized spending responses, Colarieti et al. (2024) use a large survey to classify individuals into consumer types based on self-reported responses to hypothetical income shocks.

[14] Appendix Table A.5 provides estimates for individual events.

variation in individual MPCs is more likely to produce occasional assignments to a non-modal cluster, mechanically compressing the purity measure even when the underlying type is stable. Importantly, purity remains well above zero across all values of k, indicating that cluster membership is meaningfully persistent within person. This is remarkable because individual MPCs contain a lot of "noise", being influenced by other lumpy spending decisions unrelated to income changes and other observable household characteristics, and this noise remains substantial for average MPCs within individuals even if we observe multiple income events per person.[15] This stability, alongside the within-person MPC persistence documented in Section 3.3.3, supports the interpretation that the four clusters capture durable behavioral types rather than temporary circumstances.

In addition to these main results, columns 5–6 separately examine characteristics and average consumption responses for individuals in the top quartile of liquidity and those in the bottom quartile. Here, we see many similarities with prior literature which have extensively documented differences in MPCs along the dimension of liquidity. We note that individuals with high levels of liquidity have post-MPCs of near zero, while those with low levels of liquidity have post-MPCs almost an order of magnitude larger.

Crucially, however, the high and low liquidity groups do not significantly differ from one another in their pre-MPC values for any event. Both have small average preemptive responses, again consistent with most of the prior literature. In contrast, we saw many more significant differences along both dimensions in the clustered groups in columns 1–4.

Finally, row 3 separately tests for announcement effects on non-durable spending for the subset of events where we can identify an announcement date (COVID Stimulus 1, 2, 3, and the Alaska PF Dividends for 2020 and 2021 which had 'bonus' components), as not all events have a unique and cleanly identifiable announcement date. Here, we measure an post-announcement MPC in the same manner as the 1-week post-payment MPCs, calculating the excess spending in the week following the events' announcements. Comparing the post-announcement results to the estimates of pre-payment MPCs in row 2, we find that average spending impacts upon announcement are typically of lower magnitude than the pre-payment MPCs. This may partly be the result of attenuation bias from measurement error in identifying the precise announcement date. The measurement includes empirically identifying the official announcement date as well as when the announcement of each event enters into the information set of each individual. Importantly, we find that the spending responses at announcement are highly correlated with the pre-MPC estimates across clusters. That is, the pre-payment MPCs tend to be good predictors of the announcement spending effects and are available for all events rather than just the subset with well-defined announcement dates.

We now discuss how the four clusters map to well-known theoretical models of consumer behavior. We have sorted the clusters by the pre-payment MPC to highlight the value added of that metric compared to the few previous papers who study MPCs heterogeneity solely based

[15] We use the term "noise" to emphasize that because we have financial transaction data instead of survey responses, this variation in individual pre- and post-MPCs and the corresponding cluster assignments are not due to measurement error.

on post-payment MPCs.

4.1 Type 1: Always Spenders (∼9%)

The consumption behavior of individuals in the first cluster is largely consistent with that of highly impatient or present-biased consumers such as hyperbolic discounters. Table 8 shows that both their pre- and post-payment MPCs are large, suggesting that they prefer to consume a lot and quickly. Table A.5 confirms that these patterns of high pre- and post-MPCs appear across all events independent of whether they are one-time, repeated, or regular payment events. The pre-payment MPC is typically an order of magnitude larger than that of other consumers, including that of the canonical rational consumers in cluster 2.

Turning to auxiliary variables that were not used to form the clusters, we observe that consumers in this cluster exhibit financial distress when using different measures of financial constraints, such as a low credit ratios, high levels of overdrafts, and low liquidity in the form of low bank balances and low available credit. Furthermore, they also exhibit high values for the APC and excess sensitivity (post-MPCs), which has been shown to be associated with present bias and hyperbolic discounting (Shea 1995; Angeletos et al. 2001; Aguiar et al. 2025). Despite this, these consumers are able to pull spending forward, indicating that they find other ways of financing their anticipatory spending.

This cluster makes up about 9% of individuals in our sample. Consumers in this cluster are attentive and forward-looking as they seem aware of upcoming events. They have a desire to smooth consumption by pulling spending forward, resulting in a positive pre-MPC. Table 9 highlights the tendency of preemptive consumption among individuals in cluster 1 to be highly responsive to liquidity. However, their pre-MPC is very large, suggesting a high subjective discount rate or hyperbolic discounting behavior. At the same time, they continue to exhibit elevated consumption following the arrival of the payment, possibly reflecting time-inconsistency in their consumption plans. The time-inconsistency or difficulty in following plans more generally is reflected in the very high number of overdraft fees they incur. They also score low on measures of financial sophistication, displaying relatively low levels of investment deposits and a lower likelihood of being an employee of a financial firm.

In addition, we are able to test for differential responses of total spending and non-durable spending in the weeks surrounding these income shocks, as noted in rows 4 and 5 of Table 8. Interestingly, the present-biased group has the highest ratio of non-durable response to total consumption response, consistent with impatience and a preference for immediate gratification (Angeletos et al., 2001; Laibson et al., 2022).

4.2 Type 2: Consumption Smoothers (∼45%)

The second cluster exhibits post- and pre-payment MPCs that are both positive but small. This consumption pattern with a small one-week pre-MPC of about 0.03 on average and a small one-week post-MPC of about 0.05 is in line with that predicted by canonical macroeconomic

models (Kaplan and Violante, 2022). This cluster makes up about 45% of individuals in our sample. Consumers in this cluster are attentive, forward-looking planners who successfully smooth consumption over these predictable income changes. In Table 9, we see that individuals in this cluster tend to bring forward spending when they have sufficient liquidity to do so.

Turning again to auxiliary characteristics that were not part of the clustering algorithm, we observe that consumers in this cluster have high levels of liquidity and low levels of financial distress. This is consistent with other studies that find that consumption behavior appears more consistent with the canonical model after conditioning on individuals with high liquidity.

While they have substantial liquidity, they are less well-off than consumers in cluster 4, and are likely borrowing constrained at some times either because of low total assets, as in models with just one asset such as Zeldes (1989) and Deaton (1991), or because of low liquid assets as in two-asset models such as Kaplan and Violante (2014), leading to some moderate 'excess sensitivity' of consumption when receiving the predictable payment relative to canonical models of consumption under certainty, such as the permanent income hypothesis (Friedman, 1957) or the life cycle hypothesis (Ando and Modigliani, 1963).

They also show evidence of financial sophistication, including a relatively low APC and incurring few overdrafts. Table 8 also displays the fact that this cluster features a high ratio of nondurable to total expenditures in the pre period and a low ratio in the post period. This is consistent with these households engaging in optimal timing of large lumpy expenditures to coincide with income receipt (Attanasio and Weber, 2010; Aaronson et al., 2012).

4.3 Type 3: Late Spenders (∼24%)

The third cluster exhibits a high post-payment MPC. However, unlike the first cluster, the average pre-payment MPC is close to zero. In fact, Table 9 demonstrates that this group's pre-MPCs are non-responsive to the presence of liquidity or credit. This pattern is consistent with mental accounting (Thaler 1985; Shefrin and Thaler 1988), an explanation for consumption behavior that has recently grown in prominence and has argued that consumer spending is influence by when income arrives or how it is 'labeled' (Hastings and Shapiro, 2013).

This cluster makes up about 24% of individuals and also exhibits signs of low liquidity and high financial distress, albeit to a lesser degree than cluster 1. This could explain why many studies have typically grouped these two clusters together as liquidity constrained while more recent studies have concluded that they exhibit mental accounting (Bernard, 2023; Mijakovic, 2023). Our earlier results in Section 3.4 show that when using only within individual variation, liquidity constraints do not appear to explain consumption behavior to any great degree. These results are consistent with the recent popularity of mental accounting as an explanation of excess sensitivity compared to liquidity constraints.

We see that these individuals tend to be the least responsive to the salience of upcoming income shocks (see Table 10). While other clusters of individuals tend to increase anticipatory spending when exposed to greater amounts of news about future income, these mental accountants have near zero response.

4.4 Type 4: The Unresponsive (∼21%)

The fourth cluster exhibits post-payment and pre-payment MPCs that are both small or negative. This cluster makes up about 21% of individuals in our sample and possesses very high levels of liquidity and income, and very low levels of financial distress. In terms of observable characteristics, they look similar to the rational consumers of cluster 2.

While this cluster looks similar to cluster 2 in terms of observables, the MPC estimates suggest an important difference. Upon closer inspection, we can see that the MPC estimates for this cluster are much noisier. We can see this more clearly in the event-level responses in Table A.5. Tables 9 and 10 further highlight the fact that these individuals seem fairly unresponsive both to liquidity and to salience of the income shock itself.

Lastly, cluster 4 has the lowest payment-to-income ratio. Consistent with Kueng (2014), this group likely has low welfare losses from not smoothing and we note that any post-MPC is largely made up of durable expenditures rather than non-durables. The noisy estimates centered around zero suggest that these individuals exhibit near rational or inattentive behavior. The top panel of Table 10 suggests that only shocks of large sizes tend to prompt any significant consumption response among this group.

4.5 Discussion

We can further test whether our cluster classifications make sense by studying the behavior of MPCs with other measures across clusters.

As one example, Table 9 explores the effect of individual-level financial constraints on consumption response surrounding income payments. In Panel A, we test whether the difference between pre-payment and post-payment MPCs is significantly affected by having high liquidity or available credit. We see a greater degree of smoothing present for both the present-biased (cluster 1) and the rational (cluster 2) users. In Panel B, we narrow our focus to specifically the Pre-MPC for each cluster, finding that spending is pulled forward into the pre-period for these same users.

Table 10 illustrates two other dimensions of heterogeneity across these clusters. Panel A displays the extent to which larger payments drive differences in Post-MPCs across individuals in the entire sample and in each cluster, individually. We find that, while larger payments do not affect the MPC of the 'rational' households in cluster 2, larger payments tend to be associated with lower MPCs for clusters 1 and 3, consistent with models in Lee and Maxted (2023) and Beraja and Zorzi (2024). In addition, cluster 4 (the rationally inattentive cluster), responds positively to larger payments. Intuitively, this cluster is less likely to change their consumption plan for small payments, consistent with theory and evidence in Kueng (2018) and Andreolli and Surico (2021).

Panel B of Table 10 then examines differences across clusters in the response to news coverage of the various payments. Taken across the entire sample, households tend to respond preemptively in response to increases in news coverage of a given payment. However, this response differs substantially across clusters. Cluster 2 (rational users) tends to be most responsive to increases in news and salience. In contrast, cluster 3 (mental accountants) tends to have a response

close to zero for increases in news, consistent with a behavioral desire to only spend following income arrival.

Our clustering results highlight the importance of the pre-payment MPC in separately identifying groups of individuals. Without using pre-payment MPCs, cluster 1 and cluster 3, as well as cluster 2 and cluster 4, look similar enough that they could be combined into one group. However, including the pre-payment MPC allows us to separate these groups. Upon separation, the characteristics of each group are consistent with theory. Cluster 1 has a higher post-payment MPC, a positive and large pre-payment MPC and the highest values of APC and excess sensitivity which have been shown to reflect present-bias. One important point to note is that these pairs of clusters have relatively similar measures of typical observables such as income, liquidity, and financial distress. Therefore, the pre-payment MPC plays an important role in identifying those that exhibit mental accounting from those that exhibit more present-biased behavior and classic rational users from more inattentive users. We are unaware of any recent studies that have highlighted which observable characteristics are correlated with mental accounting.

5 Conclusion

This paper leverages a rich set of comprehensive and high-frequency financial transaction data to examine how individuals respond to a wide range of federal and state transfer programs in the United States. In this setting, we aim to describe how spending responds both prior to and following the arrival of these income shocks. We document three main findings that highlight the importance of mapping out heterogeneity across individuals in their marginal propensities to consume (MPCs).

