NBER WORKING PAPER SERIES

LATENT HETEROGENEITY IN THE MARGINAL PROPENSITY TO CONSUME

Daniel Lewis
Davide Melcangi
Laura Pilossoph

Working Paper 32523
http://www.nber.org/papers/w32523

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2024

The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Daniel Lewis, Davide Melcangi, and Laura Pilossoph. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Latent Heterogeneity in the Marginal Propensity to Consume
Daniel Lewis, Davide Melcangi, and Laura Pilossoph
NBER Working Paper No. 32523
May 2024
JEL No. D12,E21,E62

ABSTRACT

We estimate the unconditional distribution of the marginal propensity to consume (MPC) using clustering regression applied to the 2008 economic stimulus payments. By deviating from the standard approach of estimating MPC heterogeneity using interactions with observables, we can recover the full distribution of MPCs. We find households spent between 4 and 133% of the rebate within a quarter, and individual households used rebates for different goods. While many observable characteristics correlate individually with our estimated MPCs, these relationships disappear when tested jointly, except for income and the average propensity to consume. Household observable characteristics explain only 8% of MPC variation, highlighting the role of latent heterogeneity.

Daniel Lewis
Department of Economics
University College London
Drayton House
30 Gordon Street
London
WC1H 0AX
United Kingdom
daniel.lewis920@gmail.com

Davide Melcangi
Federal Reserve Bank of New York
33 Liberty Street
New York, NY 10045
Davide.Melcangi@ny.frb.org

Laura Pilossoph
Department of Economics
Duke University
Social Sciences Building
Durham, NC 27708
and NBER
pilossoph@gmail.com

1 Introduction

Recent research has highlighted the importance of heterogeneity in the marginal propensity to consume (MPC) out of transitory income shocks for fiscal policy and the transmission of monetary policy.[1] Despite the importance of the *distribution* of MPCs, estimates of this object have been largely elusive. Even with plausibly identified transitory income shocks, estimating individual-level MPCs requires panel data with long horizons, which is typically not available; when available, it also usually requires the unappealing assumption that an individual's marginal propensity to consume is time invariant.[2] The existing literature, therefore, has followed one of two avenues: estimating a fully structural model and simulating a distribution of MPCs, or grouping observations by some presupposed observable characteristics and estimating group-specific MPCs out of transitory income shocks thereafter.[3] However, because both of these approaches require taking a stance on the source of MPC heterogeneity, they may fail to uncover the true degree of heterogeneity, miss other relevant dimensions of heterogeneity that predict an individual's MPC, or both.

In this paper, we estimate the distribution of the MPC directly. We adopt a Gaussian mixture linear regression (GMLR) (e.g., Quandt (1972)), which jointly (i) groups households together that have similar latent consumption responses to the 2008 tax rebate and (ii) provides estimates of the MPCs within these groups. Specifically, the algorithm takes a standard regression of consumption changes on the tax rebate received and basic controls originally studied in Johnson, Parker, and Souleles (2006) and Parker, Souleles, Johnson, and McClelland (2013) (updated according to recommendations from Borusyak et al. (2024)), but allows the coefficient on the rebate to be heterogeneous across unknown groups; the groups as well as their rebate coefficients are then jointly estimated. The group structure is parametrically identified from a Gaussianity assumption on the error terms in the household's consumption equation, which allows the MPCs to be consistently estimated.

This approach offers four advantages over existing efforts to recover the distribution of MPCs. First, it allows us to estimate the full unconditional distribution of MPCs, which

[1]The MPC distribution is a crucial object in Heterogeneous Agent New Keynesian (HANK) models of monetary policy (see e.g., Kaplan, Moll, and Violante (2018) and Auclert (2019)).

[2]Nearly all theories of MPC heterogeneity have some form of state dependence; for example, in Carroll (1992) the MPC is a declining function of gross household wealth.

[3]For the former, see for instance Kaplan and Violante (2014) and Carroll, Slacalek, Tokuoka, and White (2017). For the latter, Fagereng, Holm, and Natvik (2016) exploit randomized lottery winnings to identify transitory income shocks, and subsequently group observations on observables to estimate group-level MPCs. See also Johnson et al. (2006), Blundell, Pistaferri, and Preston (2008), Parker et al. (2013), Kaplan, Violante, and Weidner (2014), and Crawley and Kuchler (2018).

may be driven both by latent factors and observable characteristics, broadly defined; understanding the range of such a distribution sheds light on whether there is potential value, in principle, in attempting to target fiscal transfers to households more likely to spend the funds. Standard methods that rely on sample splitting by observable characteristics can only ever recover the extent to which MPCs co-vary with chosen household characteristics (as opposed to a marginal distribution) and cannot recover heterogeneity in MPCs associated with latent factors (by definition) or different observables. Second, because our approach does not require taking an *ex ante* stance on what observables correlate with MPC heterogeneity, we can "let the data speak" by investigating *ex post* which observables predict the recovered individual MPCs. Third, we find a number of statistically significant relationships between MPCs and observables that sample-splitting methodologies using the same data failed to uncover, suggesting that our approach may have a statistical power advantage over sample-splitting. Finally, by estimating household-level MPCs we are able to project them on various explanatory variables *jointly.* By the same token, we can quantify the share of MPC variation explained jointly by observables.

Our estimation strategy hinges on the fact that clustering algorithms like the one we adopt assign individuals to groups not based on observable characteristics, but based on how well each set of estimated group-specific MPCs describe the observed consumption patterns within the group. This feature allows us to bypass the *ex ante* decision of which observables matter for MPC heterogeneity, and instead estimate the heterogeneity directly. GMLR specifies a linear regression model with different regression parameters for each group or "cluster". It is a probabilistic clustering approach, in which individuals are not assigned to groups in a binary fashion, but instead have posterior weights derived from a Gaussian distribution of regression errors. Conditional on these weights, GMLR simply represents a weighted least squares (WLS) regression. When the panel dimension present in "hard clustering" approaches like that in Bonhomme and Manresa (2015) or Bonhomme et al. (2019) is absent, as is the case in our empirical setting, it is unrealistic to think that group assignment can be determined binarily in the presence of noise, so such continuous posterior weights are required to represent the level of uncertainty that exists in the assignment.[4] Despite this uncertain assignment, GMLR provides estimates for MPCs and unconditional group membership probabilities that are consistent and asymptotically normal as the sample size grows (e.g., Desarbo and Cron (1988)).

Applying our estimator to study the MPC distribution using the 2008 Economic Stimulus Act, we uncover a substantial degree of heterogeneity. In particular, households

[4] We describe in later sections why a panel structure is neither possible nor desirable in our application.

spend at least 4% of the rebate within one quarter, with some households displaying an MPC above one. Generally, the share of households with a particular MPC declines as the MPC increases. We show that the range of estimated MPCs is larger than what one can recover by interacting the rebate with household observables.

We next estimate the distribution of MPCs for specific spending categories. For nondurable expenditures, we find a lower bound for the MPC of 9 cents on the dollar. On the other hand, spending on durables is dichotomous, with a nontrivial fraction of households not spending any of the rebate on such goods, and a significant share spending the majority of the rebate on these items. Finally, since our approach provides household level, good-specific MPCs, we compute their cross-correlation. We find that households with higher nondurable MPCs also display higher durable MPCs, although the correlation is weak (0.13).

Having characterized the distribution of marginal propensities to consume, we recover its observable drivers. Historically, the literature has found mixed empirical evidence and generally weak relationships between MPCs and observable household characteristics, with the possible exception of liquid wealth.[5] Our results suggest that this may be due to a loss of statistical power when re-estimating the MPC with interactions or sample splits. We find that our estimated MPCs are significantly correlated, individually, with many observable drivers, despite the fact that we use the same dataset and a similar identification strategy that previously delivered insignificant relationships. For example, we find that homeowners have significantly higher MPCs than renters, and households with a mortgage display even greater marginal propensities to consume than outright homeowners. These correlations hold for all expenditure categories that we consider.

Our estimates for household-level MPCs also allow us to study multivariate relationships without further losses of power. We find that only two observables are robust to the inclusion of additional regressors and positively correlate with MPCs: households' income and their APC. The income relationship holds for both salary and non-salary (including transfer, retirement, and business or investment income) components. The existing empirical evidence on this relationship is mixed; in line with our results, Kueng (2018) also finds that high-income households have higher MPCs in Alaska Permanent Fund data, and suggests that this behavior could be explained by theories of near-rationality. Examining how MPCs vary jointly with income and the APC, we uncover three groups of households. "Poor-savers", with low total income and a low APC, have the lowest MPCs. Households with high total income and a low APC, or vice versa, display intermediate

[5] Parker et al. (2013) find statistically insignificant differences by age, income, and liquid wealth. Broda and Parker (2014) and Fagereng et al. (2016) find significant relationships for the latter.

marginal propensities to consume. The greatest MPCs are found among "rich-spenders", who not only have high total income, but also typically spend a large portion of it. This group of households has not received much attention in models of consumption and savings. Since we do not have reliable data on liquidity, one possibility is that this group is made up of wealthy-hand-to-mouth as in Kaplan and Violante (2014); however, our results suggest that those households would also need to have high APCs.

Importantly, our best array of observable predictors is able to explain at most 8% of the variation in estimated MPCs. With the vast majority of heterogeneity unexplained by standard observables, our results suggest that a substantial portion of MPC heterogeneity may be driven by latent household traits. For example, heterogeneity in discount rates and/or intertemporal elasticities of substitution (as in, e.g., Aguiar, Boar, and Bils (2019)) would deliver heterogeneity in MPCs, and is further supported by the aforementioned significance of APCs in predicting MPCs, as APCs can also be a function of the same unobserved traits. This type of unobserved heterogeneity could never be recovered by simply splitting the sample on observable characteristics and estimating within-subsample homogeneous MPCs, as is typically done in the literature.[6]

Our results have several important policy implications. First, we find that the 2008 Economic Stimulus payments increased spending for all households, at least in partial equilibrium. Second, the fact that we uncover considerable MPC heterogeneity suggests that, in principle, aggregate spending could be further increased by targeting fiscal transfers to high MPC households. Our significant correlations suggest that it may be desirable to target relatively higher-income households to maximize the aggregate consumption effects of stimulus checks. This is true whether salary or non-salary income is considered; in the latter case, the higher income group includes those receiving transfers, retirees, and entrepreneurs.[7] Nevertheless, such a strategy may imply a tension between the stimulus and relief/insurance motives of lump-sum transfers.[8] However, since we find that observable characteristics predict little of the variation in MPCs, it is likely that any attempt at targeting will only exploit a limited share of the overall variation in MPCs, given the

[6]This is true unless preference heterogeneity is explicitly elicited in survey questions so that it can be used as an observable control. Using Nielsen panel data, Parker (2017) finds that the MPC out of the tax rebate is indeed strongly correlated with a self-reported measure of impatience.

[7]Stimulus checks were phased out for households whose income was above $150,000 ($75,000 for single filers), implying that higher earners did not receive the rebate. Thus, our findings on the positive correlation of MPCs with total income are limited to households within the income range of stimulus checks recipients.

[8]Shapiro and Slemrod (2009) find that low-income individuals were more likely to use the 2008 rebates to pay off debt. Similar patterns have been observed for the CARES act transfers in 2020, see Armantier et al. (2020). Our analysis focuses on the consumption effects of fiscal transfers but is consistent with a potential distributional trade-off. See Koşar et al. (2023) for a theoretical framework which rationalizes this finding.

information available to fiscal authorities.

We also believe our findings can be used to discipline heterogeneous-agent consumption-savings models. We provide empirical measures for (i) the extent of heterogeneity in MPCs, (ii) their correlations with observables, and (iii) the fraction of their variation explained by observables. Virtually any heterogeneous-agent macro model of consumption can be tested vis-a-vis these three implications. While preference heterogeneity can generate more dispersion in MPCs and limit the share of variation explained by observables, it cannot readily produce MPCs that positive correlate with income. We return to these points when we conclude.

Our paper is related to an extensive literature estimating the marginal propensity to consume out of transitory income shocks, and a smaller, complementary literature examining how it varies across households. In addition to the aforementioned strategies, some papers have used the "reported preference" approach, eliciting MPC heterogeneity directly from responses to survey questions. Recent examples include Sahm, Shapiro, and Slemrod (2010), Jappelli and Pistaferri (2014) and Fuster, Kaplan, and Zafar (2018). Complementary to this approach, we flexibly estimate observed and unobserved MPC heterogeneity instead using realized spending. In this respect, Misra and Surico (2014) is closest in spirit to our work. They estimate a quantile regression of consumption responses to the 2008 tax rebate using the same data, and find substantial heterogeneity. However, quantile regression estimates the role of regressors at specific points in the overall conditional distribution of the dependent variable. In Supplement B, we show how this approach is sensitive to the correlation of MPC heterogeneity with other forms of heterogeneity, since other factors may be quantitatively larger drivers of the conditional distribution of consumption than the tax rebate. Our results also relate to a burgeoning literature that has turned its attention to unobserved household traits and preference heterogeneity. Our findings corroborate the importance of this dimension, recently highlighted by Alan, Browning, and Ejrnaes (2018), Parker (2017), Aguiar et al. (2019), and Gelman (2019).

The paper proceeds as follows. In Section 2 we outline our empirical strategy based on the 2008 tax rebate. In Section 3, we describe in detail our clustering approach, which allows us to identify and estimate our empirical specification. Our results are outlined in Section 4, where we provide estimates of the distribution of MPCs for various consumption categories. Section 4.3 discusses observable characteristics that are correlated with the estimated MPCs. Section 4.4 shows the longer-run consumption responses to stimulus checks. Section 5 concludes.

2 Empirical methodology

In order to estimate the marginal propensity to consume, and how it varies across households, we consider an off-the-shelf, well-identified, quasi-natural experiment: the 2008 Economic Stimulus Act (ESA), as studied originally by Parker et al. (2013). Between April and July of 2008, $100 billion in tax rebates was sent to approximately 130 million US tax filers.[9] Importantly, the timing of rebate receipt was determined by the last two digits of the recipient's Social Security Number (SSN), making the timing of receipt random. As in Parker et al. (2013), we exploit the randomized timing of the rebate receipt, but instead estimate heterogeneous marginal propensities to consume rather than a homogeneous marginal propensity to consume.

Our data come from the 2008 Consumer Expenditure Survey (CEX), which contains comprehensive and detailed measures of household-level consumption expenditures. Since we observe expenditures, rather than consumption, our MPCs are in fact marginal propensities to spend, as is common in the empirical literature and as recently discussed by Laibson, Maxted, and Moll (2022). The 2008 CEX wave also has supplemental questions on the ESA, including the amount of each stimulus payment received. While CEX expenditures are reported at the quarterly frequency, new households enter the survey at each month, making the frequency of our data monthly. Each household is interviewed up to four times and reports consumption expenditures during the three preceding months. We defer to Johnson et al. (2006) and Parker et al. (2013) for a detailed description of the data. Since we depart from Parker et al. (2013) by allowing for treatment heterogeneity and incorporating recommendations from Borusyak et al. (2024), we present their homogeneous specification first as a useful benchmark, introducing our generalizations thereafter.

2.1 Homogeneous MPC

Parker et al. (2013) consider the following specification:

$$\Delta C_j = \beta' \omega_j + \lambda R_j + \alpha + u_j, j = 1, \ldots, N, \tag{1}$$

where ΔC_j is the first difference of consumption expenditures for $j = (i, t)$, corresponding to household i in quarter t, and $j = 1, \ldots, N$.[10] ω_j is a set of controls including month

[9] We defer to Parker et al. (2013) and Sahm et al. (2010) for an exhaustive discussion of the Economic Stimulus Act.

[10] To maintain consistent notation throughout the paper, we refer to j as the (i, t) combination of household i in quarter t. While we have information on the same households i in different periods t, as in Parker et al. (2013), identification is not obtained by comparing individual responses over time, but rather by comparing those who do and do not receive a rebate within a given period. We do not exploit any limited panel

dummies aimed at absorbing common time effects such as aggregate shocks, as well as seasonal factors.[11] α is a constant, and u_j the error term. The independent variable of interest is R_j, which denotes the amount of the tax rebate received by each household. λ is then interpreted as the causal effect of the rebate on expenditures. Identification is achieved by instrumenting the rebate value with an indicator for whether the rebate was received during the relevant quarter, so that the estimator implicitly compares expenditure changes of households that received the rebate in a certain period to expenditure changes of households that did not receive the rebate in the same period. Instrumentation avoids possible correlation between rebate value and the error term (since payments depended on income and household size, for instance).

2.2 Heterogeneous MPCs

We depart from the homogeneous specification in Equation (1) and allow for heterogeneity in expenditure responses to the tax rebate across households. In particular, we augment the specification in Parker et al. (2013) as follows:

$$\Delta C_j = \beta' \omega_j + \sum_{g=1}^{G} d_{jg} \left(\lambda_g R_j + \alpha_g \right) + \varepsilon_j, j = 1, \ldots, N, \tag{2}$$

where $d_{jg} = \mathbf{1}\left[j \in g \right]$ is an indicator that takes a value of 1 if household j (household i in period t) belongs to a certain group $g = 1, \ldots, G$ and ε_j is the error term. That is, we assume that heterogeneity in responses to the rebate can be summarized with G distinct groups, characterized by the pair of coefficients $\left\{ \alpha_g, \lambda_g \right\}$. We include group-specific intercepts, α_g, to correctly interpret λ_g as a marginal propensity to consume. For example, since we do not observe quarterly changes in income, failing to include group-specific level effects may bias MPC estimates due to heterogeneity in income changes unrelated to the tax rebate.

Recently, Borusyak et al. (2024) have argued that specifications like (1) may fail to identify (functions of) treatment effects in event studies with staggered treatment, particularly in the presence of heterogeneity, due to "forbidden comparisons" between newly treated and previously treated units. In practice, we modify (2) following their recommendations,

structure, except to construct consumption changes for the left-hand-side variable ΔC_j and other objects. We return to this point below.

[11] In Parker et al. (2013), the other controls are age, change in number of adults in the household, and change in the number of children in the household. The controls we use are the same, additionally including age squared, as in Misra and Surico (2014).

and instead estimate

$$\Delta C_j = \beta' \omega_j + \sum_{g=1}^{G} d_{jg} \left(\sum_{l=0}^{2} \iota_{jl} \lambda_{lg} + \alpha_g \right) + \epsilon_j, j = 1, \ldots, N, \tag{3}$$
$$\forall g = 1, \ldots, G, \; E\left[\epsilon_j \mid \omega_j, \iota_{jl}, j \in g\right] = 0,$$

where ι_{jl} is an observable indicator that takes a value of 1 if household j received the rebate exactly l periods ago. In contrast to (2), this specification regresses consumption changes directly on the rebate indicators, ι_{jl} (interacted with group membership), and includes lagged rebate indicators to account for lagged effects of the rebate, ensuring valid comparisons. We include the maximum number of lags, which in our data is 2, since household expenditure changes are observed at most for 2 quarters after rebate receipt. We further assume that the usual conditional mean zero assumption holds separately within each group for the error term, ϵ_j, and later will also assume that ϵ_j is normally distributed within each group. (3) is a special case of equation (4) in Borusyak et al. (2024), where heterogeneity takes the grouped structure described above. Borusyak et al. (2024) show that if there are no anticipation effects and (3) is correctly specified, then the associated OLS estimator is the unique efficient unbiased estimator for the treatment effects. As in Borusyak et al. (2024), we report coefficients λ_{lg} re-scaled by the average rebate amount, $\bar{R}$, to measure MPCs in per-dollar terms.

