Latent Heterogeneity in the Marginal Propensity to Consume[*]

Daniel Lewis — FRBNY
Davide Melcangi — FRBNY
Laura Pilossoph[†] — FRBNY

February 11, 2021

Abstract

We estimate the distribution of marginal propensities to consume (MPCs) using a novel clustering approach, generalizing the fuzzy C-means algorithm to regression settings. We find that households spent at least one fourth of their 2008 stimulus payments, with considerable MPC heterogeneity which varies by consumption good. We document observable determinants of this heterogeneity, without imposing *ex ante* assumptions on such relationships, and evaluate potential determinants jointly. MPCs correlate positively with income and the average propensity to consume, but much heterogeneity remains unexplained. The partial equilibrium response to the stimulus is twice the homogeneous estimate, highlighting the importance of heterogeneity.

Keywords: Marginal Propensity to Consume, Consumption, Tax Rebate, Heterogeneous Treatment Effects, Clustering, C-means
JEL Codes: D12, D91, E21, E32, E62

[*]This paper was previously circulated as "A New Approach to Estimating Heterogeneous Effects and the Distribution of the Marginal Propensity to Consume". The views expressed herein are solely those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System. We thank René Chalom and Meghana Gaur for excellent research assistance. For helpful comments, we also thank Michael Boutros, Sara Casella (discussant), Richard Crump, Marco Del Negro, Keshav Dogra, Domenico Giannone, Simon Gilchrist, Greg Kaplan, Ulrich Müller, Eric Nielsen (discussant), Mikkel Plagborg-Møller, Isaac Sorkin, Mary Wootters, and various seminar and conference participants.
[†]daniel.lewis@ny.frb.org, davide.melcangi@ny.frb.org, laura.pilossoph@ny.frb.org.

1 Introduction

Recent work highlights the importance of heterogeneity in the marginal propensity to consume (MPC) out of transitory income shocks for fiscal policy, the transmission of monetary policy, and welfare.[1] Despite their importance, estimates of the *distribution* of MPCs are largely elusive. Even with plausibly identified transitory income shocks, estimating individual-level MPCs requires panel data with long horizons, which are typically not available; it also usually requires the unappealing assumption that an individual's marginal propensity to consume (MPC) 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.[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 propose a novel methodology to estimate the distribution of MPCs directly. We develop an estimator based on the fuzzy C-means (FCM) clustering algorithm (Dunn (1973), Bezdek (1973)), 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 receipt and basic controls (Johnson, Parker, and Souleles (2006), Parker, Souleles, Johnson, and McClelland (2013)), but allows the coefficient on the rebate to be heterogeneous across unknown groups; the groups as well as their rebate coefficients are jointly estimated.

This approach offers several advantages over existing efforts to recover the distribution of MPCs. First, it allows us to estimate the full unconditional distribution of MPCs, which can be driven both by latent factors and observable characteristics, broadly defined. Standard methods that rely on sample splitting by the latter cannot recover heterogeneity

[1] The MPC distribution is a crucial object in Heterogeneous Agent New Keynesian (HANK) models of monetary policy (see Kaplan, Moll, and Violante (2018)). For example, Auclert (2019) shows that the response of aggregate consumption to monetary policy shocks depends on the covariance of the distribution of MPCs with the cyclicality of income, net nominal position, and unhedged interest rate exposure.

[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).

in MPCs associated with the former (by definition). Yet, we find that the majority of MPC heterogeneity can be attributed to such latent characteristics. Second, because our approach does not require taking an ex-ante stand on what observables correlate with MPC heterogeneity, we can "let the data speak" by investigating *ex post* which observables predict the recovered individual MPCs. We find that a household's MPC is individually correlated with various observable characteristics, and that these relationships are statistically significant. Hence, our approach overcomes the loss of statistical power which may have plagued the sample-splitting approach in existing studies. Importantly, however, we further show that the significance of most of these univariate relationships with the MPC disappears when we instead examine predictors *jointly*, with the exception of non-salary income and the average propensity to consume (APC). Finally, we formally show that household observable characteristics explain only a small portion of the variance in MPCs.

Our contribution 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 parameters describe the observations within the group. This feature allows us to bypass the decision on which observables matter for MPC heterogeneity, and instead estimate the heterogeneity directly first. FCM is a "fuzzy" clustering approach, in which individuals are not assigned to groups in a binary fashion, but instead have continuous weights. When the panel dimension present in Bonhomme and Manresa (2015) 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 continuous weights are desirable to represent the level of uncertainty that exists in the assignment. If the econometrician knows the true distribution of the data, a likelihood-based approach like a Gaussian mixture model can be used, but, absent any information on the distribution, the form of weights used by FCM has optimality properties (Bezdek (1973)).

We therefore generalize FCM from the cluster means case to the regression setting, where it constitutes a weighted least squares (WLS) problem, and call the resulting generalized estimators fuzzy C-parameters (FCP).[4] We establish the consistency and asymptotic normality of a GMM formulation for the FCP problem. In a simulation study calibrated to our empirical setting, we find that FCP is able to accurately recover the distribution of MPCs, more so than quantile regressions and binary classification or likelihood-based competitors. We further show how FCP can be used in the second stage of two-stage least squares (TSLS) to accommodate the use of instrumental variables, which we

[4]Further generalizations to nonlinear models, while outside the scope of this paper, appear trivial under suitable assumptions.

exploit when we instrument the rebate value with its receipt. These characteristics mean that FCP is well-suited to studying heterogeneity in a wide variety of economic settings with cross-sectional or short-panel data, and can be used in contexts well beyond our current focus.