First, we show that consumers exhibit small but statistically significant spending increases in the days leading up to anticipated income payments. These pre-payment responses – pre-MPCs – are much smaller than the substantial post-payment MPCs that follow the arrival of income, but are robust across a wide range of income events.

Second, by estimating individual-level MPCs across multiple events per person, we find that both pre- and post-MPCs exhibit considerable heterogeneity, but are positively correlated within individuals over time. This within-person correlation is inconsistent with standard models that attribute excess sensitivity entirely to binding liquidity constraints. Moreover, after conditioning on individual-level fixed effects, we note that directly observable measures of income and liquidity no longer exhibit the "predicted" relationship with post-MPCs. These findings suggest that persistent behavioral traits play a central role in shaping consumption dynamics, and that financial observables such as liquidity or income alone cannot fully explain observed heterogeneity in consumption responses.

Finally, we use the joint distribution of individual-level pre- and post-MPCs to classify consumers into meaningful behavioral types. These empirically derived types align closely with four theoretical archetypes: present-biased consumers, rational agents, mental accountants, and rationally inattentive consumers. The assignment of individuals to these clusters correlates with a wide range of untargeted outcomes – including average levels of liquidity, income, and financial

sophistication – and also produces groups that respond distinctly to news, payment size, and liquidity. These results highlight the importance of incorporating anticipatory behavior when distinguishing between behavioral mechanisms. In particular, we show that anticipatory MPCs are crucial to separating present biasedness from mental accounting, two mechanisms that yield similarly high post-MPCs but differ sharply in their pre-payment behavior.

Together, these findings have implications for both macroeconomic theory and policy design. From a theoretical perspective, our results challenge models that assume a one-dimensional source of heterogeneity, such as liquidity constraints, and instead favor frameworks that incorporate behavioral frictions, limited attention, and present bias. From a policy perspective, our results suggest that anticipation effects, although modest in size, are pervasive and systematically related to both the salience of the policy and the behavioral traits of the targeted population. As such, failing to account for these effects may lead to biased estimates of the full consumption response and thus mischaracterize the effectiveness and distributional impact of fiscal transfer programs.

References

Aaronson, Daniel, Sumit Agarwal, and Eric French, "The Spending and Debt Response to Minimum Wage Hikes," *The American Economic Review*, 2012, *102* (7).

Agarwal, Sumit and Wenlan Qian, "Consumption and debt response to unanticipated income shocks: Evidence from a natural experiment in Singapore," *American Economic Review*, 2014, *104* (12), 4205–4230.

Aguiar, Mark, Mark Bils, and Corina Boar, "Who are the Hand-to-Mouth?," *Review of Economic Studies*, 2025, *92* (3), 1293–1340.

Aiello, Darren, Scott R. Baker, Tetyana Balyuk, Marco Di Maggio, Mark J. Johnson, and Jason D. Kotter, "The effects of cryptocurrency wealth on household consumption and investment," *NBER Working Paper*, 2023.

_ , _ , _ , _ , _ , and _ , "Who Invests in Crypto? Wealth, Financial Constraints, and Risk Attitudes," *NBER Working Paper*, 2023.

Akerlof, George A. and Janet L. Yellen, "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?," *American Economic Review*, 1985, *75* (4), 708–720.

Ando, Albert and Franco Modigliani, "The "Life Cycle" Hypothesis of Saving: Aggregate Implications and Tests," *American Economic Review*, 1963, *53* (1), 55–84.

Andreolli, Michele and Paolo Surico, "Less is more: Consumer spending and the size of economic stimulus payments," *CEPR Discussion Paper No. DP15918*, 2021.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 2001, *15*, 47–68.

Angelucci, Manuela, Carlos Chiapa, Silvia Prina, and Irvin Rojas, "Transitory income changes and consumption smoothing: Evidence from Mexico," *Journal of Public Economics*, 2024, *230*, 105013.

Attanasio, Orazio P. and Guglielmo Weber, "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy," *Journal of Economic Literature*, 2010, *48.*

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," *American Economic Review*,

2019, *109* (6), 2333–2367.

_ , Bence Bardóczy, and Matthew Rognlie, "MPCs, MPEs, and Multipliers: A Trilemma for New Keynesian Models," *Review of Economics and Statistics*, 2023, *105* (3), 700–712.

_ , Matthew Rognlie, and Ludwig Straub, "The Intertemporal Keynesian Cross," *Journal of Political Economy*, 2024, *132* (12), 4068–4121.

Baker, Scott R., "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data," *Journal of Political Economy*, 2018, *126* (4), 1504–1557.

_ and Andrey Fradkin, "The Impact of Unemployment Insurance on Job Search: Evidence from Google Search Data," *Review of Economics and Statistics*, 2017, *99* (5), 756–768.

Baker, Scott R and Constantine Yannelis, "Income changes and consumption: Evidence from the 2013 federal government shutdown," *Review of Economic Dynamics*, 2017, *23*, 99–124.

Baker, Scott R. and Lorenz Kueng, "Household financial transaction data," *Annual Review of Economics*, 2022, *14* (1), 47–67.

_ , _ , Brian Melzer, and Leslie McGranahan, "Do Household Finances Constrain Unconventional Fiscal Policy?," *NBER Tax Policy and the Economy*, 2018, *3* (1), 1–32.

_ , Robert A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis, "Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments," *Review of Finance*, 2023, *27* (6), 2271–2304.

_ , Stephanie Johnson, and Lorenz Kueng, "Shopping for Lower Sales Tax Rates," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 209–250.

Balyuk, Tetyana and Emily Williams, "Friends and family money: P2p transfers and financially fragile consumers," *SSRN Working Paper No. 3974749*, 2021.

Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker, "Asymmetric consumption smoothing," *American Economic Review*, 2021, *111* (1), 192–230.

Beraja, Martin and Nathan Zorzi, "Durables and size-dependence in the marginal propensity to spend," *NBER Working Paper*, 2024.

Bernard, René, "Mental accounting and the marginal propensity to consume," *Deutsche Bundesbank Discussion Paper 13/2023*, 2023.

Blinder, Alan S., "Temporary income taxes and consumer spending," *Journal of Political Economy*, 1981, *89* (1), 26–53.

Boehm, Johannes, Etienne Fize, and Xavier Jaravel, "Five facts about MPCs: Evidence from a Randomized Experiment," *American Economic Review*, 2025, *115* (1), 1–42.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa, "Discretizing unobserved heterogeneity," *Econometrica*, 2022, *90* (2), 625–643.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and efficient Estimation," *Review of Economic Studies*, 2024.

Broda, Christian and Jonathan A. Parker, "The economic stimulus payments of 2008 and the aggregate demand for consumption," *Journal of Monetary Economics*, 2014, *68*, S20–S36.

Caldwell, Sydnee, Scott Nelson, and Daniel Waldinger, "Tax refund uncertainty: Evidence and welfare implications," *American Economic Journal: Applied Economics*, 2023, *15* (2), 352–376.

Campbell, John Y. and Angus Deaton, "Why is consumption so smooth?," *Review of Economic Studies*, 1989, *56* (3), 357–373.

_ and N. Gregory Mankiw, "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," *NBER Macroeconomics Annual*, 1989, *4*, 185–216.

Carlstrom, Charles T., Timothy S. Fuerst, and Matthias Paustian, "Inflation and Output in New Keynesian Models with a Transient Interest Rate Peg," *Journal of Monetary Economics*, 2015, *76*, 230–243.

Carroll, Christopher D., "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 1997, *112* (1), 1–55.

_ , "A Theory of the Consumption Function, With and Without Liquidity Constraints," *Journal of Economic Perspectives*, 2001, *15* (3), 23–45.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White, "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 2017, *8* (3), 977–1020.

Choi, Hyunyoung and Hal Varian, "Predicting the Present with Google Trends," *Economic Record*, 2012, *88* (s1), 2–9.

Cochrane, John H., "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *American Economic Review*, 1989, pp. 319–337.

Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva, "The How and Why of Household Reactions to Income Shocks," *NBER Working Paper no. 32191*, 2024.

Commault, Jeanne, "Heterogeneity in MPC Beyond Liquidity Constraints: The Role of Permanent Earnings," *SciencesPo Working Paper No. 03870685v3*, 2024.

Crawley, Edmund and Andreas Kuchler, "Consumption heterogeneity: Micro drivers and macro implications," *American Economic Journal: Macroeconomics*, 2023, *15* (1), 314–341.

Da, Zhi, Joseph Engelberg, and Pengjie Gao, "In search of attention," *Journal of Finance*, 2011, *66* (5), 1461–1499.

Deaton, Angus, "Life-cycle models of consumption: is the evidence consistent with the theory?," 1986.

_ , "Saving and Liquidity Constraints," *Econometrica*, 1991, pp. 1221–1248.

Del Negro, Marco, Marc P. Giannoni, and Christina Patterson, "The Forward Guidance Puzzle," *Journal of Political Economy Macroeconomics*, 2023, *1* (1), 43–79.

Di Maggio, Marco, Emily Williams, and Justin Katz, "Buy now, pay later credit: User characteristics and effects on spending patterns," *NBER Working Paper*, 2022.

Flavin, Marjorie A., "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy*, 1981, *89* (5), 974–1009.

Forgy, Edward W., "Cluster analysis of Multivariate Data: Efficiency versus Interpretability of Classifications," *Biometrics*, 1965, *21*, 768–769.

Friedman, Milton, "A Theory of the Consumption Function," *NBER Books*, 1957.

Fuchs-Schündeln, Nicola and Tarek A. Hassan, "Natural Experiments in mMcroeconomics," in "Handbook of Macroeconomics," Vol. 2, Elsevier, 2016, pp. 923–1012.

Fuster, Andreas, Greg Kaplan, and Basit Zafar, "What would you do with $500? Spending responses to gains, losses, news, and loans," *Review of Economic Studies*, 2021, *88* (4), 1760–1795.

Gelman, Michael, "What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics," *Journal of Monetary Economics*, 2021, *117*, 521–542.

_ , "The Self-Constrained Hand-to-Mouth," *Review of Economics and Statistics*, 2022, *104* (5), 1096–1109.

_ and Nikolai Roussanov, "Managing Mental Accounts: Payment Cards and Consumption Expenditures," *Review of Financial Studies*, 2024, *37* (8), 2586–2624.

_ , Shachar Kariv, Matthew D. Shapiro, and Dan Silverman, "Rational Illiquidity and Consumption: Theory and Evidence from Income Tax Withholding and Refunds," *American Economic Review*, 2022, *112* (9), 2959–2991.

_ , _ , _ , _ , and Steven Tadelis, "Harnessing naturally occurring data to measure the response of spending to income," *Science*, 2014, *345* (6193), 212–215.

_ , _ , Matthew D Shapiro, Dan Silverman, and Steven Tadelis, "How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown," *Journal of Public Economics*, 2020, *189*, 103917.

Graham, James and Robert McDowall, "Mental Accounts and Consumption Sensitivity Across the Distribution of Liquid Assets," *American Economic Journal: Macroeconomics*, 2025.

Hamilton, Steven, Geoffrey Liu, Jorge Miranda-Pinto, and Tristram Sainsbury, "A $100,000 marshmallow experiment: Withdrawal and spending responses to early retirement-savings access," *SSRN Working Paper No. 4389699*, 2024.

Hastings, Justine S. and Jesse M. Shapiro, "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks," *Quarterly Journal of Economics*, 2013, *128*, 1449–98.

Havranek, Tomas and Anna Sokolova, "Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not"," *Review of Economic Dynamics*, 2020, *35*, 97–122.

Indarte, Sasha, Raymond Kluender, Ulrike Malmendier, and Michael Stepner, "What Explains the Consumption Decisions of Low-Income Households?," *NBER Working Paper*, 2024.

Jappelli, Tullio and Luigi Pistaferri, "The consumption response to income changes," *Annual Review of Economics*, 2010, *2* (1), 479–506.

_ and _ , "Fiscal policy and MPC heterogeneity," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 107–136.