Crucially, we treat group membership, $d_{jg} = \mathbf{1}\left[j \in g\right]$, as an unknown object that is to be estimated, as explained in the next section. Group membership is binary and predetermined, meaning that groups, and thus MPCs, are discrete. We do not force group membership for household i to be fixed across t; even if individuals' preferences are constant, the MPC may be time-varying, due, for instance, to changes in state variables such as income and wealth. The model in (3) imposes no restrictions on the relationship between heterogeneity in λ_{lg} and the variables in ω_j, or any other excluded variables. Indeed, we avoid specifying such relationships (nor do we need to) for the purposes of estimating $\bar{\lambda} = \left\{\lambda_{lg}, l = 0, \ldots, 2, g = 1, \ldots, G\right\}$ for two reasons. First, excluding potential determinants of the MPCs from (3) means that when we estimate the distribution of MPCs, the heterogeneity we recover represents the full unconditional distribution; this includes variation that may be correlated only with unobservable characteristics, as well as that which may be correlated with excluded observables. Second, leaving relationships with observables unspecified in (3) allows us to explore such relationships more flexibly *ex post*. In particular, we can easily project a single baseline MPC distribution on a range of arrays of linear and nonlinear functions of observables. In the next section we discuss our clustering methodology to jointly estimate $\bar{\lambda}$ and $d_{jg} = \mathbf{1}[j \in g]$.

3 A clustering approach to MPC estimation

To estimate group-specific MPCs according to (3), we appeal to Gaussian mixture linear regression, which is a form of clustering regression. In general, clustering methods have the advantage of assigning individuals to groups not based on observable characteristics or the econometrician's judgment, but based on how well each set of estimated group-specific parameters (intercepts and MPCs in our case) describe the observed outcomes (consumption changes in our case) for each individual. This feature allows us to bypass the *ex ante* decision of which observables matter for MPC heterogeneity, and instead estimate the heterogeneity directly first. In other words, it allows us to estimate the full unconditional distribution of MPCs, including heterogeneity associated with unobservable characteristics.[12]

Gaussian Mixture Linear Regression The model (3) can be rewritten more compactly as

$$\Delta C_j = \sum_{g=1}^{G} d_{jg} \psi_g^{G\prime} x_j + \epsilon_j, \tag{4}$$

where $x_j = \begin{pmatrix} 1 & \iota_{j0} & \iota_{j1} & \iota_{j2} & \omega_j{}' \end{pmatrix}'$, ψ_g^G collects α_g, λ_{0g}, λ_{1g}, λ_{2g}, and β, and the elements of ψ_g^G corresponding to ω_j (β) are restricted to be constant across g. In this section, we include the G superscript on ψ_g^G and other parameters to make explicit the dependence of the parameter values on the specified number of groups; we omit the G superscript subsequently for compactness. We assume that the errors ϵ_j are Gaussian i.i.d. with mean zero and variance σ_g^2 for group $g \in G$. Let the $G \times 1$ vector σ^G collect σ_g^2 across groups. Let π_g^G denote the unconditional probability that any observation belongs to group g. More specifically, π_g^G is the share of individuals in the population belonging to each group, or $E\left[d_{jg}\right]$. As is standard, we assume that these prior probabilities, π_g, are independent of

[12] As an alternative to clustering approaches, quantile regression is used by Misra and Surico (2014) to characterize heterogeneous responses to the 2008 tax rebate. Quantile regression differs from clustering; because quantile regression computes relationships at percentiles of the overall conditional distribution, the estimated MPC distribution depends on the correlation of MPCs with other forms of heterogeneity. If the "ranking" of the conditional distribution is mostly driven by factors other than responsiveness to the rebate (like fixed effects or other covariates), and these factors are uncorrelated with the rebate, heterogeneity of the MPC distribution will be underestimated in the presence of noise. We provide a simple example in Supplement B.

observables in x.[13] Finally, $\phi\left(\Delta C_j; \psi_g^{G'} x_j, \sigma_g^2\right)$ is the normal probability density function evaluated for mean $\psi_g^{G'} x_j$ and variance σ_g^2. Then the complete-data likelihood (i.e. where group membership d_{jg} is known with certainty) is

$$L\left(\Delta C, X, D; \theta^G\right) = \prod_{j=1}^{N} \prod_{g=1}^{G} \left(\pi_g^G\right)^{d_{jg}} \phi\left(\Delta C_j; \psi_g^{G'} x_j, \sigma_g^2\right)^{d_{jg}}, \tag{5}$$

where D collects d_{jg}, and θ^G collects Ψ^G, σ^G, and π_g^G across groups. Since d_{jg} are latent variables, L cannot be maximized directly. Intuitively, this means that if a household belongs to group g, then ΔC_j is normally distributed with mean $\psi_g^{G'} x_j$ (which is equal to the untreated potential outcome, $\alpha_g + \beta' \omega_j$, plus the appropriate-horizon MPC λ_{lg}if treatment has already occurred) and variance σ_g^2. We are assuming that there are essentially G different populations (each normally distributed) with different MPCs and fixed effects, and the econometrician does not know *ex ante* to which population any observation belongs. As a result, we maximize the *expected* log-likelihood. Integrating $\mathcal{L} = \log L$ over d_{jg} (conditional on $(\Delta C_j, x_j)$) yields

$$l_\phi\left(\theta^G\right) \equiv E_{D|\Delta C, X}\left[\mathcal{L}\left(\Delta C, X, D; \theta^G\right)\right] = \sum_{j=1}^{N} \sum_{g=1}^{G} \gamma_{jg} \left(\log \pi_g^G + \log \phi\left(\Delta C_j; \psi_g^{G'} x_j, \sigma_g^2\right)\right), \tag{6}$$

where

$$\gamma_{jg} = \Pr\left(d_{jg} = 1 \mid \Delta C_j, x_j\right) = \frac{\pi_g^G \phi\left(\Delta C_j; \psi_g^{G'} x_j, \sigma_g^2\right)}{\sum_{h=1}^{G} \pi_h^G \phi\left(\Delta C_j; \psi_h^{G'} x_j, \sigma_h^2\right)} \tag{7}$$

are the posterior weights. These weights, and in particular their deviations from binary values, represent the econometrician's uncertainty about an individual's group membership due to noise in the consumption equation; latent true group membership is discrete and fixed for each observation, not probabilistic. The posterior probabilities are conditional on the information contained in ΔC_j and x_j; while additional covariates may be available (as we consider in subsequent regressions on observables), and including those observables in an expanded conditioning set might lead to different posterior probabilities, γ_{jg} provides valid posterior probabilities conditional on the information that is in-

[13] As a robustness check, we additionally consider a group membership function that depends explicitly on observable characteristics in Supplement A.5; the results are nearly identical to our baseline.

cluded in the model. Thus, Ψ^G can be obtained by maximizing

$$\begin{aligned}\tilde{l}_{\phi}\left(\Psi^G,\sigma^G\right) &= \sum_{j=1}^{N}\sum_{g=1}^{G}\gamma_{jg}\log\phi\left(\Delta C_j;\psi_g^{G\prime}x_j,\sigma_g^2\right)\\ &= \sum_{j=1}^{N}\sum_{g=1}^{G}\gamma_{jg}\left[\log\left(\frac{1}{\sqrt{2\pi\sigma_g^2}}\right)-\frac{\left(\Delta C_j-\psi_g^{G\prime}x_j\right)^2}{2\sigma_g^2}\right],\end{aligned}$$

which omits the portion of the expected log-likelihood that depends only on $\pi_1^G,\ldots\pi_g^G$. Conditional on γ_g and σ^G , this implies minimizing

$$Q_{\phi}\left(\Psi^G\right) = \sum_{j=1}^{N}\sum_{g=1}^{G}\gamma_{jg}\left(\Delta C_j-\psi_g^{G\prime}x_j\right)^2/\sigma_g^2, \tag{8}$$

with respect to ψ_g^G. When $\sigma_g^2\equiv\sigma^2$, and γ_{jg} is left unrestricted (as opposed to the Gaussian posterior probability in (7)), minimizing jointly over $\gamma_{jg}(\cdot)$ and Ψ^G delivers the "hard K-means" algorithm considered by Bonhomme and Manresa (2015).[14] In this algorithm, each observation would have binary weights $\left(\gamma_{jg}\in\{0,1\}\right)$, equal to one for whichever group minimizes its residual. This approach has two drawbacks in our empirical application. First, our setting is effectively cross-sectional as in Parker et al. (2013), because households receive the rebate only once and identification of MPCs is obtained by comparing those who do and do not receive a rebate in a given period. As such, we do not have a panel dimension to reduce (and asymptotically eliminate) estimation error in the group assignment. Therefore, the econometrician ought to treat the assignment of households to groups as probabilistic – representing the uncertainty introduced by ϵ_j – even though the data-generating process entails binary assignment. Second, even if it was possible to exploit a panel dimension, it is overly restrictive to force group membership for household i to be fixed across t (as required by Bonhomme and Manresa (2015)); indeed, MPCs may be time-varying, due, for instance, to changes in state variables such as income and wealth.

In practice, given that the expected log-likelihood in (6), and the subsequent objective in (8), involve γ_{jg}, which depends on $\pi_1^G,\ldots\pi_g^G,\Psi^G$, and σ^G, the model is estimated using the Expectation-Maximization (EM) algorithm (e.g., Dempster et al. (1977)), where the E-step updates the posterior weights γ_{jg} conditional on a set of parameters and the M-step updates $\pi_1^G,\ldots\pi_g^G,\Psi^G$, and σ^G as in WLS. For a detailed discussion of the GMLR

[14]See Lewis et al. (2023) for a "possibilistic" "fuzzy" approximation to the hard K-means assignment.

problem and its implementation via the EM algorithm, see Desarbo and Cron (1988) or Jones and McLachlan (1992). To minimize reliance on initial conditions, we maximize the objective for 2,500 different start values and then keep the solution associated with the highest log-likelihood.

An advantage of the GMLR approach to regression-based clustering is that asymptotic properties of the estimator (consistency and asymptotic normality) follow immediately under regularity conditions from standard maximum likelihood results. This means that analytical inference on Ψ^G (and other parameters) is straightforward, even though inference on an individual's posterior group membership probabilities γ_{jg} – which are a function of the realized error ϵ_j – is not possible. Desarbo and Cron (1988) provide a discussion of these inference results; we use the Fisher Information for inference on Ψ^G.[15] Identification of the mixture components follows from the parametric specification of the model, which, conditional on the number of mixture components, allows the econometrician to distinguish Gaussian variation in the errors, which lack first-order group structure, from variation in the conditional mean, which possesses group structure. Unfortunately, we are not aware of any identification results for mixture models without a panel dimension, as in our dataset, that do not leverage parametric structure on the errors as GMLR does. However, in addition to our empirical analysis below, we have also conducted a robustness exercise based on a non-parametric estimation approach, fuzzy clustering regression, that does not impose parametric assumptions, but is only consistent for pseudo-true parameters as a result (see Lewis et al. (2023) for a discussion). The results are similar to our baseline, and we conclude that our findings are not particularly fragile with respect to our parametric assumptions.

Choosing the number of groups In all clustering models, it is necessary to choose G, the number of groups; we use the Bayesian Information Criterion (BIC). In particular, the BIC for a candidate number of groups, $\tilde{G}$, is given by

$$BIC\left(\tilde{G}\right) = k_{\theta^{\tilde{G}}} \log N - 2l_{\phi}\left(\hat{\theta}^{\tilde{G}}\right),$$

[15] For a more detailed discussion of inference in Gaussian mixture models in general, see McLachlan and Basford (1988).

where $k_{\theta^{\tilde{G}}}$ is the number of unique parameters in $\theta^{\tilde{G}}$, and $l_\phi\left(\hat{\theta}^{\tilde{G}}\right)$ is the maximized incomplete-data log-likelihood for $\tilde{G}$ groups,

$$l_\phi\left(\hat{\theta}^{\tilde{G}}\right) = \sum_{i=1}^{N} \log \sum_{g=1}^{\tilde{G}} \hat{\pi}_g \phi\left(\Delta C_j; \hat{\psi}_g^{\tilde{G}'} x_j, \hat{\sigma}_g^2\right).$$

Under regularity conditions, the BIC is consistent for the true value of G (see, e.g., Celeux et al. (2018)). To confirm the BIC's estimate, we compare the selected G to that obtained from $K-$fold cross validation, and ensure that the chosen model is compatible with both criteria.[16]

4 Results

We estimate the distribution of marginal propensities to consume out of the 2008 tax rebate checks using our clustering methodology. As discussed in the previous section, our approach allows us to consistently estimate MPC heterogeneity described by the discrete MPCs $\bar{\lambda}$ and the population shares of households belonging to each group, π_g. The combination of these two objects describes the unconditional MPC distribution in the population, which constitutes our core results in the following subsections. Unless explicitly noted, by the MPC we refer to the contemporaneous MPC, measured by λ_{0g}. Additionally, we compute household-specific *predicted* MPCs using the estimated posterior group-membership probabilities, γ_{jg}. The distribution of predicted MPCs is useful for two reasons. First, it is the ideal input for any policymaker designing targeted fiscal transfers aimed at maximizing overall spending, as we describe below. Second, this is the required distribution for characterizing the observable determinants of heterogeneity in the marginal propensity to consume, as we discuss in Section 4.3.

We estimate a considerable degree of MPC heterogeneity in both distributions, whose extent varies depending on the consumption category considered. In Section 4.1, we present the distribution of MPCs for total expenditures and discuss the stability of our predicted MPCs. In Section 4.2 we report the MPC distribution for nondurable and durable goods. Importantly, our approach also allows us to directly test whether households display similar spending propensities for different consumption goods, or instead display differential responses across expenditure types when they receive a transitory

[16] $K-$fold cross-validation is a model validation approach that splits the data into K different subsamples. For each $k = 1, \ldots, K$, the model is estimated on the $K-1$ subsamples as training data and the fit assessed on the holdout sample, k.

income shock such as a tax rebate. Finally, in Section 4.3 we explore which observable household characteristics are correlated with the predicted marginal propensities to consume, both individually and jointly, and in Section 4.4 we analyze the longer-run spending effects of the 2008 tax rebates.

4.1 The distribution of marginal propensitixes to consume

Our headline specification considers total expenditures, defined as in Parker et al. (2013). Following Kaplan and Violante (2014), who show that properly accounting for outliers reduces the homogeneous rebate coefficient, while increasing precision, we drop the top and bottom 1.5% of consumption changes.[17] We select the number of groups, G, using the BIC, as discussed above, indicating $G = 3$; cross validation selects the same model. The MPC distributions, as well as the observable determinants of MPCs (Section 4.3), are broadly unchanged for alternative $G = 2$.[18] Reassuringly, we confirm that the heterogeneity we recover is not spurious: to do so, we generate 250 Monte Carlo samples of data from estimates for the homogeneous model and compute the BIC for each; we find that the BIC selects $G = 1$ in all samples.

Table 1 reports the points of the unconditional MPC distribution for total expenditures, $\{\lambda_{0g}\}$, following the discrete data-generating process (DGP) in (3). To enable interpretation as MPCs, we rescale all estimates λ_{lg} by the average rebate payment across rebate recipients in the sample, as in Borusyak et al. (2024). We uncover a large degree of heterogeneity, ranging from 4 cents to 1.33 dollars spent for every dollar of rebate. More than three quarters of households are associated with an MPC of 23 cents on the dollar or less, but a substantial share of households display sizable spending responses. Aggregating λ_{0g} by π_g, we find that the average marginal propensity to consume, 0.42, is larger than the MPC in the homogeneous specification, 0.24.

We report the standard errors arising from the maximum likelihood estimation in parentheses for both MPCs, λ_{0g}, and group shares, π_g. We additionally report standard errors for the MPCs that take posterior group probabilities γ_{jg} as known, to better parallel the existing literature, which assumes that group membership is known *ex ante* (for exam-

[17] This is the only way in which our sample departs from the sample used in Panel B of Table 3 in Parker et al. (2013). In Supplement A.4 we repeat our main analysis without trimming; in this case we typically uncover a larger degree of MPC heterogeneity, although our estimates are less precise. We thus see our baseline results as potentially conservative with respect to the degree of heterogeneity. Our results on the drivers of MPC heterogeneity are broadly unchanged without trimming.

[18] The BIC has a kink at $G = 2$ and then flattens at $G = 3$. While the results are similar across the two specifications, there are economically important differences in the two distributions, as we show in Supplement A.2. For instance, our baseline results feature a maximum MPC of 1.33 compared to 1.14 under $G = 2$, as well as an additional mode at 0.23.

Table 1: MPC heterogeneity

	λ_{01}	π_1	λ_{02}	π_2	λ_{03}	π_3
Estimate	0.04	0.30	0.23	0.48	1.33	0.23
Std. Err	(0.07)	(0.01)	(0.12)	(0.01)	(0.47)	(0.01)
Cond. Std. Err	[0.03]	–	[0.07]	–	[0.42]	–

Notes: The first row reports estimates for group-specific MPCs, λ_{0g}, and population shares, π_g. Standard errors, reported in the second row in parentheses $(\cdot)$, account for all uncertainty;. conditional standard errors, reported in the third row in square brackets $[\cdot]$, instead take GMLR weights as given, to parallel the way that group assignment is imposed *ex ante* in the existing literature.

ple, by splitting households on quantiles of wealth and assuming each quantile features a homogeneous MPC). The group shares, π_g, are precisely estimated. The largest MPC is statistically significant at the 1% level using either standard errors. The smallest MPC is instead not statistically different from zero even when using the second set of standard errors that does not incorporate all estimation uncertainty. We therefore cannot rule out that there exists a fraction of households adhering to the permanent income hypothesis when examining total expenditures.

The distribution described by Table 1 is consistently estimated in the population and characterizes unconditional MPC heterogeneity. However, there are several reasons to also consider the predicted MPC for each household. For example, a policymaker aiming to maximize the overall spending effect of stimulus checks would want to know which households are likely to have high or low MPCs. The posterior predicted MPC for each household provides exactly this estimate, as this prediction minimizes the mean squared error of each households' spending response.[19] The government could then use this distribution to target specific individuals with high predicted marginal propensities to consume. A household's posterior MPC is also the natural candidate to consider as a dependent variable in projections of MPCs on household characteristics. Our DGP postulates the MPC distribution as being discrete, but households' posterior predicted MPCs will generally be continuous.