Applying our estimator to study the MPC distribution using the 2008 Economic Stimulus Act, we uncover a substantial degree of heterogeneity. In particular, households spend at least one fourth of the rebate within one quarter, with some households displaying an MPC above one. While the share of households with different MPCs generally decreases with the magnitude of the MPC, we estimate sizable MPCs even at the bottom of the distribution. Fagereng et al. (2016) and Olafsson and Pagel (2018) find evidence of similar behavior in Norway and Iceland respectively, and recent papers have proposed models of limited cognitive perception to rationalize such behavior (e.g., Ilut and Valchev (2020)). However, we also find that, while all households spend a non-trivial portion of the rebate overall, they often smooth nondurable consumption. Specifically, a third of the households do not use the rebate to buy additional nondurable goods. Our results are consistent across different specifications and sample restrictions. For instance, instrumenting the rebate with an indicator for its receipt, as in Parker et al. (2013), leaves the results qualitatively unchanged and in fact increases the estimated heterogeneity in MPCs.[5]

Having characterized the distribution of marginal propensities to consume, we describe its main drivers. Historically, the literature has found mixed empirical evidence and generally weak relationships between MPC and observable household characteristics, with the possible exception of liquid wealth.[6] This is likely due to a significant loss of statistical power when re-estimating the MPC with interactions or sample splits. Our approach allows us to sidestep these issues. Indeed, we find that our estimated MPCs are significantly correlated, individually, with many observable drivers, despite the fact that we use the same identification strategy and dataset 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.

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 controls and positively correlate with MPCs: households' non-

[5]The same is true when we exclude from the sample households that never received a rebate, or when we include lagged values of the rebate to control for persistent effects of receiving the rebate.

[6]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.

salary income and their APC. Kueng (2018) also finds that high-income households have higher MPCs in Alaskan data. We highlight how our result crucially hinges on the non-salary component of income, such as business and financial income. 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 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.

Importantly, our best array of observable predictors is able to explain only about 11% of the variation in estimated MPCs. With the vast majority of heterogeneity unexplained by standard observables, our results suggest that a relevant portion of MPC heterogeneity is driven by latent household traits. For example, heterogeneity in discount rates and/or intertemporal elasticities of substitution (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.[7]

Finally, correctly accounting for MPC heterogeneity also matters for the aggregated consumption effects of the fiscal stimulus. We show that the sample average of our estimated heterogeneous responses is larger than the homogeneous marginal propensity to consume. When considering the cumulated heterogeneous responses over two quarters, the aggregated response from our distribution is almost twice as large as its homogeneous counterpart. While still a partial-equilibrium object in nature, this result suggests that correctly accounting for heterogeneity is important in order to correctly evaluate the impact of the 2008 fiscal stimulus.

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. As previously mentioned, the vast majority of existing papers study *observable* drivers of the MPC; we relate our findings to this literature. A burgeoning literature has turned its attention to unobserved household traits and preference heterogeneity. Our findings corroborate the importance of this dimension,

[7] 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.

recently highlighted by Alan, Browning, and Ejrnaes (2018), Parker (2017), Aguiar et al. (2019), and Gelman (2019).

Our approach allows us to flexibly and non-parametrically combine observed and unobserved MPC heterogeneity. In this respect, Misra and Surico (2014) is closest in spirit to our work.[8] 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. We discuss 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 changes than the tax rebate. Indeed, we find in simulations that quantile regression overstates the dispersion of MPCs, particularly in the left tail. We not only estimate a different MPC distribution relative to Misra and Surico (2014) (in particular, one featuring a lower bound on MPCs that is considerably higher), but we also find new relationships with observables.

Finally, we believe our findings can be used to discipline heterogeneous agent models, for three main reasons. First, our estimated full distribution of MPCs is an agnostic target, regardless of a model's characteristics. Second, individual correlations between MPCs and observable characteristics are crucial objects for many heterogeneous agent models, and we are able to estimate them with statistical precision. Third, we provide researchers with an explicit number for the joint importance of observable and unobservable drivers for the distribution of MPCs.

The paper proceeds as follows. In Section 2 we describe our empirical strategy based on the 2008 tax rebate. In Section 3, we formulate the problem at hand and present the FCP estimator. We extend FCM from the cluster means case to a fully-general regression setting, as well as instrumental variables regression, and derive asymptotic properties of the corresponding estimators. 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 4.5 aggregates the estimated household MPCs to arrive at a partial equilibrium effect on aggregate consumption. Section 5 concludes.

[8]Other 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).

2 Empirical methodology

In order to estimate the marginal propensity to consume, and how it varies across households, we look at an off-the-shelf well-identified quasi-natural experiment: the 2008 Economic Stimulus Act (ESA), as studied by Parker et al. (2013), among others. Between April and July of 2008, $100 billion in tax rebates was sent to approximately 130 million US tax filers.[9] 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 also exploit the randomized timing of the rebate receipt, but instead estimate heterogeneous (and unobserved) marginal propensities to consume rather than a homogeneous marginal propensity to consume. Our data come from the Consumer Expenditure Survey (CEX), which contains comprehensive and detailed measures of household-level consumption expenditures. The 2008 CEX wave also includes 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. Since we depart from Parker et al. (2013) by allowing for treatment heterogeneity, 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:

$$\Delta C_j = \beta' W_j + \phi R_j + \alpha + \epsilon_j, j = 1, \ldots, N, \tag{1}$$

where ΔC_j is the first difference of consumption expenditure of household i in quarter t.[10] W_j is a set of controls including month dummies aimed at absorbing common time effects such as aggregate shocks, as well as seasonal factors.[11] The independent variable of interest is R_j, which denotes the amount of the tax rebate received by each household. ϕ

[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. We wish to emphasize that while we have information on the same households i in different periods t, identification is not obtained by comparing individual responses over time. We do not exploit any limited panel structure, except to construct consumption changes for the left-hand-side variable. 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 will use are the same, but additionally include age squared.