_ and _ , "Reported MPC and unobserved heterogeneity," *American Economic Journal: Economic Policy*, 2020, *12* (4), 275–297.

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ and _ , "The Marginal Propensity to Consume in Heterogeneous Agent Models," *Annual Review of Economics*, 2022, *14* (1), 747–775.

_ , Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, 2018, *108* (3), 697–743.

Karger, Ezra and Aastha Rajan, "Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Payments," *FRB of Chicago Working Paper No. WP 2020-15*, 2021.

Kiley, Michael T., "Policy Paradoxes in the New Keynesian Model," *Review of Economic Dynamics*, 2016, *21*, 1–15.

Koşar, Gizem, Davide Melcangi, Laura Pilossoph, and David G. Wiczer, "Stimulus through insurance: The marginal propensity to repay debt," *SSRN Working Paper No. 4477986*, 2024.

Kuchler, Theresa and Michaela Pagel, "Sticking to Your Plan: The Role of Present Bias for Credit Card Paydown," *Journal of Financial Economics*, 2021, *139* (2), 359–388.

Kueng, Lorenz, "Tax News: The Response of Household Spending to Changes in Expected Taxes," *NBER Working Paper no. 20437*, 2014.

_ , "Excess Sensitivity of High-Income Consumers," *Quarterly Journal of Economics*, 2018, *133* (4), 1693–1751.

Laibson, David, Peter Maxted, and Benjamin Moll, "A Simple Mapping from MPCs to MPXs," *NBER Working Paper*, 2022.

Lee, Sean Chanwook and Peter Maxted, "Credit card Borrowing in Heterogeneous-agent Models: Reconciling Theory and Data," *SSRN Working Paper No. 4389878*, 2023.

Lewis, Daniel, Davide Melcangi, and Laura Pilossoph, "Latent heterogeneity in the marginal propensity to consume," *NBER Working Paper*, 2024.

Mankiw, N. Gregory, "The Savers–Spenders Theory of Fiscal Policy," *AEA Papers and Proceedings*, 2000, *90* (2), 120–125.

Maxted, Peter, "Present Bias Unconstrained: Consumption, Welfare, and the Present-bias Dilemma," *Quarterly Journal of Economics*, forthcoming.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson, "The Power of Forward Guidance Revisited," *American Economic Review*, 2016, *106* (10), 3133–3158.

Mijakovic, Andrej, "Marginal Propensities to Consume with Behavioural Agents," *SSRN Working Paper No. 4603292*, 2023.

Misra, Kanishka and Paolo Surico, "Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 84–106.

Olafsson, Arna and Michaela Pagel, "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software," *Review of Financial Studies*, 2018, *31* (11), 4398–4446.

Parker, Jonathan A., "Why don't households smooth consumption? Evidence from a $25 million experiment," *American Economic Journal: Macroeconomics*, 2017, *9* (4), 153–183.

_ , Jake Schild, Laura Erhard, and David S. Johnson, "Economic Impact Payments and Household Spending during the Pandemic," *Brookings Papers on Economic Activity*, 2022, *2022* (2), 81–156.

Poterba, James M., "Are consumers forward looking? Evidence from fiscal experiments," *American Economic Review*, 1988, *78* (2), 413–418.

Schild, Jake, Sophie M. Collyer, Thesia Garner, Neeraj Kaushal, Jiwan Lee, Jane Waldfogel, and Christopher T. Wimer, "Effects of the Expanded Child Tax Credit on Household Spending: Estimates Based on US Consumer Expenditure Survey Data," *NBER Working Paper*, 2023.

Shea, John, "Union contracts and the life-cycle/permanent-income hypothesis," *American Economic Review*, 1995, *85*, 186–200.

Shefrin, Hersh M. and Richard H. Thaler, "The Behavioral Life-cycle Hypothesis," *Economic inquiry*, 1988, *26* (4), 609–643.

Sokolova, Anna, "Marginal propensity to consume and unemployment: A meta-analysis," *Review of Economic Dynamics*, 2023, *51*, 813–846.

Thakral, Neil and Linh Tô, "Anticipation and consumption," *SSRN Working Paper No. 3756188*, 2024.

Thaler, Richard, "Mental Accounting and Consumer Choice," *Marketing Science*, 1985, *4* (3), 199–214.

Toczynski, Jan, "Spend or invest? Analyzing MPC heterogeneity across three stimulus programs," *Swiss Finance Institute Research Paper*, 2022.

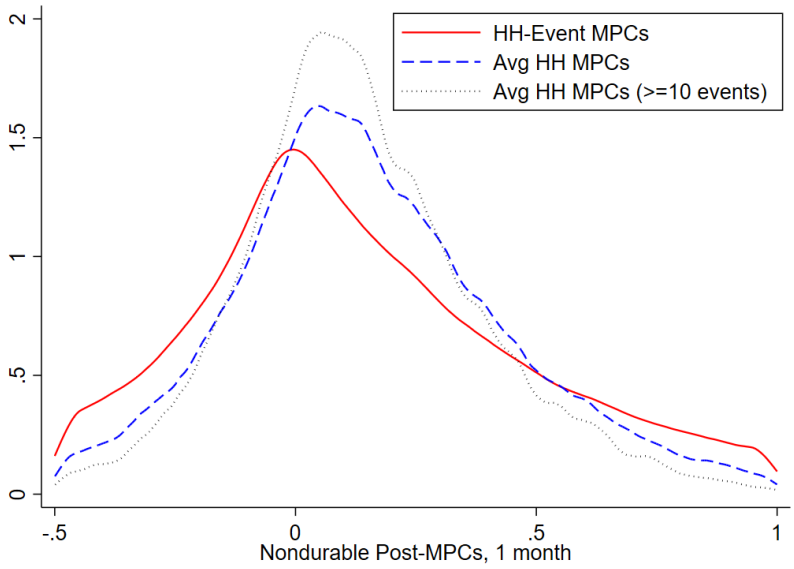
West, Kenneth D., "The insensitivity of consumption to news about income," *Journal of Monetary Economics*, 1988, *21* (1), 17–33.

Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty

Equivalence," *Quarterly Journal of Economics*, 1989, *104* (2), 275–298.

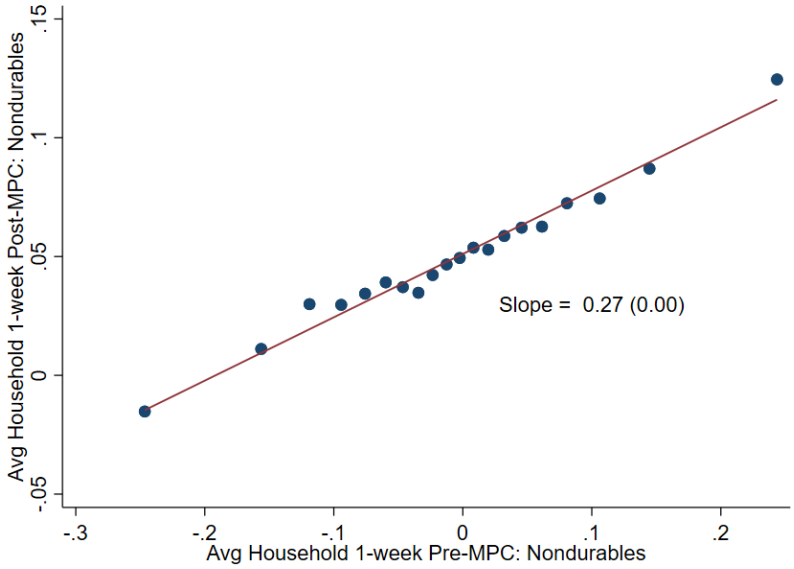
Copy of Figures and Tables from Main Text

Figure 1: Distribution of Post-Payment MPCs



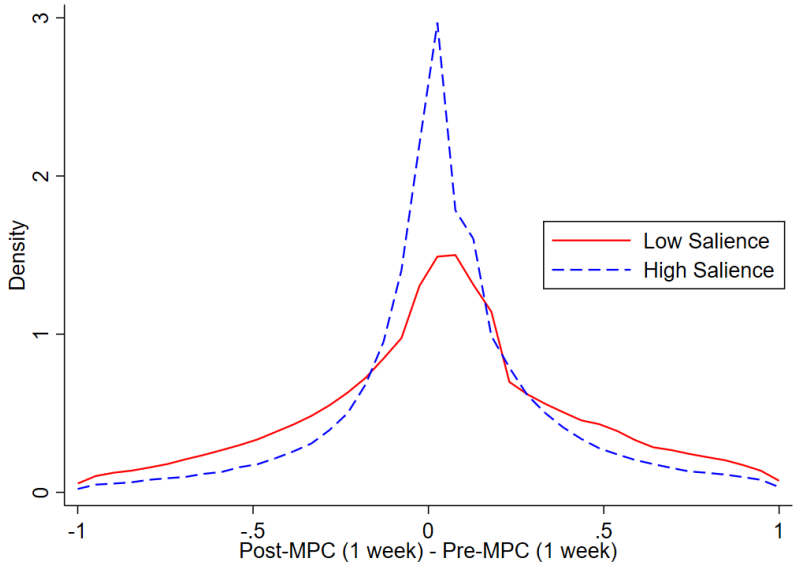
Notes. This figure plots the distribution of post-payment nondurables MPCs computed over a one-month period after each event's payment date. The red solid line shows kernel densities of the marginal propensity to consume out of pre-announced one-time income changes in the month after the arrival of income at a user-event level. The blue dashed line then plots the same distribution after *averaging* these user-events across all events within a user. Finally, the black dotted line represents the distribution of averaged user-level MPCs for only those users where we observe 10 or more income events. All distributions utilize MPCs derived from expenditures on nondurables and services only. Each observation represents a user-event (red solid) or user (blue dashed and black dotted).

Figure 2: Relationship Between Within-Person Average Pre- and Post-MPCs



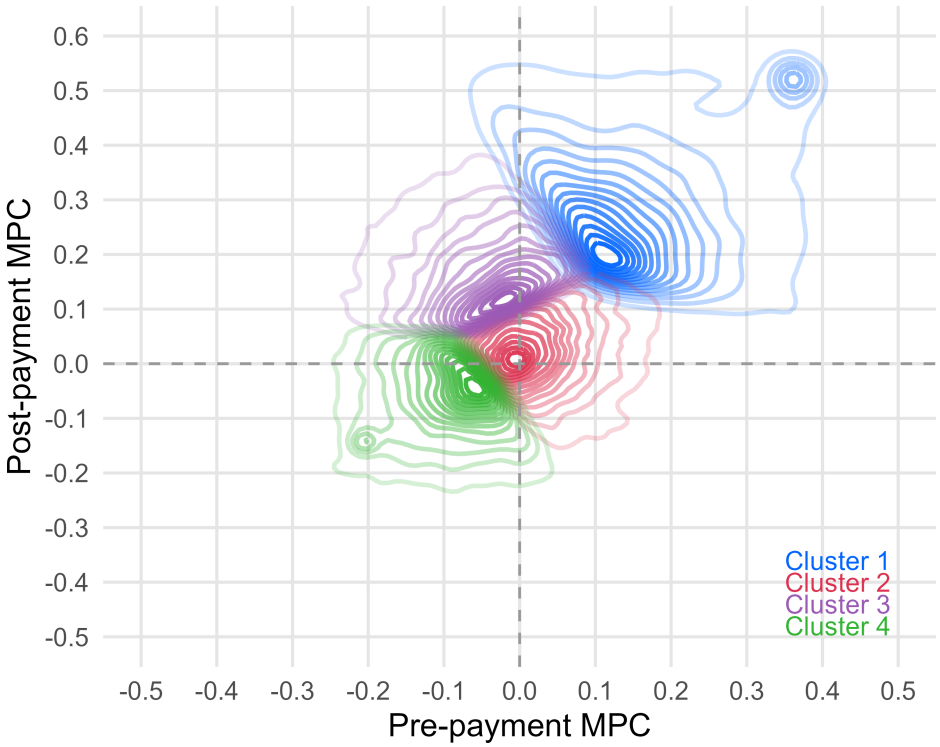
Notes. This figure displays a bin-scatter plot of the one-week within-user average pre-payment MPC against the one-week average post-payment MPC across all events for each user, where each underlying observation represents a user.