Figure 1 shows the distribution of households' predicted MPCs, formally defined as $\tilde{\lambda}_{0,j} = \sum_{g=1}^{G} \gamma_{jg}\lambda_{0g}$. The vast majority of households display a relatively low (but non-negligible) predicted MPC, and the share of households with a given predicted MPC slowly decays as the value increases. No household is predicted to consume fewer than

[19]Alternatively, one could predict household-specific MPCs by assigning the λ_{0g} associated with their most likely group (ie, $\lambda^*_{0,j} = \lambda_{0g^*_j}$, where $g^*_j = \arg\max_g \gamma_{jg}$). While this approach minimizes classification error, it is less relevant from the perspective of a government interested in maximizing the overall spending response. Nevertheless, we show this distribution in Supplement A.2.

Figure 1: Predicting individual MPCs out of the tax rebate

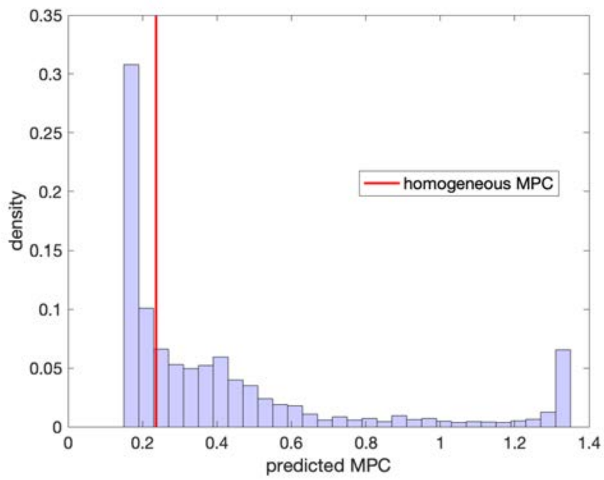

Notes: The histogram (light blue bars) plots the individual MPCs constructed using the GMLR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}. For each household we compute the posterior-weighted MPC across the discrete group-specific MPCs. Results are for total expenditures and we only plot households that received the rebate. The sample is defined as in the text. The BIC selects $G = 3$. The homogeneous MPC (red vertical line) is estimated imposing $G = 1$ in our baseline specification.

15 cents for every rebate dollar: this is equivalent to saying that no household is estimated to have a very high posterior weight, γ_{jg}, associated with the lowest MPC group. Therefore, while we have previously shown that there is a group-specific MPC as low as 4 cents for each dollar rebate, associated with a substantial share of the population, the households in our sample are all *individually* predicted to consume at least 15% of their rebate. Non-zero spending responses can be dispersed if agents face bounded rationality. Ilut and Valchev (2020), for instance, develop a model in which MPCs can be high for all households, even those with slack liquidity constraints. Due to limited cognitive perception, households can find themselves in the midst of a "learning trap", "which makes the high MPC behavior the norm, rather than exception."

At the opposite end of the predicted MPC distribution, a sizable fraction (11%) of households are predicted to consume more than the entirety of the rebate. This suggests that many households have a posterior weight near 1 for the highest-MPC group. Finally, we confirm the stability of the predicted MPC distribution. While we cannot formally conduct inference on a household's weights – since they are a function of a single realization of a random error, not a parameter – we assess the stability of our findings via bootstrapping. In particular, we repeat the GMLR estimation of the distribution of MPCs for total expenditures, with 3 groups, over 250 samples obtained by bootstrap with replacement. For each bootstrapped sample, we construct the cumulative density function of predicted household MPCs, $\tilde{\lambda}_{0,j}$. The dash-dotted blue line in Figure 2 shows the median

Figure 2: Bootstrapped individual MPCs out of the tax rebate

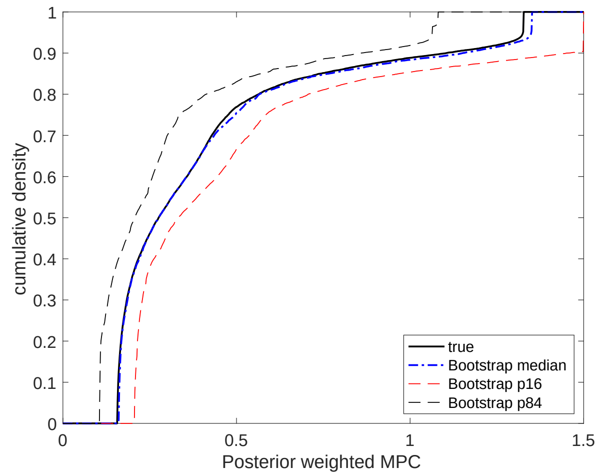

Notes: The black solid line plots CDF of the distribution of individual MPCs for total expenditures, calculated as shown in Figure 1. We predict individual responses using posterior-weighted MPCs. The blue dash-dotted line shows the median CDF of the estimated distribution of MPCs, across 250 bootstraps with replacement. The dashed red and black lines correspond to the centered 68% confidence interval.

CDF across bootstraps, which tracks the CDF of the distribution shown in Figure 1 (here depicted in solid black) remarkably well. Moreover, there is reasonably little variation across bootstraps, as evidenced by the centered 68% confidence interval. The predicted distribution is found to be quite stable even when we instead assign individuals to their most likely (maximum posterior weight) MPCs for each bootstrapped sample.[20]

Taken together, both the unconditional MPC distribution and the distribution of predicted MPCs show that there is indeed considerable variation in spending responses across households. From a policy perspective, this implies a significant benefit in targeting transfers; for a given dollar value of transfer, those households with a higher MPC will spend more and save less, leading to a greater increase in consumption and stimulatory effect on aggregate demand. We return to the question of whether such targeting is feasible in practice in Section 4.3 when we evaluate the observable correlates of MPCs.

4.2 The MPC distribution for different consumption goods

We have shown how households differ with respect to their propensity to consume out of their tax rebate. How does the distribution of these propensities differ across consumption goods? The granularity of the CEX data allows us to tackle this question, while our methodology allows us to explore how good-specific MPCs vary at the household level.

[20] In the vast majority of the bootstrap samples the BIC also flattens at $G = 3$, further confirming our selection for the number of groups.

Table 2: MPCs out of the tax rebate: nondurables and durables

(a) Nondurables

	λ_{0g}	π_g
$g = 1$	0.09	0.48
	(0.04)	(0.01)
	[0.02]	–
$g = 2$	0.18	0.52
	(0.07)	(0.01)
	[0.07]	–

(b) Durables

	λ_{0g}	π_g
$g = 1$	0.03	0.29
	(0.03)	(0.01)
	[0.01]	–
$g = 2$	0.15	0.50
	(0.07)	(0.01)
	[0.04]	–
$g = 3$	0.67	0.21
	(0.42)	(0.01)
	[0.42]	–

***Notes*:** The tables show estimates for group-specific MPCs, λ_{0g}, and population shares, π_g. Standard errors, in parentheses $(\cdot)$, account for all uncertainty;. conditional standard errors, in square brackets $[\cdot]$, instead take GMLR weights as given, to parallel the way that group assignment is imposed *ex ante* in the existing literature. Nondurable expenditures (panel (a)) are defined, following Parker et al. (2013), as strictly nondurables (Lusardi (1996)) plus apparel goods and services, health care expenditures (excluding payments by employers or insurers), and reading material (excluding education). As in Parker et al. (2013), we define durable expenditures (panel (b)) as the difference between total and nondurable expenditures. The homogeneous MPCs ($G = 1$) are 0.13 for nondurables and 0.05 for durables. For nondurables the BIC selects $G = 2$ and for durables $G = 3$.

The left panel of Table 2 reports the estimated MPCs, λ_{0g}, and population shares, π_g, for nondurable goods consumption. As expected, households consume a lower fraction of the rebate in nondurables than estimated for total expenditures, as nondurable goods account for, on average, less than two thirds of household total expenditures. Roughly half of households spend 9 cents for each dollar of rebate in nondurables; this coefficient is statistically different from zero at the 5% level. Thus, nondurable MPCs exhibit a lower bound above zero, suggesting no household strictly follows the Permanent Income Hypothesis (Friedman (1957)). The existence of a lower bound for the MPC is a topic of ongoing debate in the literature. For example, Fagereng et al. (2016) and Olafsson and Pagel (2018) find sizable spending responses even for households with high liquid wealth in Norwegian administrative data and Icelandic application user data, respectively. We likewise find evidence that even the smallest nondurable MPCs are larger than zero, although not substantially so, even when estimating the full unconditional MPC distribution, in standard U.S. survey data. In contrast, in this same data, Misra and Surico (2014) use quantile regression to estimate a substantial share of MPCs at or *below* zero; we discuss in Supplement B how our approach differs from theirs.

The remaining half of the distribution consumes just less than a fifth of the rebate in nondurables. Households' predicted MPCs are distributed within similar bounds, reported in Supplement A.2. As before, we confirm the stability of the predicted MPC distribution via bootstrapping. The minimum predicted MPC is higher than 5 cents in 94%

of the bootstraps, providing further evidence of a positive lower bound on nondurables consumption.

The right panel of Table 2 shows the estimated MPC distribution for durable goods. About 30% of households have a durable MPC that is not statistically different from zero when we incorporate all estimation uncertainty. The highest MPC is noisily estimated at 0.67. This uncertainty at the top is also evident in our bootstrapping exercise (see Supplement A.2).[21] The upper bound of the durable MPC distribution is sensitive to the estimation sample. In particular, our baseline sample excludes the top and bottom 1.5% of consumption changes: without the trimming, the highest durable MPC increases – far exceeding 1 – and becomes statistically significant, while the lower bound does not change (see Supplement A.4). This suggests that some of the trimming at the top excludes large purchases that are indeed a direct response to rebate receipt. Overall, the dichotomous nature of this MPC distribution accords with the discrete nature of durable goods purchases.

Finally, we assess whether households with high propensities to consume nondurable goods after receiving the rebate are also more likely to consume durable goods. The correlation between household-level predicted MPCs for nondurable goods with those for durables is 0.13 (significant at the 1% level). We can therefore rule out substitution between goods, but the estimated complementarity at the margin is quantitatively small. Nevertheless, this weak complementarity might signal the presence of heterogeneous preferences or a small share of "spender" types, who are more prone to adjust *any* type of consumption in response to transitory income shocks. While the structure of our data does not allow us to draw conclusions regarding whether the heterogeneity we measure is permanent or transitory, we can investigate what observable characteristics explain the estimated MPC distributions that we recover. We tackle this issue in the next section.

4.3 What drives MPC heterogeneity?

Our empirical strategy uncovers the unconditional distribution of marginal propensities to consume without taking a stance, *ex ante*, on its observable determinants. Consequently, we can use the estimated distribution to understand how predicted MPCs correlate, *ex post*, with observable characteristics. As such, we contribute to the literature

[21] The bootstrapping exercise is also informative for the selection of the number of groups, G. While the BIC in our estimation sample flattens at $G = 3$, before dropping again at $G = 5$, it remains flat beyond $G = 3$ in the majority of bootstrapped samples. We take this as a signal of the sensitivity of durable MPCs to a small number of observations. While we show results for $G = 3$, the estimated MPC distribution is quite similar for $G = 5$.

in three ways. First, using our methodology, we uncover a large number of statistically significant individual correlations between MPCs and observable drivers. This is true despite the fact that we use a dataset and an identification strategy that previously failed to find statistically significant relationships (e.g., Parker et al. (2013)). Second, we show that the distribution of MPCs is correlated with observable characteristics *jointly*, and, unlike in previous methodologies requiring progressively smaller interacted groups, we can be confident that any lack of significant correlations is not due to loss of statistical power. When performing this joint analysis, we highlight the most important drivers of our estimated propensities. Third, we quantitatively assess the share of MPC heterogeneity explained by observables. This metric is important for gauging the potential the government has for targeting payments explicitly. Throughout the section we mostly focus on total expenditures, since results are largely consistent across consumption goods; we highlight the few instances in which results differ for nondurables or durables.

We first examine how relevant household characteristics individually correlate with MPCs. To do so, we replicate each household's data G times, and assign to each of household j's G replications one of the λ_{0g} MPCs and its associated posterior weight, γ_{jg}. We then estimate weighted least-squares (WLS) regressions, using the household posterior probabilities for each MPC, γ_{jg}, as weights. In columns (1)-(4) of Table 3 we report how different sets of predictors are individually related with the estimated MPCs and then show, in column (5), how these results change when we regress our estimated MPCs on all household characteristics jointly. Due to the low response rate for liquid wealth, and the associated potential non-response bias, we do not include it as an explanatory variable here, but do so in Supplement A.3. Our results are robust across specifications and, even more notably, are unchanged when considering the MPC distribution estimated with a different number of groups, as shown in Supplement A.3. Moreover, we confirm that all the results throughout the remainder of Section 4 are unlikely to arise if the estimated heterogeneity in the MPC distribution is spurious. To do so, we generate 250 Monte Carlo samples from the estimated *homogeneous* model, impose a spuriously large number of groups ($G = 3$) and repeat our regressions on observables using the estimated MPC distributions. Doing so, we find significant relationships with observable characteristics in virtually none of the Monte Carlo samples.

Table 3: Explanatory variables of MPCs

	(1)	(2)	(3)	(4)	(5)
Mortgage interest/income	0.043				-0.112
	(0.070)				(0.095)
Homeowner dummy	0.090***				0.076***
	(0.021)				(0.026)
Mortgage dummy	0.022				-0.011
	(0.021)				(0.029)
Zero salary		0.002			-0.087***
		(0.022)			(0.031)
Salary: middle tercile		0.054**			0.112***
		(0.024)			(0.027)
Salary: top tercile		0.130***			0.188***
		(0.025)			(0.035)
Zero non-salary			-0.011		0.031
			(0.022)		(0.027)
Non-salary: middle tercile			-0.050*		0.049
			(0.026)		(0.033)
Non-salary: top tercile			0.015		0.166***
			(0.027)		(0.038)
APC				0.075***	0.189***
				(0.014)	(0.023)
adj. R^2	0.008	0.011	0.001	0.009	0.059

Notes: Regressions are estimated on the sample that includes only households receiving a rebate in the current period. Robust standard errors are reported in parentheses. In column (5) we also control for five education dummies, a married dummy, number of children, age and age squared; these coefficients are not reported. *, **, and *** denote significance at the 10, 5, and 1% levels respectively.

Housing-related variables

As shown in column (1), homeowners have significantly larger MPCs, a result that aligns with Parker et al. (2013). Moreover, having a mortgage is associated with a still higher propensity to consume, and mortgagors that face larger interest payments relative to their income spend a larger fraction of their rebate, although both effects are statistically insignificant. When we control for all household characteristics jointly in column (5), the homeownership effect persists, although it is statistically and economically weaker. Interestingly, the effects of mortgages flip signs (while remaining insignificant), along both the extensive and intensive margin. Recent papers, such as Cloyne, Ferreira, and Surico (2020) and Wong (2021), have shown that MPC heterogeneity by housing and mortgage status can translate into heterogeneous consumption effects of monetary policy and, in turn, different aggregate effects of interest rate adjustments. Our results suggest that these relationships are complex and might change when conditioning on other household characteristics. In particular, these results imply that income – rather than having a mortgage or not – is the more important predictor for the MPC, as we elaborate on next.

Salary and nonsalary income

We split income between its salary and non-salary components, see columns (2) and (3).[22] When considered individually, both margins are broadly positively correlated with the marginal propensity to consume, but there are interesting non-monotonicities. In particular, MPCs are highest for top earners, and U-shaped in non-salary income.[23]

These findings are broadly maintained when controlling for other household characteristics in column (5). The differences are also economically meaningful, as households whose earnings are in the top tercile spend 17 cents more per each rebate dollar, while having a high stream of non-salary income implies 19 additional cents spent for each rebate dollar. Moreover, households with no salary have slightly lower MPCs, although the effect is small and insignificant for nondurable goods, as we show in Supplement A.3. The relationship with income, especially non-salary income, is generally weaker for nondurable MPCs, although it remains positive and mildly significant. In order to shed further light on our results, we split non-salary income between the sum of its financial and business components, and the remaining sources (mainly social security income, besides unemployment compensation and other transfers). Households in the top tercile of either component display statistically higher MPCs. Hence, it seems that households whose level of income is higher are more likely to spend a larger proportion of their rebate check.

While perhaps surprising, the positive correlation between income and MPCs that we

[22] Income in the CEX is measured in the first interview and relates to income over the prior 12 months. Non-salary income consists of farm and business income, financial income (e.g., income from interest, dividends, pensions and annuities) and all other income except foodstamps (e.g., retirement, supplemental security, unemployment compensation), following the categorization in Coibion et al. (2017). We deflate income variables by CPI.

[23] Results are broadly unchanged when we split by quintiles, as we show in Supplement A.3. We also find that these relationships hold jointly: households with highest salary *and* non-salary income have the highest marginal propensities to consume.

find is consistent with various antecedents in the empirical literature. For example, Kueng (2018), who studies consumption responses to regular and predetermined payments from the Alaska Permanent Fund, also finds that spending propensities were higher for relatively high income households. Results of this kind have been often rationalized by the relative size of the payment. High-income households may perceive, ex ante, a small cost of deviating from consumption smoothing, as rebates are smaller relative to their income. This behavior could be explained by theories of near-rationality or mental accounting; for instance Boutros (2022) proposes a model of bounded intertemporal rationality, in which the smaller the relative size of the payment, the more planning costs dominate the benefits of consumption smoothing.

Many other papers have found weak or inconclusive relationships between MPCs and income. Misra and Surico (2014) also find that median income is higher at the top of the conditional distribution of consumption changes, which they find to be associated with higher propensities to consume, although the overall relationship is U-shaped. While Broda and Parker (2014) find that low-income households had larger propensities to spend in the month of the 2008 rebate receipt than households in the top income tercile, this difference "becomes indistinguishable by the end of the quarter". Shapiro and Slemrod (2009) use data on self-reported propensities to spend the 2008 rebate and show that low-income individuals were more likely to pay off debt rather than consume. They also find that 21% of households making more than $75,000 of total annual income reported to spend most of the rebate, compared to 18% for households with total income below $20,000. Miranda-Pinto, Murphy, Walsh, and Young (2020) develop a model that can rationalize these findings via time-varying consumption thresholds. Koşar et al. (2023) provide another explanation, supported by empirical evidence: poor, and especially high–debt, households are more likely to use fiscal transfers to pay down their debts. On the other hand, some studies have found that low-income households have a higher marginal propensity to spend: see, for instance, Johnson et al. (2006) for the 2001 tax rebate and Jappelli and Pistaferri (2014), with respect to cash on hand, for Italian data on self-reported MPCs. Importantly, however, while we find that low-income households have relatively lower MPCs than high earners, they still exhibit large marginal propensities to consume, suggesting sizable deviations from consumption smoothing at the bottom of the income distribution.