is then interpreted as the causal effect of the rebate on expenditures, where identification is achieved by comparing 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.[12]

2.2 Heterogeneous MPCs

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

$$\Delta C_j = \beta' W_j + \sum_{g \in G} \left(\phi_g \mathbf{1}\left[j \in g\right] R_j + \alpha_g \mathbf{1}\left[j \in g\right] \right) + \epsilon_j, j = 1, \ldots, N, \tag{2}$$

where $\mathbf{1}\left[j \in g\right]$ is an indicator that takes a value of 1 if household i in period t belongs to a certain group $g = 1, \ldots, G$. That is, we assume that heterogeneity in responses to the rebate can be summarized with G groups, characterized by the vector of coefficients $\{\alpha_g, \phi_g\}$. We include group-specific intercepts α_g to correctly interpret ϕ_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. Our object of interest is $\phi = \begin{pmatrix} \phi_1 & \ldots & \phi_G \end{pmatrix}'$, which describes MPC heterogeneity, while $\mathbf{1}\left[j \in g\right]$ tells us the group membership of each household. The vector of coefficients, combined with $1\left[j \in g\right]$, gives an approximation of the MPC distribution. In the next section we introduce a new methodology to jointly estimate ϕ and $\mathbf{1}\left[j \in g\right]$. We also discuss how we choose the total number of groups, G.

3 A fuzzy clustering approach to MPC estimation

In this section, we describe our approach to recovering the distribution of MPCs. We first outline the FCM algorithm for the cluster means case, before generalizing it to our FCP methodology, and outlining simulation results. Finally, we compare the FCP approach to alternative methodologies.

[12] Kaplan and Violante (2014) discuss why ϕ may not correctly measure the marginal propensity to consume out of a transitory income shock, but is instead better thought of as a "rebate coefficient". We address these issues in Supplement D.5.

3.1 Fuzzy clustering

To estimate group-specific MPCs we must first assign individuals to groups. Previous papers have grouped individuals based on observable characteristics, but doing so presupposes the determinants of the MPCs *a priori*. However, we do not suppose to know these determinants in advance (there is indeed considerable empirical and theoretical disagreement on this point).[13] Moreover, we aim to investigate correlates of the MPC *ex post*, requiring us to remain agnostic while recovering the MPCs. For these reasons, we group individuals on the basis of their heterogeneous MPCs themselves (and potentially other group-specific parameters). Clustering methods are tailored to this goal.

To build intuition for our approach, we start with the cluster means case, where the econometrician simply wishes to characterize the distribution of means across groups. This problem consists of minimizing

$$L_1 (P, \psi) = \int \sum_{g=1}^{G} w_g (y; \psi) \left\| y - \psi_g \right\|^2 P (dy), \tag{3}$$

where $y \in \mathbb{R}$ is the observed variable with probability measure P on $\mathbb{R}$, $g = 1, \ldots, G$ indexes groups, and $\psi \in \mathbb{R}^G$ is the vector of cluster means.[14] The weights, w_g, satisfy $\sum_{g=1}^{G} w_g (y; \psi) = 1 \forall y$. Equation (3) constitutes a weighted least-squares (WLS) problem, where the weights are unknown. Dunn (1973) shows that if the econometrician does not have any additional information about the distribution of y, the optimal solution to Equation (3) consists of binary weights $w_g^* = \mathbf{1} \left[\left\| y - \psi_g \right\|^2 \leq \left\| y - \psi_h \right\|^2 \ \forall\, h \neq g \right]$. This "hard" classification corresponds to the so-called hard K-means (HKM) approach (e.g., Bonhomme and Manresa (2015)).[15] If the econometrician has information to provide a likelihood for y, then w_g can be derived from that likelihood, as in a Gaussian mixture model, for instance. However, performance can deteriorate (e.g., Winkler et al. (2011)) if the likelihood is misspecified or there are outliers, which motivates the use of the (optimal) non-parametric weights, w_g^*.

Moreover, the binary nature of w_g^* is not well-suited to empirical settings like ours. Some clustering applications focus on cases where y is a $T \times 1$ vector of outcomes (a panel structure), sometimes even assuming $T \to \infty$. In contrast, in our data, $T = 1$. Given a

[13] We cover the extensive literature on MPC heterogeneity in Sections 1 and 4.3.

[14] For simplicity, we present our theory in the text in terms of scalar-valued y, but our theoretical results in Supplement A are presented in full generality for vector-valued $y \in \mathbb{R}^T$ to accommodate a panel structure, for instance.

[15] Bonhomme et al. (2019) provide a recent extension tailored to settings where true heterogeneity is not discrete, a 2-step procedure using HKM for classification in the first step and estimating group-specific heterogeneity in the second.

single observation for each individual, in the presence of noise – including measurement error, which is common in survey data like the CEX – definitive group assignment is an unrealistic representation of the econometrician's information. Assigning individuals for whom $\left\|y-\psi_g\right\|^2 \approx \left\|y-\psi_h\right\|^2, g \neq h$ discretely to a single group has the potential to mischaracterize group membership and to distort WLS estimates. Indeed, in a simple example with two Gaussian clusters, the cluster means are shifted outwards, as we show analytically in Proposition 1 in Supplement A.1. The "fuzzy clustering" literature recognizes that group assignments may be better represented in practice as uncertain, with observations belonging to multiple groups with some weight in $(0,1)$. To obtain a set of weights that is both optimal and non-binary, we must modify Equation (3) so that the derivative of the objective function with respect to w_g is positive at $w_g = 1$ for all g (see Rousseeuw et al. (1995)). Dunn (1973) and Bezdek (1973) propose a simple solution:

$$L_m(P,\rho) = \int \sum_{g=1}^{G} w_g^m(y;\rho) \left\|y-\rho_g\right\|^2 P(dy), \tag{4}$$

where m is a "fuzziness parameter" that determines the deviation from binary assignment. If $m = 1$, then $L_m = L_1$, with binary weights, but if $m > 1$, the optimal weights are "fuzzy", with $w^*_{g,m} \to 1/G$ as $m \to \infty$. In particular, Bezdek (1973) shows that the minimization of Equation (4) results in the optimal weights:

$$\mu_g(y;\rho) \equiv w^*_{g,m}(y;\rho) = \left(\sum_{h=1}^{G} \frac{\left\|y-\rho_g\right\|^{2/(m-1)}}{\left\|y-\rho_h\right\|^{2/(m-1)}}\right)^{-1}, g = 1, \ldots, G.$$

The resulting objective function,

$$J_m(P,\rho) = \int \sum_{g=1}^{G} \mu_g^m(y;\rho) \left\|y-\rho_g\right\|^2 P(dy) \tag{5}$$

is known as "fuzzy C-means" (FCM). Without specifying the distribution of y, and given the formulation in Equation (4), FCM provides the optimal non-parametric solution to clustering data in cross-sectional or small-T settings like the one we face empirically (Bezdek (1973)).

Asymptotic properties for FCM in this simple cluster means case are derived by Yang and Yu (1992) and Yang (1994), who show that sample estimates of ρ_g are consistent for their population counterparts with normal limiting distributions. FCM has frequently been found to perform very well in practice, sometimes even outperforming well-specified mixture models (e.g., Chapter 6 of Bezdek (1981)).

3.2 Fuzzy clustering regression

While FCM has been widely adopted in practice as a tool to recover cluster means, the MPC, which we wish to uncover, is a regression coefficient on the tax rebate as outlined in Section 2. We therefore extend the FCM approach to a regression setting by considering the model

$$y_i = \sum_{g=1}^{G} \mathbf{1}\left[i \in g\right] \theta_g' x_i + \varepsilon_i, \quad i = 1, \ldots, N, \tag{6}$$

where group membership indicators $1\left[i \in g\right]$ are unobserved and x is a $K \times 1$ vector. We propose fuzzy C-parameters (FCP) as a novel tool to recover the distribution of heterogeneous treatment effects θ_g in Equation (6). FCP replaces cluster means ρ_g in Equation (5) with the conditional mean, $\theta_g' x$:

$$J_m^{reg}\left(\Pi, \theta\right) = \int \int \sum_{g=1}^{G} \mu_g^{reg}\left(y, x; \theta\right)^m \left\|y - \theta_g' x\right\|^2 \Pi_{y|x}\left(dy \mid x\right) \Pi_x\left(dx\right). \tag{7}$$

where $\theta \in \Theta \subset \mathbb{R}^{K \times G}$, Π denotes the distribution of (y, x) and

$$\mu_g^{reg}\left(y \mid x; \theta\right) \equiv \left(\sum_{h=1}^{G} \frac{\left\|y - \theta_g' x\right\|^{2/(m-1)}}{\left\|y - \theta_h' x\right\|^{2/(m-1)}} \right)^{-1}, g = 1, \ldots, G. \tag{8}$$

As before, Equation (7) has the interpretation of a WLS regression, where the weights are unknown. We define θ^* as the population minimizer of the objective in Equation (7). Our approach determines group weights based on the data alone, without requiring a classification *a priori*.

Asymptotic properties of fuzzy C-parameters

We now characterize the asymptotic properties of an estimator, $\hat{\theta}$, based on the sample counterpart of Equation (7). In Supplement A.2, we provide detailed proofs of these properties, generalized to allow y to be vector valued (as in panel data).

Substituting μ_g^{reg} in Equation (8) into Equation (7) and taking first-order conditions with respect to θ yields a set of moment conditions

$$E\left[\frac{\partial J_m^{reg}}{\partial vec\left(\theta\right)}\right] \equiv E\left[\eta\left(\theta, y_i, x_i\right)\right] = E\left[\left(\sum_{h=1}^{G} \frac{\left\|y_i - \theta_g' x_i\right\|^{2/(m-1)}}{\left\|y_i - \theta_h' x_i\right\|^{2/(m-1)}}\right)^{-m} \left(y_{it} - \theta_g' x_i\right) x_i\right] = 0$$
$$\text{for } g = 1, \ldots, G,$$

which are satisfied by θ^*. These sample counterparts of these moment conditions can be used to estimate θ^* by (just-identified) GMM using the objective function

$$S_N\left(\theta\right) = \frac{1}{N} \sum_{i=1}^{N} \eta\left(\theta, y_i, x_i\right)' \sum_{i=1}^{N} \eta\left(\theta, y_i, x_i\right). \tag{9}$$

Define $\hat{\theta}$ as the solution to (9). To conserve space, we detail all assumptions in Supplement (A.2), but sketch them here. Assumption 1 presents standard OLS assumptions, extended to require exogeneity of x_i for each group, as well as finite G. Assumption 2 adds standard additional assumptions for consistency of GMM, in particular the uniqueness of θ^* (up to group ordering) and compactness of Θ. Theorem 1 shows that $\hat{\theta}$ is consistent for θ^*, the population solution to Equation 7:

Theorem 1. *(Consistency) Under Assumptions 1-2, $\hat{\theta} \xrightarrow{p} \theta^*$ as $N \to \infty$.*

To characterize the asymptotic distribution, we require further standard technical conditions for the asymptotic normality of GMM estimators, outlined in Assumption 3. The limiting distribution of $\hat{\theta}$ is given by Theorem (2):

Theorem 2. *(Asymptotic Normality) Under Assumptions 1 - 3,*

$$\sqrt{N}\left(vec\left(\hat{\theta}\right) - vec\left(\theta^*\right)\right) \xrightarrow{d} \mathcal{N}\left(0, H^{-1} V H^{-1}\right),$$

where

$$V = E\left[\eta\left(\theta^*, y_i, x_i\right) \eta\left(\theta^*, y_i, x_i\right)'\right],$$

and H is the Hessian of Equation (7).