Figure 3: Payment Sensitivity (Post-MPC − Pre-MPC) by Event Salience



Notes. Displayed are two kernel densities of the 'dynamic slope', i.e., the difference (Post-MPC − Pre-MPC) at a user-event level, which reflects an individual's payment or cash flow sensitivity of consumption. We measure the salience of each event based on the fraction of newspaper articles written about the event in the local media market. We then split the sample into a top and bottom half of salience and plot the two sub-samples separately.

Figure 4: Clustering of Consumer Types Based on Pre- and Post-MPCs



Notes. This figure shows a density contour plot of individual-level average one week pre- and post-payment MPCs. The individual-level averages are calculated by taking the average of each measure for all events observed within each individual. Each individual is assigned to a cluster based on our k-means clustering algorithm and each cluster is represented by a different color. The darker color contours represent higher density areas.

Table 1: Summary Statistics

	Obs.	Mean	Std. Dev	P5	P25	P50	P75	P95
Panel A: Spending characteristics								
Total Spending: Monthly	6,801,149	5,174	4,367	92	2,116	4,165	7,115	13,486
Daily	205,747,911	171	420	0	0	20	149	809
Nondur. Spending: Monthly	6,801,149	2,881	2,443	19	1,136	2,322	4,001	7,563
Daily	205,747,911	95	210	0	0	9	97	453
Panel B: Household characteristics								
Total Monthly Income	6,801,149	6,026	5,987	0	2,127	4,499	8,051	17,604
Bank Balances	87,187	20,123	71,673	0	388	1,567	5,960	100,023
Available Credit	70,202	14,509	12,215	228	5,112	12,476	20,617	37,108
Credit Ratio	70,184	0.7190	0.2985	0.0549	0.5583	0.8441	0.9533	0.9987
Overdrafts	363,348	0.5804	3.9763	0	0	0	0	2
Panel C: Payment transaction characteristics								
COVID Stimulus 1	29,684	2,222	1,018	1,200	1,200	2,400	2,900	3,900
COVID Stimulus 2	26,968	1,352	864	600	600	1,200	1,800	3,000
COVID Stimulus 3	30,091	3,275	2,115	1,400	1,400	2,800	4,200	7,000
Child Tax Credit	82,510	464	262	167	250	450	550	1,000
Alaska PF Dividend	15,467	3,538	3,077	992	1,600	2,750	4,456	9,852
Tax Refund	329,413	2,714	3,771	63	560	1,549	3,780	8,647
Paychecks	11,197,958	1,883	1,648	208	743	1,482	2,490	4,955

Notes. Panel A presents the monthly and daily income and spending amounts. Panel B presents financial characteristics of each user across the events. Note that the bank account balances and available credit amounts are not available for all individuals. Panel C presents the changes in income from fiscal events. Child Tax Credit transactions include monthly transactions for six months. Alaska PF Dividend is the annual transfer to Alaskan residents from the Alaska Permanent Fund.

Table 2: Event Dates

Event	Date of Payment	Date of Announcement	Predicted Announcement: Google Search	Predicted Announcement: Newspapers	Salience Measures: G-Search	Salience Measures: Newsp.
Panel A: One-off events						
COVID Stimulus 1	4/11/2020	3/27/2020	3/25/2020	3/27/2020	21	57
COVID Stimulus 2	12/29/2020	12/21/2020	12/21/2020	12/23/2020	33	60
COVID Stimulus 3	3/17/2021	3/11/2021	3/13/2021	3/11/2021	86	80
Panel B: One-time monthly payments for 6 months						
Child Tax Credit	7/15/2021	3/11/2021	7/13/2021	5/17/2021	43	10
	8/13/2021					
	9/15/2021					
	10/15/2021					
	11/15/2021					
	12/15/2021					
Panel C: Once every year						
Alaska PF Dividend	10/04/2012	09/18/2012	9/18/2012	9/19/2012	21	20
	10/03/2013	09/17/2013	9/18/2013	9/5/2013	20	20
	10/02/2014	09/17/2014	9/17/2014	9/18/2014	19	19
	10/01/2015	09/21/2015	9/21/2015	9/21/2015	29	29
	10/06/2016	09/23/2016	9/23/2016	9/29/2016	39	26
	10/05/2017	09/04/2017	10/3/2017	9/27/2017	49	28
	10/04/2018	09/14/2018	10/2/2018	8/15/2018	64	6
	10/03/2019	09/27/2019	9/27/2019	9/27/2019	53	25
	07/01/2020	06/12/2020	6/29/2020	5/29/2020	49	20
	10/11/2021	09/30/2021	10/8/2021	9/30/2021	25	40
	09/20/2022	09/08/2022	9/8/2022	9/8/2022	35	36
Tax Refund	Various	Filing date				

Notes. This table presents the announcement and payment dates of each income change used that is induced by fiscal policy events and measures of event salience. The date of payment reflects when the actual payment was made. The date of announcement is a subjective measure based on our reading of news media. The predicted announcement dates are based off of our two salience measure. For each one, the date reflects the highest salience measure recorded before the date of payment. For our salience measures, both of the columns G-Search (i.e., Google Search) and Newspaper reflect the average salience measure over the period that starts from the predicted announcement date and ends at the date of payment. The search measure is a relative score between 0 and 100 while the newspaper value is the number of daily articles related to the event divided by the total number of articles published in the relevant state for the relevant event, scaled by 100.

Table 3: Replication of Previous Studies' Post-Payment MPCs out of Fiscal Transfers

	Original study			
Event	Paper	Period	Main result	Replication
Panel A: Fiscal transfers				
COVID Stimulus 1	Baker et al. (2023)	2020.01–2020.08	0.22	0.22
Child Tax Credit	Schild et al. (2023)	2019.01–2022.03	0.41	0.37
Alaska PF Dividend	Kueng (2018)	2010–2014	0.11	0.10
Tax Refund	Graham et al. (2025)	2014–2015	0.42	0.37
Panel B: Salary receipts				
Paycheck Income	Gelman et al. (2014)	2012–2013	0.07	0.06

Notes. This table presents the replication results of existing papers. The column 'main result' reports the 1-month MPX in Baker et al. (2023), Table 3; the 3-months MPC in Schild et al. (2023), Table 10; the 1-month MPX in Graham and McDowall (2025), Figure 2; the 1-month MPC in Kueng (2018), Figure 3; and the 1-day MPX in Gelman et al. (2014), Table S3. All estimated responses are what we label *post-payment* MPCs respectively MPXs, which measure the additional consumer spending after receiving the fiscal transfer relative to some baseline spending before the transfer.

Table 4: Nondurable Spending Responses to Income Changes Around the Payment Date

			Post-MPC			Pre-MPC
Event	Period	Observations	1-week	1-month	3-month	1-week
Panel A: All events pooled						
Anticipated Payment	2012–2023	205,743,872	0.0357***	0.1033***	0.1853***	0.0064***
			(0.0013)	(0.0042)	(0.0084)	(0.0004)
Panel B: One-off events						
COVID Stimulus 1	2020	19,294,172	0.0478***	0.1368***	0.2446***	0.0053***
			(0.0018)	(0.0041)	(0.0095)	(0.0014)
COVID Stimulus 2	2020–2021	37,379,306	0.0815***	0.1038***	–	0.0039*
			(0.0029)	(0.0063)		(0.0023)
COVID Stimulus 3	2021	18,082,437	0.0493***	0.1351***	0.1616***	0.0081***
			(0.0014)	(0.0033)	(0.0067)	(0.0011)
Panel C: One-time monthly payments for 6 months						
Child Tax Credit	2021–2022	33,251,181	0.1193***	–	–	0.0328***
			(0.0063)			(0.0062)
Panel D: Once every year						
Alaska PF Dividend	2012–2023	205,747,911	0.0760***	0.2058***	0.3992***	0.0194***
			(0.0025)	(0.0061)	(0.0152)	(0.0015)
Tax Refund	2012–2023	205,747,911	0.0302***	0.0867***	0.1506***	0.0037***
			(0.0012)	(0.0038)	(0.0074)	(0.0003)
Panel E: Periodically						
Monthly Paycheck	2012–2023	98,787,012	0.0709***	–	–	0.0220***
			(0.0009)			(0.0007)

Notes. This table presents the responses of spending on nondurables and services to predictable income changes around the payments dates. The corresponding total spending responses (MPXs) are reported in Appendix Table A.2. The 3-month post-MPC for COVID Stimulus 2 is omitted because it overlaps with the pre-MPC of COVID Stimulus 3, which occurred 11 weeks later; see Appendix Figure A.2. Similarly, we report only the one-week post-MPC of the Child Tax Credit because of the recurring payments two to six months after the initial payment. Date and individual by day-of-month fixed effects are included in all specifications. Standard errors clustered at the individual level are reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 5: Persistence of Pre- and Post-Payment MPCs within Person

	Pre-Payment MPC		Post-Payment MPC	
	(1)	(2)	(3)	(4)
Lagged Pre-MPC	0.0249***	0.0182***		
	(0.0036)	(0.0047)		
Lagged Pre-MPC × High News		0.0104		
		(0.0281)		
Lagged Pre-MPC × Repeat Event		0.0161**		
		(0.0074)		
Lagged Post-MPC			0.116***	0.0858***
			(0.0035)	(0.0045)
Lagged Post-MPC × High News				0.0811***
				(0.0275)
Lagged Post-MPC × Repeat Event				0.0685***
				(0.0068)
Observations	313,285	313,285	313,285	313,285
R^2	0.009	0.010	0.021	0.022
Event FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Financial Controls	YES	YES	YES	YES

Notes. Dependent variables are the marginal propensity to consume out of payments in the week before (*Pre-MPC*) or the week after (*Post-MPC*) the income's arrival. *Lagged values* represent the relevant MPC value from the most recent prior event for a given user. For example, for an individual the pre- and post-MPCs out of say the April 2021 tax refund might be the lagged pre- and post-MPCs for the pre- and post-MPCs out of the Child Tax Credit in July 2021. *High News* refers to events where both the event and lagged event were in the top quartile of news coverage. *Repeat Event* refers to events in which both the event and lagged event were not the first of an event type. *Financial controls* include a moving average of bank balances, available credit card credit, income, and incurred overdrafts. Standard errors clustered at the individual level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 6: Liquidity and Pre- and Post-Payment MPCs

	Post-Payment MPC (in %)			Pre-Payment MPC (in %)		
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Bank Balance)	-0.312***	-0.290***	-0.0013	0.0091	0.0269	0.215***
	(0.0330)	(0.0328)	(0.0843)	(0.0293)	(0.0294)	(0.0743)
ln(Available Credit)	-0.374***	-0.376***	0.178**	-0.120***	-0.116***	0.145**
	(0.0308)	(0.0308)	(0.0710)	(0.0270)	(0.0271)	(0.0632)
Numb. Overdrafts	0.304***	0.307***	0.122***	0.0195	0.0188	-0.0201
	(0.0188)	(0.0190)	(0.0226)	(0.0150)	(0.0150)	(0.0215)
ln(Income)	-2.468***	-2.596***	0.584***	0.285***	0.332***	2.442***
	(0.0950)	(0.0961)	(0.191)	(0.0783)	(0.0793)	(0.171)
Observations	359,209	359,209	359,209	359,209	359,209	359,209
R^2	0.007	0.010	0.184	0.000	0.002	0.147
Event FE		YES	YES		YES	YES
User FE			YES			YES

