Anticipatory Spending: Identifying Consumer Types with Pre-Payment MPCs*

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

July 17, 2026

Abstract

How much of the heterogeneity in marginal propensities to consume (MPCs) reflects transitory financial circumstances versus persistent behavioral traits? We address this question by jointly estimating novel "pre-payment" MPCs—spending increases prior to payment arrival—and traditional post-payment MPCs. Using high-frequency transaction data, we leverage multiple pre-announced income changes per person across a range of fiscal policies. We find a robust positive correlation between pre- and post-payment responses, which exhibits substantially higher within-person persistence than predicted by canonical models with liquidity constraints. Building on this evidence, we use k-means clustering to classify individuals into four distinct types: *Active Savers* (43%), *Passive Savers* (23%), *Active Spenders* (9%), and *Passive Spenders* (25%). By examining auxiliary dimensions such as liquidity, payment size, and event salience, we map these empirical types to established consumption models such as rational consumption smoothing, inattention, mental accounting, and present bias. Our findings suggest that persistent behavioral traits are important drivers of consumption dynamics and help explain the empirical attenuation of forward guidance.

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, Chicago Fed, FGV-Sao Paulo, INSPER, PUC-Rio, the Universities of Copenhagen, New South Wales, St. Gallen, Sydney, Vienna, WHU, 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) is central to modern macroeconomics. It determines the transmission of fiscal transfers, governs monetary policy potency in heterogeneous-agent (HANK) models, and serves as a sufficient statistic for the welfare effects of stabilization (Kaplan and Violante, 2014; Auclert, 2019). While a large empirical literature confirms that MPCs vary substantially across individuals (Jappelli and Pistaferri, 2010; Havranek and Sokolova, 2020), standard financial observables such as income and liquid wealth leave most of this variation unexplained (Kaplan and Violante, 2022).

This residual variation raises a fundamental question: are individuals' consumption responses driven by *transitory economic circumstances* or *persistent behavioral traits*? In canonical models, high MPCs arise from temporary binding liquidity constraints due to negative income shocks (Zeldes, 1989; Deaton, 1991). In this "circumstance-based" view, the population of high-MPC agents shifts with the business cycle, rendering fiscal multipliers state-dependent and policy impacts less predictable. Alternatively, if high MPCs reflect "spender-saver" traits (Campbell and Mankiw, 1989), consumption responses are more stable across environments, expectation management is less powerful, and policy effectiveness depends more on features like salience and payment size than on the timing of cash flows. Distinguishing between these mechanisms requires a source of variation that complements the static financial snapshots used in previous research, such as income, liquidity, or credit access.

This paper introduces a new diagnostic for consumption dynamics: the *pre-payment MPC* (pre-MPC). We define the pre-MPC as the spending response in the week immediately preceding a predictable income change. While the literature typically focuses on *post-MPCs* (excess spending after income arrives), canonical models impose sharp restrictions on both. In standard buffer-stock models, unconstrained consumers smooth consumption by responding when they learn of future income, while constrained consumers respond only upon arrival. Theoretically, the sets of "anticipatory spenders" and "post-payment spenders" should be distinct. By estimating pre- and post-MPCs together, we can quantify the extent to which standard models explain observed dynamics and, crucially, use the joint distribution to classify individuals into behavioral types where they cannot.

Identifying pre-MPCs is empirically demanding: it requires high-frequency data to capture brief anticipatory windows and large samples to detect the relatively small spending changes predicted by consumption smoothing in canonical models. We utilize transaction-level data for over 50,000 U.S. individuals from 2012 to 2023, leveraging a suite of exogenous, pre-announced income increases, including COVID-19 stimulus rounds, Child Tax Credit expansions, Alaska Permanent Fund Dividends, and state and federal tax refunds. These events vary in frequency, size, and salience, allowing us to disentangle liquidity from behavioral effects. Most importantly, our panel allows us to observe the same individual across multiple events, a feature essential for separating persistent consumer traits from transitory financial shocks. We validate our data by replicating the post-MPC estimates of five major prior studies that utilized different data sources for these same events.

Utilizing this panel, we establish two main findings. First, we document that heterogeneity in both pre- and post-MPCs is pervasive, but that single-event MPCs mix persistent heterogeneity with substantial event-level consumption noise. While high-frequency transaction data minimize the measurement error common in expenditure surveys (Baker and Kueng, 2022), they do not eliminate idiosyncratic consumption movements unrelated to income, such as emergency repairs or one-off purchases. We use repeated income events per person to separate these components. Among consumers observed in at least ten income events, the standard deviation of one-month post-MPCs falls from 79% for a single event to 30% after averaging over ten events, yet it remains well above the homogeneous-consumer benchmark. A simple variance decomposition implies that the persistent component has a standard deviation of 18%, while the event-level noise component has a standard deviation of 76%. Hence, most of the dispersion in single-event MPCs is noise that averages out, but a substantial residual remains. This residual dispersion is the signal we use to classify consumers into persistent types. This persistence suggests that person-level differences, rather than purely random shocks, drive a significant portion of the observed variation in consumption responses.

Second, we analyze the *joint* distribution of pre- and post-MPCs, providing a sharp test of canonical theory. Standard models – including the permanent-income hypothesis and buffer-stock frameworks – predict a *negative* correlation between these two measures: unconstrained smoothers should exhibit small, positive pre- and post-MPCs, while constrained agents should show a zero pre-MPC followed by a large post-MPC. Contradicting this benchmark, we find a robust and significant *positive* relationship. On average, consumers who increase spending in anticipation of a transfer also spend most heavily upon its arrival.

Moreover, these behaviors are persistent: an individual's pre-MPC and post-MPC for one event predicts their responses surrounding the next, even after conditioning on income and liquid assets. In the cross-section, high liquidity produces the standard association with low MPCs, but this relationship decays considerably when including individual-specific fixed effects and relying on variation *within* individual across events. Together, these patterns suggest that persistent behavioral "consumer types," rather than transitory liquidity positions, are the primary drivers of consumption dynamics.

Motivated by this evidence of persistent individual dynamics, we use the joint distribution of pre- and post-MPCs to group individuals into four distinct consumer types. While the pioneering work of Campbell and Mankiw (1989) and Mankiw (2000) uses a one-dimensional "spender-saver" split to describe post-MPCs, our two-dimensional approach reveals critical heterogeneity within those groups. We identify: (i) *Active Savers* (43% of the sample), who exhibit small, symmetric responses consistent with rational smoothing; (ii) *Passive Savers* (23%), who are largely unresponsive to all but the most salient events; (iii) *Passive Spenders* (25%), who exhibit high sensitivity only upon arrival, tracking the classic hand-to-mouth rule; and (iv) *Active Spenders* (9%), who spend both in anticipation and upon receipt. We use three different tests to demonstrate that these classifications contain a persistent component. First, among consumers with at least five events, the share of a consumer's events assigned to their own modal type is higher than a

random benchmark. Second, event-to-event type transitions exhibit above-chance persistence for every type. Third, in a fully out-of-sample test, a user's spending response around regular paycheck receipts places a user within their own measured type more often that would be expected under a random null. This stability reinforces the conclusion that persistent behavioral traits, rather than fluctuating liquidity positions, are the primary drivers of MPC heterogeneity.

We then map these empirical types to established consumption models by examining auxiliary dimensions not targeted by the clustering algorithm, such as liquidity, income, average propensity to consume (APC), transfer size, and event salience. We find that while *Active* and *Passive Savers* share similar high-liquidity profiles, they differ sharply in their attention to news: Active Savers pull spending forward specifically for salient events, behaving as *rational unconstrained agents*, whereas Passive Savers appear *inattentive or near-rational* (Akerlof and Yellen, 1985). Similarly, both spender types exhibit low liquidity and frequent overdrafts, but their anticipatory behavior distinguishes them. *Active Spenders* behave like *present-biased consumers* (Laibson, 1997), impatiently pulling spending forward and continuing to spend when the cash arrives. In contrast, *Passive Spenders* appear to follow *mental accounting* rules (Thaler, 1985), remaining unresponsive to upcoming payments even when they have the financial capacity to spend, exhibiting a behavioral reluctance to consume before income is 'in hand.'

Heterogeneity in MPCs driven by consumer types has direct, first-order implications for macroeconomic modeling. Current HANK models typically generate MPC heterogeneity through a single channel: the endogenous distribution of liquid wealth (Kaplan et al., 2018; McKay and Wolf, 2023; Auclert et al., 2025). Our results suggest that while liquidity is an essential constraint, it does not fully describe consumer behavior. The positive within-person correlation between pre- and post-MPCs, together with the existence of distinct behavioral types within identical liquidity groups, indicates that next-generation models should incorporate ex-ante behavioral heterogeneity alongside standard ex-post heterogeneity from endogenous asset positions. Specifically, variation in discount factors, attention, or mental accounting rules is required to match the observed two-dimensional joint distribution of spending dynamics. The population shares, combined with the persistence of the classification of individual consumers into behavioral types we document, provide an empirical benchmark to discipline the calibration of these behavioral macro frameworks (Farhi and Werning, 2019; Gabaix, 2019; Pfäuti and Seyrich, 2026).

For example, this behavioral taxonomy provides a direct micro-foundation for the *forward guidance puzzle*: the stark empirical disconnect between the powerful aggregate demand effects of future-dated policies predicted by canonical theory and the more modest responses observed in practice (Carlstrom et al., 2015; Kiley, 2016; McKay et al., 2016; Del Negro et al., 2023). Our classification reveals the mechanism driving this attenuation: roughly 25% of individuals behave as 'mental accountants' who exhibit a behavioral reluctance to pull spending forward despite having the financial capacity to do so, while an additional 23% are largely inattentive to upcoming policy events. Consequently, while we detect anticipatory responses across all income events that are sizable when viewed through the lens of unconstrained rational benchmarks such as the PIH, the aggregate pre-payment effects are macroeconomically small – typically an order of

magnitude smaller than traditional post-MPCs. While these estimates of anticipatory spending responses may represent a lower bound due to the inherent difficulty of pinning down exact announcement dates and heterogeneity in consumers' policy awareness, their small aggregate magnitude suggests they contribute negligibly to aggregate demand. For policy design, this implies that neglecting the pre-announced nature of fiscal stimulus introduces only minimal downward bias when evaluating total aggregate demand effects.

Distinguishing between these underlying mechanisms is not merely a classification exercise; it fundamentally alters the optimal design, distributional consequences, and welfare implications of fiscal and monetary stabilization policies. If MPC heterogeneity were purely circumstance-based, the population of high-MPC agents would shift mechanically with transitory shocks and business cycle fluctuations. However, as mentioned above, we find that event-level type assignments are persistently above chance and that deviations are structured along the active–passive margin within the spender and saver families rather than across the spender–saver divide. Decomposing these persistent behavioral traits is in turn crucial for policy targeting. For instance, incorporating present bias into an otherwise standard HANK model amplifies both fiscal and monetary policy shocks substantially (Lee and Maxted, 2023; Maxted et al., 2025), whereas the presence of fixed mental accountants dampens the effectiveness of expectation management (Lian, 2019, 2023). More generally, identifying these specific behavioral frictions is essential in an environment where financial decisions are increasingly self-directed. Households today directly manage complex retirement portfolios, navigate tax credit schedules, and utilize FinTech products. Heterogeneity in discount factors, attention, and mental accounting rules therefore has profound implications for how households respond to these policies – implications that remain hidden in models focusing solely on liquid assets or income.

Literature.— This paper contributes to three strands of literature. First, we extend the literature estimating consumption responses to unanticipated (Hall and Mishkin, 1982; Blundell et al., 2008) and predictable income changes (Altonji and Siow, 1987; Johnson et al., 2006), as well as the more recent work documenting their wide distribution (Misra and Surico, 2014; Lewis et al., 2026; Boehm et al., 2025).[1] While prior work often focuses on single events, our multi-event approach allows us to document how responses vary across a wide range of policy environments within the same individual.[2]

Second, we provide evidence on anticipatory spending, relating to the classic "excess smoothness" literature which documented that aggregate consumption responds too little to new information about future income (Deaton, 1986; Campbell and Deaton, 1989).[3] While survey-based

[1]More recent studies of post-MPCs include Fagereng et al. 2021; Ganong et al. 2025 for unanticipated shocks and Andreolli and Surico 2021; Crawley and Kuchler 2023; Commault 2024; Koşar et al. 2024 for anticipated income changes.

[2]We also find that spending responses to regular paychecks yield qualitatively similar but quantitatively larger pre- and post-MPCs, consistent with prior research (Olafsson and Pagel, 2018; Gelman, 2022). However, as documented by Gelman et al. (2014), Baker and Yannelis (2017), and Gelman et al. (2020, 2022), these large MPCs often reflect the timing of regular bill payments rather than consumption. The income payments studied in this paper are less frequent and do not lend themselves to this type of liquidity management.

[3]Early tests of anticipation effects using aggregate time-series data include Blinder (1981), Flavin (1981),

evidence suggests individuals react to hypothetical future income (Fuster et al., 2021; Colarieti et al., 2024), recent empirical work using individual transaction data provides a more nuanced picture. For instance, Bräuer et al. (2022) use high-frequency bank data to show that even investors who clearly anticipate dividend payments increase their spending precisely upon receipt. Similarly, Thakral and Tô (2024) find that the consumption response to anticipated transfers depends on the length of the planning horizon, with households exhibiting "excess sensitivity" to receipt despite the pre-announced nature of the income. We reconcile these mixed findings, where studies using micro-data—including our own—often reveal null average anticipation despite large post-payment responses,[4] by showing that anticipation is concentrated among specific consumer types and is highly sensitive to event salience.[5]

Finally, we join the debate on whether MPC heterogeneity reflects transitory *circumstances*—shocks to liquidity or balance sheets (Carroll, 1997; Kaplan et al., 2014)—or persistent *behavioral traits* such as patience, attention, or financial sophistication (Parker, 2017; Gelman, 2021; Indarte et al., 2024).[6] Our findings align with evidence from Scandinavian administrative data suggesting that financial behavior has a stable, potentially hard-wired component (Cesarini et al., 2010; Cronqvist and Siegel, 2015). We provide the first empirical quantification of these population shares based on revealed consumption dynamics, rather than model-imposed assumptions. Grounded in the joint distribution of pre- and post-MPCs, this classification offers a new empirical benchmark for calibrating structural models of consumption with behavioral heterogeneity.