The average propensity to consume and rich-spenders

We find that the average propensity to consume (APC) is an important predictor of marginal propensities to spend the rebate, and this positive association is even stronger when controlling for all household characteristics, as can be seen by comparing columns (4) and (5) of Table 3. Through the APC, we aim to capture persistent preference heterogeneity, by identifying something like a "spender" type.[24] Households that spent 1 percentage point more of their income before receiving the rebate spent 19 additional cents out of each rebate dollar. In the theoretical literature, the relationship between MPC and APC has been extensively studied, and we discuss later what different models of household behavior predict regarding this relationship. However, to the best of our knowledge, we are the first to provide empirical evidence on the cross-sectional relationship between the APC and the MPC to well-identified transitory income shocks.

We show in Figure 3 how the MPC varies jointly with the APC and total income. We separately compute quintiles of the APC and total income, and calculate the average MPC for each quintile pair using households' posterior weights. The MPC is increasing in both income, conditional on the APC, and vice versa. We uncover three main groups. First, households with low total income and a low APC display the lowest marginal propensity to consume. We label these households "poor savers". Second, households with a high APC and low total income, or vice versa, display intermediate MPCs. The largest marginal propensity to consume is found among households with a high APC and high total income. We label this group "rich spenders".[25]

The results presented thus far are particularly relevant for disciplining macro models of household consumption. For example, the relationship between the MPC, the APC, and income can be directly tested in even the simplest of consumption-savings models. In a standard one-asset incomplete markets model with borrowing constraints, the unconditional relationship between the MPC and the APC is ambiguous. On the one hand, households with temporarily low income will also typically display high APCs and MPCs. On the other hand, typically high income agents – who have reached their target level of wealth – will have a low MPC and a high APC as they no longer save.[26] In general, the canonical consumption-savings model typically predicts a negative correlation between

[24] Empirically, we define the APC as average lagged consumption divided by total income, which, as previously mentioned, relates to the prior 12 months. We lag expenditures to avoid the possibility of a mechanical positive correlation with the MPC. To ensure stability of APCs, and avoid sensitivity to a single shock in the previous period and instead reflect persistent heterogeneity, we average expenditures over all the available lagged quarters at the household-level, but the results are virtually unchanged if we only consider the first lag. We consider income as measured in the first interview for each household, which refers to the previous 12 months. We winsorize the APC upwards at 3, which is about 5% of the observations.

[25] We find similar relationships for the MPC for nondurables and durables, especially the presence of "rich spenders", as we show in Supplement A.3.

[26] In addition, these models typically fail to generate savings rates (APCs) that increase (decrease) with wealth and permanent income, at odds with what is observed in the data and documented by Dynan, Skinner, and Zeldes (2004) and Straub (2017).

Figure 3: The relationship between MPCs, APCs, and income

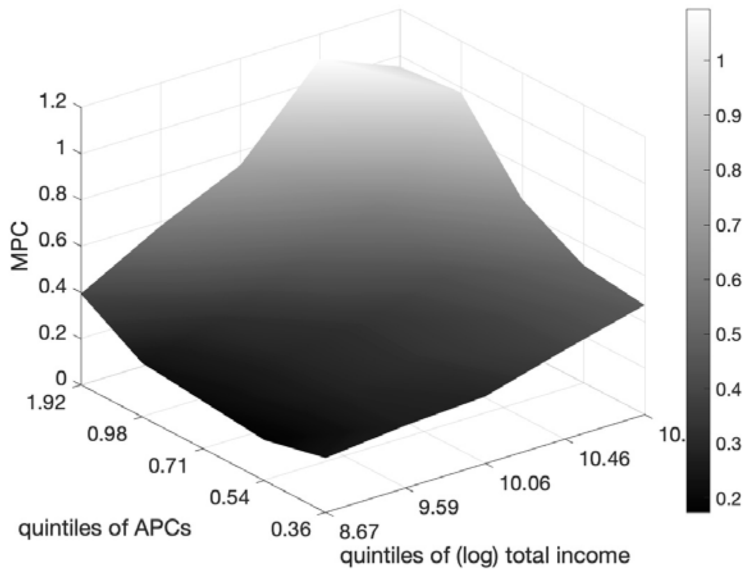

Notes: For each pair of quintiles of APC and log total income, we compute the average value for λ_{0g}, weighted by γ_{jg}, and plot it as a surface. The color bar on the right represents the MPC.

the MPC and the APC conditional on income, thus at odds with our empirical findings. Two-asset versions of these models such as Kaplan and Violante (2014) generate wealthy hand-to-mouth households with high income and high MPCs, but these households do not necessarily have high APCs.

The preceding characterizations are conditional on homogeneous preferences.[27] Preference heterogeneity, in contrast, can help rationalize our findings. Aguiar et al. (2019), for instance, highlight the importance of heterogeneity in the intertemporal elasticity of substitution (IES) in order to generate heterogeneous target levels of wealth in this same class of models; high-IES households have both high MPCs and high APCs.[28] Empirically, and consistent with this idea, Parker (2017) finds that the majority of consumption responsiveness to the tax rebate in the Nielsen data is driven by a measure of impatience, defined as households reporting to be "the sort of people who would rather spend their money and enjoy today rather than save more for the future." Our results suggest that this preference heterogeneity may be positively correlated with income. For example, if entrepreneurs or investors have a high IES, they may display high MPCs, APCs, and income. Next, we present additional suggestive evidence on the importance of latent unobserved heterogeneity.

[27] Another dimension of (latent) heterogeneity may come from differences in income processes.
[28] For a more in-depth discussion of MPCs in heterogeneous-agent models see Kaplan and Violante (2021).

Demographics and unobserved heterogeneity

Besides the main drivers of MPC heterogeneity discussed so far, we investigated whether many other household characteristics are correlated with the MPC. We report some of these results in Supplement A.3; MPCs increase with education, are non-monotonically related with age, and display little difference by race and gender of households' reference persons. Most importantly, all these variables are insignificantly associated with the MPC when tested jointly with income, the APC, and housing status, with the sole exception of marital status.

Our results therefore suggest that only a handful of household observable characteristics matter for estimated MPC heterogeneity. Moreover, these main drivers – as well as other household characteristics that do not strongly correlate with the MPC – explain a relatively small portion of the variance of the weighted MPC distribution. Indeed, our linear regression framework that predicts MPCs using observable characteristics delivers an adjusted R^2 of 6%, as shown in Table 3, and the explanatory power is even lower for nondurables. Unlike previous studies, we obtain a statistical measure of the portion of the variance in the MPC distribution explained by observable characteristics through the R^2. Technically, the reported R^2 is a lower bound on the true R^2 due to measurement error in the estimated MPCs. We discuss this issue in detail in Supplement A.6 and propose a back-of-the-envelope adjustment to the R^2 to account for measurement error in recovering the MPCs in our clustering approach. The correction increases the R^2 for total expenditures to 8%, which still indicates that only a small fraction of the MPC heterogeneity can be explained by observables.

A low R^2 could also be partly explained by non-linear relationships that are either difficult to parametrize or simply not captured by variables in our dataset.[29] For example, the CEX contains only sparsely populated information on wealth. In Supplement A.3, we report the relationship between the MPC and liquid wealth, aware of the potential nonresponse bias highlighted by Parker et al. (2013); we cannot do the same for total wealth, though, given the lack of reliable data. While such "unobserved" characteristics could potentially explain some variation in MPCs, our results strongly suggest the presence of latent drivers, since some observables we do consider may proxy for such unobserved characteristics. The finding that relatively little MPC variation can be explained by observable characteristics is not only useful for disciplining heterogeneous agent models, but is also informative about the degree to which fiscal policy can target high MPC households, as we discuss later.

[29] Our results are robust to different sets of explanatory variables, including an array of predictors selected by a linear Lasso.

Table 4: MPC heterogeneity: full vs. observable distribution

	(1) GMLR	(2) age	(3) total income	(4) salary	(5) non-salary	(6) APC
g = 1	0.04	0.13	0.26	0.25	0.47**	0.09
	(0.03)	(0.21)	(0.18)	(0.23)	(0.20)	(0.20)
g = 2	0.23***	0.23	0.18	0.31	0.34	0.21
	(0.07)	(0.22)	(0.23)	(0.23)	(0.31)	(0.21)
g = 3	1.33***	0.34*	0.45	0.36	0.27	0.06
	(0.42)	(0.19)	(0.27)	(0.30)	(0.23)	(0.26)

Notes: Table 4 reports estimated total expenditure MPCs for different groupings of households. Robust standard errors are reported in parentheses. In the first column, we repeat the results from our baseline GMLR estimation, which were shown in Table 1. In the other columns, we report MPCs obtained by estimating Equation (3) using terciles of age, total income, salary and non-salary components, and APC to determine $d_{jg} = 1[j \in g]$, respectively. Groups are ordered from the lowest to the highest MPCs in the first column and by terile in the other columns.

A comparison with interacted regression models

We next conduct a comparison of our results with those arising from the approach typically taken in the literature. Previous papers interact the rebate with group membership dummies constructed by partitioning the sample based on some household observable characteristic. We report the MPCs recovered using our estimated group membership function against those the econometrician would recover by splitting the sample into groups using terciles of commonly-studied observable characteristics.[30] Table 4 displays the results, starting with the MPCs using our probabilistic membership function, λ_{0g}; the remaining columns use groups based on terciles of age, total income, salary and non-salary income, and the APC, in tercile order. MPCs mildly increase with age. Importantly, if a researcher used only age to characterize the extent of MPC heterogeneity, she would obtain estimates between 13 and 34 cents, much narrower than the range we uncover. Splitting by either total income or the APC, which we show above to be the most robust drivers of MPC heterogeneity, would allow a researcher to uncover higher MPCs, but never above 0.5, and would thus deliver a much lower degree of heterogeneity than using our approach. These conclusions are broadly unchanged if we interact the rebate by income components. Therefore, the existing literature, by splitting on individual observables that are likely noisy in practice, and correlated with only a portion of MPC heterogeneity, would under-estimate the true extent of MPC heterogeneity. Moreover, the MPCs estimated with our approach are much more precisely estimated, as shown by

[30]Formally, we estimate $\Delta C_j = \beta' \omega_j + \sum_{g=1}^{G} d_{jg} \left(\sum_{l=0}^{2} \iota_{jl} \lambda_{lg} + \alpha_g \right) + \epsilon_j$, in which $d_{jg} = \mathbf{1}[j \in g]$ is defined by terciles of a certain characteristic such as age. Table 4 reports estimates for λ_{0g}.

the smaller p-values. These results corroborate earlier statements that our approach may deliver improvements in statistical power.

Implications for Policy and Consumption-Savings Models

From a policy perspective, the results in this section have two implications. Only three observable household traits are robustly correlated with the MPC in a statistically significant manner: salary income, non-salary income, and the average propensity to consume. Hence, the first implication is that fiscal authorities might consider targeting those households who receive relatively more income as recipients of lump-sum transfers, in an attempt to maximize the effect on aggregate consumption. While we cannot speak to the MPCs of higher earners ineligible for the rebate in this natural experiment, among the eligible population in this experiment, this high-income group might include either top earners, households receiving high flows of business or financial income, or households with large social security payouts. This type of policy of course poses a potential tradeoff between the stimulus and relief/insurance motives for lump-sum transfers, since it suggests targeting higher-income households who are less in need of insurance. However, this tension is consistent with the empirical finding that low-income households are more likely to use stimulus checks to pay down debt, both in 2008 (Shapiro and Slemrod (2009)) and in 2020 (Armantier et al. (2020)). Therefore, whom the government should target may depend on the primary goal of the program: stimulus or insurance. Koşar et al. (2023) discuss this tradeoff in detail.

Second, the finding of a small R^2 in the regression of the MPC on observable characteristics suggests that attempts to target transfers based on factors observable by policymakers will ultimately exploit only a small fraction of the variation in households' MPCs. This means that feasible targeted transfers can harness only a small share of the possible gains in terms of consumption responses that would otherwise be available if policymakers could observe the identity of high MPC households directly, or if such MPCs were more strongly associated with observable characteristics.

The empirical results we have presented offer a new set of testable implications for consumption savings models. First, rather than simply targeting a wealth distribution, we provide a target distribution for the MPC itself. Second, our estimates offer guidance on how much of the variation in MPCs should be tied to observables like income. Finally, our estimates imply that MPCs should be positively correlated with the APC and income. Preference heterogeneity can generate substantial MPC heterogeneity that is by definition unexplained by observables; however, it needs to be quite rich in order to generate a positive correlation between MPCs and APCs, and cannot readily produce MPCs that

Figure 4: Predicting the longer-run effect of the tax rebate

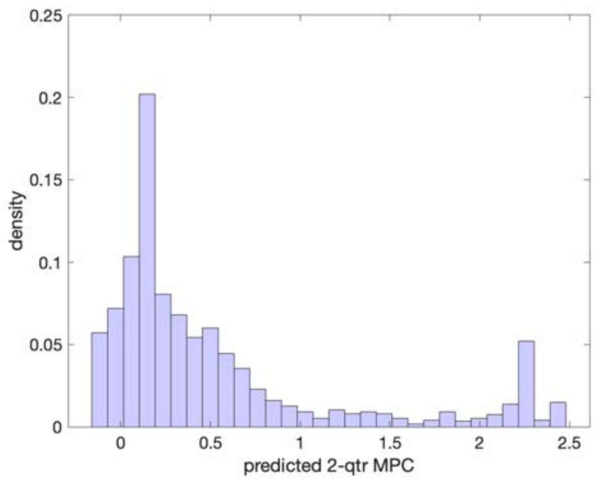

Notes: The histogram plots the individual cumulative MPCs constructed using the GMLR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}. For each household, we compute the weighted contemporaneous and lagged MPCs across groups and plot the total response as discussed in the text.

positively correlate with income. Alternative mechanisms, such as debt repayment or bounded rationality, appearing promising options to replicate this empirical finding.

4.4 The longer-run effects of the 2008 ESA

Finally, we analyze the household-level long-run spending effects of the 2008 tax rebates, using the dynamic MPCs found in λ_{1g} and λ_{2g}.[31] These long-run effects are also studied by Parker et al. (2013). In our estimation, we do not force group membership for household i to be fixed across t; even if individuals' preferences are constant, the MPC may be time-varying, due, for instance, to changes in state variables such as income and wealth.[32] Hence, we use the expected group-level MPC associated with household i at time t and observation (i, t)-specific posterior weights to construct each household's 2-quarter total expected effect of the rebate as $2\tilde{\lambda}_{0,j} + \tilde{\lambda}_{1,j}$.

Relative to the baseline results depicted in Figure 1, the distribution depicted in Figure 4 spreads out, with some households having a total effect near zero, but with many cumulated effects being larger than responses within the quarter. Moreover, all predicted

[31] We focus on the 2-quarter cumulative effect since there are too few observations present in the sample with sufficiently populated observable characteristics to reliably estimate projections on observables using 3-quarter cumulative effects. We show the 3-quarter distribution in Appendix A.2.

[32] Even in the homogenous case, λ_0 can be different from λ_1 because they measure two different objects; the coefficient on the lagged rebate value is an inter-temporal MPC which can be different from the contemporaneous MPC. See Auclert et al. (2018) for a theoretical discussion of intertemporal MPCs.

household lagged responses are negative, implying that households consume a smaller fraction of the rebate in the second (and third) period after receipt than in the first (since a value of zero indicates a constant consumption response), or mean reversion in consumption changes.

Finally, in Table 5 we show that our previous analysis regarding the drivers of MPC heterogeneity is confirmed when looking at longer-run spending responses. Top salary earners, as well as households with the highest non-salary income, have the highest longer-run marginal propensities to consume. This set of results on cumulative MPCs suggest that the spending effects of the rebate are persistent for most households. Moreover, the relationship between the MPC and its determinants does not appear to be the result of short-lived effects that could be erased by intertemporal substitution, at least in the first two quarters following rebate receipt.[33] In addition, the R^2 remains low, suggesting that unobserved heterogeneity remains important at this longer horizon.

5 Conclusion

We exploit a flexible clustering regression to uncover the unconditional distribution of the marginal propensity to consume. Our strategy improves on existing approaches by recovering the unconditional (marginal) distribution of MPCs and not simply estimating how the MPC co-varies with selected observable characteristics. Applying this methodology to consumption data following the 2008 Economic Stimulus Payments, households display a considerable degree of heterogeneity in their MPCs. A non-negligible share of households spent the checks in their entirety, and all households spent at least 8% of the rebate within one quarter, although this lower bound is subject to some statistical uncertainty. Nondurable consumption is also characterized by a lower bound that is significantly larger than zero, while durable consumption features two distinct groups with MPCs close to zero and one.

Examining which observables predict different portions of the MPC distribution, we obtain various statistically significant relationships with the MPC, but only those associated with income (both salary and non-salary components) and the APC survive the inclusion of additional drivers. These results hinge on the fact that our approach proves statistically more powerful than existing methodologies. Moreover, the R^2 from such regressions is a natural measure of the share of the unconditional MPC heterogeneity

[33] We found in Section 4.3 that marital and homeowner status were mildly significant predictors of higher 1-quarter MPCs. These relationships are no longer statistically significant when considering 2-quarter MPCs, although they preserve the same sign.

Table 5: Explanatory variables of longer-horizon MPCs

	(1) 1-qtr MPC	(2) 2-qtr MPC
Zero salary	-0.099***	-0.167**
	(0.037)	(0.071)
Salary: middle tercile	0.163***	0.299***
	(0.033)	(0.064)
Salary: top tercile	0.231***	0.412***
	(0.042)	(0.083)
Zero non-salary	0.033	0.079
	(0.030)	(0.060)
Non-salary: middle tercile	0.053	0.109
	(0.039)	(0.076)
Non-salary: top tercile	0.212***	0.383***
	(0.048)	(0.093)
APC	0.227***	0.435***
	(0.027)	(0.055)
adj. R^2	0.143	0.123

Notes: The table reports posterior-weighted predicted MPCs for total expenditures. For both columns, we consider the sub-sample of households that are observed in the quarter of rebate receipt and in the following quarter. We also control for, but do not report, five education dummies, the number of children, age and age squared, a married dummy, a homeowner dummy, a mortgage dummy, and mortgage interest to income ratio. Robust standard errors are in parentheses. *, **, and *** denote significance at the 10, 5, and 1% levels respectively.

that can be explained by observables. Since observable characteristics explain a minor portion of the estimated MPC heterogeneity, we posit that other latent factors, such as preference heterogeneity, are likely important in determining marginal propensities to consume. Taken together our results provide a range of facts useful for disciplining an emerging literature of macroeconomic models of consumption as well as significant policy implications, particularly for the targeting of transfers.

Finally, two caveats help to highlight possible avenues for future work. First, we estimate the MPC distribution using a single cross-section of data during a recession; if an individual's MPC is a function of the aggregate state, extrapolating our estimates requires caution. Second, because our empirical setting is one in which individuals only experience positive transitory shocks, we cannot speak to income losses, to which households may respond differently (Fuster et al. (2018)). However, clustering approaches like the one we use can easily be applied to other datasets with suitably identified transitory income shocks, making comparisons straightforward. We leave such exercises for future work.