Notes. Dependent variables are the marginal propensity to consume out of income payments in the week after (*Post-Payment MPC*) in columns 1–3 or the week before (*Pre-Payment MPC*) the income's arrival in columns 4–6. Bank balances, available credit (credit card utilization), overdraft counts, and logged income are all measured as the average of the stated variable across the four quarters prior to the arrival of the payment. Coefficients multiplied by 100 for readability. Standard errors clustered at the individual level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 7: Event Salience and Pre- and Post-Payment MPCs

	Pre-Payment MPC		Post-Payment MPC		Difference: Post−Pre MPC	
	(1)	(2)	(3)	(4)	(5)	(6)
Google Search	0.136***		0.0412		-0.0950	
	(0.0450)		(0.0500)		(0.0626)	
State News Coverage		0.393***		-0.147**		-0.539***
		(0.0660)		(0.0724)		(0.0938)
Observations	102,346	359,209	102,346	359,209	102,346	359,209
R^2	0.287	0.147	0.329	0.184	0.320	0.162
User FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES
Financial Controls	YES	YES	YES	YES	YES	YES

Notes. This table presents regressions of pre- and post-MPCs and their difference on measures of event salience. Financial controls include a moving average of bank balances, available credit card credit, income, and incurred overdrafts Standard errors clustered at the individual level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 8: Individual Characteristics by Cluster and Liquidity

	Cluster				Liquidity	
Sub-samples:	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Low	High
Panel A: Consumption						
Nondurable Post-MPC, 1-week	0.2422***	0.0475***	0.1457***	0.0315***	0.0814***	0.0447***
	(0.0075)	(0.0014)	(0.0027)	(0.0027)	(0.0036)	(0.0033)
Nondurable Pre-MPC, 1-week	0.1073***	0.0327***	0.0216***	0.0033	0.0245***	0.0249***
	(0.0046)	(0.0012)	(0.0016)	(0.0022)	(0.0022)	(0.0027)
Post-Announcement MPC	0.0543***	0.0143***	0.0292***	0.0021	0.0158***	0.0116***
	(0.0042)	(0.0012)	(0.0020)	(0.0022)	(0.0024)	(0.0029)
Total Post-MPX, 1-week	0.379***	0.095***	0.229***	0.070***	0.142***	0.085***
	(0.013)	(0.003)	(0.004)	(0.005)	(0.006)	(0.006)
Total Pre-MPX, 1-week	0.143***	0.045***	0.032***	0.005	0.030***	0.036***
	(0.008)	(0.002)	(0.003)	(0.005)	(0.004)	(0.005)
Average Prop. to Cons. (APC)	1.00	0.91	0.97	0.91	0.98	0.90
Panel B: Liquidity						
Bank Balances	6,197	11,523	6,891	15,011	275	33,894
Available Credit	10,392	14,049	11,232	15,863	2,599	25,160
Credit Ratio	0.58	0.74	0.63	0.76	0.36	0.88
Number of Overdrafts	1.55	0.39	0.90	0.36	1.61	0.32
Panel C: Income						
Income	66,540	81,927	72,767	86,727	61,470	103,815
Payment-to-Income Ratio	0.04	0.05	0.05	0.03	0.05	0.04
Income Variation	0.04	0.04	0.04	0.04	0.03	0.04
Panel D: Financial Sophistication						
Investment Deposits	2,401	3,678	2,568	4,320	2,084	5,517
Finance Employee	0.05	0.07	0.04	0.09	0.03	0.10
Number of Observations	22,739	180,587	87,037	73,213	37,906	41,389
Number of Users	4,646	22,757	12,111	10,776	9,028	8,371

Notes. This table reports the means of variables for different sub-samples. The MPCs in the first three rows do not use predictable income changes from tax refunds and regular salary paychecks. MPCs from these events, which are observed much more frequently, are reported separately in the following rows. *Income Variation* is measured as the volatility of monthly income within a year divided by average income in the same year. *Finance Employee* is an indicator for an individual who received a salary from a financial institution during the sample period. *Low Liquidity* refers to the bottom quartile of both bank balances and available borrowing capacity on the credit card. *High Liquidity* refers to users in the top quartile along the same metric.

Table 9: Liquidity, Credit, and Differences in Pre- and Post-MPCs by Cluster

	Cluster 1	Cluster 2	Cluster 3	Cluster 4
Panel A. Difference: Post-MPC − Pre-MPC (in %)				
ln(Bank Balance)	-0.703***	-0.190***	-0.158*	0.0110
	(0.191)	(0.0447)	(0.0835)	(0.0687)
ln(Available Credit)	-0.917***	-0.202***	-0.0913	-0.0438
	(0.173)	(0.0407)	(0.0732)	(0.0641)
R^2	0.002	0.003	0.013	0.001
Panel B. Pre-payment MPC (in %)				
ln(Bank Balance)	0.228*	0.0955***	-0.0962*	0.0321
	(0.122)	(0.0294)	(0.0575)	(0.0463)
ln(Available Credit)	0.519***	0.0553**	-0.256***	0.0353
	(0.110)	(0.0271)	(0.0506)	(0.0431)
R^2	0.026	0.005	0.009	0.020
Observations	22,739	180,587	87,036	73,213
Event FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES

Notes. Dependent variables are the marginal propensity to consume (MPC) out of income payments in the week before (*Pre MPC*, panel B) and the difference of this pre MPC and the post MPC the week after the income's arrival (panel A). Dependent variables multiplied by 100 for display purposes. Bank balances and available credit are measured as the average of the stated variable across the four quarters prior to the income's arrival. *** p<0.01, ** p<0.05, * p<0.1.

Table 10: Payment Size, Salience and MPCs, by Cluster

	All Clusters	Cluster 1	Cluster 2	Cluster 3	Cluster 4
Panel A. Post-MPC					
ln(Shock Amount)	-0.0135***	-0.164***	-0.0003	-0.0949***	0.0816***
	(0.00245)	(0.0090)	(0.0036)	(0.0051)	(0.0048)
R^2	0.193	0.244	0.143	0.165	0.172
Panel B. Pre-MPC					
State News Coverage	0.248***	0.269	0.422***	-0.0277	0.230
	(0.091)	(0.421)	(0.122)	(0.180)	(0.207)
R^2	0.147	0.143	0.088	0.105	0.111
Observations	345,325	21,050	172,043	82,448	69,784
User FE	YES	YES	YES	YES	YES
Event FE	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES
Financial Controls	YES	YES	YES	YES	YES

Notes. Panel A presents results of regressions of Post-MPCs on the ratio between the payment amount and an individual's annual income. Panel B displays results of regressions of Pre-MPCs on the amount of news coverage for a particular event in the state of residence for an individual. Financial controls include a moving average of bank balances, available credit card credit, income, and incurred overdrafts. Standard errors clustered at the individual level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Online Appendix: Anticipatory Spending

Scott R. Baker, Michael Gelman, Lorenz Kueng, Seung Hyeong Lee

A Data Appendix

A.1 Choosing the number of clusters k

With a two-dimensional joint distribution, using $k = 4$ clusters is a natural extension of the two type grouping in Campbell and Mankiw (1989) and Mankiw (2000). In this section, we provide some metrics beyond this economic approch of selecting the number of clusters. In general, choosing the number of clusters k can be challenging as there is no universally agreed upon method. We use three different measures to evaluate our choice of k.

A.1.1 Within-cluster sum of squares (WCSS)

The first method we use is WCSS. It's an intuitive metric to use because this is the object that the clustering algorithm is minimizing.

The downside is that WCSS is declining in k. Therefore choosing the smallest WCSS would in the limit choose $k = n$. In practice, practitioners use the "elbow method" which attempts to find an obvious inflection point. This "elbow" is where there is the most obvious gain from including an extra cluster.

Figure A.5 plots the WCSS by the number of clusters. The figure shows that gains from including more clusters starts to level out around 4 or 5 clusters. By this measure, our choice of 4 is not unreasonable.

A.1.2 Silhouette

Our second measure is the average Silhouette score. The silhouette score is a measure of how close any given point is to all the other points in its assigned cluster compared to the next closest cluster. Under ideal conditions, most points would be much closer to each other in its assigned cluster compared to the other clusters.

Figure A.6 shows the results. The highest score is for $k = 2$ while the rest of the clusters have comparable numbers. Although a cluster of 2 might maximize the silhouette score, we believe that there is economic value in including more clusters as they are able to identify behavior that maps to theoretic models popular in the literature.

A.1.3 R^2 measures

Our third measure is the adjusted R^2 in a regression of the slope of the MPC (post-MPC - pre-MPC) on indicator variables for each cluster k. We also analyze this measure for other alternative

clustering variables post-MPC, pre-MPC, bank balance, and income. Figure A.7 plots the results. When using both the pre- and post- MPC, the adjusted R^2 levels out after a k of 5. While increasing the number of clusters can improve the fit of the model, our goal is to identify the maximum number of clusters that uniquely map to well known theoretical models. Therefore, we believe that this method also shows that our choice of k=4 is reasonable.

The results also show that using only the post-MPC to cluster explains more variation for each k relative to only using the pre-MPC. This is not surprising, as the pre-MPC is likely estimated with more noise. Combining the two measures leads to sustained increases in the R^2 as k increases while the gains level out when only using each variable alone. This implies that there is some benefit to combining the two metrics that improves the fit beyond each single measure alone. This highlights the benefit of our novel pre-MPC measure.

Lastly, our analysis shows that clustering on bank balance and income explain very little of the variance of the slope of the MPC.

A.1.4 Summary

Figure A.8 plots the clusters classifications for $k = 2$, 3, 4, and 5. Starting at $k = 2$ generates the typical high and low post-MPC clusters which tend to have the same average pre-MPC. This is a typical split you would see when using some split that reflects liquidity constraints (liquidity, income, age, etc). Going from 2 to 3 splits up the lower post-MPC cluster into 2 clusters which we call canonical rational and near rational or inattentive consumers. Going to 3 to 4 allows us to split the high MPC cluster into hyperbolic discounters and mental accountants. Moving from 4 to 5 takes the near rational or inattentive consumers and splits them into two groups based on their pre-payment MPC. However, we think this isn't a meaningful split as this cluster has a lot of noise that is a result of their non-optimization. Going from 5 to 6 mainly cleaves off a portion fo the rational cluster into a higher MPC cluster that is somewhat of a scaled down version of the hyperbolic discounters. Using $k = 7$ and above (not shown here) starts splitting qualitatively similar clusters into more extreme splits.

In terms of matching the clusters to theoretical models, the biggest gains come from going from 2 to 3 and then 3 to 4. After $k = 4$, the algorithm is creating further quantitative refinements of these categories which we think are not qualitatively important enough to warrant another split. As discussed in the main text, we arrive at our preferred value of 4 because we believe that it creates groups that map best to consumption models that are widely used in the literature.

A.2 Within-Person Persistence of Cluster Assignments

Let income transfer events be indexed by e. E_i counts the number of income transfer events that individual i is exposed to during the sample period. We restrict the sample to individuals who face at least 5 income events, i.e., $E_i \geqslant 5$. For each event, we run the k-means clustering algorithm, say for $k = 4$, assigning each treated individual to a cluster $g_{i,e} \in \{1, 2, 3, 4 = k\}$. This yields an unbalanced panel of cluster assignments by individual and event, with panel dimension E_i. g_i^{mod} denotes the number of observations in individual i's modal cluster across all events.

We measure the within-person persistence of cluster assignments – which we call purity – by the ratio of the number of observations in an individual's modal cluster to the total number of observations:

$$purity_i = \frac{g_i^{mod}}{E_i}. \tag{A.1}$$

This measure of persistence ranges from $\frac{1}{k} = \frac{1}{4}$ to 1. With $k = 4$ clusters, g_i^{mod} has a maximum of E_i and a minimum of $ceil(\frac{E_i}{k})$. For example, if cluster assignment has no persistence and occurs randomly, then the purity is $\frac{2}{5}$ for individuals with $E_i = 5$ and $\frac{2}{6}$ if $E_i = 6$, $\frac{2}{7}$ if $E_i = 7$, $\frac{2}{8}$ if $E_i = 8$, $\frac{3}{9}$ if $E_i = 9$, $\frac{3}{10}$ if $E_i = 10$, and so forth.