The rest of the paper proceeds as follows. Section 2 describes the data and fiscal interventions. Section 3 discusses the empirical strategy and documents the heterogeneity in pre- and post-MPCs. Section 4 leverages the joint distribution of pre- and post-MPCs to document the underlying persistence and Section 5 uses it to classify consumers into types. Section 6 uses additional data moments not targeted by the clustering algorithm to map these types to established consumption models. Section 7 concludes.

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,

Poterba (1988), and West (1988).

[4]See Broda and Parker 2014; Kueng 2018; Baugh et al. 2021; Graham and McDowall 2025.

[5]Other work has documented anticipatory responses to non-income events, such as sales tax increases (Baker et al., 2018, 2021; Angelucci et al., 2024) and federal income tax changes (Kueng, 2014). We contribute to this line of inquiry by focusing on the anticipatory dynamics of transfer payments.

[6]Further evidence supporting the role of persistent behavioral frictions includes Kuchler and Pagel (2021), Maxted (forthcoming), Mijakovic (2023), Gelman and Roussanov (2024), and Hamilton et al. (2024).

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' 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 users from January 2012 to September 2023.[7]

Table 1, panel A, reports the summary statistics for the users 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 users. We define income as including not only salary but also other types of income, such as pensions, bonuses, and interest income. In addition, it presents several household financial characteristics (at the user-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.

[7]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.

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). 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 2.3.

2.2 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.

2.2.1 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.[8] 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.

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.

As mentioned in the introduction, the announcement date does not necessarily coincide with individuals receiving and processing information about the policy and the upcoming transfer payments, the so-called 'news shocks' (Beaudry and Portier 2004, 2006; Kueng 2014; Ramey 2016). Some people might not follow the news while others might receive information already before the official announcement, especially for repeated events. We therefore do not report spending responses around these announcement dates. Instead, we use them to define a common pre-period window of one week that we can apply across all events without pushing the start of the pre-period window before the announcement date and without having the post-period window of one event overlap with the pre-period window of the next event.

2.2.2 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. This approach 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.[9] Prior work, such as Da et al. (2011), Choi

[8]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.

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

and Varian (2012), Baker and Fradkin (2017), and Baker et al. (2021) demonstrate the utility of using Google Trends to measure contemporaneous attention for a wide range of economic and financial topics. 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. 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 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.

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).

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 a consumer-specific level.

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.

2.3 Data Validation: Replication of Five Prior Post-MPC Estimates

While the focus in this paper is on understanding heterogeneity in individuals' consumption responses to predictable income changes by leveraging novel pre-MPCs, the response *following* the particular income payments – what we call the post-MPC – has been previously estimated

by researchers in a range of studies. 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 spending across the liquidity distribution. We mirror their approach, calculating the MPC across all expenditures in the one month 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 a three-month MPX of 0.37 relative to their estimate of 0.41.

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).[10] Despite using

[10]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, respectively, 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-month MPCs reported in Table A.2 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.

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 Average 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 reported in the online appendix, 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).

Following the recent literature on two-way fixed effects models with staggered treatment (e.g., Borusyak et al. 2024), we estimate pre- and post-payment MPCs using the following dynamic distributed lag specification of daily consumption expenditures of individual i on calendar date t controlling for date (α_t) and individual-by-day of month fixed effects ($\alpha_{i,dom(t)}$) to account for common shocks across all consumers and for regular spending patterns that occur on a specific day of the month for different individuals (e.g., regular shopping trips, work lunches):

$$c_{i,t} = \sum_{s=-28}^{90} \beta_{s,e} \, \mathbf{1}[t = T_{i,e} + s] \times Amount_{i,e} + \alpha_{i,dom(t)} + \alpha_t + u_{i,t} \tag{1}$$

Standard errors are clustered by individuals. $Amount_{i,e}$ is the amount of the income a treated individual receives at date $T_{i,e}$ from fiscal transfer event e. This specification employs several sources of identifying variation in both the receipt ($\mathbf{1}[t = T_{i,e} + s]$), timing ($T_{i,e}$), and amount of income transfers conditional on receipt ($Amount_{i,e}$). 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.2 for convenience). We then sum daily responses to obtain cumulative pre- and post MPCs over different horizons before and after receiving an anticipated income transfer: $\beta^{pre}_{e,\underline{h}} = \sum_{s=-\underline{h}}^{-1} \beta_{s,e}$ for horizon $\underline{h} = 7$ days and $\beta^{post}_{e,\bar{h}} = \sum_{s=0}^{\bar{h}} \beta_{s,e}$ for horizons $\bar{h} \in \{7, 28, 90\}$ days.[11]

[11] For the monthly events, we include the 7 days before and the 14 days after the events.

3.1 Post-Payment MPCs: Excess Sensitivity

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 responses across payment events. MPCs generally range from 0.03 to 0.13 in the week following the income's arrival to approximately 0.15 to 0.38 during the first three months after the transfer.[12] 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 Pre-Payment MPCs: Anticipatory Spending Behavior

Table 4 also examines the pre-payment MPCs for one week prior to income arrival. We do find that users increase spending in the week prior to the arrival of income, but that magnitudes for most events that 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).[13] 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 users, 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.3 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.

[12] 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.

[13] 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.

4 The Distribution of Individual Pre- and Post-MPCs

4.1 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 the vast majority of individuals experience multiple events. Because person-level daily spending responses are too noisy and computationally costly, we estimate individual event-level cumulative pre- and post-payment MPCs directly using the following regression specification previously used in Baker et al. (2023):

$$c_{i,t} = \beta^{pre}_{i,e,\underline{h}} \, Pre_{i,t,e,\underline{h}} \times Amount_{i,e} + \beta^{post}_{i,e,\bar{h}} \, Post_{i,t,e,\bar{h}} \times Amount_{i,e} + \alpha_{i,dom(t)} + \alpha_t + u_{i,t}, \tag{2}$$

where $\beta^{pre}_{e,i,\underline{h}}$ and $\beta^{post}_{e,i,\bar{h}}$ are allowed to vary across individuals and payment events and thus estimate heterogeneous cumulative MPCs for each individual and each event separately. $Pre_{i,t,e,\underline{h}} = \mathbf{1}[T_{i,e} - \underline{h} \leqslant t < T_{i,e}]/\underline{h}$ is an indicator that equals one for treated individuals for $\underline{h}$ days before the payment date $T_{i,e}$, scaled by the estimation window length $\underline{h}$. Similarly, $Post_{i,t,e,\bar{h}} = \mathbf{1}[T_{i,e} \leqslant t \leqslant T_{i,e} + \bar{h}]/\bar{h}$ is the scaled indicator for the post period. To illustrate why β corresponds to the cumulative MPC over the period of length h, imagine that a \$1 transfer leads to \$1 dollar of additional spending in the day immediately after receipt, a one-day post-MPC of 100%. Thus, if we estimated the effect over one day, we would scale by $\bar{h} = 1$ and $\beta^{post}_{i,e,1} = 1$. If we estimate the effect over 10 days, the average effect each day is 0.1, which would be the coefficient on a regression of $Post_{i,t,e,\bar{h}} \times Amount_{i,e}$. Dividing by $\bar{h} = 10$ scales the coefficient by 10 so that the cumulative post-MPC is again $\beta^{post}_{i,e,10} = 1$.[14]

4.1.1 Unconditional and Within-Person Distributions of MPCs

Figure 1 plots several distributions of one month post-payment nondurables MPCs across our sample. 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. (2026), 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

[14] Appendix Table A.2 shows that estimating the equivalent specification for average MPCs imposing treatment effect homogeneity, $c_{i,t} = \beta^{pre}_{e} \, Pre_{e,i,t,\underline{h}} \cdot Amount_{i,e} + \beta^{post}_{e} \, Post_{e,i,t,\bar{h}} \cdot Amount_{i,e} + \alpha_{i,dom(t)} + \alpha_t + u_{i,t}$ yields similar results as specification (1).

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 an 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.

4.2 Persistent Behavior versus Event-Level Consumption Noise

The wide dispersion of individual-by-event MPCs raises a natural question: how much of the residual dispersion in within-person averages reflects persistent heterogeneity across consumers, and how much is event-level noise that has not yet averaged out? We address this with a within-person design that holds the population fixed. We keep the $N = 13{,}490$ consumers who experience at least $K = 10$ income events, and for each $E = 1, \ldots, 10$ we draw E of that consumer's events at random and average their post-MPCs. Because the same consumers enter at every E, any decline in dispersion reflects the averaging of noise rather than a change in *who* is being averaged. This design therefore does not require the assumption that consumer types are unrelated to the number of events observed and thus is free from potential biases stemming from changes in the composition of the sample population as we vary E.

Volatility test. Panel (b) of Figure 1 plots the standard deviation of within-person average post-MPCs against E. Under the null that individual MPCs equal a common population MPC plus i.i.d. draws of "consumption noise," the within-person average has standard deviation $\sigma_{\text{pop}}/\sqrt{E}$, which falls to zero as E grows. The data decline steeply, from 78.6% at $E = 1$ to 30.0% at $E = 10$, but they stay above this benchmark at every $E \geqslant 2$ and flatten toward a positive floor rather than collapsing toward zero. Fitting $\text{sd}(E)^2 = \sigma^2_{\text{true}} + \sigma^2_{\text{noise}}/E$ gives a persistent component of $\sigma_{\text{true}} =$ 17.8% (95% CI 16.4–18.9) and an event-level noise component of $\sigma_{\text{noise}} = 76.1\%$. Hence, most of the dispersion in single-event MPCs is noise that averages out, but a substantial and precisely estimated dispersion survives averaging. It reflects ex-ante heterogeneity of consumer types together with ex-post heterogeneity stemming from differences in temporary circumstances.

Share-negative test. Appendix Figure A.5 applies the same design to the share of consumers whose within-person average post-MPC falls below zero, -5%, and -10%. The share below zero

declines from 41.4% at $E = 1$ to 29.0% at $E = 10$, so averaging over ten events removes less than a third of the negative averages. The shares below -5% and -10% decline in parallel, from 35.8% to 22.3% and from 30.9% to 16.8%. Both nulls apply the central limit theorem to the within-person average and predict a negative share of $\Phi\big(-\mu_{\text{pop}}/\text{sd}(E)\big)$, with $\mu_{\text{pop}} = 15.6\%$ and $\text{sd}(E)$ taken from the corresponding null in Panel (b). They coincide at $E = 1$ and separate as E grows. The homogeneous null falls to 26.4% at $E = 10$ and lies below the bootstrap interval from $E = 6$ onward, so this test rejects it as well. The heterogeneous null, which uses the σ_{true} and σ_{noise} estimated in Panel (b) without any re-fitting, tracks the data far more closely and stays inside the interval through $E = 8$, at 42.1% versus 41.4% at $E = 1$ and 30.1% versus 29.0% at $E = 10$. It overshoots slightly at $E = 9$ and $E = 10$, which is what a right-skewed distribution of true MPCs produces: once the noise has largely averaged out, the normal approximation underlying $\Phi(\cdot)$ places too much mass in the left tail.

The dispersion of true MPCs. Our $\sigma_{\text{true}} = 17.8\%$ measures the dispersion of *persistent* MPCs, purged of event-level noise, and few comparable numbers exist. Most papers that document MPC heterogeneity report quantiles or group-specific means rather than a second moment. Jappelli and Pistaferri (2014) are the exception, reporting a standard deviation of stated MPCs of 35.7% (their Table 1, mean 47.6%), and Jappelli and Pistaferri (2020) report 36% and 35% in the two waves of the same survey. That number is not comparable to ours in either direction. The survey enforces that the spend and save shares sum to 100, which truncates dispersion from below, while 62% of the mass sits on the three focal answers 0, 0.5, and 1, which inflates the second moment through rounding rather than through heterogeneity. Two further estimates can be recovered from published tables rather than read off them.[15] Fuster et al. (2021) do not report a standard deviation, but their Table 3 implies one of roughly 28% for the \$500 gain. Lewis et al. (2026) do not either, but their estimated three-point latent distribution (MPCs of 0.04, 0.23, and 1.33 with population shares of 0.30, 0.48, and 0.23) implies a mean of 0.42 and a standard deviation of about 50%. The latter is the closest comparison available, since their latent distribution is also meant to be free of noise, though it refers to quarterly total expenditure rather than to one-month nondurable spending and so includes durable purchases, which are both larger and more dispersed. Hence, the comparisons are loose, and the useful statement is a methodological one: because we observe the same consumer over many pre-announced events, σ_{true} is estimated rather than assumed, and it does not rest on a rank-invariance assumption, a parametric mixture,

[15] Neither paper reports a standard deviation; we compute both from their published tables. For Lewis et al. (2026) we use the estimated latent distribution in their Table 1, which places MPCs of $\lambda = (0.04, 0.23, 1.33)$ on population shares of $\pi = (0.30, 0.48, 0.23)$. Normalizing the shares, which sum to 1.01 in their table because of rounding, the implied mean $\sum_g \pi_g \lambda_g = 0.42$ reproduces the average MPC they report, and the implied standard deviation is $\big[\sum_g \pi_g (\lambda_g - \bar{\lambda})^2\big]^{1/2} = 50\%$. We use the latent distribution rather than the posterior-weighted individual MPCs plotted in their Figure 1, because posterior weighting shrinks each household toward the mean and would understate the dispersion of the underlying types. The calculation treats the estimated distribution as known and hence ignores estimation uncertainty in λ_g and π_g. For Fuster et al. (2021) we invert the standard error of the mean reported in their Table 3: with SE $= 0.007$ and $N = 1{,}638$ for the \$500 gain scenario, SD $\approx$ SE$\sqrt{N} = 28\%$. This is approximate, since their statistics are weighted and their reported means are winsorized at the 2.5th and 97.5th percentiles, both of which break the exact relation.

or the accuracy of a stated plan.