We provide expressions for H (including several extensions) in Supplement A.2. With these results, we are able to conduct inference for the population parameter θ^*.

Practical considerations

In practice, the econometrician must choose both the number of groups, G, and the smoothing parameter, m; we adopt data-dependent rules to select them. For G, we adapt the "gap statistic" of Tibshirani et al. (2001) to the regression setting. This approach runs the clustering algorithm on data with no cluster structure (homogeneous coefficients) and compares the FCP objective function in this reference distribution to the value obtained for the empirical data. The gap statistic normalizes declines in (7) as G increases by expected declines in (7) that would arise mechanically even without a group structure; the optimal G maximizes the gap between the reference and empirical FCP objective function.

We choose m to minimize the loss function

$$Q\left(m\right) = \left(\bar{J}_m^{reg}\left(\hat{\theta}\right)\right)^2 + \left(\bar{E}_m\left(\hat{\theta}\right)\right)^2,$$

where $\bar{J}_m^{reg}$ is the sample analog of Equation (7), normalized to the range $[0,1]$, and $\bar{E}_m$ is the sample cluster entropy (see e.g., Bezdek et al. (1984)), measuring dispersion within clusters, normalized to the range $[0,1]$. We trade off the objective function against the similarity of observations within each group. This approach is related to the method of Cui et al. (2010) for cluster means. In our data, we select $m = 2.35$, in line with the range of values found in the FCM literature.[16] Since the optimal value for G is invariant to m in our data, we choose G and m sequentially.

We implement FCP by numerically minimizing the GMM objective function, as detailed in Section (3.2). The FCM literature has solved the minimization problem iteratively, but we find GMM to be both considerably faster and more accurate in our setting than the iterative approach. As such, we view casting the problem in a GMM formulation as another contribution of the paper. Supplement C.1 contains details.

FCP instrumental variables regression

In our empirical setting, the value of the rebate an individual receives is potentially endogenous, so we propose an instrumental variables extension of FCP (FCP-IV). In particular, we consider a two-stage least squares (TSLS) estimator, where we denote the rebate as x^e, additional controls as ω, and the instrument as z (in our case, an indicator for the

[16]The literature generally suggests $1.5 \leq m \leq 2.5$ (e.g., Bezdek et al. (1984); Pal and Bezdek (1995); Yu et al. (2004); Jing et al. (2014); Wu (2012)). Optimality results are not generally available (theoretically or numerically).

rebate receipt). We estimate the first stage

$$x_i^e = \gamma z_i + \tau' \omega_i + u_i,$$

via OLS, and then generate $\tilde{x}_i^e = \gamma z_i + \tau' \omega_i$. We define $\tilde{x} = \begin{pmatrix} \tilde{x}^e & \omega' \end{pmatrix}'$, and estimate Equation (7), replacing x with $\tilde{x}$ as the second stage. In Supplement A.3, we establish the consistency and asymptotic normality of this estimator.

3.3 Simulation results

To evaluate the performance of FCP in our setting, we conduct a simulation study calibrated to our empirical data. We summarize the main results below, with full details reported in Supplement C.2. We simulate data according to Equation (6) using the empirical parameter estimates, taking an individual's true group membership as the group with maximal weight. We consider three specifications: $G = 5$ (our empirical baseline), with Gaussian errors ε_i (with group-specific variance); $G = 5$, with ε_i resampled from each group's empirical errors; and $G = 10$, with Gaussian errors (with group-specific variance). For each specification, we estimate the model using FCP ($m = 2.35$, our empirical choice), HKM extended to a regression framework (as in Bonhomme and Manresa (2015)), and a Gaussian mixture model (using the EM algorithm). For each estimator, we report the mean and RMSE for each MPC, the share of observations misclassified (by maximal weight), and the median CDF of MPCs (by maximal weight), across 500 simulations.

For $G = 5$ and Gaussian errors, both FCP and HKM perform quite well by all metrics: the average point estimates are accurate, the RMSEs are fairly low, only 9% of individuals are misclassified, and the median CDF recovers the true distribution of MPCs very closely. The Gaussian mixture model recovers a spuriously wide distribution of MPCs. For $G = 5$ and empirical errors, FCP continues to perform well, demonstrating a clear advantage over HKM in terms of average point estimates and continuing to match the true CDF of MPCs. The Gaussian mixture model recovers a sizable mass of negative MPCs. Finally, for $G = 10$, a much more challenging clustering problem, FCP performs remarkably well. Average point estimates remain reasonably close to the truth, the RMSEs generally fall, the share misclassified rises only slightly, and the median CDF matches the true distribution of MPCs very closely. HKM recovers a large mass of negative MPCs, while the Gaussian mixture model recovers MPCs greater than one. Additionally, for FCP, tests of the null of the true MPCs are generally well-sized. Together, these results support the ability of FCP to recover the distribution of MPCs in our empirical setting.

Apart from these 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; rather than recovering groupings of the data that share similar parameters and estimating those parameters, quantile regressions investigate the relationship between the dependent variable and the regressors at different points of the conditional distribution of the dependent variable. 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 the 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 C.3.