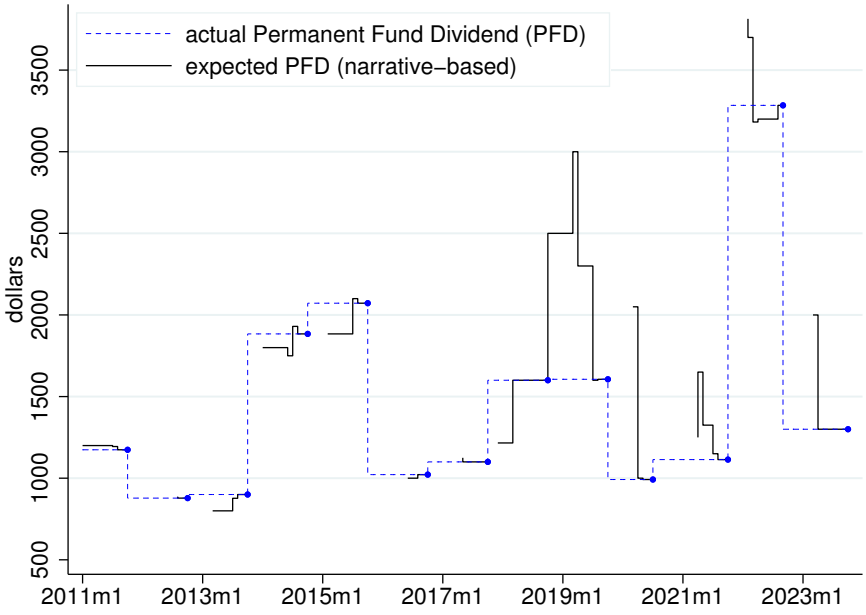
We therefore use the normalized purity measure

$$purity_i^* = \frac{g_i^{mod} - ceil(\frac{E_i}{k})}{E_i - ceil(\frac{E_i}{k})} \tag{A.2}$$

as our main measure of cluster persistence. This normalized measure ranges from 0 for random assignment with no persistence to 1 for individuals that are assigned to the same cluster for each event.

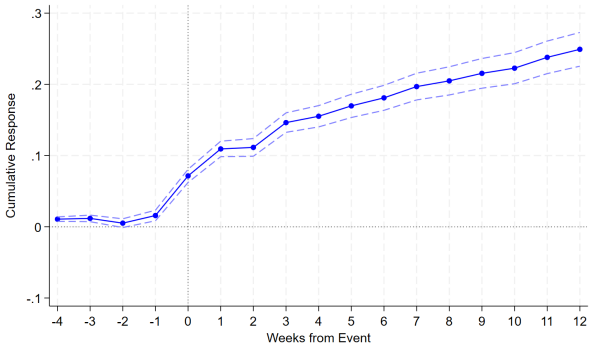
B Appendix Figures and Tables

Figure A.1: Alaska Permanent Fund Dividend News

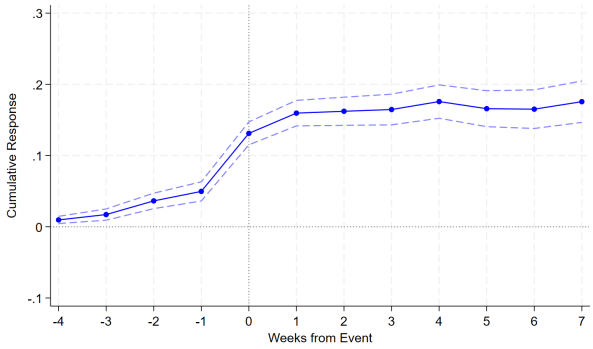


Notes. This figures shows expected amounts of the next Alaska Permanent Fund Dividend (black line) compared to the realized dividend (blue dots). The dividend expectation series is based on a narrative analysis of Alaskan newspaper and extends the series in Kueng (2018) to 2023.

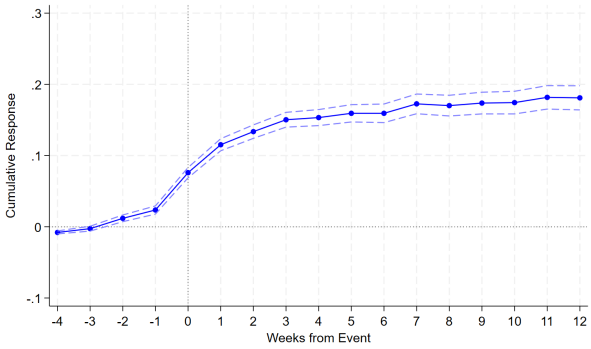
Figure A.2: Nondurable Spending Responses Relative to Payment Date



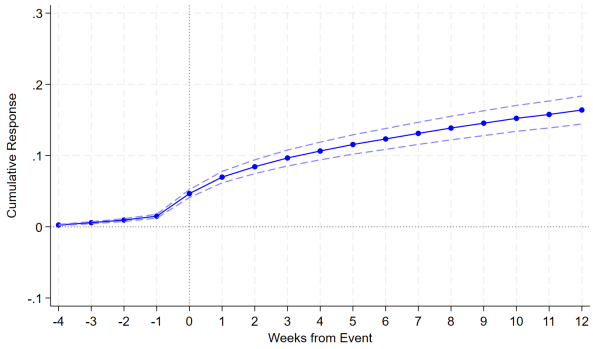
(a) COVID Stimulus 1: 4/11/2020



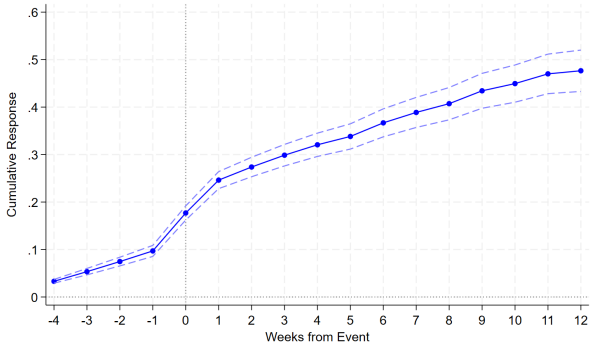
(b) COVID Stimulus 2: 12/29/2020



(c) COVID Stimulus 3: 3/17/2021



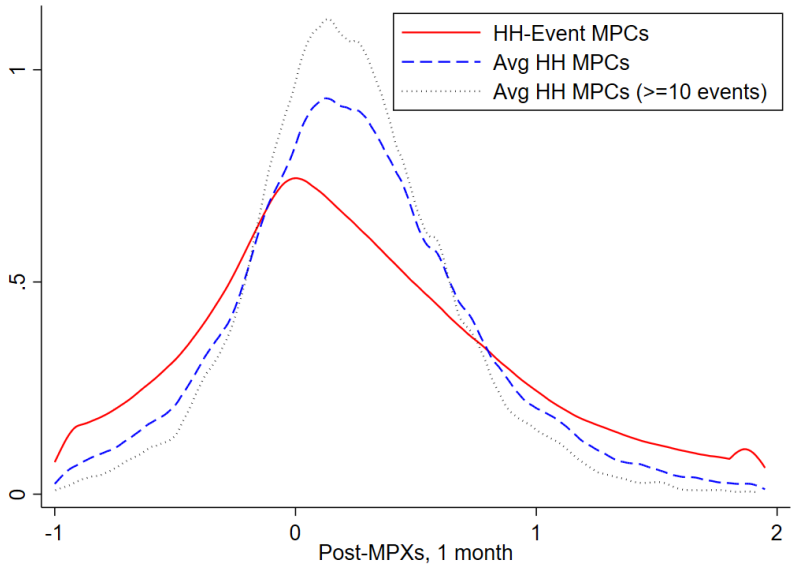
(d) Tax Refund: 2012-2023



(e) Alaska Permanent Fund Dividend: 2012–2023

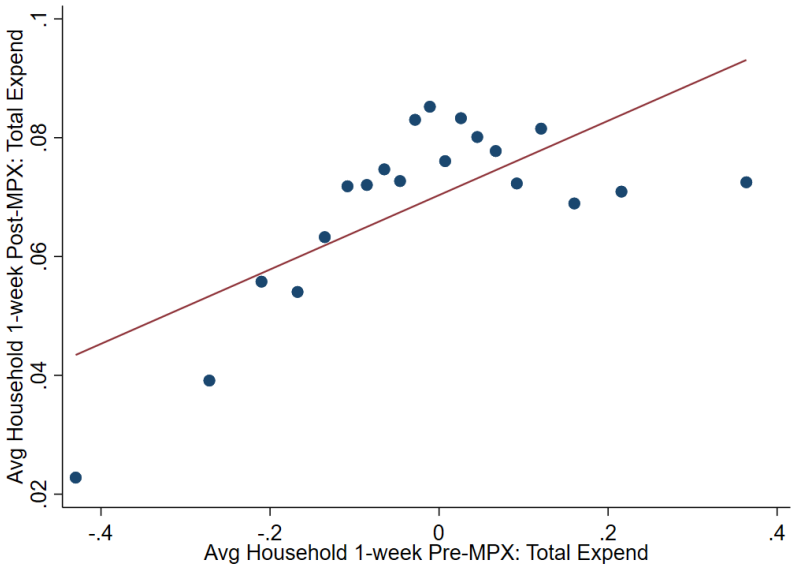
Notes. This figure plots the nondruable spending responses relative to payment date. We first estimate daily responses with the following distributed lead and lag specification: $c_{i,t} = \sum_{s=-28}^{90} \beta_s \; Amount_i \times \mathbf{1}[t = s]_{it} + \alpha_{i,d(t)} + \alpha_t + u_{i,t}$. We then combine the estimates into weeks relative to the payment date. The solid line shows the aggregated point estimates of β_s relative to one week prior to the payment (i.e., week -1 is normalized to zero). The vertical bars show the 95% confidence interval. Standard errors are clustered at the individual level. Date and individual-by-'day of month' fixed effects are included. Time relative to the payment date is equal to zero on the day of receiving the transfer. Since the event windows of COVID Stimulus 2 and 3 partially overlap, panel (b) is displaced by up to 7 weeks post-event.

Figure A.3: Distribution of Post-Payment MPXs



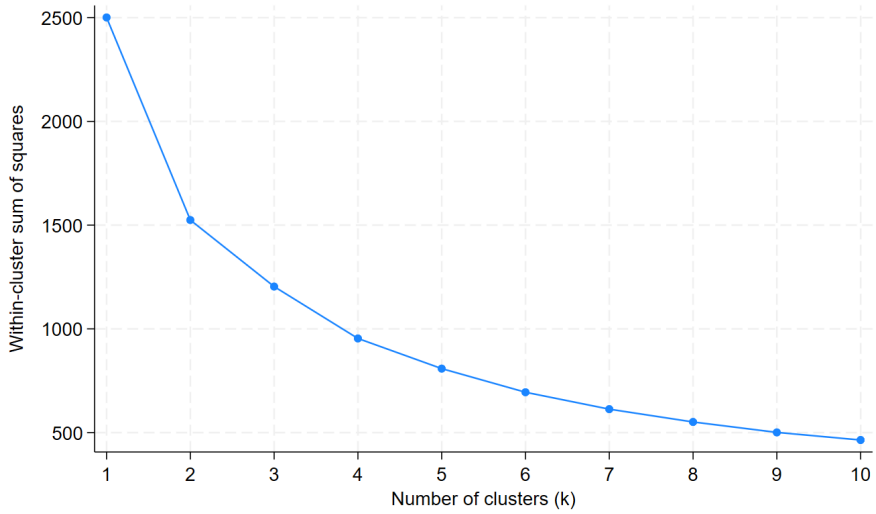
Notes. This figure is the equivalent of Figure 1 except for using total spending MPXs instead of nondurable spending MPCs.