Negative MPCs. Our estimates of the share of individual-level MPCs that are negative help reconcile a disagreement in the literature about whether negative MPCs are real. Two papers use the same CEX rebate data and reach opposite conclusions. Misra and Surico (2014) estimate quantile regressions of the change in spending on the rebate; in their Figure 1 the nondurable response crosses zero at about the 35th percentile for the 2008 rebate and the 40th for 2001, which implies a large negative share provided that households keep their rank in the conditional spending distribution. Lewis et al. (2026) instead fit a three-component Gaussian mixture to the same data and find essentially no negative mass. That result is close to mechanical: their household-level MPCs are posterior-weighted averages of the three estimated group MPCs, all of which are positive, so the entire left tail is absorbed by the within-group error rather than by a negative-MPC type. The two designs differ not in their data but in how much of the left tail each assigns to noise, and neither can settle the question, because each household is observed for a single event.

Our design does not need that assumption, because we observe the same consumer across many events, and it delivers both numbers. At $E = 1$ the measured share of negative post-MPCs is 41.4%, close to what the quantile-regression reading implies. This is arithmetic rather than evidence of widespread dissaving: single-event responses have mean $\mu_{\text{pop}} = 15.6\%$ against a standard deviation of $\sigma_{\text{pop}} = 78.1\%$, and a small mean on a wide distribution mechanically places much mass below zero. As events are averaged the noise falls away and the share declines toward $\Phi\left(-\mu_{\text{pop}}/\sigma_{\text{true}}\right) = 19.0\%$. Hence, most of the measured negative mass is event-level noise, but not all of it, which is why both prior readings can look right on the same data.

Designs that avoid consumption noise point in the same direction. Boehm et al. (2025) recover the MPC distribution by deconvolution in a randomized transfer experiment. Their baseline constrains the support to be positive, but their unconstrained estimate places the bottom 10 to 15% of the unrestricted-card arm below zero (their Figure A25). Fuster et al. (2021) instead elicit stated spending responses: respondents report the causal response directly, so it never has to be inferred from observed spending, which moves for many reasons other than the transfer. Their design also does not restrict answers to $[0, 1]$, and 8% report a negative response to a \$500 windfall. Both estimates lie below our 19.0%, and both comparisons are bounded in the same direction: our figure overstates the true share because the normal approximation over-weights the left tail of a right-skewed distribution, while stated responses remain subject to reporting error, which inflates both tails if it is symmetric. Note that eliciting the response directly removes consumption noise, not measurement error. Taken together, the evidence supports a small share of consumers with a persistently negative MPC, well below what single-event estimates suggest. Appendix A.1 records the design differences behind these comparisons.

Hence, the volatility and share-negative tests support the same conclusion. Most of the wide cross-sectional dispersion in individual-event MPCs is event-level noise that averages out, but a substantial residual persists: the dispersion of within-person averages settles at $\sigma_{\text{true}} = 17.8\%$, and 29% of consumers still show a negative average post-MPC after ten events. Both statistics

reject the homogeneous-MPC null once enough events are averaged, and the same $(\sigma_{\text{true}}, \sigma_{\text{noise}})$ pair matches both. This residual motivates the cluster analysis in Section 5, where we ask whether consumers separate into a small number of stable behavioral types.

4.3 Persistence of Consumption Responses

4.3.1 The Joint Distribution of Pre- and Post-MPCs

The large average post-payment MPCs and the small but statistically significant average pre-payment MPCs shown in Table 4 across several fiscal transfer events are 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 $\beta_{pre,i} \approx \beta_{post,i}$ 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. 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 various post-payment MPCs. In particular, we include one-week post-payment MPCs, one-month MPCs, and one-month MPCs while excluding spending in the first week after payment (to mitigate any mechanical correlation in spending immediately before and after payment from a single shock like a vacation timed to span payment arrival). We confirm the graphical relationship while accounting for event and user 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.

4.3.2 Salience and the Predictive Power of Pre- & Post-MPCs Across Events

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

In columns 1 and 3, we find that there is a significant persistence in the size of a users' consumption response, both prior to income arrival and following arrival. That is, users 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 articles in the user's state that discussed the payment. 'Repeat Event' denotes an event type with which the user 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, users need to be aware of the income shock in order for persistent characteristics to drive similarity in spending responses.

While higher levels of salience tend to produce more consistent responses over time within individual, they also shift spending across time around a given event. 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 6 explores this shifting in a regression format. In column (1), we find that higher levels of salience are associated with higher levels of spending in the month following payment arrival. However, columns (2) and (3) show that, despite spending being higher across a month, higher levels of salience tend to reallocate some spending from the week after payment arrival to the week before arrival. That is, users pull forward spending into the pre-period when they are likely more aware of the upcoming payments. This also implies a smaller jump in spending around payment arrival (column (4)).

4.3.3 Liquidity and MPC Heterogeneity

We now turn to a more in depth examination of financial 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 users 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 7 examines the relationship between both pre- and post-payment MPC and a primary driver of MPC heterogeneity in the canonical models: liquidity. 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 MPCs.

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 (3) show strong effects of 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.

However, this finding is not robust to controlling for user fixed effects in columns (2) and (4). Focusing on within-individual variation by adding user fixed effects decreases or even flips the sign for these measures of liquidity. Within user, more credit is associated with *larger* post-payment MPCs.

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 (5)–(6) repeat the analysis for pre-payment MPCs. The cross-sectional relationships in column (5) are weaker when compared to those of the post-payment MPC. High credit and low overdrafts also associated with lower MPCs without user fixed effects. Adding user 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 users

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 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.

5 Consumer 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 4.3.3. 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.

5.1 Cluster Algorithm

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.[16] 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

[16] 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), Colarieti et al. (2024), Andre et al. (2025), Boehm et al. (2025), Lewis et al. (2026), and Andre et al. (2025). An early and robust finding of this literature – which our analysis corroborates and sharpens – 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 two exogenous income events: the 2001 and 2008 fiscal stimulus programs. Lewis et al. (2026) extend this approach by employing a parametric Gaussian mixture model to address the substantial measurement error in CEX spending data. Because each household is observed for a single event, however, both designs must *assume* how much of the dispersion is measurement error rather than measure it, and the assumption drives the answer: the two papers use the same 2008 CEX rebate data and reach opposite conclusions about the share of consumers with a negative MPC, as we discuss in Section 4.2. Their reliance on cross-sectional variation and a single income event also 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. 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. Finally, Andre et al. (2025) use within-household survey responses to hypothetical income gains and losses of varying sizes to identify four near-rational "quick-fixing" types, which explain substantial MPC heterogeneity despite being weakly predicted by financial and demographic observables.

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 descriptive names to the resulting groups based on their location in the joint distribution of pre- and post-payment MPCs. Finally, we use additional observables not used by the algorithm, such as socio-economic demographics, to assess whether the descriptive labels capture broader empirical differences across consumers.

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 a simple two-by-two classification that separates the saver-spender dimension from the active-passive timing dimension. This structure naturally extends the two-type approach of Campbell and Mankiw (1989) and Mankiw (2000) by adding the pre-payment MPC as a second dimension.[17]

5.2 Four Consumer Types

The four clusters have a simple interpretation. The post-payment MPC separates savers from spenders, while the pre-payment MPC separates active from passive consumers. We therefore label the groups as *Active Savers*, *Passive Savers*, *Active Spenders*, and *Passive Spenders*. These labels are descriptive: they summarize the joint dynamics of consumption around predictable payments, before we map the groups to models in Section 6.

Table 8 reports the average MPCs that define these types. The first two rows of Panel A report one-week post- and pre-payment MPCs for users assigned to each type. For each type, we separately re-estimate the pre- and post-payment MPCs within the group in order to compute correct within-type standard errors. This re-estimation implies that the type-level MPCs in Table 8 differ slightly from the average individual MPCs used to form the k-means clusters in Figure 4.

Columns 5–6 split individuals by liquidity and show a familiar result from the literature: low-liquidity consumers have much larger post-payment MPCs than high-liquidity consumers. However, the same liquidity split generates little difference in pre-payment MPCs. In contrast, the four consumer types differ along both dimensions. This is the main empirical value of the pre-payment MPC: it separates consumers with similar post-payment MPCs into groups with different timing behavior, splitting the usual saver and spender groups into active and passive types.

Type 1: Active Savers (43% of the sample).— Active Savers have small but positive pre- and post-payment MPCs, consistent with consumers who smooth spending over the payment date.

Type 2: Passive Savers (23%).— Passive Savers have near-zero pre-payment MPCs and small post-payment MPCs. They look similar to Active Savers if one only observes post-payment responses, but they do not pull spending forward before predictable payments arrive.

[17] Appendix A.2.4 shows the effect of using different numbers of clusters, with k ranging from 2 to 6 (see Appendix Figure A.6).

Type 3: Active Spenders (9%).— Active Spenders have large pre-payment MPCs and large post-payment MPCs, indicating that they both pull spending forward and continue to spend upon arrival.

Type 4: Passive Spenders (25%).— Passive Spenders have high post-payment MPCs but much smaller pre-payment MPCs.

5.3 Cluster Persistence

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 an individual's events assigned to their modal type, where each event is assigned to the nearest user-level cluster centroid. We benchmark observed purity against its expectation under random reassignment of event labels, holding fixed the aggregate assignment distribution and each individual's event count (Appendix A.2.1). At our preferred $k = 4$, observed purity lies far outside the null distribution (permutation standard error ≈ 0.0005, $p < 0.01$; Appendix Table A.5), and the same pattern holds at every k from 2 to 10.

Purity is a coarse statistic: most of its level reflects cluster sizes and event-count mechanics rather than behavior. The structure of transitions provides the sharper diagnostic because individual MPCs contain a lot of "noise," being influenced by other lumpy spending decisions unrelated to income changes and this noise remains substantial for average MPCs within individuals even with multiple income events per person.[18] Our second test therefore examines event-to-event transitions: we assign each individual-event observation to the nearest user-level centroid and pair it with the same individual's next observed event.

Appendix A.2.2 details the construction and the permutation inference, and Appendix Table A.6 reports the full matrix of normalized transition probabilities, where a value of one corresponds to no persistence. Aggregate persistence is modest but decisively non-random: Cohen's κ—the observed same-type share of transitions relative to the share expected from the marginal type distributions alone—is 0.060, with a permutation null tightly concentrated at zero ($p < 0.01$). Three patterns emerge. First, all four own-type values exceed one (1.12–1.40). Second, when assignments do change, they move disproportionately *within* the saver or spender family: in every row the second-largest value is the other type in the same family, with flows between Active and Passive Spenders of 1.15 in both directions. Third, every transition that crosses the spender–saver boundary has a normalized probability below one. Event-level noise thus perturbs assignments along the active–passive margin rather than across the spender–saver divide, indicating that the classification captures a hierarchy: a highly persistent spender–saver trait modified by a second active–passive margin.

Both purity and the transition structure, however, are computed on the same events that define the types. Our third test is therefore fully out-of-sample: we ask whether a user's spending

[18] 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.

response around regular paycheck receipt—an income event excluded from the type definitions—recovers the user's type. We assign each user's salary response to the nearest cluster centroid and compare this assignment to the type based on fiscal transfers (Appendix A.2.3). The salary response places 40.3% of users in their own type ($N = 40{,}109$). Against the random assignment benchmark, Cohen's $\kappa = 0.078$, with a permutation null tightly concentrated at zero (standard error ≈ 0.003, $p < 0.01$).

Overall, this stability of cluster assignments alongside the within-person MPC persistence documented in Section 4.3.2, supports the interpretation that the four clusters capture durable behavioral types, measured with considerable idiosyncratic consumption noise, rather than temporary circumstances.

6 Mapping Consumer Types to Models

Having defined the four types descriptively, we now map them to well-known models of consumer behavior. This mapping is informal and interpretive. The clustering algorithm uses only pre- and post-payment MPCs; it does not target liquidity, credit access, payment size, salience, APC, overdrafts, finance employment, or the composition of spending.

The auxiliary moments therefore provide an over-identifying check on whether the descriptive labels correspond to mechanisms suggested by theory. No single model should be expected to match every spending pattern and auxiliary moment perfectly. The goal is instead to ask which model provides the most useful interpretation of each empirical type.

6.1 Active Savers → Canonical Macro/PIH

Active Savers have small positive pre- and post-payment MPCs. Their one-week pre-MPC is about 0.04, while their one-week post-MPC is about 0.05. The similar magnitudes of the two MPCs are the pattern predicted by canonical forward-looking models: once a predictable income change is incorporated into permanent income, consumption rises before the payment arrives and remains at approximately the same higher level afterward (Friedman, 1957; Ando and Modigliani, 1963; Zeldes, 1989; Deaton, 1991; Carroll, 1997). At the same time, the one-week pre-MPC of 0.04 may still be large relative to the response predicted by standard calibrated models. Occasional borrowing constraints can help explain why the average post-payment MPC exceeds the payment's annuity value predicted by the frictionless PIH, but they do not readily explain spending before the payment arrives: households whose borrowing constraints bind cannot spend against the future transfer. Active Savers are therefore the closest empirical counterpart to canonical smoothers among the four types, not literal frictionless PIH agents or calibrated buffer-stock model agents. The mapping remains imperfect because observed individual MPCs combine persistent behavior with idiosyncratic consumption noise, and because true frictionless PIH responses may be too small to distinguish empirically from near-rational inattention, so such consumers may be assigned to the Passive Saver group.

The auxiliary moments support this interpretation. Active Savers have relatively high liquid-

ity, few overdrafts, and indicators of financial sophistication. They also pull spending forward when liquidity or available credit is high (Table 9), and their pre-MPC responds strongly to news coverage (Table 10). These patterns suggest consumers who are attentive and able to adjust the timing of spending when it is useful to do so, including through the timing of larger expenditures around income receipt (Attanasio and Weber, 2010; Aaronson et al., 2012).