References

AGUIAR, M., C. BOAR, AND M. BILS (2019): "Who Are the Hand-to-Mouth?" in *2019 Meeting Papers*, Society for Economic Dynamics, 525.

ALAN, S., M. BROWNING, AND M. EJRNAES (2018): "Income and Consumption: A Micro Semistructural Analysis with Pervasive Heterogeneity," *Journal of Political Economy*, 126, 1827–1864.

ARMANTIER, O., L. GOLDMAN, G. KOŞAR, J. LU, R. POMERANTZ, W. VAN DER KLAAUW, ET AL. (2020): "How Have Households Used Their Stimulus Payments and How Would They Spend the Next?" Tech. rep., Federal Reserve Bank of New York.

AUCLERT, A. (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109, 2333–67.

AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2018): "The Intertemporal Keynesian Cross," Working Paper 25020, National Bureau of Economic Research.

BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption inequality and partial insurance," *American Economic Review*, 98, 1887–1921.

BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): "A Distributional Framework for Matched Employer Employee Data," *Econometrica*, 87, 699–739.

BONHOMME, S. AND E. MANRESA (2015): "Grouped Patterns of Heterogeneity in Panel Data," *Econometrica*, 83, 1147–1184.

BORUSYAK, K., X. JARAVEL, AND J. SPIESS (2024): "Revisiting Event-Study Designs: Robust and Efficient Estimation," *The Review of Economic Studies*, rdae007.

BOUTROS, M. (2022): "Windfall income shocks with finite planning horizons," Tech. rep., Bank of Canada.

BRODA, C. AND J. A. PARKER (2014): "The economic stimulus payments of 2008 and the aggregate demand for consumption," *Journal of Monetary Economics*, 68, S20–S36.

CARROLL, C., J. SLACALEK, K. TOKUOKA, AND M. N. WHITE (2017): "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 8, 977–1020.

CARROLL, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 23, 61–156.

CELEUX, G., S. FRUEWIRTH-SCHNATTER, AND C. ROBERT (2018): "Model Selection for Mixture Models - Perspectives and Strategies," *arXiv: Methodology*.

CLOYNE, J., C. FERREIRA, AND P. SURICO (2020): "Monetary policy when households have debt: new evidence on the transmission mechanism," *The Review of Economic Studies*, 87, 102–129.

COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): "Innocent Bystanders? Monetary policy and inequality," *Journal of Monetary Economics*, 88, 70–89.

CRAWLEY, E. AND A. KUCHLER (2018): "Consumption Heterogeneity: Micro Drivers and Macro Implications," *Danish National Bank Working Paper 129*.

DEMPSTER, A. P., N. M. LAIRD, AND D. B. RUBIN (1977): "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society. Series B (Methodological)*, 39, 1–38.

DESARBO, W. AND W. CRON (1988): "A Conditional Mixture Maximum Likelihood Methodology for Clusterwise Linear Regression," *Journal of Classification*, 5, 249–282.

DYNAN, K. E., J. SKINNER, AND S. P. ZELDES (2004): "Do the Rich Save More?" *Journal of Political Economy*, 112, 397–444.

FAGERENG, A., M. B. HOLM, AND G. J. NATVIK (2016): "MPC Heterogeneity and Household Balance Sheets," Discussion Papers 852, Statistics Norway, Research Department.

FRIEDMAN, M. (1957): *A Theory of the Consumption Function*, Princeton University Press.

FUSTER, A., G. KAPLAN, AND B. ZAFAR (2018): "What Would You Do With $500? Spending Responses to Gains, Losses, News, and Loans," Staff Reports 843, Federal Reserve Bank of New York.

GELMAN, M. (2019): "The Self-Constrained Hand to Mouth," .

GUDICHA, D. W. AND J. K. VERMUNT (2013): "Mixture Model Clustering with Covariates Using Adjusted Three-Step Approaches," in *Algorithms from and for Nature and Life*, ed. by B. Lausen, D. Van den Poel, and A. Ultsch, Cham: Springer International Publishing, 87–94.

ILUT, C. L. AND R. VALCHEV (2020): "Economic agents as imperfect problem solvers," Tech. rep., National Bureau of Economic Research.

JAPPELLI, T. AND L. PISTAFERRI (2014): "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics*, 6, 107–36.

JOHNSON, D. S., J. A. PARKER, AND N. S. SOULELES (2006): "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 96, 1589–1610.

JONES, P. AND G. J. MCLACHLAN (1992): "Fitting Finite Mixture Models in a Regression Context," *Australian Journal of Statistics*, 34, 233–240.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

KAPLAN, G. AND G. L. VIOLANTE (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82, 1199–1239.

——— (2021): "The Marginal Propensity to Consume in Heterogeneous Agents Models," Tech. rep.

KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 45, 77–153.

KOŞAR, G., D. MELCANGI, L. PILOSSOPH, AND D. G. WICZER (2023): "Stimulus through Insurance: The Marginal Propensity to Repay Debt," *FRB of New York Staff Report*.

KUENG, L. (2018): "Excess sensitivity of high-income consumers," *The Quarterly Journal of Economics*, 133, 1693–1751.

LAIBSON, D., P. MAXTED, AND B. MOLL (2022): "A Simple Mapping from MPCs to MPXs," Tech. rep., National Bureau of Economic Research.

LEWIS, D. J., D. MELCANGI, L. PILOSSOPH, AND A. TONER-RODGERS (2023): "Approximating grouped fixed effects estimation via fuzzy clustering regression," *Journal of Applied Econometrics*, 38, 1077–1084.

LUSARDI, A. (1996): "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets," *Journal of Business & Economic Statistics*, 14, 81–90.

MAJESKE, K. D., T. LYNCH-CARIS, AND J. BRELIN-FORNARI (2010): "Quantifying R 2 bias in the presence of measurement error," *Journal of Applied Statistics*, 37, 667–677.

MCLACHLAN, G. AND D. PEEL (2004): *Finite Mixture Models*, Wiley Series in Probability and Statistics, Wiley.

MCLACHLAN, G. J. AND K. E. BASFORD (1988): *Mixture models. Inference and applications* *to clustering*.

MIRANDA-PINTO, J., D. MURPHY, K. J. WALSH, AND E. R. YOUNG (2020): "A model of expenditure shocks," .

MISRA, K. AND P. SURICO (2014): "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs," *American Economic Journal: Macroeconomics*, 6, 84–106.

OLAFSSON, A. AND M. PAGEL (2018): "The liquid hand-to-mouth: Evidence from personal finance management software," *The Review of Financial Studies*, 31, 4398–4446.

ORCHARD, J., V. RAMEY, AND J. WIELAND (2023): "Using Macro Counterfactuals to Assess Plausibility: An Illustration using the 2001 Rebate MPCs," NBER Working Papers 31808, National Bureau of Economic Research, Inc.

PARKER, J. A. (2017): "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment," *American Economic Journal: Macroeconomics*, 9, 153–83.

PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103, 2530–53.

QUANDT, R. E. (1972): "A New Approach to Estimating Switching Regressions," *Journal of the American Statistical Association*, 67, 306–310.

SAHM, C. R., M. D. SHAPIRO, AND J. SLEMROD (2010): "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications," in *Tax Policy and the Economy, Volume 24*, National Bureau of Economic Research, Inc, NBER Chapters, 69–110.

SHAPIRO, M. D. AND J. SLEMROD (2009): "Did the 2008 Tax Rebates Stimulate Spending?" *American Economic Review*, 99, 374–79.

STRAUB, L. (2017): "Consumption, Savings, and the Distribution of Permanent Income," Working Paper.

WONG, A. (2021): "Refinancing and the transmission of monetary policy to consumption," .

Supplemental Appendix to "Latent Heterogeneity in the Marginal Propensity to Consume"

Daniel Lewis Davide Melcangi Laura Pilossoph

Abstract

This supplement has two parts. Section A contains additional empirical results and robustness checks. These include alternative specifications, further results on the MPC distribution and the correlation of MPCs and observables, the impact of sample trimming, group membership as a function of covariates, and adjustments to the R^2 for measurement error. Section B contains a discussion of the differences between our approach and quantile regression.

A Empirical Appendix

In this section we report additional empirical results. First, we allow for heterogeneity in all coefficients in Equation (3), which leaves the estimated distribution of MPCs largely unchanged. Second, we report additional details on the MPC distribution for several specifications. Third, we document further results on the relationship between MPC heterogeneity and observable characteristics. Fourth, we report results for a sample that includes individuals with far left- and right-tail expenditure changes. Fifth, we allow group membership to be an explicit function of observable characteristics. Finally, we describe how we adjust the R^2 in Table 3 to account for measurement error.

Table A.1: The MPC distribution: heterogeneous coefficients on controls

	Baseline		Heterogenous controls	
	(I)		(II)	
	λ_{0g}	π_g	λ_{0g}	π_g
$g=1$	0.04	0.30	0.11	0.29
$g=2$	0.23	0.48	0.28	0.48
$g=3$	1.33	0.22	1.42	0.23

Notes: The table shows estimates for group-specific MPCs, λ_{0g}, and population shares, π_g. Column (I) is the baseline shown in Table 1. In column (II) coefficients on all controls ω_j are group specific. In this specification, the BIC also selects $G=3$.

A.1 Heterogeneous coefficients on controls

Our baseline specification assumes common coefficients on time dummies and household-level controls. It is natural to wonder if there is also a role for heterogeneity with respect to those covariates. In Equation (3), this amounts to interacting ω_j with the group dummies and allowing β to vary across group. In Table A.1 we compare the estimation results in this alternative specification to our baseline estimates. The estimates lie within a very similar range. The smallest group-specific MPC increases from 4 to 11 cents, and the highest from 1.33 to 1.42. The similarity of the two distributions is also evident when looking at Figure A.1, which plots the histograms of predicted posterior-weighted MPCs. As previously mentioned, we observe only a slight shift to the right. The average MPC differs only by 7 cents in the two specifications. Very similar findings are obtained if we allow heterogeneity in all household characteristics, but maintain homogeneous coefficients on time dummies: in this case, the distribution of predicted MPCs is almost identical to our baseline.

Table A.2: Estimated MPCs with different number of groups G

	$G = 1$	$G = 2$	$G = 3$
λ_{01}		0.09	0.04
λ_{02}	0.24		0.23
λ_{03}		1.14	1.33
π_1		0.67	0.22
π_2	1		0.48
π_3		0.33	0.30

Notes: The table shows estimates for group-specific MPCs, λ_{0g}, and population shares, π_g, estimated with a different number of groups G. Baseline sample and specification for total expenditures.

Figure A.1: Predicting individual MPCs: heterogeneous coefficients on controls

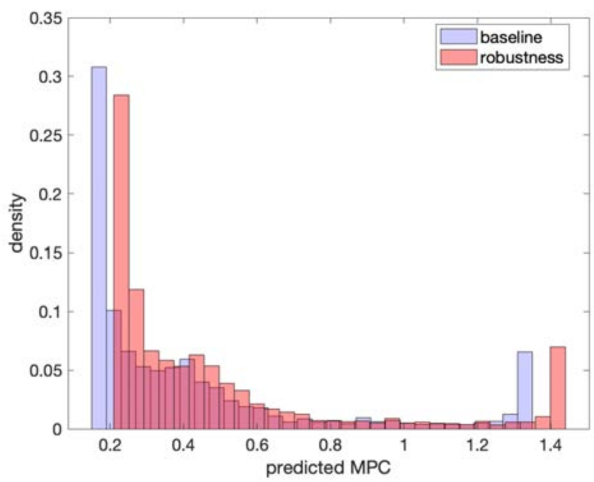

Notes: The histograms plot the individual MPCs constructed using the GMLR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}. Light blue bars refer to the baseline total expenditures specification as in Figure 1. The red bars refer to a specification that allows coefficients on all controls ω_j to be group specific. In this specification, the BIC selects $G = 3$.

A.2 The MPC distribution: additional results

We next show that the distribution of marginal propensities to consume is robust to the choice of G. As noted in the text, the BIC has a kink at $G = 2$ and meaningfully flattens at $G = 3$, and hence we compare results for these groups. Table A.2 reports the estimates, λ_{0g} and π_g, that underlie the MPC distribution. Adding more groups clearly implies the appearance of new MPC estimates. However, the range of MPCs is similar across G, although the highest MPC is 19 cents larger in our baseline.

The posterior predicted MPCs, plotted in Figure A.2, highlight similarities and differences depending on the number of groups G. As expected, we see a more pronounced bimodality with 2 groups. However, the average MPC is quite similar for $G = 2$ and

Figure A.2: Predicting individual MPCs out of the tax rebate: different G

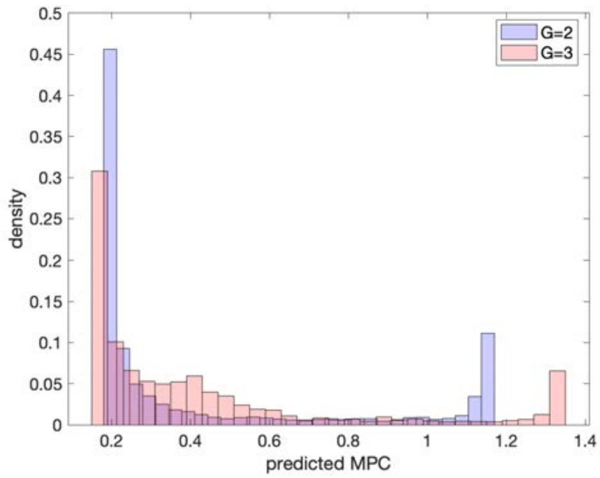

***Notes*:** As in Figure 1, the figure plots a histogram of the individual MPCs constructed using the GMLR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}. Results are for total expenditures and we only plot households that received the rebate. Light blue bars show results for $G = 2$, while red bars for our baseline results with $G = 3$.

$G = 3$, around 45 cents per rebate dollar. In both distributions, about 75% of households consume less than half of the rebate.

In Figure A.3 we predict individuals' MPCs minimizing the classification error. To do so, we assign each individual the estimated MPC λ_{0g} associated with their highest posterior probability group (i.e., $\lambda^*_{j,0g} = \lambda_{0g^*_j}$, where $g^*_j = \arg\max_g \gamma_{jg}$). In line with the unconditional MPC distribution shown in Table 1, most households are predicted to spend less than 25 cents for each rebate dollar, but about 15% of the rebate recipients spend more than the entirety of the rebate.

Figure A.3: Predicting individual MPCs out of the tax rebate: highest probability MPCs

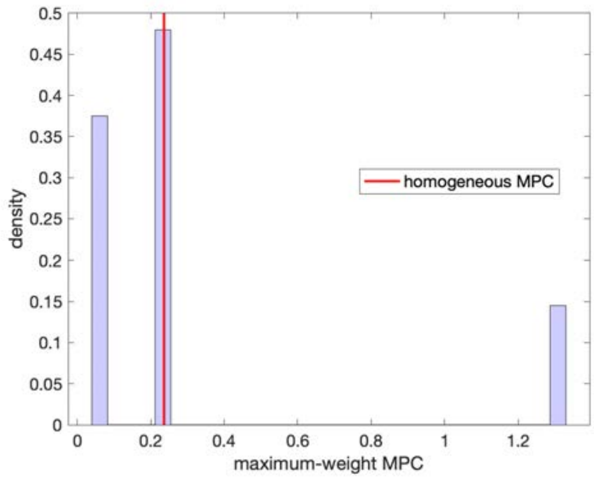

Notes: The histogram (light blue bars) plots the individual MPCs constructed using the GMLR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}. We assign individuals to the group – and the associated λ_{0g} – associated with the maximum γ_{jg} across groups. Results are for total expenditures and we plot only households that received the rebate. The sample is defined as in the text. The BIC selects $G = 3$. The homogeneous MPC (red vertical line) is estimated imposing $G = 1$ in our baseline specification.

Figure A.4 shows the stability of the distribution of predicted most likely MPCs, repeating the bootstrapping exercise of Section 4.1. Also in this setting, the median CDF across boostraps, plotted in dash-dotted blue, reassuringly tracks the estimated distribution, in solid black. There is also limited variation across bootstraps, especially for higher MPCs.

Table A.3 reports the statistical significance of the point estimates for the MPCs, in the baseline specification under GMLR. In the left panel, we make use of the analytical formulas outlined in Section 3 to compute Wald tests of pairwise equality across MPCs (accounting for uncertainty in individual weights). The right panel shows the same tests, taking the estimated weights as given; we report these results to parallel tests typically conducted in the literature, where group membership is taken as known (based on assumed observable relationships). In Table A.4, we repeat the same analysis for nondurables and durables. For compactness, we only show tests accounting for uncertainty in the weights.

In Figure A.5 we show how the posterior predicted MPCs, constructed using γ_{jg}, vary across individuals, for nondurable and durable expenditures.

Figure A.4: Bootstrapped individual MPCs out of the tax rebate: highest probability MPCs

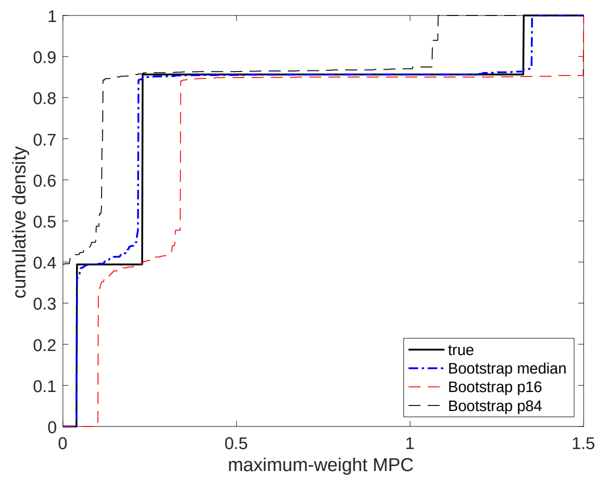

Notes: The black solid line plots the CDF of individual MPCs for total expenditures, calculated as shown in Figure A.3. In each bootstrap, we assign individuals to the group - and the associated λ_{0g} - associated with the maximum γ_{jg} across groups. The blue dash-dotted line shows the median CDF of the estimated distribution of MPCs, across 250 bootstraps with replacement. The dashed red and black lines correspond to the centered 68% confidence interval.