Figure A.4: Relationship Between Within-Person Average Pre- and Post-MPXs



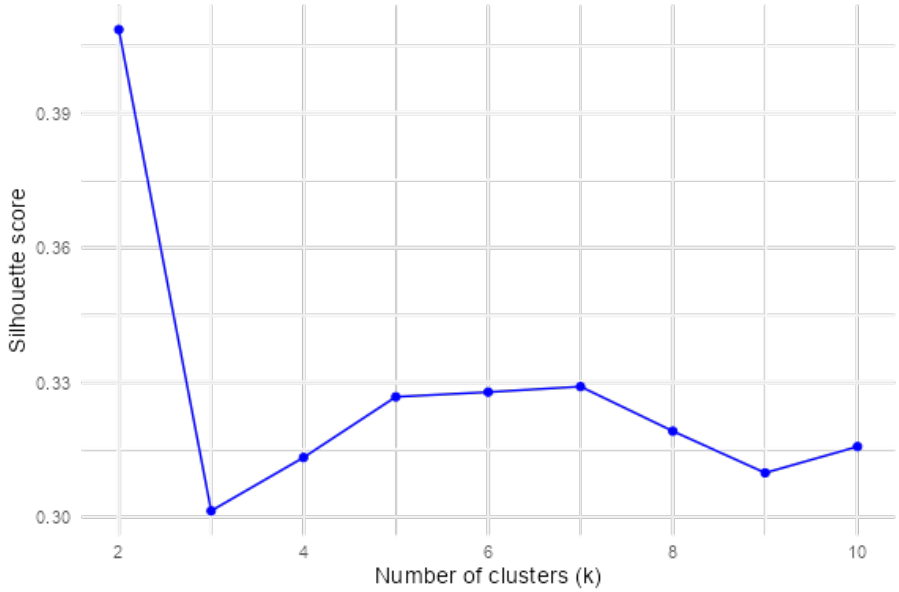
Notes. This figure is the equivalent of Figure 2 except for using total spending MPXs instead of nondurable spending MPCs.

Figure A.5: Within-Cluster Sum of Squares by Number of Clusters



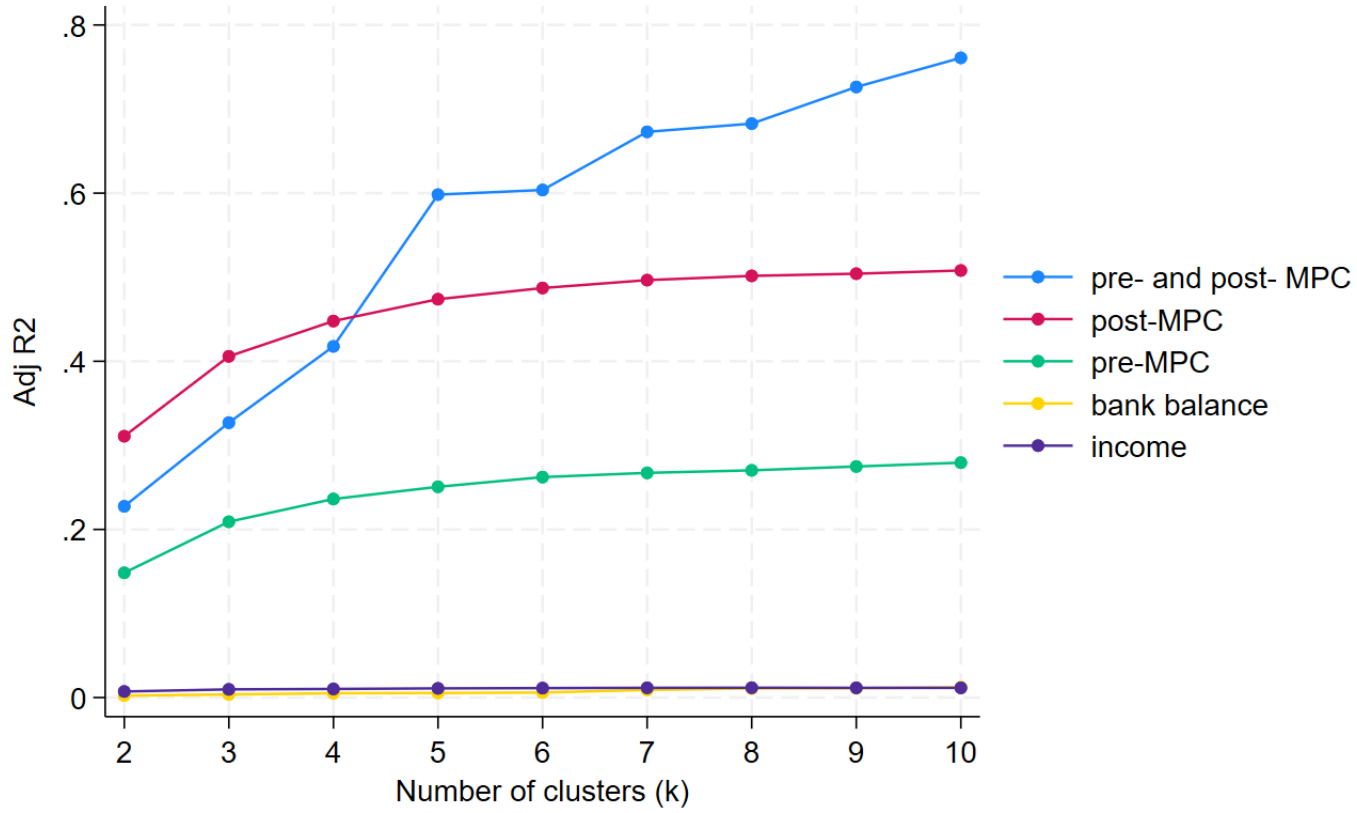
Notes. This figure plots the within-cluster sum of squares for each value of k in our k means clustering algorithm.

Figure A.6: Silhouette score by Number of Clusters



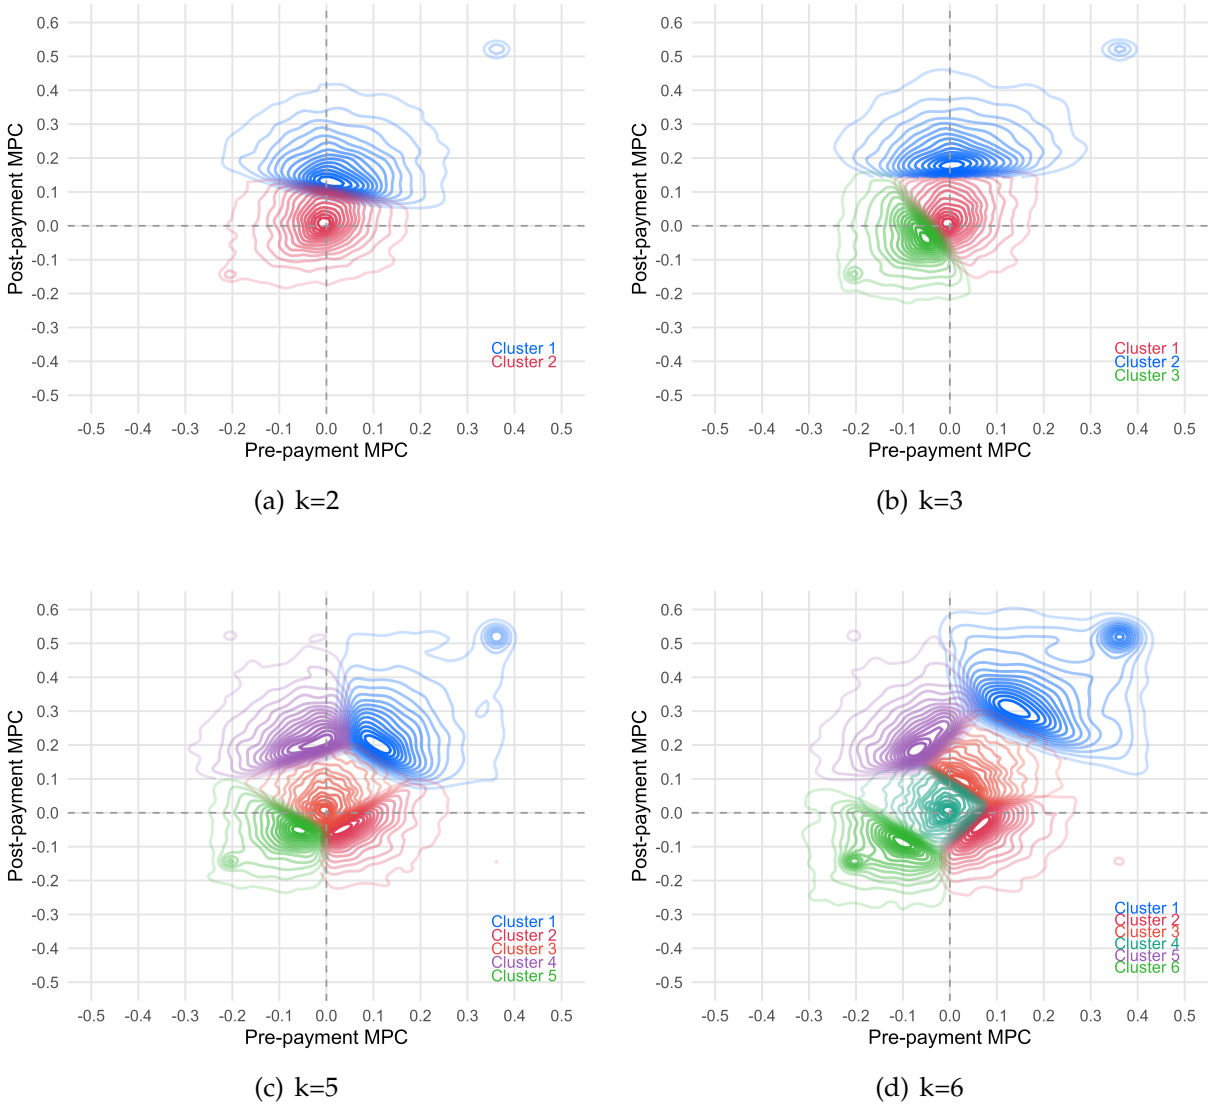
Notes. This figure plots the average silhouette score for each k (number of clusters). The silhouette score is a measure how how well each point in the cluster fits in its assigned cluster compared to the next closest cluster.

Figure A.7: Adjusted R^2 by number of clusters (k)



Notes. This figure shows a the R^2 of a regression of the MPC slope (post-MPC - pre-MPC) on indicator variables for each cluster k for various k.

Figure A.8: Alternative Number of Clusters

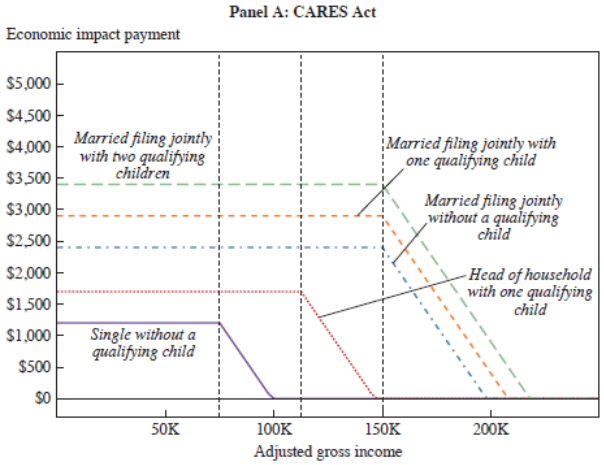


(a) k=2 (b) k=3

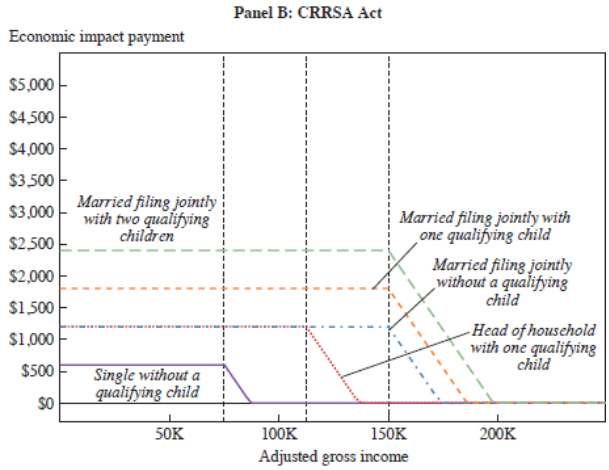
(c) k=5 (d) k=6

Notes. This figure shows a density contour plot of individual-level average one week pre-payment and post-payment MPCs. The individual-level averages are calculated by taking the average of each measure for all events observed within each individual. Each individual is assigned to a cluster based on our k means clustering algorithm and each cluster is represented by a different color. The darker color contours represent higher density areas.