6.2 Passive Savers → Near-Rational Inattention

Passive Savers have near-zero pre-payment MPCs and small post-payment MPCs. They look similar to Active Savers on several observables: they have high liquidity, high income, few overdrafts, and low financial distress. However, their consumption responses are much less active. Their pre-MPC is close to zero, and Tables 9 and 10 show little response to liquidity or salience.

This pattern is consistent with near-rational inattention. Their low payment-to-income ratio supports this interpretation. These consumers have the resources to smooth, and the welfare gains from reacting to small predictable payments may be low (Akerlof and Yellen, 1985; Cochrane, 1989; Kueng, 2018). Hence, the muted response need not reflect binding constraints; it can reflect limited attention or small perceived benefits from re-optimizing around each payment.

6.3 Active Spenders → Present Biased

Active Spenders have large pre- and post-payment MPCs. Their pre-payment MPC is more than twice as large as the Passive Spenders' pre-MPC and about three times as large as the Active Savers' pre-MPC. They therefore appear to anticipate future payments, but they do not merely shift spending from after receipt to before receipt; they continue to display high post-payment MPCs after the payment arrives. This combination is consistent with present bias, high impatience, or hyperbolic discounting (Shea, 1995; Angeletos et al., 2001; Aguiar et al., 2025).

The auxiliary moments point in the same direction. Active Spenders have low liquidity, low available credit, many overdrafts, high APC, and low indicators of financial sophistication. Their pre-MPC is also highly responsive to liquidity (Table 9). These consumers are attentive enough to pull spending forward when they can, but their high post-payment MPC and financial distress are consistent with difficulty smoothing consumption over time. At the same time, they continue to exhibit elevated consumption following the arrival of the payment, possibly reflecting time-inconsistency in their consumption plans. This time-inconsistency, or difficulty following plans more generally, is also reflected in the very high number of overdraft fees they incur.

The composition of the spending response further supports this interpretation. Rows 4 and 5 of Table 8 show that Active Spenders have the highest ratio of nondurable response to total consumption response, consistent with impatience and a preference for immediate gratification (Angeletos et al., 2001; Laibson et al., 2022).

6.4 Passive Spenders → Mental Accounting

Passive Spenders have high post-payment MPCs but much smaller pre-payment MPCs. They spend when income arrives, but they do not pull much spending forward. This timing pattern

is consistent with mental accounting: the payment matters when it is received or labeled as available income, even if the payment was predictable beforehand (Thaler, 1985; Shefrin and Thaler, 1988; Hastings and Shapiro, 2013).

The auxiliary moments again help distinguish this group from Active Spenders. Passive Spenders experience financial distress and have low liquidity, but less extreme values than Active Spenders. This may explain why many studies would typically group the two spender types together as liquidity constrained, while more recent work has emphasized mental accounting as a separate mechanism (Bernard, 2023; Mijakovic, 2023). More importantly, their pre-MPC is not very responsive to liquidity, credit, or salience. Thus, the high post-payment MPC does not appear to reflect only temporary borrowing constraints. It instead looks like a stable tendency to spend when the income arrives.

6.5 Discussion

The mapping is not one-to-one. No single model should be expected to match every auxiliary moment in Table 8, Table 9, and Table 10. Active and Passive Savers both have small post-payment MPCs, but only Active Savers pull spending forward when payments are salient. Active and Passive Spenders both have high post-payment MPCs, but only Active Spenders display large anticipatory responses.

This distinction matters for modeling. A one-dimensional saver-spender split can capture the post-payment MPC distribution, but it misses the timing dimension revealed by predictable payments. The joint distribution of pre- and post-payment MPCs instead points to two margins: the saver-spender margin, which governs the level of excess sensitivity to income arrival, and the active-passive margin, which governs whether consumers respond before the payment arrives. This is why the four types are useful even though individual MPCs are noisy, type assignment is imperfect, and the model mapping remains interpretive.

7 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.

Altonji, Joseph G. and Aloysius Siow, "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data," *Quarterly Journal of Economics*, 1987, *102* (2), 293–328.

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

Andre, Peter, Joel P Flynn, Georgios Nikolakoudis, and Karthik Sastry, "Quick-fixing: Near-rationality in consumption and savings behavior," Technical Report, National Bureau of Economic Research 2025.

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.

_ , _ , and _ , "Fiscal and Monetary Policy with Heterogeneous Agents," *Annual Review of Economics*, 2025, *17.*

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.

Beaudry, Paul and Franck Portier, "An exploration into Pigou's theory of cycles," *Journal of Monetary Economics*, 2004, *51* (6), 1183–1216.

_ and _ , "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 2006, *96* (4), 1293–1307.

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.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, 2008, *98* (5), 1887–1921.

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.

Bräuer, Konstantin, Andreas Hackethal, and Tobin Hanspal, "Consuming Dividends," *Review of Financial Studies*, 2022, *35* (10), 4802–4857.

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.

Cesarini, David, Magnus Johannesson, Paul Lichtenstein, Örjan Sandewall, and Björn Wallace, "Genetic variation in financial decision-making," *Journal of Finance*, 2010, *65* (5), 1725–1754.

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.

Cronqvist, Henrik and Stephan Siegel, "The Origins of Savings Behavior," *Journal of Political Economy*, 2015, *123* (1), 123–169.

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.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik, "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 2021, *13* (4), 1–54.

Farhi, Emmanuel and Iván Werning, "Monetary Policy, Bounded Rationality, and Incomplete Markets," *American Economic Review*, 2019, *109* (11), 3887–3928.

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 Macroeconomics," 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.

Gabaix, Xavier, "Behavioral Inattention," in "Handbook of Behavioral Economics," Vol. 2, Else-

vier, 2019, pp. 261–343.

Ganong, Peter, Damon Jones, Pascal Noel, Fiona Greig, Diana Farrell, and Chris Wheat, "Liquid Wealth and Consumption Smoothing of Typical Labor Income Shocks," *Working Paper*, 2025.

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.

Hall, Robert E. and Frederic S. Mishkin, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, 1982, *50* (2), 461–481.

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.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

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.

_ , Giovanni L. Violante, and Justin Weidner, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, *2014* (1), 77–138.

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, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–478.

_ , 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," *Review of Economic Studies*, 2026.

Lian, Chen, "Consumption with Imperfect Perception of Wealth," *Working paper*, 2019.

_ , "Mistakes in Future Consumption, High MPCs Now," *American Economic Review: Insights*, 2023, *5* (4), 563–581.

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.

_ , David Laibson, and Benjamin Moll, "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy," *Quarterly Journal of Economics*, 2025, *140* (1), 691–743.

McKay, Alisdair and Christian K. Wolf, "Monetary Policy and Inequality," *Journal of Economic Perspectives*, 2023, *37* (1), 121–144.

_ , 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.

Pfäuti, Oliver and Fabian Seyrich, "A Behavioral Heterogeneous Agent New Keynesian Model," *Working Paper*, 2026.

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

Ramey, Valerie A., "Macroeconomic Shocks and Their Propagation," *Handbook of Macroeconomics*, 2016, *2*, 71–162.

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.

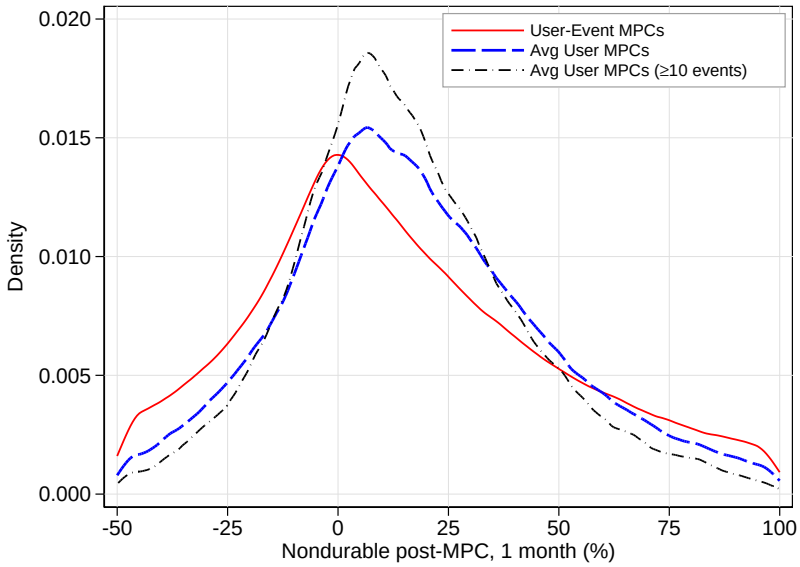
Toczynski, Jan, "Spend or invest? Analyzing MPC heterogeneity across three stimulus programs," *Swiss Finance Institute Research Paper*, 2022.

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

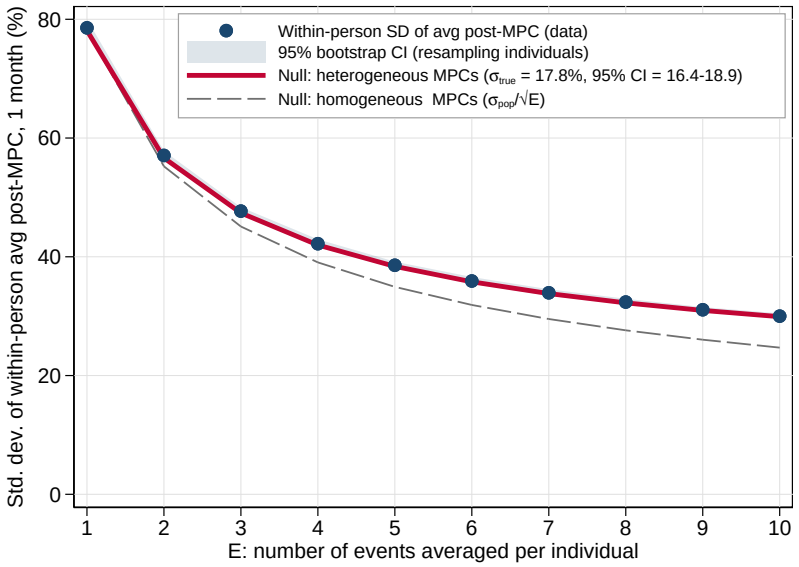
Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 1989, *104* (2), 275–298.

Copy of Figures and Tables from Main Text

Figure 1: Distribution of Post-Payment MPCs



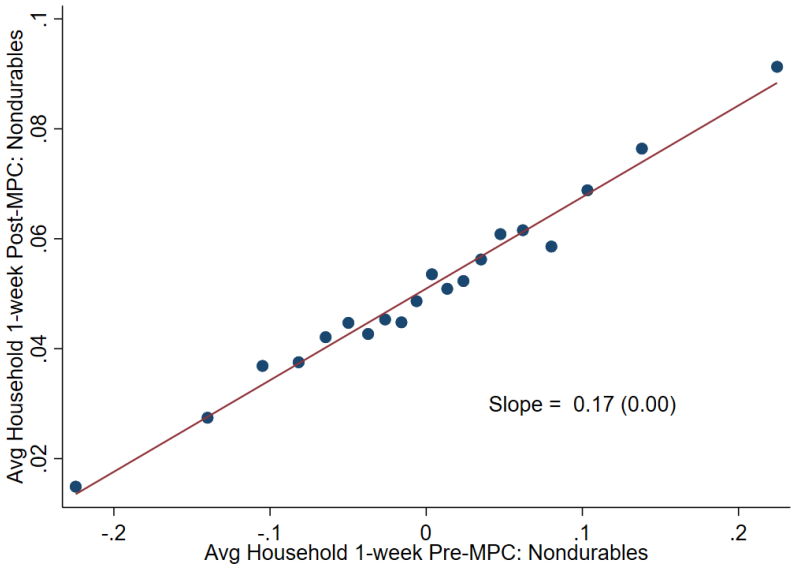
(a) Distribution of Individual Post-MPCs



(b) Volatility of Within-Person Average Post-MPCs (Composition-Free)

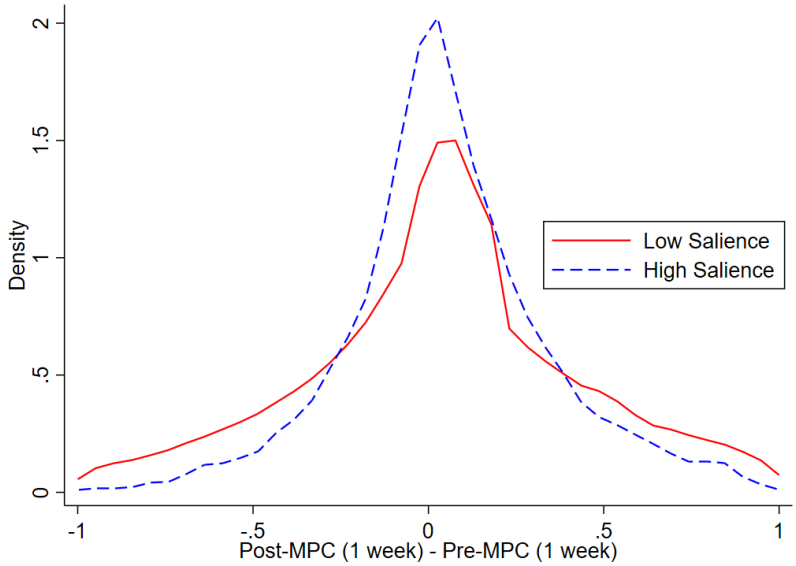
Notes. Panel (a) plots the distribution of one-month post-payment MPCs of nondurables. The red solid line shows the kernel density of the distribution of single user-by-event MPCs. The blue dashed line then plots the same distribution after *averaging* these user-events across all events within a user. 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. Panel (b) traces how the dispersion of within-person average post-MPCs declines as more events are averaged per individual and addresses a composition concern: individuals observed for many events might differ from those observed for few events. We therefore fix the population to the $N = 13{,}490$ individuals observed in at least $K = 10$ income events. For each $E = 1, \ldots, 10$ we draw E of that individual's events at random, compute the within-person average post-MPC, and take its standard deviation across individuals (solid dots, averaged over 200 Monte Carlo draws; shaded band is bootstrapped 95% CI resampling individuals). Because the *same* individuals enter at every E, the decline reflects pure averaging, not composition. We compare this decline to two nulls. Under the homogeneous null (dashed line), all individuals share one true MPC and all dispersion is i.i.d. event-level noise, so the within-person average has standard deviation $\text{sd}(E) = \sigma_{\text{pop}}/\sqrt{E}$. $\sigma_{\text{pop}} = 78.1\%$ is the standard deviation of *individual event-level* post-MPCs in the same fixed sample, before averaging. Under the heterogeneous null (solid line), MPCs differ persistently across individuals with standard deviation σ_{true} and each event adds independent noise with standard deviation σ_{noise}: $\text{sd}(E)^2 = \sigma^2_{\text{true}} + \sigma^2_{\text{noise}}/E$. σ^2_{true} is the intercept of an OLS regression of $\text{sd}(E)^2$ on $1/E$ over $E = 1, \ldots, 10$, and $\sigma^2_{\text{noise}} = \sigma^2_{\text{pop}} - \sigma^2_{\text{true}}$. This yields $\sigma_{\text{true}} = 17.8\%$ and $\sigma_{\text{noise}} = 76.1\%$. Appendix Figure A.4 shows that σ_{true} lies between 17.4% and 19.0% for thresholds $K = 4, \ldots, 11$ and declines at higher thresholds, where the sample becomes very small ($N = 726$ at $K = 16$).