We confirm quantile regression may provide a poor approximation of the MPC distribution using the empirically-motivated simulation setup described above. We find that, in this setting, quantile regression estimates spurious dispersion in the MPC, particularly in the left tail, computing a sizable mass of responses below the lowest MPC and even below zero. As we argue below, the lower tail of the MPC distribution is indeed a crucial distinction between our estimates and those of Misra and Surico (2014). Quantile regression also less accurately assigns individuals to MPCs in our simulations. This latter point is important since we examine determinants of MPC heterogeneity *ex post*; if individual MPCs are not accurate, results of these regressions will be spurious. Such exercises also pose a conceptual problem for quantile regression. Following quantile regression, it is possible to regress *percentiles* on the characteristics of their members, similar to how Misra and Surico (2014) plot observables across percentiles. However, it is not possible to examine the relationship between MPCs and covariates directly. Thus, this exercise relies on percentiles and MPCs being well-correlated for validity.

4 Results

We apply our FCP approach to the 2008 tax rebate. In particular, our heterogeneous formulation in Equation (2) is a special case of Equation (6), where the coefficients on W_j are assumed to be common across groups; we then estimate (6) using FCP with $m = 2.35$, in particular minimizing

$$\frac{1}{N}\sum_{j=1}^{N}\left(\sum_{g=1}^{G}\left\|\Delta C_j - \alpha_g - \phi_g R_j - \beta' W_j\right\|^{-2/(m-1)}\right)^{1-m}$$

with respect to $\{\alpha_g, \phi_g\}_{g=1}^{G}$.[17,18] Our findings highlight a considerable degree of MPC heterogeneity whose extent varies depending on the consumption category considered. We first show the distribution of marginal propensities to consume for total expenditures and illustrate how our results are robust to different specifications and sample selection procedures. We then investigate how the MPC distribution changes as we consider non-durable and durable goods as the dependent variables. Importantly, our approach also allows us to directly test whether households display similar propensities for different consumption goods, or instead substitute across expenditure types when they receive a transitory income shock such as a tax rebate. Finally, we explore which observable household characteristics are correlated with the estimated marginal propensities to consume, both individually and jointly, and discuss the partial equilibrium response of expenditures to the 2008 tax rebates in our heterogeneous setting relative to the homogeneous one.

4.1 The distribution of marginal propensities to consume

We start by considering 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.[19] The gap statistic indicates that 5 is the optimal number of groups. For each household that receives the rebate, we compute the weighted average MPC, using the household-specific weights $\mu_{i,g}$ and the group-specific MPCs ϕ_g estimated by the algorithm.[20] Figure 1 shows the distribution of this object for the (i,t) pairs receiving the rebate.

The vast majority of households display a relatively low (but certainly non-negligible) MPC between 0.25 and 0.4, and the share of households with a given MPC slowly decays as the MPC increases. While under this specification no household can be strictly defined as having an MPC of one, the majority of the sample exhibits a sizable propensity to consume; our findings suggest that most households consume at least part of the rebate.[21]

[17] The weights are given by Equation 8, which are a function of the estimates themselves.

[18] We show in Appendix D.1 that the estimated MPC distribution is robust to also allowing for heterogeneity in the household level controls.

[19] This is the only way in which our sample departs from Parker et al. (2013), and explains why the homogeneous MPC we estimate for total consumption differs from theirs.

[20] Supplement D.3 shows the distribution of the MPC associated with the maximal weight at the individual level. Because we find that many such weights are close to one, the average of MPCs associated with maximal weights is very close to the average of weighted MPCs.

[21] In Section D.2, we report the MPC distribution estimated by HKM and a Gaussian mixture model. The distribution is qualitatively robust to the estimator used, although there is variation in the degree of dis-

Figure 1: Estimated distribution of MPCs out of the tax rebate

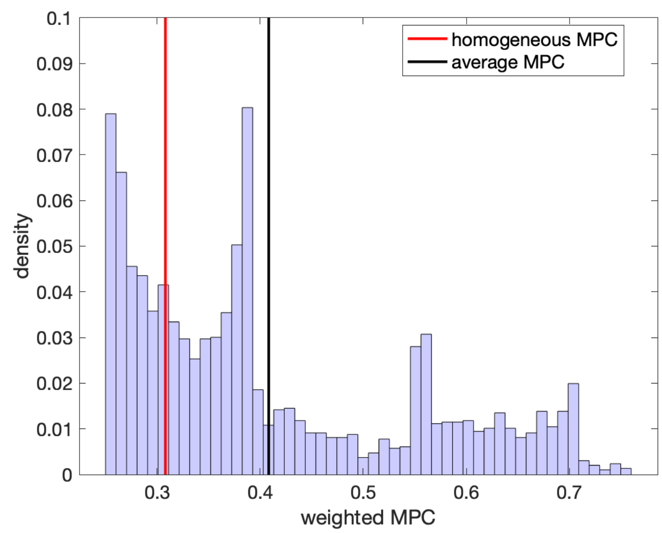

Notes: Figure 1 plots a histogram (light blue bars) of the estimated distribution of MPCs for total expenditures among households that received the rebate, defined as in Parker et al. (2013). The homogeneous MPC (red vertical line) is estimated assuming a homogeneous response to the tax rebate as in Parker et al. (2013), following Equation (1). For each household we compute the weighted MPC across groups. The black vertical line shows the average weighted MPC in our sample.

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 app user data, respectively. We confirm that even the smallest MPCs are substantially larger than zero, 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. As we discuss in Section 3.3 and Appendix C.3, our simulations suggest that an extended left tail could be a spurious feature of the quantile regression approach in this setting, which would explain the difference in our recovered MPC distributions. A lower bound of the MPC distribution above zero can be explained by 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."