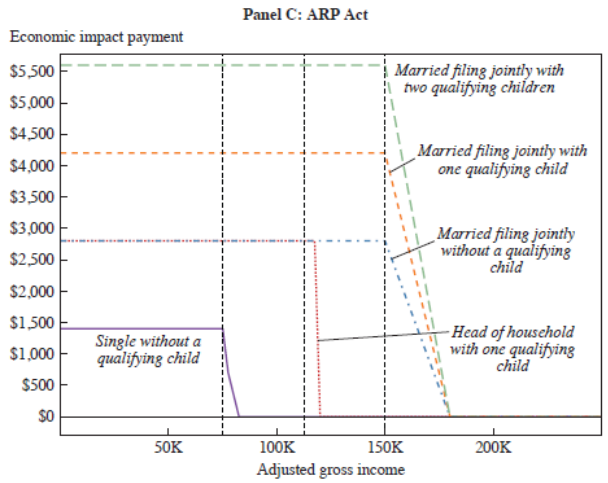
Figure A.9: Sources of Identifying Variation from the COVID Stimulus Programs



(a) COVID Stimulus 1: CARES Act of March 27, 2020



(b) COVID Stimulus 2: CRRSA Act of Dec. 29, 2020



(c) COVID Stimulus 3: ARP Act of March 11, 2021

Notes. These Figures are reproduced for convenience from Parker et al. (2022), Figure 2. CARES refers to the *Coronavirus Aid, Relief, and Economic Security* Act of 2020; CRRSA to the *Coronavirus Response and Relief Supplemental Appropriations* respectively the *Consolidated Appropriations* Act of 2021; and ARP to the *American Rescue Plan* Act of 2021.

Table A.1: Transaction Category Classification

Transaction Category	Total Spending	Nondurable Spending	Total Income
Panel A: Spending transactions			
ATM/Cash Withdrawals	Yes	No	No
Automotive/Fuel	Yes	Yes	No
Cable/Satellite/Telecom	Yes	Yes	No
Charitable Giving	Yes	No	No
Check Payment	Yes	No	No
Education	Yes	Yes	No
Electronics/General Merchandise	Yes	Yes	No
Entertainment/Recreation	Yes	Yes	No
Gifts	Yes	No	No
Groceries	Yes	Yes	No
Healthcare/Medical	Yes	Yes	No
Home Improvement	Yes	No	No
Insurance	Yes	Yes	No
Loans	Yes	No	No
Mortgage	Yes	No	No
Office Expenses	Yes	Yes	No
Other Expenses	Yes	No	No
Personal/Family	Yes	Yes	No
Pets/Pet Care	Yes	Yes	No
Postage/Shipping	Yes	Yes	No
Rent	Yes	No	No
Restaurants	Yes	Yes	No
Service Charges/Fees	Yes	No	No
Services/Supplies	Yes	Yes	No
Subscriptions/Renewals	Yes	Yes	No
Travel	Yes	Yes	No
Utilities	Yes	Yes	No
Panel B: Other transactions			
Credit Card Payments	No	No	No
Deposits	No	No	Yes
Expense Reimbursement	No	No	No
Interest Income	No	No	Yes
Investment/Retirement Income	No	No	Yes
Other Income	No	No	Yes
Refunds/Adjustments	No	No	No
Retirement Contributions	No	No	No
Rewards	No	No	No
Salary/Regular Income	No	No	Yes
Sales/Services Income	No	No	Yes
Savings	No	No	No
Securities Trades	No	No	No
Taxes	No	No	No
Financial Account Transfers	No	No	No

Notes. This table presents the transaction category classification.

Table A.2: Responses to Income Changes Around the Payment Date: Total Spending (MPX)

Event	Period	Observations	Post-MPX 1-week	Post-MPX 1-month	Post-MPX 3-month	Pre-MPX 1-week
Panel A: All events pooled						
Anticipated Payment	2012–2023	205,743,872	0.0761***	0.2066***	0.3586***	0.0142***
			(0.0028)	(0.0079)	(0.0153)	(0.0008)
Panel B: One-off events						
COVID Stimulus 1	2020	19,294,172	0.1023***	0.2639***	0.4485***	0.0120***
			(0.0033)	(0.0073)	(0.0167)	(0.0028)
COVID Stimulus 2	2020–2021	37,379,306	0.1398***	0.1708***	–	0.0135***
			(0.0053)	(0.0111)		(0.0042)
COVID Stimulus 3	2021	18,082,437	0.0984***	0.2311***	0.2746***	0.0141***
			(0.0025)	(0.0055)	(0.0113)	(0.0021)
Panel C: One-time monthly payments for 6 months						
Child Tax Credit	2021–2022	33,251,181	0.1761***	–	–	0.0264**
			(0.0107)			(0.0112)
Panel D: Once every year						
Alaska PF Dividend	2012–2023	205,747,911	0.1293***	0.3108***	0.5243***	0.0217***
			(0.0042)	(0.0098)	(0.0241)	(0.0025)
Tax Refund	2012–2023	205,747,911	0.0682***	0.1844***	0.3085***	0.0110***
			(0.0029)	(0.0076)	(0.0140)	(0.0008)
Panel E: Periodically						
Monthly Paycheck	2012–2023	98,787,012	0.1573***	–	–	0.0432***
			(0.0019)			(0.0013)

Notes. This table reports the same statistics as Table 4 but for total expenditures (MPXs) instead of spending on nondurables and services (MPCs). See Table 4 for details.

Table A.3: One Week Spending Responses Around the Announcement Date

Event	Period	Observations	Announcement Date	
			Pre-MPC	Post-MPC
Panel A: Nondurable spending (MPCs)				
COVID Stimulus 1	2020	19,294,172	-0.0069***	-0.0055***
			(0.0016)	(0.0017)
COVID Stimulus 2	2020–2021	37,379,306	0.0004	0.0241***
			(0.0028)	(0.0029)
COVID Stimulus 3	2021	18,082,437	0.0049***	0.0252***
			(0.0012)	(0.0013)
Alaska PF Dividend	2012–2023	205,747,911	0.0151***	0.0236***
			(0.0014)	(0.0016)
Panel B: Total spending (MPXs)				
COVID Stimulus 1	2020	19,294,172	-0.0025	0.0027
			(0.0030)	(0.0034)
COVID Stimulus 2	2020–2021	37,379,306	-0.0030	0.0289***
			(0.0048)	(0.0049)
COVID Stimulus 3	2021	18,082,437	0.0049**	0.0500***
			(0.0023)	(0.0024)
Alaska PF Dividend	2012–2023	205,747,911	0.0278***	0.0378***
			(0.0026)	(0.0027)

Notes. This table presents the one-week spending responses to future fiscal income transfers around the time of announcement. For the COVID Stimulus 3, the post-announcement-MPC may contain the payment dates for several individuals due to the close time gap between the announcement and payment. Date and Individual-by-day of month fixed effects are applied. Standard errors clustered at the individual level are reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A.4: Relationship between Pre- and Post-Payment MPCs

	Dependent variable: Pre-Payment MPC, one-week					
	(1)	(2)	(3)	(4)	(5)	(6)
Post-MPC, one-week	0.0633***	0.0563***	0.0581***			
	(0.0015)	(0.0017)	(0.0019)			
Post-MPC, one-month				0.0423***	0.0410***	0.0424***
				(0.0006)	(0.0007)	(0.0008)
Observations	348,785	348,785	272,514	348,785	348,785	272,514
R^2	0.007	0.169	0.173	0.015	0.176	0.181
Event FE	YES	YES	YES	YES	YES	YES
User FE		YES	YES		YES	YES
Financial Controls			YES			YES

Notes. This table presents regressions of one-week Pre-MPCs on Post-MPCs, where post-MPCs are measured either also over a one-week periiod or over one-month period for increased precision. Financial controls include bank balances, available credit card credit, income, and the count of incurred overdrafts as well as a user's average MPC out of paychecks and a user's average propensity to consume (APC) across the sample period. Standard errors clustered at the individual level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A.5: One-Week Nondurable Spending Responses by Cluster

Event	Cluster 1	Cluster 2	Cluster 3	Cluster 4
Panel A: One-off events				
COVID Stimulus 1: Post-MPC	0.1966***	0.0153***	0.1095***	-0.0016
	(0.0090)	(0.0021)	(0.0039)	(0.0036)
Pre-MPC	0.0614***	0.0146***	-0.0110***	-0.0219***
	(0.0064)	(0.0020)	(0.0028)	(0.0032)
COVID Stimulus 2: Post-MPC	0.3251***	0.0229***	0.2067***	-0.0227***
	(0.0135)	(0.0034)	(0.0059)	(0.0065)
Pre-MPC	0.1039***	0.0316***	-0.0599***	-0.0454***
	(0.0125)	(0.0032)	(0.0044)	(0.0057)
COVID Stimulus 3: Post-MPC	0.1532***	0.0233***	0.0986***	0.0151***
	(0.0067)	(0.0016)	(0.0031)	(0.0029)
Pre-MPC	0.0403***	0.0096***	0.0073***	-0.0102***
	(0.0051)	(0.0015)	(0.0024)	(0.0026)
Panel B: One-time monthly payments for 6 months				
Child Tax Credit: Post-MPC	0.6621***	0.0546***	0.3224***	-0.2527***
	(0.0281)	(0.0079)	(0.0122)	(0.0164)
Pre-MPC	0.4693***	0.1238***	-0.0883***	-0.3263***
	(0.0300)	(0.0081)	(0.0124)	(0.0157)
Panel C: Once every year				
Alaska PF Dividend: Post-MPC	0.2038***	0.0481***	0.1349***	0.0517***
	(0.0235)	(0.0026)	(0.0049)	(0.0058)
Pre-MPC	0.0599***	0.0214***	0.0119***	0.0175***
	(0.0094)	(0.0020)	(0.0027)	(0.0044)
Tax Refund: Post-MPC	0.1129***	0.0179***	0.0831***	0.0076***
	(0.0170)	(0.0010)	(0.0026)	(0.0012)
Pre-MPC	0.0314***	0.0063***	-0.0038***	-0.0038***
	(0.0045)	(0.0005)	(0.0006)	(0.0011)
Panel D: Periodically				
Monthly Paycheck: Post-MPC	0.1844***	0.055***	0.1103***	0.0322***
	(0.0046)	(0.0011)	(0.0022)	(0.0017)
Pre-MPC	0.0872***	0.0275***	0.0168***	0.0041***
	(0.0032)	(0.0008)	(0.0014)	(0.0014)

Notes. This table presents the one-week nondurable spending responses around the payment date by cluster. Date and individual-by-day of month fixed effects are applied. Standard errors clustered at the individual level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A.6: Mean Normalized Purity of Cluster Assignments by Number of Clusters

Number of Clusters (k)	Mean Normalized Purity
2	0.553
3	0.542
4	0.554
5	0.540
6	0.434
7	0.436
8	0.396
9	0.392
10	0.368

Notes. This table reports mean normalized purity of within-person cluster assignments for each value of k. Normalized purity equals zero under random assignment and one under perfect within-person persistence across events; see Appendix Section A.2.