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



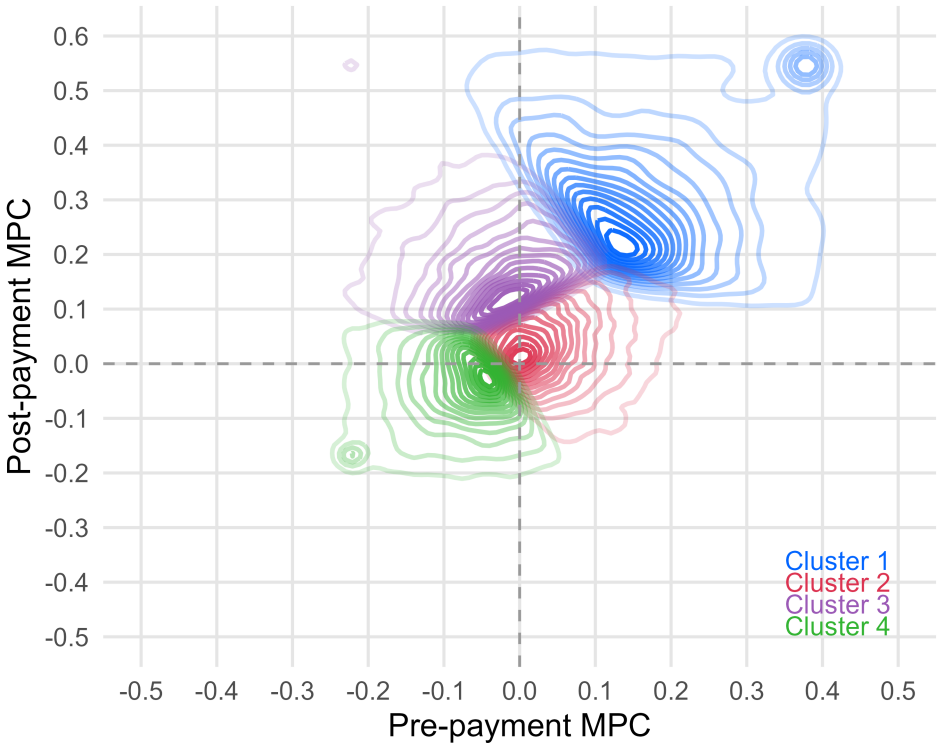
Notes. This figure displays a bin-scatter plot of the one-week within-user average pre-payment MPC against the one-week average post-payment MPC across all events for each user, where each underlying observation represents a user.

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
Nondurable 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								
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
Paycheck	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	09/15/2017	10/3/2017	9/27/2017	49	28
	10/04/2018	05/14/2018	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 on 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	Main Estimand	Result	Replication
Panel A: Fiscal transfers				
COVID Stimulus 1	Baker et al. (2023)	1-month MPX	0.22	0.22
Child Tax Credit	Schild et al. (2023)	3-month MPC	0.41	0.37
Alaska PF Dividend	Kueng (2018)	1-month MPC	0.11	0.10
Tax Refund	Graham et al. (2025)	1-month MPX	0.42	0.37
Panel B: Salary receipts				
Paycheck Income	Gelman et al. (2014)	1-day MPX	0.07	0.06

Notes. This table presents the replication results of existing papers. The column 'Main Estimand' reports the 1-month MPX in Baker et al. (2023), Table 3 (estimation period 2020.01–2020.08); the 3-month MPC in Schild et al. (2023), Table 10 (2019.01–2022.03); the 1-month MPC in Kueng (2018), Figure 3 (2010–2014); the 1-month MPX in Graham and McDowall (2025), Figure 2 (2014–2015); and the 1-day MPX in Gelman et al. (2014), Table S3 (2012–2013). Alaska PF Dividend is the annual transfer to Alaskan residents from the Alaska Permanent Fund. 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	1-week	1-month	3-month	1-week
Panel A: One-off events				
COVID Stimulus 1	0.0556***	0.1499***	0.2334***	0.0107***
	(0.0019)	(0.0046)	(0.0099)	(0.0015)
COVID Stimulus 2	0.0816***	0.1168***		0.0133***
	(0.0030)	(0.0068)		(0.0027)
COVID Stimulus 3	0.0524***	0.1406***	0.1575***	0.0117***
	(0.0015)	(0.0035)	(0.0069)	(0.0012)
Panel B: One-time monthly payments for 6 months				
Child Tax Credit	0.1258***	–	–	0.0313***
	(0.0064)			(0.0062)
Panel C: Once every year				
Alaska PF Dividend	0.0799***	0.2116***	0.3795***	0.0224***
	(0.0031)	(0.0075)	(0.0180)	(0.0019)
Tax Refund	0.0317***	0.0894***	0.1491***	0.0054***
	(0.0018)	(0.0051)	(0.0090)	(0.0005)
Panel D: Periodically				
Monthly Paycheck	0.0769***	–	–	0.0214***
	(0.0010)			(0.0006)

Notes. This table presents the responses of spending on nondurables and services to predictable income changes around the payments dates. 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.3. 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	0.0703***	0.0487***		
	(0.0123)	(0.0158)		
Lagged Pre-MPC × High News		0.0708		
		(0.0955)		
Lagged Pre-MPC × Repeat Event		0.0518**		
		(0.0251)		
Lagged Post-MPC			0.349***	0.257***
			(0.0108)	(0.0139)
Lagged Post-MPC × High News				0.247***
				(0.0840)
Lagged Post-MPC × Repeat Event				0.211***
				(0.0206)
Observations	313,285	313,285	313,285	313,285
R^2	0.002	0.003	0.016	0.017
Date and Event 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: Event Salience and Pre- and Post-Payment MPCs

	Post-MPC, 1m	Post-MPC, 1w	Pre-MPC, 1w	Post MPC - Pre MPC, 1w
	(1)	(2)	(3)	(4)
State News Coverage	2.157***	-0.271**	0.377***	-0.649***
	(0.253)	(0.106)	(0.099)	(0.130)
Observations	344,308	359,209	359,209	359,209
R^2	0.190	0.184	0.147	0.161
User FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Financial Controls	YES	YES	YES	YES

Notes. This table presents regressions of pre- and post-MPCs on 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 7: Liquidity and Pre- and Post-Payment MPCs (in %)

	Post-MPC, 1m		Post-MPC, 1w		Pre-MPC, 1w	
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Bank Balance)	-0.862***	0.0325	-0.322***	0.00474	0.0146	0.216***
	(0.0873)	(0.219)	(0.0332)	(0.0851)	(0.0295)	(0.0759)
ln(Available Credit)	-1.201***	0.372**	-0.382***	0.177**	-0.148***	0.180***
	(0.0812)	(0.189)	(0.0311)	(0.0717)	(0.0272)	(0.0643)
Num Overdrafts	0.249***	0.0816***	0.104***	0.0420***	0.0102**	-0.00384
	(0.0189)	(0.0207)	(0.00637)	(0.00767)	(0.00512)	(0.00731)
Observations	345,005	345,005	345,005	345,005	345,005	345,005
R^2	0.018	0.194	0.020	0.202	0.014	0.166
Date and Event FE	YES	YES	YES	YES	YES	YES
User FE		YES		YES		YES

Notes. Dependent variables are the marginal propensity to consume out of income payments in the month after (*Post-Payment MPC*) in columns 1–2, the week after (*Post-Payment MPC*) in columns 3–4 or the week before (*Pre-Payment MPC*) the income's arrival in columns 5–6. Bank balances, available credit (credit card utilization), and overdraft counts 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 8: Individual Characteristics by Consumer Type and by Liquidity

	By Pre- and Post-MPCs					
	Savers:		Spenders:		By Liquidity:	
	Active	Passive	Active	Passive	Low (Q1)	High (Q4)
Panel A: Consumption						
Nondurable Post-MPC, 1-week	0.0528***	0.0329***	0.2624***	0.1490***	0.0860***	0.0472***
	(0.0015)	(0.0024)	(0.0067)	(0.0028)	(0.0038)	(0.0034)
Nondurable Pre-MPC, 1-week	0.0381***	0.0052**	0.1162***	0.0294***	0.0293***	0.0270***
	(0.0013)	(0.0021)	(0.0050)	(0.0017)	(0.0023)	(0.0027)
Total Post-MPX, 1-week	0.1036***	0.0710***	0.4063***	0.2348***	0.1487***	0.0881***
	(0.0028)	(0.0044)	(0.0108)	(0.0046)	(0.0064)	(0.0059)
Total Pre-MPX, 1-week	0.0543***	0.0064*	0.1557***	0.0433***	0.0379***	0.0393***
	(0.0023)	(0.0036)	(0.0079)	(0.0028)	(0.0038)	(0.0052)
Salary Post-MPC	0.0619***	0.0296***	0.1895***	0.1118***	0.0862***	0.0417***
	(0.0012)	(0.0017)	(0.0050)	(0.0023)	(0.0027)	(0.0023)
Salary Pre-MPC	0.0358***	0.0061***	0.0970***	0.0296***	0.0207***	0.0212***
	(0.0009)	(0.0015)	(0.0034)	(0.0015)	(0.0019)	(0.0019)
Average Prop. to Cons. (APC)	0.91	0.91	1.00	0.97	0.98	0.90
Panel B: Liquidity						
Bank Balances	11489	15263	5498	6693	276	33895
Available Credit	14020	15768	10069	11188	2599	25161
Credit Ratio	0.74	0.76	0.59	0.63	0.36	0.88
Number of Overdrafts	1.25	0.99	4.51	2.67	4.83	0.96
Panel C: Income						
Income	81349	87610	65182	72995	61471	103815
Payment-to-Income Ratio	0.05	0.04	0.04	0.05	0.05	0.04
Income Variation	0.04	0.04	0.04	0.04	0.03	0.04
Panel D: Financial Sophistication						
Investment Deposits	3692	4267	2332	2552	2085	5518
Finance Employee	0.07	0.09	0.04	0.04	0.03	0.10
Panel E: Mapping to Models						
Model Class	Canonical Macro/PIH	Near-Rational Inattention	Present Biased	Mental Accounting	Canonical Macro: Constrained	Canonical Macro: Unconstrained
Number of Observations	169,574	82,543	21,911	89,548	37,906	41,389
Number of Users	21,540	11,851	4,538	12,361	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 Type

	Savers:		Spenders:	
	Active	Passive	Active	Passive
Panel A. Difference: Post-MPC − Pre-MPC (in %)				
ln(Bank Balance)	-0.161***	0.0509	-0.830***	-0.130
	(0.0470)	(0.0640)	(0.213)	(0.0832)
ln(Available Credit)	-0.268***	-0.0130	-1.156***	-0.0433
	(0.0428)	(0.0601)	(0.195)	(0.0729)
R^2	0.009	0.002	0.006	0.022
Panel B. Pre-payment MPC (in %)				
ln(Bank Balance)	0.0692**	0.0007	0.347***	-0.0766
	(0.0313)	(0.0427)	(0.133)	(0.0583)
ln(Available Credit)	0.0723**	-0.00215	0.662***	-0.258***
	(0.0286)	(0.0402)	(0.120)	(0.0511)
R^2	0.010	0.025	0.033	0.013
Observations	169,574	82,543	21,911	89,547
Event FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Model Class	Canonical Macro/PIH	Near-Rational Inattention	Present Biased	Mental Accounting

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 Consumer Type

	Savers:		Spenders:	
	Active	Passive	Active	Passive
Panel A. Post-MPC				
ln(Shock Amount)	-0.030***	0.096***	-0.224***	-0.107***
	(0.004)	(0.005)	(0.010)	(0.006)
R^2	0.145	0.167	0.258	0.164
Panel B. Pre-MPC				
State News Coverage	0.486***	0.227	0.409	0.028
	(0.130)	(0.197)	(0.426)	(0.181)
R^2	0.092	0.109	0.149	0.100
Observations	168,055	81,503	21,155	88,496
Event FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Financial Controls	YES	YES	YES	YES
Model Class	Canonical Macro/PIH	Near-Rational Inattention	Present Biased	Mental Accounting

Notes. Panel A presents results of regressions of Post-MPCs on the logged payment amount. 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 Individual MPCs and Consumption Noise

This appendix expands on the within-person design behind Panel (b) of Figure 1 and on the two appendix figures that support it, Figures A.4 and A.5.