Figure A.5: Predicting individual MPCs out of the tax rebate: nondurables and durables

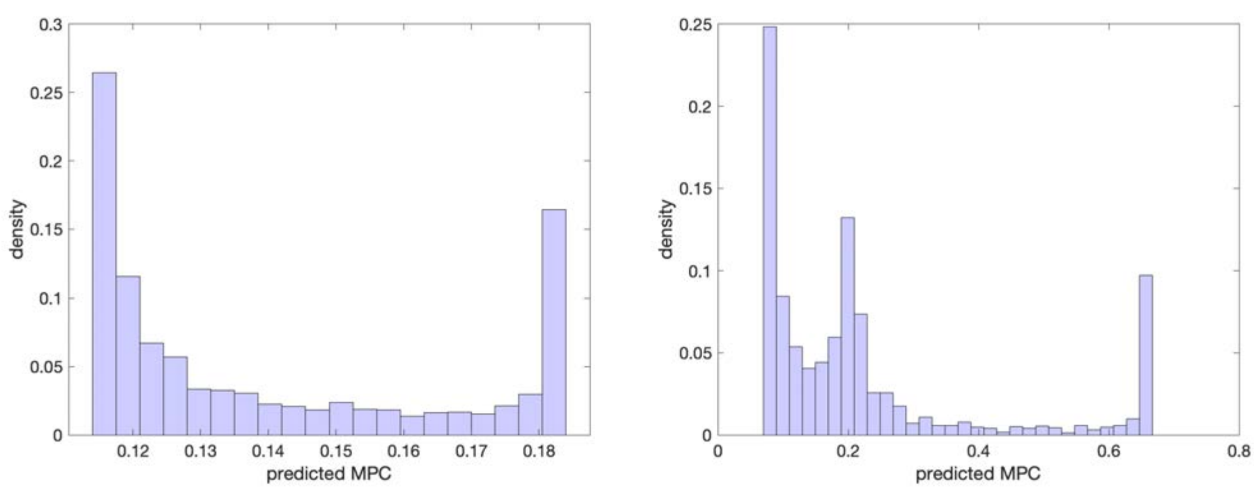

(a) Nondurables (b) Durables

***Notes*:** The histograms (light blue bars) plot the individual MPCs constructed using the GMLR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}. The left panel reports nondurable and the right durable expenditures, defined as in Parker et al. (2013), among households that received the rebate. The sample is defined as in the text. The homogeneous MPCs ($G = 1$) are 0.13 for nondurables and 0.05 for durables. For each household we compute the posterior-weighted MPC across groups. For nondurables the BIC selects $G = 2$ and for durables $G = 3$.

We also show that nondurable predicted MPCs are meaningfully above zero across bootstrap samples, parallel to the exercise of Figure 2. Results are reported below in Figure A.6a. We also perform the same bootstrap exercise for durables, reported in Figure

Table A.3: Statistical tests of MPCs: total expenditures

(a) Analytical standard errors

MPC	0.04	0.23	1.33
0.04	0.32		
	(0.57)		
0.23	1.55	3.45	
	(0.21)	(0.06)	
1.33	0.32	1.55	7.59
	(0.01)	(0.03)	(0.00)

(b) Conditional on weights

MPC	0.04	0.23	1.33
0.04	2.48		
	(0.12)		
0.23	6.95	10.68	
	(0.01)	(0.00)	
1.33	9.27	6.59	9.88
	(0.00)	(0.01)	(0.00)

Notes: The table shows $F-$statistics for pairwise tests of equality of MPCs (the diagonals show tests of equality with zero) for the baseline total expenditures specification estimated under GMLR. In the left panel, standard errors account for all estimation uncertainty. In the right panel, GMLR weights are taken as given, to parallel the way that group assignment is taken as known in the existing literature. $p-$values are reported in parentheses.

Table A.4: Test for MPC equality: nondurables and durables

(a) Nondurables

MPC	0.09	0.18
0.09	4.88	
	(0.03)	
0.18	1.11	6.19
	(0.29)	(0.01)

(b) Durables

MPC	0.03	0.15	0.67
0.03	1.69		
	(0.19)		
0.15	2.81	5.36	
	(0.09)	(0.02)	
0.67	1.69	2.81	2.22
	(0.14)	(0.24)	(0.12)

Notes: The table shows $F-$statistics for pairwise tests of equality of MPCs (the diagonals show tests of equality with zero) for the baseline total expenditures specification estimated under GMLR. $p-$values are reported in parentheses.

A.6b.

Figure A.6: Bootstrapped distribution of MPCs out of the tax rebate

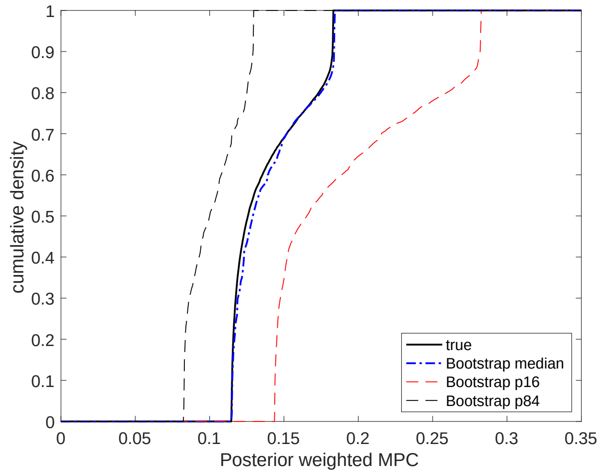

(a) Nondurable MPCs

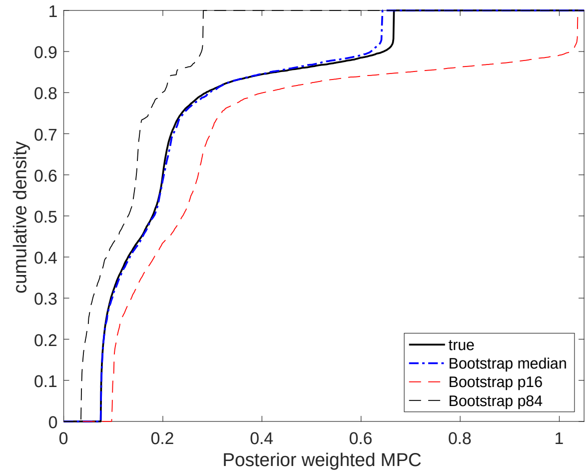

(b) Durable MPCs

Notes: The black solid lines plot CDF of the distribution of individual MPCs, calculated as shown in Figure A.5. We predict individual responses using posterior-weighted MPCs. The blue dash-dotted lines show the median CDF of posterior-weighted MPCs across 250 bootstraps. The dashed black and red lines denote the centered 68% confidence interval.

The results shown in Section 4 follow the approach by Borusyak et al. (2024) and thus include households that never receive the rebate in the control group. An alternative approach is to exclude these households from the sample, and implement a IV strategy using only the contemporaneous rebate receipt. Many papers in the literature have estimated a homogeneous MPC using this approach, including Parker et al. (2013). We estimate our model using this specification and sample and repeat the same predictive exercise for households' posterior MPCs of Figure 1. The results are shown in Figure A.7. This approach delivers higher MPCs, even when looking at the homogeneous treatment effect, a result that has been highlighted by Borusyak et al. (2024) and Orchard et al. (2023). Nevertheless, up to this level shift, MPCs are distributed in a similar fashion to our baseline results. Looking at the estimated coefficients, the lowest λ_{0g} is slightly higher, from 4 to 8 cents, whereas the highest λ_{0g} increases from 1.33 to 1.44. Importantly, the correlations with household observable characteristics that we show in Section 4.3 are little changed.

Figure A.8: Predicting the longer-run effect of the tax rebate

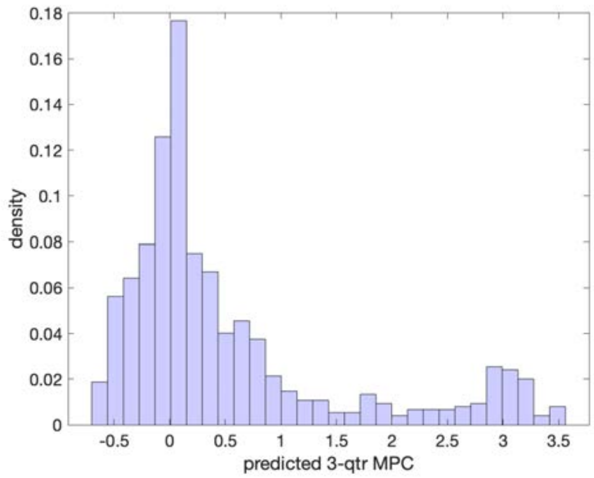

Notes: The histogram plots the individual MPCs constructed using the GMLR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}. For each household, we compute the weighted contemporaneous and lagged MPCs across groups and plot the total response as discussed in the text.

Figure A.7: Predicting individual MPCs out of the tax rebate: IV with rebate recipients only

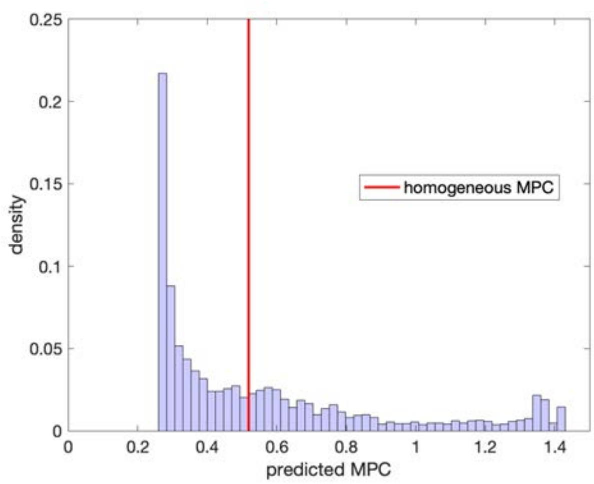

Notes: The estimation sample excludes households that never receive a rebate. The model estimated is (2), , where R_j is instrumented using the dummy for rebate receipt in a homogeneous first-stage regression where all other coefficients are restricted to be zero. The histogram (light blue bars) plots the individual MPCs constructed using the GMIVR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}. For each household we compute the posterior-weighted MPC across groups. Results are for total expenditures and we only plot households that received the rebate. The BIC selects 5 groups in this alternative specification.

In Figure A.8 we show the distribution of predicted MPCs, cumulated over 3 quarters, similarly to those in Section 4.4. Formally, the 3-period cumulative MPC is defined as $3\tilde{\lambda}_{0,j} + 2\tilde{\lambda}_{1,j} + \tilde{\lambda}_{2,j}$.

Figure A.9: The correlation of MPCs across consumption goods

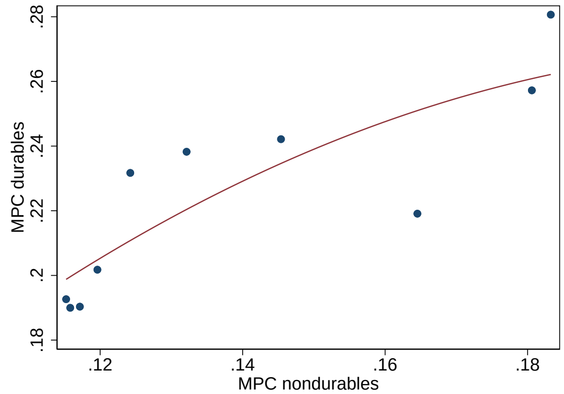

Notes: The blue dots display a binscatter of household posterior-weighted MPCs for durables against those for nondurables. Each dot shows the average posterior-weighted MPC for durable goods for each decile of the distribution of posterior-weighted MPCs for nondurable goods. The red line shows the quadratic fit.

A.3 What drives MPC heterogeneity: additional results

In this section we report additional results for the relationship between MPC heterogeneity and observable characteristics, for our baseline specification estimated with GMLR.

Figure A.9 graphically displays the correlation of posterior-weighted MPCs for durable and nondurable goods.

We have shown in Section 4.3 that some household characteristics individually correlate with the MPC distribution. Here, we investigate some of these relationships in greater detail, as well as discuss some additional findings. In Figure A.10 we explore whether there are non-monotonicities in the relationship between the MPC and income sources, considering finer quintile bins. The results are in line with those discussed in Section 4.3.

Figure A.10: MPCs by bins of income sources

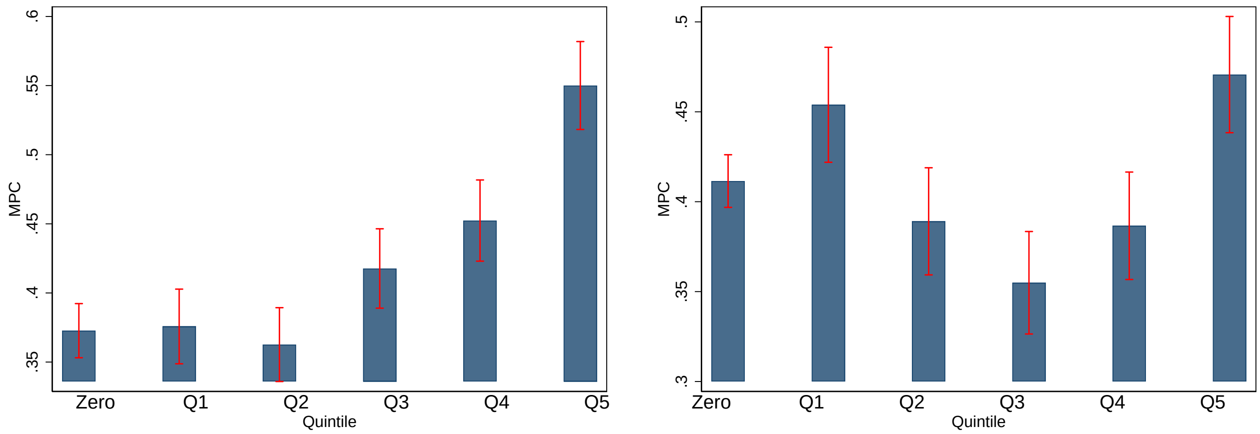

(a) Salary income

(b) Non-salary income

Notes: The bars in the left panel display the average MPC λ_{0g}, weighted by γ_{jg}, for total expenditures by quintiles of salary income and for households with zero salary income (leftmost bar). Quintiles are defined conditional on having strictly positive salary income. The sample is restricted *ex post* to households receiving a rebate, and results are for the baseline specification estimated with GMLR. Error bars show 90% confidence intervals for the average MPC within each group. The bars in the right panel repeat the same analysis for non-salary income. A similar relationship holds for both nondurables and durables.

We next explore the relationships between the MPC and age and liquid assets. Figure A.11a suggests a positive unconditional relationship between the MPC and liquid wealth. The pattern is similar when looking at the ratios of liquid assets to total income. In particular, the overall correlation between liquidity ratios and MPCs is small and typically insignificant, especially for nondurable MPCs. Turning to age, the relationship looks instead concave, as shown in Figure A.11b, but the differences are insignificant, except for the highest quintile, which consists of households whose heads are in retirement age (i.e., 65 years old or higher).

Figure A.11: Marginal propensities to consume: liquid wealth and age

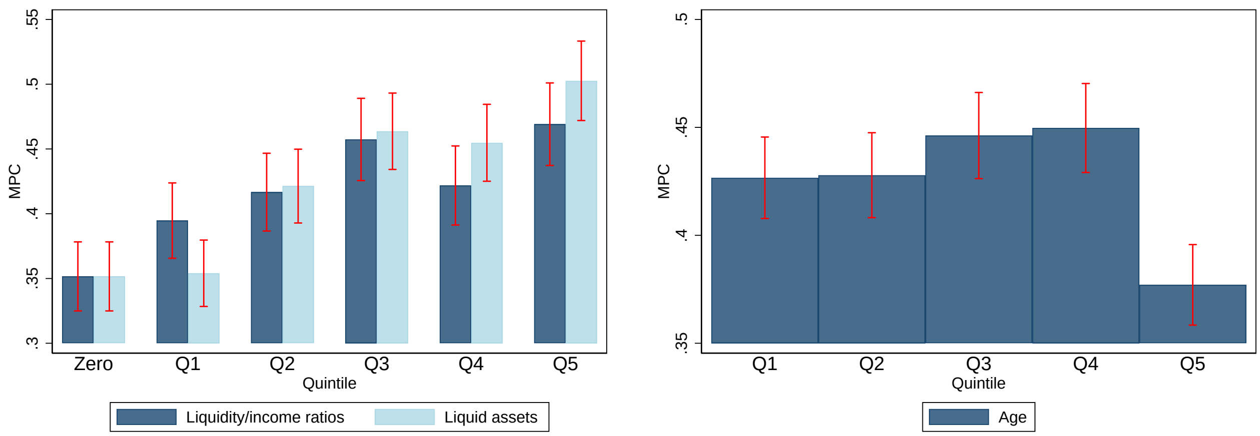

(a) Liquid wealth

(b) Age

Notes: In the left panel, the dark blue bars display the average MPC λ_{0g}, weighted by γ_{jg}, for total expenditures by quintiles of ratios of liquid assets to income and for households with zero liquid assets (leftmost bar). Quintiles are defined conditional on having strictly positive liquid assets. The sample is restricted ex-post to households receiving a rebate, and results are for the baseline specification estimated with GMLR. Error bars show 90% confidence intervals, computed using the weighted standard deviation of MPC within each group. Light-blue bars repeat the same analysis for liquid assets in levels. The right panel shows a similar bar chart for quintiles of the age of households' reference persons. Similar relationships hold for both durables and nondurables, but are weaker and even less significant for the latter.

Neither relationship is robust to the inclusion of a set of controls, as we further show in Table A.5. In the same table, we further confirm that the findings shown in Section 4.3 are robust to controlling for liquid wealth. Liquid assets to income ratios are insignificantly related to the MPCs, after controlling for the other observable characteristics. Coefficients on age, and its squared value, are statistically insignificant as in our baseline analysis, and this holds even if we include a separate dummy for retirement age, which is also insignificant.

Further additional relationships hold unconditionally, but then turn insignificant when controlling jointly for all predictors. For instance, MPCs monotonically increase in educational attainment of households' reference persons. We find that households that put money into a tax-deferred or tax-free educational savings plan have a significantly higher MPC. Households with male heads have slightly higher MPCs, but the relationship is significant only at the 10% level. There is little variation by race, with the only exception being households with Black reference persons displaying lower MPCs than households with white reference persons. Finally, married households have statistically larger MPCs (by 7 cents) even in our joint regression.

In Figure A.12, we show that "rich-spenders" (i.e. households with high APC and high total income) have high MPCs for nondurable and durable expenditures too.