Aggregating the individual-level responses, we also find a larger average marginal propensity to consume than from the homogeneous regression, as shown by the black

persion in the tails. HKM does not estimate any MPC higher than 0.45. The Gaussian mixture model, in contrast, predicts higher dispersion on both ends of the distribution, also estimating some negative MPCs; we find exaggerated dispersion to be a property of the Gaussian mixture model in simulations, see Supplement C.2.

Table 1: Test for MPC equality

(a) Analytical standard errors

	MPC				
	0.25	0.38	0.57	0.71	0.76
0.25	6.02				
	(0.01)				
0.38	0.37	1.87			
	(0.54)	(0.17)			
0.57	6.02	0.37	0.17		
	(0.68)	(0.82)	(0.48)		
0.71	8.05	1.63	0.04	10.51	
	(0.00)	(0.20)	(0.84)	(0.00)	
0.76	0.21	0.14	0.02	0.00	0.43
	(0.65)	(0.71)	(0.89)	(0.97)	(0.51)

(b) Conditional on FCP weights

	MPC				
	0.25	0.38	0.57	0.71	0.76
0.25	154.55				
	(0.00)				
0.38	8.28	78.70			
	(0.00)	(0.00)			
0.57	2.23	0.75	7.33		
	(0.14)	(0.39)	(0.01)		
0.71	115.12	32.48	0.45	304.07	
	(0.00)	(0.00)	(0.50)	(0.00)	
0.76	4.37	2.35	0.35	0.04	9.89
	(0.04)	(0.13)	(0.55)	(0.85)	(0.00)

Notes: MPCs for total expenditures. The two tables show $F-$statistics from pairwise two-sided Wald tests of equality across MPCs (the diagonals show tests of equality with zero). The left panel uses the standard errors outlined in Theorem 2. The right panel repeats the exercise, taking the weights as given. These are equivalent to weighted least squares estimates where the weights are those in Equation (8), raised to the power m, as in (12). Therefore, to run the tests in Table 1(b), we replicate the sample by the number of groups and estimate $\Delta C_j = \beta' W_j + \sum_{g \in G} \left(\phi_g \mathbf{1}\left[j \in g\right] R_j + \alpha_g \mathbf{1}\left[j \in g\right] \right) + \epsilon_j$ via weighted least squares, with standard errors corrected for heteroskedasticity, and compute the Wald tests. $p-$values are reported in parentheses.

and red vertical lines, respectively. This is not surprising; in general, estimates from a homogeneous specification like Equation (1) will not equal the average of estimates from our heterogeneous specification, unless the distribution of right-hand-side variables is invariant across groups. We discuss this point further in Supplement E, and provide a statistical decomposition to determine which variables drive this discrepancy.[22] We find that variation in the mean of regressors across groups (particularly for certain time dummies and the age variable) is principally responsible.

In Table 1 we examine whether the estimated MPCs differ from zero and are statistically different from one another. In the left panel, we make use of the analytical formulas outlined in Theorem 2 to compute Wald tests of pairwise equality across MPCs. Groups 3 and 5 (ordered from lowest to highest MPC), whose MPCs are not statistically significantly different from zero, are also those with the smallest share of households.[23] Relative to the equivalent weighted least squares regression in which the weights are treated as given, the FCP standard errors are larger because they take estimation error of the weights into account. Table 1(b) shows that – when group assignment is taken as given – all MPCs are statistically different from zero. Moreover, most MPC groups are statis-

[22]See Słoczyński (2020) for a recent discussion of this issue in a model with binary treatment.
[23]4% and 3% of rebate recipients have maximum weight on groups 3 and 5, respectively.

tically different from each other, at least at the 68% confidence level. Supplement D.3 further shows that the full *distribution* of MPCs is largely invariant when re-estimated on bootstrap samples drawn from the data.[24]

The flexibility of the FCP methodology allows us to nest instrumental variable estimation. This is particularly relevant in our framework, since the exogenous source of the transitory income shock is driven by the random timing of rebate receipt, but the value of the rebate itself may be endogenous. We therefore follow the literature and instrument the tax rebate with an indicator variable for its receipt. In this TSLS specification, we first regress the rebate value on the rebate indicator and the same controls as in Equation (2), and then use the predicted values in the second stage. Moreover, we also address potential bias arising from the inclusion of households that never get the rebate, excluding them from the analysis.[25]

Figure 2 plots the resulting distribution of weighted MPCs, and shows how it remains qualitatively unchanged relative to the OLS specification. If anything, instrumentation uncovers a non-negligible portion of households that consume the rebate in its entirety and that even display an MPC slightly larger than 1. Moreover, the gap between aggregated and homogeneous response is even larger than in OLS.

4.2 The MPC distribution for different consumption goods

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

First, in the left panel of Figure 3, we report the weighted MPC distribution for nondurable goods.[26] As expected, the distribution is shifted to the left with respect to the distribution corresponding to total expenditures in Figure 1, as nondurable goods account for, on average, only 57% of household total expenditures.

An important share of households consumes a value of nondurables consistent with the annuity value of the rebate, as suggested by the Permanent Income Hypothesis (Fried-

[24] In particular, we repeat the estimation of the distribution of MPCs for total expenditures, with 5 groups, over 250 samples obtained with bootstrap with replacement. We find that the average quantiles across bootstraps are very close to those estimated in the baseline sample, and fairly stable across bootstraps.

[25] Households may never receive the rebate because they have different characteristics, such as higher income. In Supplement D.5 we show robustness to the exclusion of this group for OLS too.

[26] Nondurable goods 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).