The decomposition.— Write the MPC estimated for individual i at event e as $\hat{\beta}_{i,e} = \beta_i + \nu_{i,e}$, where β_i is the individual's persistent MPC and $\nu_{i,e}$ is event-level consumption noise with mean zero, independent across events. Averaging E of that individual's events gives $\text{var}\big(\bar{\beta}_i(E)\big) = \sigma^2_{\text{true}} + \sigma^2_{\text{noise}}/E$, with $\sigma^2_{\text{true}} = \text{var}(\beta_i)$ and $\sigma^2_{\text{noise}} = \text{var}(\nu_{i,e})$. This is an identity in $1/E$ rather than a testable restriction. Hence, the regression of $\text{sd}(E)^2$ on $1/E$ fits almost perfectly by construction, with R^2 above 0.999 in our data, and we report σ_{true} with a bootstrap confidence interval instead of a goodness-of-fit statistic. What the data do test is whether the intercept is zero: $\sigma_{\text{true}} = 0$ is the homogeneous-MPC null, and the bootstrap rejects it decisively – the 95% CI is $[16.4, 18.9]$.

The role of the threshold K.— K is not a tuned parameter. It defines the population, since the sample consists of individuals observed in at least K income events (i.e., individuals with $E_i \geqslant K$) and at most $E = K$ events are averaged within that fixed sample. Raising K therefore trades a longer averaging path against a smaller and more selected population. Figure A.4 maps this trade-off. Panel (a) shows that the qualitative pattern − dispersion declining but flattening above the homogeneous benchmark − holds at every K, with curves shifted down for larger K. Panel (b) shows that event noise dominates the level of single-event dispersion at every K, with σ_{noise} between 70% and 81% against a σ_{true} of 10–19%. Panel (c) shows that $\sigma_{\text{true}}(K)$ lies between 17.4% and 19.0% for $K = 4, \ldots, 11$ and declines at higher thresholds, where the sample thins to 726 individuals at $K = 16$. Panel (d) shows the reason: as K rises, the sample population shifts sharply toward the Alaska Permanent Fund Dividend (APFD), from 4% of events at $K = 4$ to 38% at $K = 16$. The APFD events are fixed-date, known-amount annual transfers whose single-event MPCs are therefore cleaner to measure. We set $K = 10$ for the main figure. It sits inside the plateau of $\sigma_{\text{true}}(K)$, and its sample of 13,490 individuals is about three times the sample at $K = 13$, which tightens the confidence interval.

The implied share of persistently negative MPCs.— Figure A.5 applies the same design to the share of consumers whose within-person average post-MPC is negative. As $E \to \infty$, the heterogeneous null approaches $\Phi\big(-\mu_{\text{pop}}/\sigma_{\text{true}}\big) = 19.0\%$ at $\mu_{\text{pop}} = 15.6\%$ and $\sigma_{\text{true}} = 17.8\%$,

an estimate of the share of consumers who hold a persistently negative one-month post-MPC. Two caveats limit its interpretation. First, the null applies a normal approximation, and it over-predicts the observed share at every E, by 0.7 percentage points at $E = 1$ and 1.2 points at $E = 10$. Since the $E = 1$ gap involves no averaging and no fitting, it isolates the approximation error: the MPC distribution is right-skewed (see Figure 1(a)) and places less mass far below the mean than a normal with the same first two moments. Hence, 19.0% is likely an upper bound. Second, σ_{true} mixes ex-ante heterogeneity of consumer types with ex-post heterogeneity from temporary circumstances, so not all of this share need reflect a stable trait.

Comparison with the literature.— Section 4.2 compares our estimates to the literature. Here we record the design differences behind those comparisons, which are large enough that the numbers should be read as indicative rather than as directly comparable. Horizons differ (one month here, a quarter in Misra and Surico 2014 and Lewis et al. 2026, four weeks in Boehm et al. 2025, three months in Fuster et al. 2021, and unspecified in Jappelli and Pistaferri 2014), as do the spending measures (nondurables here, total expenditure including durables in the CEX studies, card transactions in the French experiment, and self-reported spending or donations in the surveys). Durables are both larger and more dispersed, so estimates that include them are not comparable to ours in level or in spread.

Two features separate the designs more fundamentally than horizon or category. The first is whether the response is *revealed* or *stated*. Studies that infer the response from observed spending inherit the full event-level consumption noise that we measure, whereas surveys that ask directly do not, because the respondent reports the causal response rather than leaving it to be extracted from spending that moves for many other reasons. That is not the same as being free of measurement error: stated responses remain subject to inattention, rounding, and heaping, and Fuster et al. (2021) report that 74% of respondents state an MPC of exactly zero. The second is whether the design can produce a negative MPC at all. Jappelli and Pistaferri (2014) cannot, since the interviewer enforces that the stated spend and save shares sum to 100; Fuster et al. (2021) can, and do not restrict answers to $[0, 1]$; Boehm et al. (2025) can in principle but constrain the support to be positive in their baseline, motivated by the similarity of the treated and control left tails, while conceding that the constraint is unnecessary and reporting the unconstrained estimate in their Figure A25; and Lewis et al. (2026) impose no constraint, yet cannot deliver a negative household MPC once every estimated group MPC is positive, because each household's MPC is a posterior-weighted average of them.

The one design in the literature that resembles ours is Jappelli and Pistaferri (2020), who ask the same households the same hypothetical question in two waves six years apart. Their Figure 3 plots the within-household change in the reported MPC. It has a mode at zero, but the mass is spread across the whole $[-1, 1]$ range, with secondary spikes at ± 0.5 that are simply the differences between the focal answers 0, 0.5, and 1. The authors model this variation as classical measurement error and difference it away in order to identify the sensitivity of the MPC to cash on hand. Hence, even stated MPCs are far from stable within a person, which is the same phenomenon that $\sigma_{\text{noise}} = 76.1\%$ measures in our setting, arrived at from a very different

direction. Their design, however, cannot separate the persistent and transitory components: with a single hypothetical windfall per wave, the split between f_i and ν_{it} in their notation is assumed rather than estimated.

A.2 Cluster Assignments

A.2.1 Within-Person Persistence of Cluster Assignments

We formalize this stability using a normalized purity measure that quantifies how consistently each individual is assigned to the same cluster across income events.

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., $E_i \geqslant 5$.

Cluster definitions are estimated once at the user level (k-means on users' mean pre- and post-payment MPCs across events, as in Section 5); each individual-event observation is then assigned to the type whose user-level centroid is nearest in the two-dimensional space of one-week pre- and post-payment MPCs, the same construction used for the transition analysis in Appendix Table A.6. This yields an unbalanced panel of type assignments $g_{i,e} \in \{1, \ldots, k\}$ 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:

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

The minimum possible modal count, $\lceil E_i/k \rceil$, is not the expected modal count under random assignment: the expected maximum of multinomial counts strictly exceeds this floor in small samples, and unequal cluster shares raise it further. With the largest type containing 43% of consumers, expected purity under random assignment is 0.461 at $k = 4$, far above the $\lceil E_i/k \rceil / E_i$ benchmark. We therefore benchmark observed purity against an empirical permutation null: we randomly reassign the pool of event-level type labels across all events—preserving the aggregate assignment distribution, including unequal type shares, and each individual's event count E_i—and recompute mean purity, averaging over 200 draws. Our headline persistence statistic is the chance-corrected normalized purity applied to the sample mean,

$$\text{purity}^* \;=\; \frac{\overline{\text{purity}} \;-\; \mathbb{E}\left[\overline{\text{purity}} \mid \text{random}\right]}{1 \;-\; \mathbb{E}\left[\overline{\text{purity}} \mid \text{random}\right]}, \tag{A.2}$$

which equals zero in expectation under random assignment and one under perfect persistence. The permutation distribution also delivers a standard error for observed purity under the null and an exact p-value with minimum $1/201$.

A.2.2 Event-to-Event Persistence of Type Assignments

Our second persistence test examines the structure of type assignments across consecutive income events within person. Event-level assignments follow the same construction as Appendix A.2.1: each individual-event observation is assigned to the consumer type whose user-level cluster centroid is nearest in the two-dimensional space of one-week pre- and post-payment MPCs. For each individual, we then pair every event with that individual's next observed event in calendar time, yielding 289,591 within-person event-to-event transitions.

Let π_{ij} denote the probability of assignment to type j at the next event conditional on assignment to type i at the current event, and let π_j denote the unconditional probability of assignment to type j at the next event. Appendix Table A.6 reports the normalized transition probabilities π_{ij}/π_j. A value of one indicates that the current assignment carries no information about the next; values above (below) one indicate that the transition occurs more (less) often than under independence. This normalization is essential because the marginal type shares are unequal: raw conditional probabilities would show large same-type shares for the biggest types purely mechanically.

We summarize aggregate persistence with Cohen's $\kappa = (p_o - p_e)/(1 - p_e)$, where p_o is the observed share of transitions that remain in the same type and $p_e = \sum_i \Pr(\text{current} = i) \Pr(\text{next} = i)$ is the same-type share expected from the marginal distributions alone. κ equals zero when assignments are serially independent given the marginals and one under perfect persistence.

Inference uses the same permutation null as the purity statistic in Appendix A.2.1: we randomly reassign the pool of event-level type labels across all events—preserving the aggregate assignment distribution and each individual's event count and event ordering—and recompute the transition matrix and κ for each draw. Because the reshuffle preserves the marginals, the null distribution of κ is centered at zero and that of each normalized cell at one; the draws deliver a standard error under the null and an exact p-value with minimum $1/(\text{draws} + 1)$.

In the data, $\kappa = 0.060$, far outside the permutation null (standard error ≈ 0.0011, $p < 0.01$). The cell-level structure of the deviations is discussed in Section 5.3 and reported in Table A.6.

A.2.3 Out-of-Sample Validation: Type Assignment from Salary Responses

The persistence statistics in Section 5.3 and Appendix A.2.1 are computed on the same events that define the types. As an out-of-sample check, we ask whether a user's spending response around regular paycheck receipt—estimated pre- and post-payment MPCs from salary events—assigns the user to their own type. To make the test valid, we re-estimate the user-level k-means clusters ($k = 4$) on users' mean pre- and post-payment MPCs computed *excluding* the salary observations, and assign each user's salary response to the nearest resulting centroid. This leave-out step matters: were salary responses included in the user means, a user's own salary numbers would mechanically pull their mean—and hence their cluster—toward themselves, overstating the test's success. Because the salary numbers play no role in either the user means or the centroids they are scored against, agreement between the salary-based assignment and the user's cluster reflects genuine informational content of a single income event about the user's type.

The salary response assigns 40.3% of users to their own type ($N = 40{,}109$). This exceeds the 25% implied by uniform random assignment, but the relevant benchmark must account for the uneven type distribution: given the marginal shares of types and of salary-based assignments, random pairing alone would produce 35.2% agreement. Against that benchmark, Cohen's κ is 0.078, and the permutation null is tightly concentrated at zero (standard error 0.003, exact $p < 0.01$ based on 200 draws). Cohen's κ corrects for the marginal shares in the aggregate, but it can still be disproportionately influenced by the largest type. We therefore also compute the balanced accuracy of the salary-based assignment,

$$\text{BA} \;=\; \frac{1}{k}\sum_{j=1}^{k} \Pr\big(\text{salary-based assignment} = j \,\big|\, \text{user's type} = j\big), \tag{A.3}$$

the mean per-type recall, which weights each type equally regardless of its population share. In the data, BA = 31.2%. Its benchmark under random assignment is $1/k = 25\%$: under random pairing, each type's expected recall equals the marginal share of salary-based assignments to that type, and these shares average to $1/k$ by construction. The salary response thus recovers types at above-chance rates even when all four types are weighted equally, confirming that the out-of-sample agreement in the headline 40.3% figure is not an artifact of the large Active Saver type.The magnitude of this out-of-sample signal matches the event-to-event transition persistence ($\kappa = 0.060$, Table A.6) and the excess of within-person purity over its chance benchmark (Table A.5). All three statistics deliver the same conclusion independently: any single income event carries genuine but heavily noise-attenuated information about a consumer's type, which is precisely the pattern implied by stable behavioral traits observed through large idiosyncratic consumption noise (Section 4.2), and why the classification is estimated from within-person averages across many events.

A.2.4 Varying the Number of Clusters

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 discuss the robustness of our choice of k by showing the effect of using different numbers of clusters.

Figure A.6 plots the cluster classifications for $k = 2$, 3, 5, and 6. 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 of 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

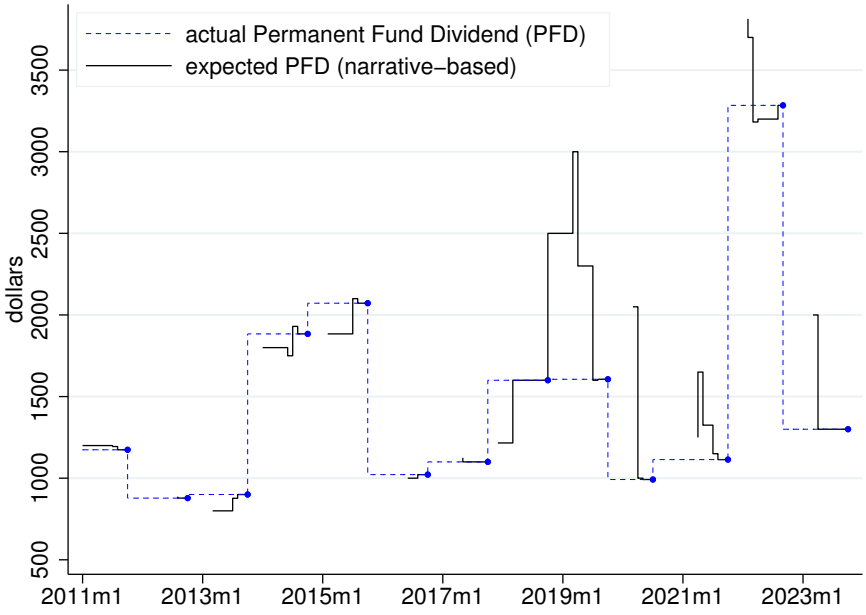
similar clusters into more extreme splits.