Table A.5: Explanatory variables of MPCs: including liquid assets

	(1) Total	(2) Total	(3) Nondurables	(4) Nondurables	(5) Durables	(6) Durables
Zero liquid assets	-0.012	-0.013	-0.001	0.002	-0.017	-0.029**
	(0.025)	(0.029)	(0.003)	(0.004)	(0.013)	(0.015)
Liquid assets: middle tercile	0.115***	0.098***	0.004	0.000	0.062***	0.045***
	(0.024)	(0.030)	(0.002)	(0.003)	(0.012)	(0.015)
Liquid assets: top tercile	0.112***	0.040	0.006**	-0.000	0.063***	0.021
	(0.024)	(0.033)	(0.002)	(0.003)	(0.012)	(0.016)
Zero salary		-0.097**		-0.007*		-0.041**
		(0.040)		(0.004)		(0.019)
Salary: middle tercile		0.077**		0.010***		0.031*
		(0.034)		(0.004)		(0.017)
Salary: top tercile		0.171***		0.018***		0.088***
		(0.047)		(0.004)		(0.023)
Zero non-salary		0.014		0.001		-0.016
		(0.032)		(0.003)		(0.016)
Non-salary: middle tercile		0.085**		0.004		0.007
		(0.040)		(0.004)		(0.020)
Non-salary: top tercile		0.159***		0.009*		0.052**
		(0.046)		(0.005)		(0.024)
APC		0.185***		0.010***		0.089***
		(0.029)		(0.002)		(0.015)
Homeowner dummy		0.060*		0.008**		0.013
		(0.032)		(0.004)		(0.016)
Married dummy		0.074***		0.003		0.032**
		(0.026)		(0.003)		(0.013)
adj. R^2	0.013	0.063	0.003	0.027	0.021	0.080

Notes: Regressions are estimated on the sample that includes only households receiving a rebate in the current period. Robust standard errors are reported in parentheses. In column (5) we also control for five education dummies, a married dummy, number of children, age and age squared; these coefficients are not reported. For formatting reasons we also do not report coefficients on a mortgage dummy and mortgage interest to income ratios, which are both statistically insignificant across all columns. *, **, and *** denote significance at the 10, 5, and 1% levels respectively.

Figure A.12: The relationship between MPCs, APCs, and income

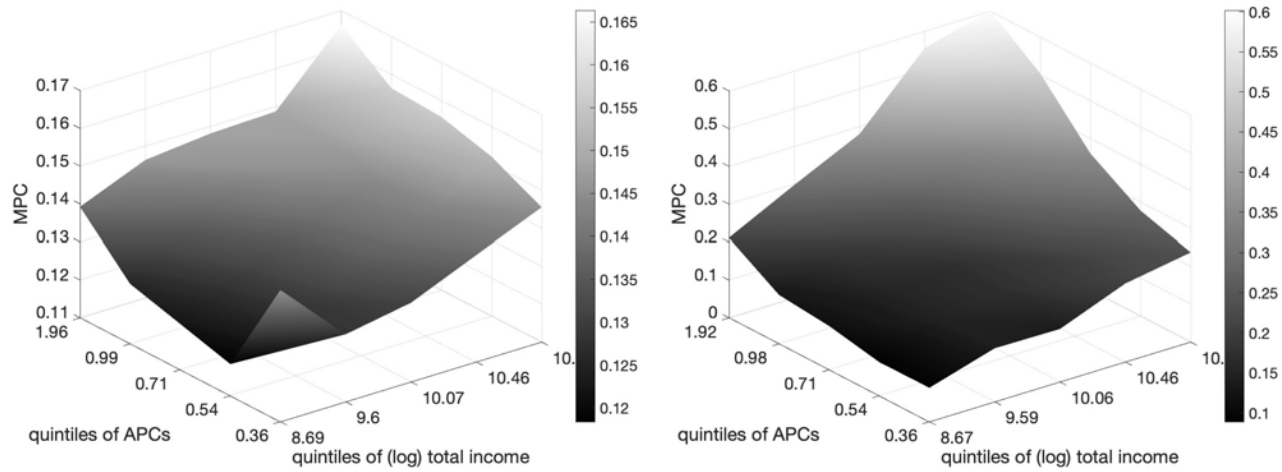

(a) Nondurable MPCs

(b) Durable MPCs

The surfaces plot average weighted MPCs for pairs of quintiles of APC and log total income. The left panel considers nondurable expenditures and the right durable expenditures. The color bar on the right represents the MPC.

Next, we show in Tables A.6 and A.7 how our results change if we consider MPCs for nondurables and durables, respectively. Most of the patterns are strikingly similar. One exception is that households with no salary income do not have statistically different MPCs for nondurable expenditures; they did in the joint regression for total expenditures. Moreover, the income effect is weaker, both economically and statistically, especially for nonsalary income. In addition, all housing variables are insignificant whereas the APC remains a very strong and significant predictor. Consistent with less significant relationship, the adjusted R^2 for nondurable MPCs is particularly low.

Finally, we show that the relationship between MPCs and observable characteristics is robust to the selection of the number of groups used in the GMLR estimation. In Table A.8 we repeat the analysis of column (5) of Table 3 for $G = 2$, where the BIC has the first kink. We report the results for total expenditure MPCs. The estimated effects, as well as their statistical significance, are barely changed across columns. In particular, the coefficients are also quantitatively stable. Moreover, the adjusted R^2 is also low when $G = 2$.

A.4 MPC results without sample trimming

We repeat our main results for a sample in which we do not drop the top and bottom 1.5% of consumption changes, a form of trimming that we adopted in our baseline specifications.

Table A.6: Explanatory variables of nondurable MPCs

	(1)	(2)	(3)	(4)	(5)
Mortgage interest/income	-0.002				-0.005
	(0.007)				(0.009)
Homeowner dummy	0.008***				0.006*
	(0.002)				(0.003)
Mortgage dummy	0.002				-0.000
	(0.002)				(0.003)
Zero salary		-0.001			-0.007*
		(0.002)			(0.003)
Salary: middle tercile		0.003			0.008***
		(0.002)			(0.003)
Salary: top tercile		0.011***			0.014***
		(0.002)			(0.004)
Zero non-salary			-0.000		0.001
			(0.002)		(0.003)
Non-salary: middle tercile			-0.003		0.004
			(0.003)		(0.003)
Non-salary: top tercile			0.002		0.010***
			(0.003)		(0.004)
APC				0.002*	0.010***
				(0.001)	(0.002)
adj. R^2	0.007	0.010	0.000	0.001	0.024

Notes: Regressions are estimated on the sample that includes only households receiving a rebate in the current period. Robust standard errors are reported in parentheses. In column (5) we also control for five education dummies, a married dummy, number of children, age and age squared; these coefficients are not reported. *, **, and *** denote significance at the 10, 5, and 1% levels respectively.

Table A.7: Explanatory variables of durable MPCs

	(1)	(2)	(3)	(4)	(5)
Mortgage interest/income	0.034				-0.031
	(0.036)				(0.052)
Homeowner dummy	0.042***				0.026**
	(0.011)				(0.013)
Mortgage dummy	0.022**				0.009
	(0.011)				(0.014)
Zero salary		0.005			-0.035**
		(0.011)			(0.015)
Salary: middle tercile		0.029**			0.052***
		(0.012)			(0.014)
Salary: top tercile		0.086***			0.104***
		(0.012)			(0.017)
Zero non-salary			-0.018		-0.004
			(0.011)		(0.014)
Non-salary: middle tercile			-0.040***		0.001
			(0.013)		(0.017)
Non-salary: top tercile			-0.011		0.063***
			(0.014)		(0.019)
APC				0.037***	0.087***
				(0.007)	(0.012)
adj. R^2	0.014	0.021	0.002	0.010	0.072

Notes: Regressions are estimated on the sample that includes only households receiving a rebate in the current period. Robust standard errors are reported in parentheses. In column (5) we also control for five education dummies, a married dummy, number of children, age and age squared; these coefficients are not reported. *, **, and *** denote significance at the 10, 5, and 1% levels respectively.

Table A.8: Explanatory variables of MPCs: robustness to G

	(1) G = 2	(2) G = 3
Zero salary	-0.082**	-0.087***
	(0.037)	(0.031)
Salary: middle tercile	0.112***	0.112***
	(0.031)	(0.027)
Salary: top tercile	0.181***	0.188***
	(0.039)	(0.035)
Zero non-salary	0.023	0.031
	(0.031)	(0.027)
Non-salary: middle tercile	0.042	0.049
	(0.037)	(0.033)
Non-salary: top tercile	0.158***	0.166***
	(0.044)	(0.038)
mortgage interest/income	-0.035	-0.112
	(0.107)	(0.095)
APC	0.183***	0.189***
	(0.024)	(0.023)
homeowner dummy	0.084***	0.076***
	(0.031)	(0.026)
dummy for mortgage	-0.025	-0.011
	(0.033)	(0.029)
married dummy	0.074***	0.074***
	(0.024)	(0.021)
adj. R^2	0.058	0.059

Notes: Regressions are estimated on the sample that includes only households receiving a rebate in the current period. Robust standard errors are reported in parentheses. We control for five education dummies, a married dummy, number of children, age and age squared; these coefficients are not reported. *, **, and *** denote significance at the 10, 5, and 1% levels respectively.

Table A.9: MPC heterogeneity: keeping tail expenditure changes

	λ_{01}	π_1	λ_{02}	π_2	λ_{03}	π_3
Estimate	0.02	0.38	0.36	0.46	2.96	0.16
Std. Err	(0.07)	(0.01)	(0.14)	(0.01)	(1.08)	(0.00)
Cond. Std. Err	[0.03]	-	[0.10]	-	[0.95]	-

Notes: The first row reports estimates for group-specific MPCs, λ_{0g}, and population shares, π_g, for the sample without trimming tail expenditure changes. Standard errors, reported in the second row in parentheses $(\cdot)$, account for all uncertainty;. conditional standard errors, reported in the third row in square brackets $[\cdot]$, instead take GMLR weights as given, to parallel the way that group assignment is imposed *ex ante* in the existing literature. The BIC selects $G = 3$ in the untrimmed data.

Figure A.13: Predicting individual MPCs: keeping tail expenditure changes

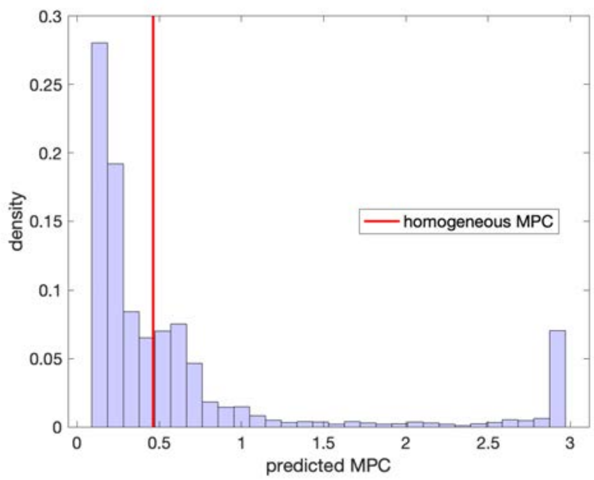

Notes: The histogram (light blue bars) plots the individual MPCs constructed using the GMLR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}, for the sample without trimming tail expenditure changes. For each household we compute the posterior-weighted MPC across the discrete group-specific MPCs. Results are for total expenditures and we only plot households that received the rebate. The sample is defined as in the text. The BIC selects $G = 3$ in the untrimmed data. The homogeneous MPC (red vertical line) is estimated imposing $G = 1$ to our baseline specification..

Table A.9 reports the points of the unconditional MPC distribution for total expenditures, $\{\lambda_{0g}\}$, the associated population shares, π_g, along with standard errors, as discussed in the main text for Table 1. The lowest MPC is broadly unchanged at 2 cents; in contrast, the largest λ_{0g} increases.

In Figure A.13 we plot the distribution of households' predicted MPCs. A non-negligible share of households are predicted to spend much more than the entirety of the rebate.

We now turn to nondurable and durable expenditures and repeat the same analysis shown in Section 4.2. Table A.10 shows our estimates and the corresponding standard errors. Including expenditure changes of large magnitude also increases the range of estimated MPCs for these subcategories. The lowest MPC for nondurables is little changed. Compared to our baseline result, 4% of households spend more than the entirety of the rebate in nondurable goods; this coefficient , however, is noisily estimated. Turning to

Table A.10: Nondurable and durable MPCs: keeping tail expenditure changes

(a) Nondurables

	λ_{0g}	π_g
$g = 1$	0.07	0.56
	(0.04)	(0.01)
	[0.02]	-
$g = 2$	0.29	0.40
	(0.10)	(0.01)
	[0.08]	-
$g = 3$	1.63	0.04
	(1.57)	(0.00)
	[0.88]	-

(b) Durables

	λ_{0g}	π_g
$g = 1$	-0.01	0.21
	(0.27)	(0.02)
	[0.13]	-
$g = 2$	-0.00	0.01
	(0.00)	(0.00)
	[0.00]	-
$g = 3$	0.01	0.08
	(0.02)	(0.01)
	[0.00]	-
$g = 4$	0.09	0.25
	(0.05)	(0.02)
	[0.01]	-
$g = 5$	0.21	0.31
	(0.10)	(0.02)
	[0.04]	-
$g = 6$	2.92	0.13
	(1.16)	(0.00)
	[1.07]	-

***Notes*:** See notes to Table 2. The homogeneous MPCs ($G = 1$) are 0.12 for nondurables and 0.34 for durables. In the untrimmed data for nondurables, the BIC selects $G = 3$ and for durables, $G = 6$. The BIC for durables flattens at $G = 4$ but then drops again before $G = 6$. The MPC distribution with 4 groups, however, is little different from the results reported here, with a largest MPC of 2.6 and most households with MPCs close to zero.

durables, the large majority of households spend a small fraction of the rebate on these goods, but about 15% of households display a durable MPC larger than 1. These right-tail responses are much larger than the greatest MPC we estimate in our baseline results, in line with the fact that upper 1.5% percent of expenditure changes may include large durable adjustments, such as car purchases.

To summarize, dropping the top and bottom expenditure changes improves precision and most likely recovers conservative estimates of the degree of MPC heterogeneity, specifically at the top. Finally, in Table A.11 we show that our findings on the drivers of MPC heterogeneity are, however, broadly unchanged. We still find that MPCs increase with salary and nonsalary income, as well as with the APC; moreover, the adjusted R^2 remains low. The fact that the estimates of the MPC distribution are somewhat sensitive to trimming (at the right tail) is not surprising, since we exploit a Gaussian mixture model. The reason for trimming is not to remove heterogeneity, but rather outliers. Outliers are often thought of as being drawn from a different distribution from the rest of a sample, i.e.

a different distribution for ϵ_j. If ϵ_j is not Gaussian for the outliers removed by trimming, leaving those observations in the sample may prevent the mixture model from recovering the correct distribution for reasons entirely unrelated to heterogeneity.

Table A.11: Explanatory variables of MPCs: keeping tail expenditure changes

	(1) Total	(2) Nondurables	(3) Durables
homeowner dummy	0.137**	0.028*	0.074
	(0.055)	(0.015)	(0.052)
mortgage dummy	-0.004	-0.010	0.011
	(0.057)	(0.015)	(0.052)
mortgage interest/income	-0.589***	-0.006	-0.433**
	(0.157)	(0.041)	(0.178)
Zero salary	-0.225***	-0.032**	-0.182***
	(0.066)	(0.015)	(0.063)
Salary: middle tercile	0.267***	0.020	0.241***
	(0.059)	(0.014)	(0.058)
Salary: top tercile	0.464***	0.073***	0.411***
	(0.078)	(0.020)	(0.075)
Zero non-salary	0.062	0.010	0.035
	(0.056)	(0.012)	(0.055)
Non-salary: middle tercile	0.121*	0.023	0.078
	(0.068)	(0.017)	(0.063)
Non-salary: top tercile	0.389***	0.056***	0.337***
	(0.079)	(0.021)	(0.077)
APC	0.452***	0.048***	0.385***
	(0.052)	(0.012)	(0.051)
adj. R^2	0.077	0.021	0.062

Notes: Regressions are estimated on the sample that includes only households receiving a rebate in the current period. Robust standard errors are reported in parentheses. We control for five education dummies, a married dummy, number of children, age and age squared; these coefficients are not reported. *, **, and *** denote significance at the 10, 5, and 1% levels respectively.

Table A.12: MPC heterogeneity: group membership as a function of ω

	λ_{01}	π_1	λ_{02}	π_2	λ_{03}	π_3
Baseline	0.04	0.30	0.23	0.48	1.33	0.23
Robustness	0.05	0.30	0.20	0.48	1.38	0.23

Notes: The first row reports estimates for group-specific MPCs, λ_{0g}, and population shares, π_g, in the baseline model. The second row reports estimates in a model in which π is a parametric function of controls, ω_j.

A.5 Group membership as a function of covariates

In this section, we allow group membership to be an explicit function of observable characteristics. In particular, following Gudicha and Vermunt (2013) and McLachlan and Peel (2004), for example, we assume that group membership is a multinomial logit function of a certain set of covariates, z. In particular, the complete-data likelihood becomes

$$\tilde{L}\left(\Delta C, X, D; \theta^G, \delta^G\right) = \prod_{j=1}^{N}\prod_{g=1}^{G} \bar{\pi}_g\left(z_j; \delta^G\right)^{d_{jg}} \phi\left(\Delta C_j; \psi_g^{G\prime} x_j, \sigma_g^2\right)^{d_{jg}},$$

where δ^G stacks the group-specific membership parameters, δ_g^G, and

$$\bar{\pi}_g\left(z_j; \delta^G\right) = \frac{\exp\left(\delta_g^{G\prime} z_j\right)}{\sum_{h=1}^{G}\exp\left(\delta_h^{G\prime} z_j\right)},$$

so the group membership function is a multinomial logistic regression on the covariates, z_j. This exercise serves to assess the sensitivity of our results results to the assumption that the prior probabilities for group membership do not depend on these variables..

We consider $z = \omega$, so group membership is allowed to depend on all the common covariates included in Equation 3, such as time dummies and age. As shown in Table A.12, the MPC distribution is very similar to that estimated in our baseline model. Predicted MPCs are also very similar in the two models, as plotted in Figure A.14. Importantly, the correlations between MPCs and observables are very similar to those shown in our baseline results (compare with Table 3), and the R^2 remains low, as we show in Table A.13. Indeed, the correlation between predicted MPCs in our baseline model and in this alternative is 0.995. Inspecting the multinomial coefficients, δ_G, those on the time dummies are mostly close to zero and statistically insignificant. Demographic coefficients are also insignificant with the exception of age, whose coefficient is small but statistically significant.

Figure A.14: Predicting individual MPCs: group membership as a function of ω

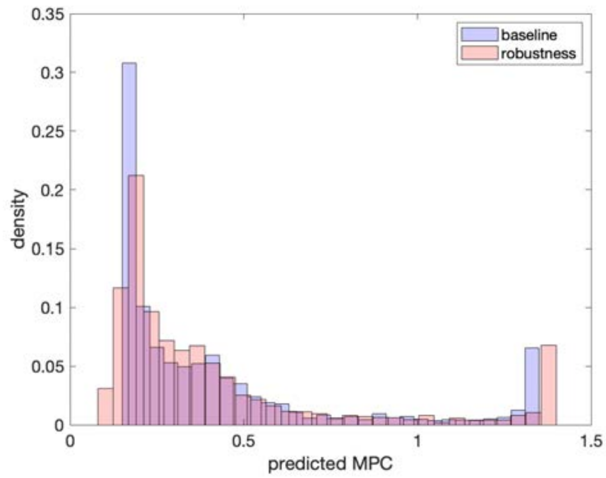

Notes: The histograms plot the individual MPCs constructed using the GMLR-estimated parameters and individuals' posterior probabilities for each group, γ_{jg}. Light blue bars refer to the baseline total expenditures specification as in Figure 1. The red bars refer to a model in which π depends on all controls ω_j. In the specification with group membership as a function of observables, the BIC selects $G = 3$.