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 further validation of this choice is that the resulting cluster assignments exhibit systematic above-chance within-person stability across events—a pattern we would not expect if the algorithm were simply partitioning noise.

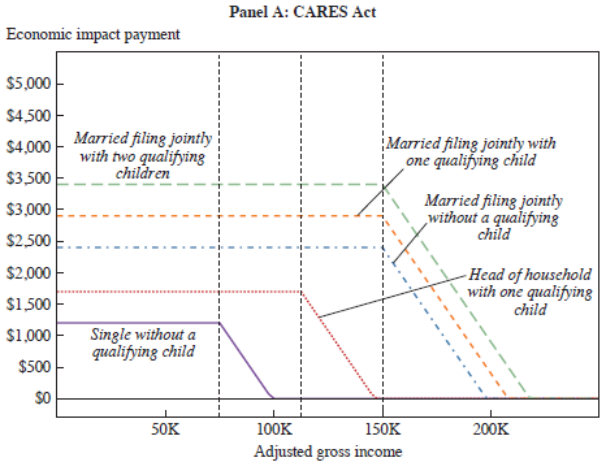
B Appendix Figures and Tables

Figure A.1: Alaska Permanent Fund Dividend News

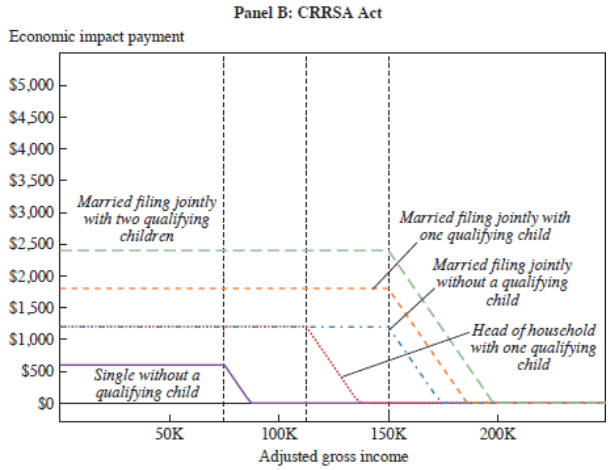


Notes. This figure shows expected amounts of the next Alaska Permanent Fund Dividend (black line) compared to the realized dividend (blue dots). The dividend expectation series is based on a narrative analysis of Alaskan newspapers and extends the series in Kueng (2018) to 2023.

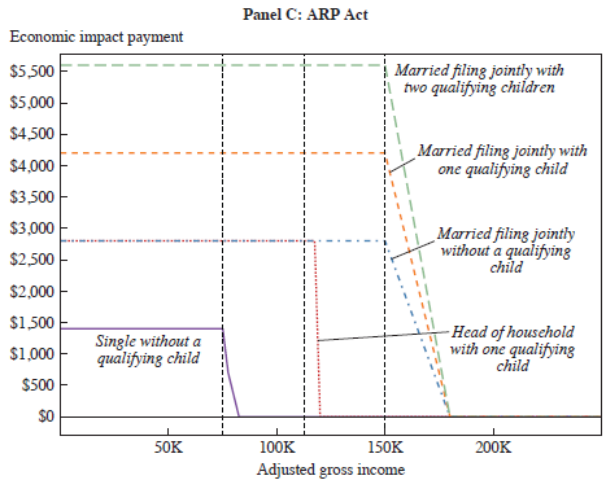
Figure A.2: Sources of Identifying Variation from the COVID Stimulus Programs



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



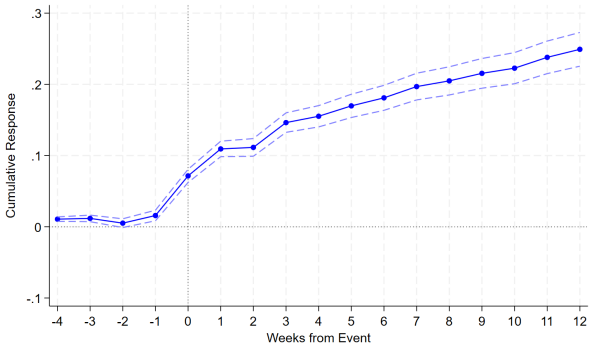
(b) COVID Stimulus 2: CRRSA Act of Dec. 29, 2020



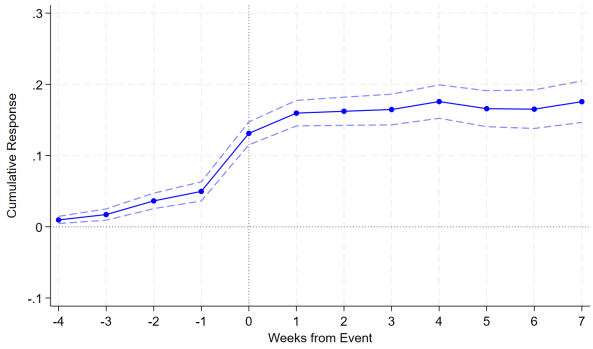
(c) COVID Stimulus 3: ARP Act of March 11, 2021

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.

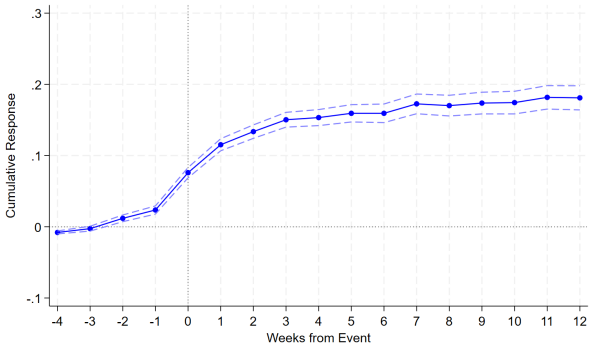
Figure A.3: Nondurable Spending Responses Relative to Payment Date



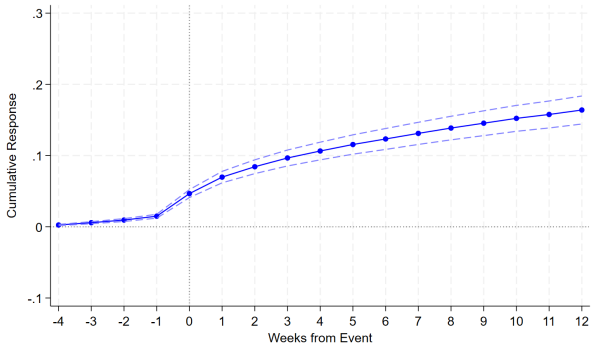
(a) COVID Stimulus 1: 4/11/2020



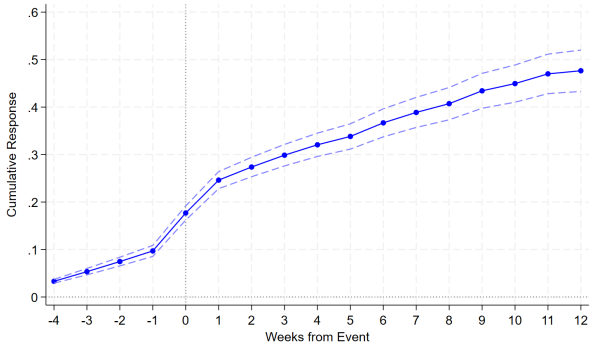
(b) COVID Stimulus 2: 12/29/2020



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



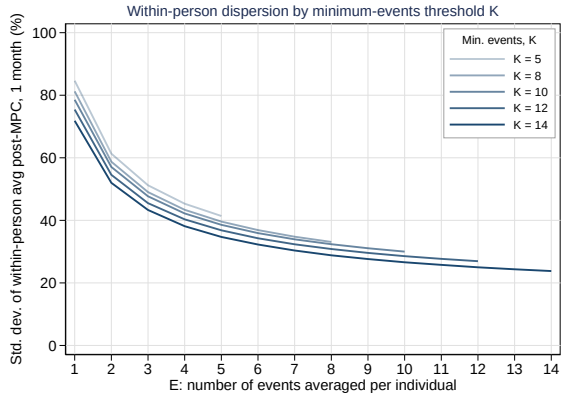
(d) Tax Refund: 2012-2023



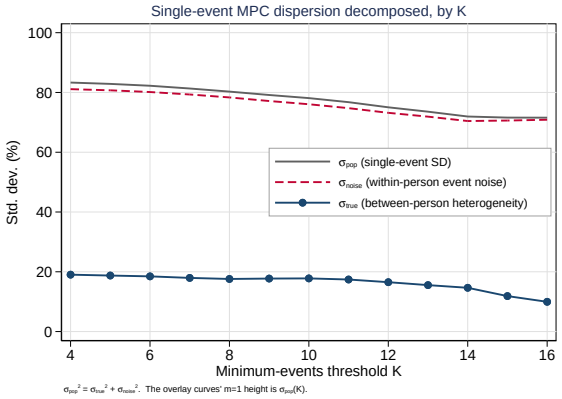
(e) Alaska Permanent Fund Dividend: 2012–2023

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.

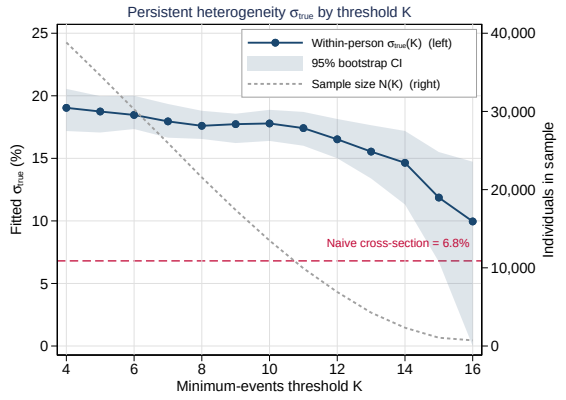
Figure A.4: Robustness of Within-Person MPC Dispersion to the Minimum-Events Threshold K