Table A.13: Explanatory variables of MPCs: group membership as a function of ω

	(1)	(2)	(3)	(4)	(5)
mortgage interest/income	0.046				-0.110
	(0.074)				(0.102)
homeowner dummy	0.086***				0.077***
	(0.022)				(0.027)
mortgage dummy	0.034				-0.016
	(0.022)				(0.030)
Zero salary		-0.019			-0.089***
		(0.023)			(0.033)
Salary: middle tercile		0.058**			0.112***
		(0.025)			(0.029)
Salary: top tercile		0.138***			0.187***
		(0.026)			(0.037)
Zero non-salary			-0.009		0.032
			(0.023)		(0.028)
Non-salary: middle tercile			-0.060**		0.049
			(0.027)		(0.035)
Non-salary: top tercile			-0.000		0.167***
			(0.028)		(0.040)
APC				0.073***	0.195***
				(0.015)	(0.025)
adj. R^2	0.009	0.014	0.001	0.008	0.062

Notes: Regressions are estimated on the sample that includes only households receiving a rebate in the current period. Robust standard errors are reported in parentheses. In column (5) we also control for five education dummies, a married dummy, number of children, age and age squared; these coefficients are not reported. *, **, and *** denote significance at the 10, 5, and 1% levels respectively.

A.6 Measuring the explanatory power of observables for heterogeneity

An advantage of our two-stage approach is that we can regress the full heterogeneity of the MPC distribution on observables. Not only does this allow us to characterize which observable variables remain significant predictors of the MPC in these joint regressions, but it also enables us to compute the share of heterogeneity that is predicted based on observables. This exercise gives a measure of what share of heterogeneity may be truly latent – driven by fundamentally unobservable factors such as preference heterogeneity – or correlated with variables that simply are not included in our dataset. One complication is that the R^2 computed from these regressions provides a lower bound on the true R^2 due to measurement error in the estimated MPCs. In this section, we describe a simple exercise to adjust the R^2 for estimation error in recovering the MPCs.

In our WLS regressions, we pair replication g of an individual's characteristics with $\hat{\lambda}_{0g} = \lambda_{0g} + \hat{e}_g$, where $\hat{e}_g$ is the estimation error for λ_{0g}, and use the estimated posterior probabilities γ_{jg}, as weights. For the purpose of this exercise, we ignore measurement error in γ_{jg}; we discuss this issue below. Suppose that regressing the true MPCs, λ_{0g}, we have

$$\lambda_{0g}\gamma_{jg}^{1/2} = \left(c + \mu' F_j + v_{jg}\right)\gamma_{jg}^{1/2}, j = 1, \ldots, N, g = 1, \ldots, G. \quad \text{(A.9)}$$

Then the regression of estimated MPCs on observables takes the form

$$\hat{\lambda}_{0g}\gamma_{jg}^{1/2} = \left(c + \mu' F_j + v_{jg} + \hat{e}_g\right)\gamma_{jg}^{1/2}, j = 1, \ldots, N, g = 1, \ldots, G. \quad \text{(A.10)}$$

Based on the infeasible (A.9),

$$R^2_{true} = 1 - \frac{E\left[v_{jg}^2 \gamma_{jg}\right]}{\text{var}\left(\lambda_{0g}\gamma_{jg}^{1/2}\right)},$$

while the value computed based on (A.10) is

$$R^2_{raw} = 1 - \frac{E\left[\left(v_{jg} + \hat{e}_g\right)^2 \gamma_{jg}\right]}{\text{var}\left(\hat{\lambda}_{0g}\gamma_{jg}^{1/2}\right)}.$$

Under the assumption that estimation error in $\hat{\lambda}_g$ is uncorrelated with v_j for any individual j, the formula simplifies to

$$R^2_{raw} = 1 - \frac{E\left[v_{jg}^2 \gamma_{jg}\right] + E\left[\hat{e}_g^2 \gamma_{jg}\right]}{\text{var}\left(\hat{\lambda}_{0g}\gamma_{jg}^{1/2}\right)}.$$

This value is biased towards zero (since $\text{var}\left(\hat{\lambda}_{0g}\gamma_{jg}^{1/2}\right) = \text{var}\left(\lambda_{0g}\gamma_{jg}^{1/2}\right) + E\left[\hat{e}_g^2 \gamma_{jg}\right]$), potentially leading us to conclude that too small a share of MPC heterogeneity can be explained by observables. As in Majeske et al. (2010), these expressions can be rearranged to show that

$$R^2_{true} = \frac{R^2_{raw}}{1 - E\left[\hat{e}_g^2 \gamma_{jg}\right] / \text{var}\left(\hat{\lambda}_{0g} \gamma_{jg}^{1/2}\right)}. \tag{A.11}$$

The methods proposed in Majeske et al. (2010) to apply this formula – based on taking repeated measurements in experimental settings – are infeasible. Instead, we use the variance of our estimator for λ_{0g} and the estimated posterior probabilities to compute $E\left[\hat{e}_g^2 \gamma_{jg}^{1/2}\right]$, since $\hat{e}_g$ is the estimation error in $\hat{\lambda}$. With this proxy in hand, we can implement (A.11) to obtain a back-of-the-envelope estimate of R^2_{true}. When we do so, the baseline R^2 rises from 6% to 8%, still indicating that the majority of heterogeneity remains unexplained by observables. This remains a back-of-the-envelope estimate of R^2_{true}, since we have not accounted for estimation error in γ_{jg}. Note, however, that γ_{jg} enters linearly in both the numerator and denominator of the R^2 formulas, so that the only way it can impact the value of the R^2 is if errors in these posterior probabilities are systematically correlated with v_{jg}, the error in the projections on observables. As a further check, we also evaluate the R^2 for the regression on observables using the MPC distributions estimated from 250 bootstrap samples. We find that the average R^2 is 6%, and the 95% confidence interval ranges from 0 to 12%. Note that while the formulas derived above are non-parametric, they do assume that a group structure is the correct model, and we rely on the asymptotic distribution of $\hat{\theta}$ being correct in order to implement them.

B The role of correlated heterogeneity in quantile regression

In this section, we demonstrate why quantile regression is ill-suited to recovering the MPC distribution in an intuitive example; the recovered MPC distribution is impacted by heterogeneity in additional parameters.

We consider a simple setting, where there are two possible fixed effect values, $\alpha_j \in \{-10,000, 10,000\}$ (the order of magnitude of our estimated fixed effects in the model with only contemporaneous MPCs), and two contemporaneous MPCs, $\lambda_j = \{0.20, 0.70\}$. We assume that there are no lagged treatment effects. We draw non-zero rebate values $R_j \sim N\left(900, 100^2\right)$, centered near the median in our data. We then generate data according to

$$\Delta C_j = \alpha_j + \lambda_j R_j + \epsilon_j, j = 1, \ldots, N,$$

where $\epsilon_j \sim N\left(0, 1000^2\right)$, somewhat lower than the estimated noise in the data. We set $N = 100,000$, with 17.5% of observations receiving a rebate, as in our data, with $R_j = 0$ for the others.

We assume α_j and λ_j take each value with 50% probability. We consider three possible relationships between α_j and λ_j. First, we assume that they are perfectly positively correlated, so $\left(\alpha_j, \lambda_j\right) \in \{(-10,000, 0.20), (10,000, 0.70)\}$, with equal probabilities. Next, we assume that fixed effects and MPCs have zero correlation. Thus, $\left(\alpha_j, \lambda_j\right) \in \{(-10,000, 0.20), (10,000, 0.70), (-10,000, 0.70), (10,000, 0.20)\}$, with equal probabilities. Finally, we assume that MPCs and fixed effects are perfectly negatively correlated, so $\left(\alpha_j, \lambda_j\right) \in \{(10,000, 0.20), (-10,000, 0.70)\}$, with equal probabilities. For each specification, we draw 10 samples, estimate the model using quantile regression for every fifth percentile, and plot the estimated MPC distributions in Figure B.15. The first panel shows that when the fixed effects and MPCs are positively correlated, the MPCs are well estimated; half of the distribution is associated with an MPC around 0.20, and half with an MPC around 0.70. Because the fixed effects dominate the conditional distribution, and the MPCs are correlated with the fixed effects, the lower MPC aligns with the lower half of the distribution. In the second panel, there is zero correlation between fixed effects and MPCs. Since the percentile of the distribution to which each observation corresponds is driven largely by the fixed effect, the two MPCs occur with approximately equal frequency at each percentile, so a value near the average MPC is estimated *at each percentile*. Finally, the third panel shows that when fixed effects and MPCs are negatively correlated, the MPCs are again well-estimated, as in the first panel. However, this time the high MPC corresponds to the lower half of the distribution, since it aligns with the lower fixed effect. These results show that if there is heterogeneity in other parameters besides the MPC, the relationship between such heterogeneity and the MPC will impact the econometrician's ability to recover the distribution of MPCs using quantile regression.

We find in our empirical results that the MPC heterogeneity estimated by Misra and Surico (2014) is in fact exaggerated relative to ours, as opposed to the compressed distribution we observe in this highly simplified example. Alternative patterns, like that one, are entirely possible depending on the precise DGP as groups and controls are added and the correlations between group-specific parameters change.

Figure B.15: The role of correlated heterogeneity in quantile regression

(a) Positive correlation (b) Zero correlation (c) Negative correlation

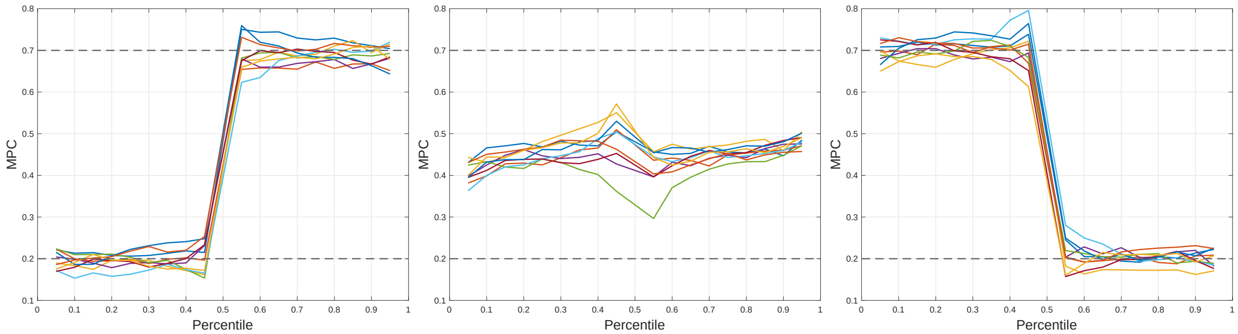

Notes: Figure B.15 plots the estimated MPCs from quantile regression for every fifth percentile for 10 samples of simulated data for three specifications. In each specification, both fixed effects and MPCs take two possible values. In the first panel, fixed effects and MPCs are perfectly positive correlated, in the second they have zero correlation, and in the third they are perfectly negatively correlated. The dashed lines represent the two true MPC values.

References

AGUIAR, M., C. BOAR, AND M. BILS (2019): "Who Are the Hand-to-Mouth?" in *2019 Meeting Papers*, Society for Economic Dynamics, 525.

ALAN, S., M. BROWNING, AND M. EJRNAES (2018): "Income and Consumption: A Micro Semistructural Analysis with Pervasive Heterogeneity," *Journal of Political Economy*, 126, 1827–1864.

ARMANTIER, O., L. GOLDMAN, G. KOŞAR, J. LU, R. POMERANTZ, W. VAN DER KLAAUW, ET AL. (2020): "How Have Households Used Their Stimulus Payments and How Would They Spend the Next?" Tech. rep., Federal Reserve Bank of New York.

AUCLERT, A. (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109, 2333–67.

AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2018): "The Intertemporal Keynesian Cross," Working Paper 25020, National Bureau of Economic Research.

BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption inequality and partial insurance," *American Economic Review*, 98, 1887–1921.

BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): "A Distributional Framework for Matched Employer Employee Data," *Econometrica*, 87, 699–739.

BONHOMME, S. AND E. MANRESA (2015): "Grouped Patterns of Heterogeneity in Panel Data," *Econometrica*, 83, 1147–1184.

BORUSYAK, K., X. JARAVEL, AND J. SPIESS (2024): "Revisiting Event-Study Designs: Robust and Efficient Estimation," *The Review of Economic Studies*, rdae007.

BOUTROS, M. (2022): "Windfall income shocks with finite planning horizons," Tech. rep., Bank of Canada.

BRODA, C. AND J. A. PARKER (2014): "The economic stimulus payments of 2008 and the aggregate demand for consumption," *Journal of Monetary Economics*, 68, S20–S36.

CARROLL, C., J. SLACALEK, K. TOKUOKA, AND M. N. WHITE (2017): "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 8, 977–1020.

CARROLL, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 23, 61–156.

CELEUX, G., S. FRUEWIRTH-SCHNATTER, AND C. ROBERT (2018): "Model Selection for Mixture Models - Perspectives and Strategies," *arXiv: Methodology*.

CLOYNE, J., C. FERREIRA, AND P. SURICO (2020): "Monetary policy when households have debt: new evidence on the transmission mechanism," *The Review of Economic Studies*, 87, 102–129.

COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): "Innocent Bystanders? Monetary policy and inequality," *Journal of Monetary Economics*, 88, 70–89.

CRAWLEY, E. AND A. KUCHLER (2018): "Consumption Heterogeneity: Micro Drivers and Macro Implications," *Danish National Bank Working Paper 129*.

DEMPSTER, A. P., N. M. LAIRD, AND D. B. RUBIN (1977): "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society. Series B (Methodological)*, 39, 1–38.

DESARBO, W. AND W. CRON (1988): "A Conditional Mixture Maximum Likelihood Methodology for Clusterwise Linear Regression," *Journal of Classification*, 5, 249–282.

DYNAN, K. E., J. SKINNER, AND S. P. ZELDES (2004): "Do the Rich Save More?" *Journal of Political Economy*, 112, 397–444.

FAGERENG, A., M. B. HOLM, AND G. J. NATVIK (2016): "MPC Heterogeneity and Household Balance Sheets," Discussion Papers 852, Statistics Norway, Research Department.

FRIEDMAN, M. (1957): *A Theory of the Consumption Function*, Princeton University Press.

FUSTER, A., G. KAPLAN, AND B. ZAFAR (2018): "What Would You Do With $500? Spending Responses to Gains, Losses, News, and Loans," Staff Reports 843, Federal Reserve Bank of New York.

GELMAN, M. (2019): "The Self-Constrained Hand to Mouth," .

GUDICHA, D. W. AND J. K. VERMUNT (2013): "Mixture Model Clustering with Covariates Using Adjusted Three-Step Approaches," in *Algorithms from and for Nature and Life*, ed. by B. Lausen, D. Van den Poel, and A. Ultsch, Cham: Springer International Publishing, 87–94.

ILUT, C. L. AND R. VALCHEV (2020): "Economic agents as imperfect problem solvers," Tech. rep., National Bureau of Economic Research.

JAPPELLI, T. AND L. PISTAFERRI (2014): "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics*, 6, 107–36.

JOHNSON, D. S., J. A. PARKER, AND N. S. SOULELES (2006): "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 96, 1589–1610.

JONES, P. AND G. J. MCLACHLAN (1992): "Fitting Finite Mixture Models in a Regression Context," *Australian Journal of Statistics*, 34, 233–240.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

KAPLAN, G. AND G. L. VIOLANTE (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82, 1199–1239.

——— (2021): "The Marginal Propensity to Consume in Heterogeneous Agents Models," Tech. rep.

KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 45, 77–153.

KOŞAR, G., D. MELCANGI, L. PILOSSOPH, AND D. G. WICZER (2023): "Stimulus through Insurance: The Marginal Propensity to Repay Debt," *FRB of New York Staff Report*.

KUENG, L. (2018): "Excess sensitivity of high-income consumers," *The Quarterly Journal of Economics*, 133, 1693–1751.

LAIBSON, D., P. MAXTED, AND B. MOLL (2022): "A Simple Mapping from MPCs to MPXs," Tech. rep., National Bureau of Economic Research.

LEWIS, D. J., D. MELCANGI, L. PILOSSOPH, AND A. TONER-RODGERS (2023): "Approximating grouped fixed effects estimation via fuzzy clustering regression," *Journal of Applied Econometrics*, 38, 1077–1084.

LUSARDI, A. (1996): "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets," *Journal of Business & Economic Statistics*, 14, 81–90.

MAJESKE, K. D., T. LYNCH-CARIS, AND J. BRELIN-FORNARI (2010): "Quantifying R 2 bias in the presence of measurement error," *Journal of Applied Statistics*, 37, 667–677.

MCLACHLAN, G. AND D. PEEL (2004): *Finite Mixture Models*, Wiley Series in Probability and Statistics, Wiley.

MCLACHLAN, G. J. AND K. E. BASFORD (1988): *Mixture models. Inference and applications to clustering*.

MIRANDA-PINTO, J., D. MURPHY, K. J. WALSH, AND E. R. YOUNG (2020): "A model of expenditure shocks," .

MISRA, K. AND P. SURICO (2014): "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs," *American Economic Journal: Macroeconomics*, 6, 84–106.

OLAFSSON, A. AND M. PAGEL (2018): "The liquid hand-to-mouth: Evidence from personal finance management software," *The Review of Financial Studies*, 31, 4398–4446.

ORCHARD, J., V. RAMEY, AND J. WIELAND (2023): "Using Macro Counterfactuals to Assess Plausibility: An Illustration using the 2001 Rebate MPCs," NBER Working Papers 31808, National Bureau of Economic Research, Inc.

PARKER, J. A. (2017): "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment," *American Economic Journal: Macroeconomics*, 9, 153–83.

PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103, 2530–53.

QUANDT, R. E. (1972): "A New Approach to Estimating Switching Regressions," *Journal of the American Statistical Association*, 67, 306–310.

SAHM, C. R., M. D. SHAPIRO, AND J. SLEMROD (2010): "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications," in *Tax Policy and the Economy, Volume 24*, National Bureau of Economic Research, Inc, NBER Chapters, 69–110.

SHAPIRO, M. D. AND J. SLEMROD (2009): "Did the 2008 Tax Rebates Stimulate Spending?" *American Economic Review*, 99, 374–79.

STRAUB, L. (2017): "Consumption, Savings, and the Distribution of Permanent Income," Working Paper.

WONG, A. (2021): "Refinancing and the transmission of monetary policy to consumption," .