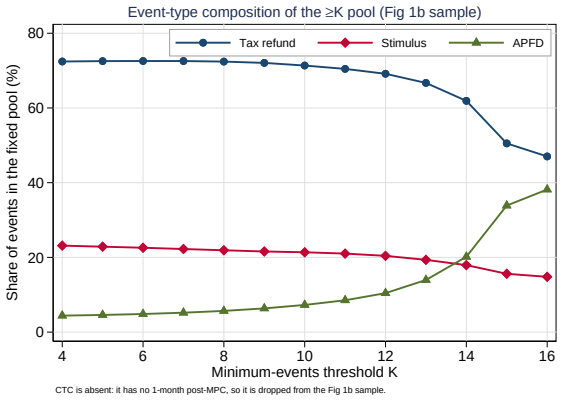
(a) Dispersion curves for several K



(b) Variance decomposition by K



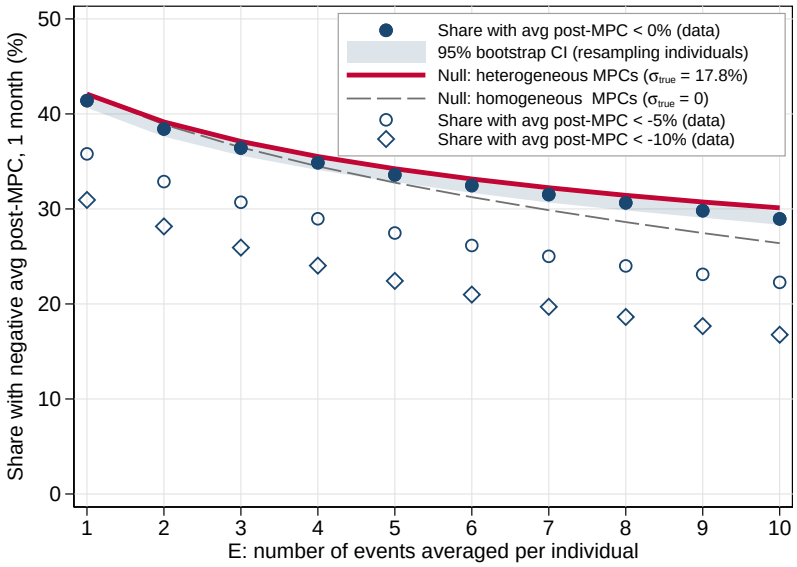
(c) σ_{true} and sample size by K



(d) Event-type mix of the sample by K

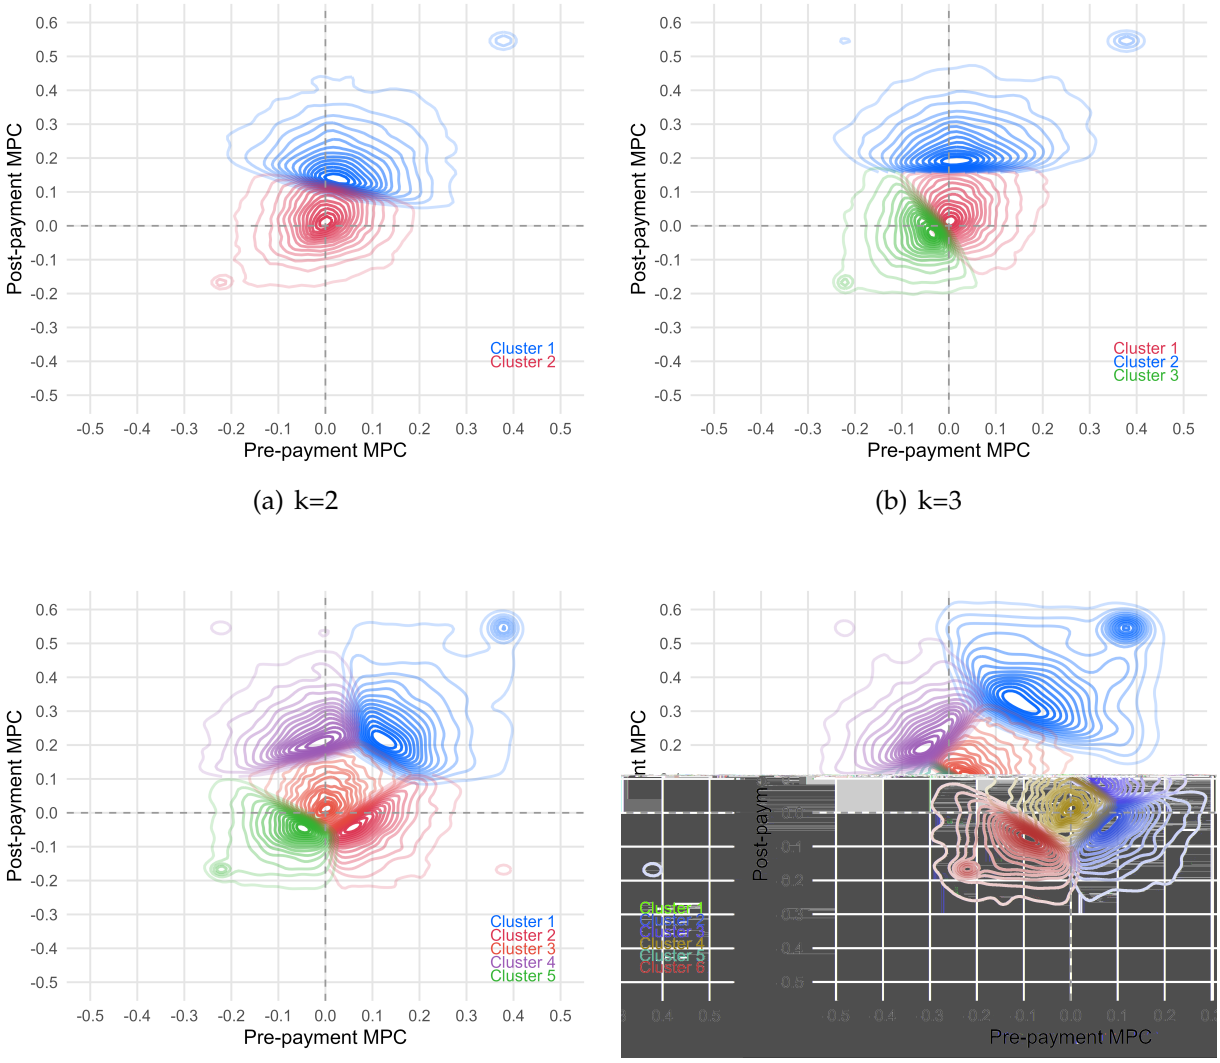
Notes. All four panels probe how the within-person dispersion of the average post-payment MPC (Panel (b) of Figure 1) depends on the minimum-events threshold K, which defines the K-threshold sample (individuals observed in at least K income events; CTC events are excluded throughout, as they have no one-month post-MPC). Panel (a) overlays the dispersion curves for several thresholds: each fixes its own K-threshold sample and plots the standard deviation of the E-event within-person average against E. The pattern − dispersion declining but flattening above the homogeneous-noise benchmark − is robust across K; curves start lower for larger K because individuals with more frequent events have less noisy single-event MPCs. Panel (b) decomposes the single-event dispersion $\sigma_{\text{pop}}(K)$ (gray solid line) into the between-person component $\sigma_{\text{true}}(K)$ (blue solid line with circle markers) and the within-person event-noise component $\sigma_{\text{noise}}(K)$ (red dashed line), with $\sigma^2_{\text{pop}} = \sigma^2_{\text{true}} + \sigma^2_{\text{noise}}$. This shows that event noise dominates the level and drives the downward shift of the curves in Panel (a). Panel (c) plots the fitted asymptote $\sigma_{\text{true}}(K)$ (blue solid line with circle markers, left axis) with its 95% individual-bootstrap band and the sample size $N(K)$ (gray dashed line, right axis). The horizontal red dashed line is the naive cross-sectional estimate (6.8%) for reference: it imposes no threshold, but instead assigns all 50,135 individuals to bins by their own event count $E_i = 1, \ldots, 22$ (dropping bins with fewer than 20 individuals) and fits $\text{sd}(E_i)^2 = \sigma^2_{\text{true}} + \sigma^2_{\text{noise}}/E_i$ by WLS weighted by bin size. Because different individuals populate each bin, it confounds within-person averaging over additional events with changes in the composition of the samples. $\sigma_{\text{true}}(K)$ lies between 17.4% and 19.0% for $K = 4, \ldots, 11$ and exceeds the naive estimate at every K, by a factor of about 2.6 across that range, declining only among the very heaviest-event individuals, where the sample also thins. Panel (d) shows the event-type composition of the K-threshold sample: as K rises, the sample shifts sharply toward the Alaska Permanent Fund Dividend − a fixed-date, known-amount annual transfer whose single-event MPC is therefore cleaner to measure − from 4% of events at $K = 4$ to 38% at $K = 16$, which explains the fall in σ_{noise} (Panel (b)) and, in part, in σ_{true} at high K.

Figure A.5: Decline of the Share of Negative Within-Person Average MPCs, by Number of Events



Notes. This figure is the composition-free (within-person) counterpart of Panel (b) of Figure 1. The population is fixed to individuals observed in at least $K = 10$ income events ($N = 13{,}490$; CTC events are excluded, as they have no one-month post-MPC). For each $E = 1, \ldots, 10$ we draw E events at random from each individual, compute the within-person average one-month post-payment MPC, and plot the share of individuals whose average falls below three thresholds $-$ 0 (solid dots), -5% (hollow circles), and -10% (hollow diamonds) $-$ averaged over 200 Monte Carlo draws; the shaded band is the bootstrapped 95% CI on the headline (< 0) share, resampling individuals. Because the *same* individuals enter at every E, the decline reflects pure averaging, not composition. Both nulls apply the central limit theorem to the within-person average, whose mean is $\mu_{\text{pop}} = 15.6\%$, and predict a negative share of $\Phi\big(-\mu_{\text{pop}}/\text{sd}(E)\big)$, with $\text{sd}(E)$ taken from the corresponding null in Panel (b) of Figure 1. Under the homogeneous null (dashed line), $\text{sd}(E) = \sigma_{\text{pop}}/\sqrt{E}$, where $\sigma_{\text{pop}} = 78.1\%$ is the standard deviation of individual event-level post-MPCs in the same fixed sample. Under the heterogeneous null (red solid line), $\text{sd}(E)^2 = \sigma^2_{\text{true}} + \sigma^2_{\text{noise}}/E$, evaluated at the $\sigma_{\text{true}} = 17.8\%$ and $\sigma_{\text{noise}} = 76.1\%$ estimated in Panel (b) of Figure 1. Importantly, these two parameters are not re-fit here, which makes this figure an out-of-sample check on them rather than a fit. The homogeneous null is the special case $\sigma_{\text{true}} = 0$, such that both nulls coincide at $E = 1$ by construction. The heterogeneous null lies slightly above the data at every E, by 0.7 percentage points at $E = 1$ and by 1.2 points at $E = 10$. This gap reflects the normal approximation rather than a mis-calibration: at $E = 1$ nothing is averaged, so the 0.7-point gap measures only the failure of the normal approximation to the right-skewed distribution of Panel (a) of Figure 1, which places less mass far below the mean than a normal with the same first two moments. Appendix A.1 discusses the implied share of persistently negative MPCs and compares it to the literature.

Figure A.6: 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.

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: 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	205,747,911	0.0376***	0.1046***	0.1737***	0.0077***
			(0.0015)	(0.0044)	(0.0078)	(0.0004)
Panel B: One-off events						
COVID Stimulus 1	2020	19,294,172	0.0519***	0.1506***	0.2280***	0.0123***
			(0.0018)	(0.0043)	(0.0094)	(0.0012)
COVID Stimulus 2	2020–2021	37,379,306	0.0873***	0.1270***	–	0.0183***
			(0.0029)	(0.0066)		(0.0024)
COVID Stimulus 3	2021	18,082,437	0.0520***	0.1362***	0.1459***	0.0128***
			(0.0014)	(0.0034)	(0.0067)	(0.0011)
Panel C: One-time monthly payments for 6 months						
Child Tax Credit	2021–2022	33,251,181	0.1204***	–	–	0.0364***
			(0.0063)			(0.0063)
Panel D: Once every year						
Alaska PF Dividend	2012–2023	205,747,911	0.0780***	0.2037***	0.3577***	0.0220***
			(0.0025)	(0.0061)	(0.0149)	(0.0015)
Tax Refund	2012–2023	205,747,911	0.0316***	0.0876***	0.1450***	0.0050***
			(0.0013)	(0.0040)	(0.0070)	(0.0004)
Panel E: Periodically						
Monthly Paycheck	2012–2023	95,695,521	0.0732***	–	–	0.0323***
			(0.0010)			(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.3. 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.3. 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 A.3: Responses to Income Changes Around the Payment Date: Total Spending (MPX)

			Post-MPX			Pre-MPX
Event	Period	Observations	1-week	1-month	3-month	1-week
Panel A: All events pooled						
Anticipated Payment	2012–2023	205,747,911	0.0798***	0.2080***	0.3356***	0.0167***
			(0.0030)	(0.0081)	(0.0142)	(0.0009)
Panel B: One-off events						
COVID Stimulus 1	2020	19,294,172	0.1098***	0.2857***	0.4316***	0.0248***
			(0.0034)	(0.0077)	(0.0167)	(0.0029)
COVID Stimulus 2	2020–2021	37,379,306	0.1510***	0.2128***	–	0.0386***
			(0.0054)	(0.0114)		(0.0043)
COVID Stimulus 3	2021	18,082,437	0.1030***	0.2342***	0.2603***	0.0227***
			(0.0026)	(0.0057)	(0.0113)	(0.0021)
Panel C: One-time monthly payments for 6 months						
Child Tax Credit	2021–2022	33,251,181	0.1771***	–	–	0.0315***
			(0.0108)			(0.0113)
Panel D: Once every year						
Alaska PF Dividend	2012–2023	205,747,911	0.1316***	0.3001***	0.4516***	0.0249***
			(0.0042)	(0.0097)	(0.0236)	(0.0026)
Tax Refund	2012–2023	205,747,911	0.0708***	0.1858***	0.2971***	0.0139***
			(0.0030)	(0.0078)	(0.0134)	(0.0009)
Panel E: Periodically						
Monthly Paycheck	2012–2023	98,787,012	0.1573***	–	–	0.0651***
			(0.0019)			(0.0014)

Notes. This table reports the same statistics as Table A.2 but for total expenditures (MPXs) instead of spending on nondurables and services (MPCs). See Table A.2 for details.

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

	Dependent variable: Pre-Payment MPC, one-week								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post-MPC, one-week	0.0751***	0.0665***	0.0698***						
	(0.00151)	(0.00167)	(0.00190)						
Post-MPC, one-month				0.0667***	0.0656***	0.0671***			
				(0.000561)	(0.000621)	(0.000706)			
Post-MPC, weeks 2-4							0.0678***	0.0666***	0.0682***
							(0.000560)	(0.000614)	(0.000699)
Observations	348,785	348,785	269,373	348,785	348,785	269,373	348,785	348,785	269,373
R^2	0.008	0.170	0.175	0.041	0.195	0.202	0.042	0.197	0.204
Event FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
User FE		YES	YES		YES	YES	YES	YES	YES
Financial Controls			YES			YES	YES	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 period, over one-month period for increased precision, or over weeks 2-4 following payment (ie. one-month post excluding week 1). 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: Within-Person Purity of Type Assignments by Number of Clusters

Number of Clusters (k)	Raw Purity	Chance Benchmark	Normalized Purity
2	0.726	0.694	0.106
3	0.549	0.513	0.074
4	0.492	0.461	0.059
5	0.449	0.409	0.067
6	0.428	0.385	0.069
7	0.406	0.361	0.070
8	0.367	0.330	0.055
9	0.357	0.320	0.054
10	0.338	0.304	0.049

Notes. Raw purity is the mean across individuals (with at least five events) of the share of events assigned to the individual's modal type, where each event is assigned to the nearest user-level cluster centroid. The chance benchmark is the expected raw purity under random reassignment of event labels, holding fixed the aggregate assignment distribution and each individual's event count (permutation null, 200 draws; see Appendix A.2.1). Normalized purity is equation (A.2). Observed purity exceeds the permutation null at every k ($p < 0.01$).

Table A.6: Persistence of Consumer-Type Assignments Across Events: Normalized Transition Probabilities

	Savers:		Spenders:	
Type at event t	Active	Passive	Active	Passive
Active Saver	**1.14**	0.98	0.81	0.92
Passive Saver	0.96	**1.12**	0.93	0.92
Active Spender	0.85	0.90	**1.40**	1.15
Passive Spender	0.93	0.91	1.15	**1.16**

Notes. This table reports normalized transition probabilities of consumer-type assignments between consecutive income events within person. Each individual-event observation is assigned to the consumer type whose user-level cluster centroid (in the two-dimensional space of one-week pre- and post-payment MPCs) is nearest to the event-level pre- and post-MPC estimates. For each individual, we pair each event with that individual's next observed event, yielding 289,591 event-to-event transitions. Rows denote the type assigned at event t and columns the type assigned at event $t+1$. The normalized transition probability is the conditional transition probability divided by the unconditional probability of the destination type, Pr(type at $t{+}1 = j$ | type at $t = i$) / Pr(type at $t{+}1 = j$). A value of one indicates that the current assignment carries no information about the next assignment; values above (below) one indicate that the transition occurs more (less) often than under independence. Diagonal elements are in bold. Aggregate persistence across all cells: Cohen's $\kappa = 0.060$, i.e., the observed same-type share of transitions exceeds the share expected from the marginal type distributions alone by 6% of the available range (permutation $p < 0.01$).