Figure 2: Estimated distribution of MPCs out of the tax rebate: two-stage least squares

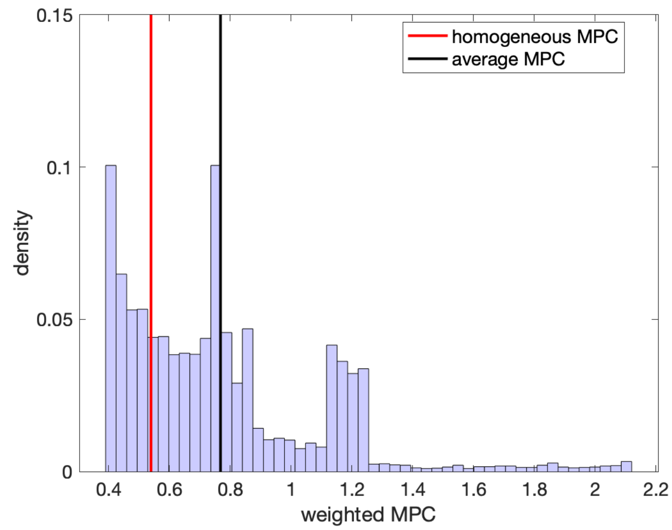

Notes: Figure 2 plots a histogram (light blue bars) of the estimated distribution of MPCs for total expenditures, defined as in Parker et al. (2013), using the two-stage least squares specification. We also drop from the sample the households that never receive a rebate. The homogeneous MPC (red vertical line) is estimated assuming a homogeneous response to the tax rebate, as in Parker et al. (2013), following the two-stage least squares equivalent of Equation (1). We estimate the model with $G = 5$. For each household we compute the weighted MPC across groups. The black vertical line shows the average weighted MPC in our sample.

man (1957)): roughly one third of households have an MPC that is not statistically distinguishable from zero. This suggests that the large lower bound on MPCs, recently found in the literature and discussed in the previous section, could depend on the type of good purchased. The majority of households have a small, but non-zero, propensity to consume nondurable goods, and around 8% of households consume nearly half of their rebate in nondurables. The heterogeneity in nondurable MPCs is not only economically meaningful, but also statistically significant. In Supplement D.3 we show that nearly all the estimated MPCs are statistically different from each other. Instrumenting the rebate with the rebate receipt indicator slightly increases both the mass and the values at the right tail, similar to the results for total expenditures.

As shown in the right panel of Figure 3, we estimate that 73% of households do not change their durable expenditures in response to the rebate; their weighted MPC is in a neighborhood of zero and the associated group-specific MPC is not statistically different from zero. A small fraction of households, however, has a durable MPC larger than one half.[27] The dichotomy of this MPC distribution is in line with the discrete nature of durable goods purchases. This discreteness implies lumpy adjustment and is consistent with the fact that most households either use most of the rebate to purchase durables, or

[27] In line with the tendency shown for other consumption categories, estimating durable MPCs with TSLS uncovers a group with larger propensity, up to 1.49.

Figure 3: MPCs out of the tax rebate: nondurables and durables

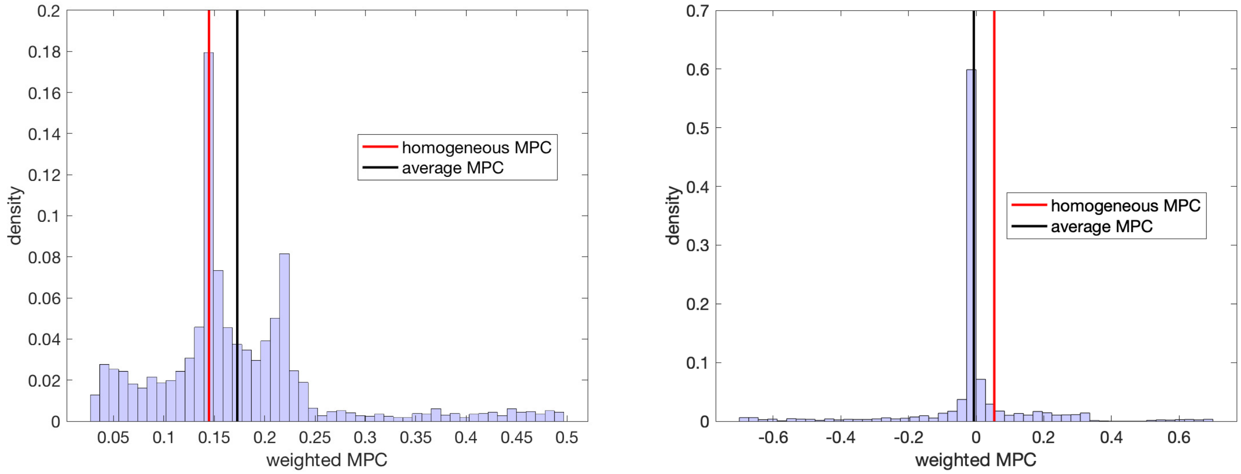

(a) Nondurables

(b) Durables

***Notes*:** Nondurable goods 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). The homogeneous MPC (red line) is estimated assuming homogeneous response to the tax rebate. For each household we compute the weighted MPC across groups. The black line shows the average weighted MPC in our sample. We estimate the models with $G = 5$ to allow direct comparability with the distribution of MPCs for total expenditures. We follow Coibion et al. (2017) and define durables as durable health expenditures, entertainment durables, furniture, jewelry, durable personal care, vehicle purchases, durable vehicle expenditures, housing durable expenditures (e.g., maintenance and repair commodities such as paint, materials.).

do not adjust their durable goods consumption at all.[28]

Finally, we assess whether households with high propensities to consume nondurable goods are also more likely to consume durable goods after receiving the rebate. While we can rule out substitution between goods, the estimated complementarity – at the margin – is, however, quantitatively small. The correlation between household-level weighted MPCs for nondurable goods with those for durables is 0.04, significant at the 5% level, while the rank correlation is also 0.04. Albeit small, the 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 permanent unobserved heterogeneity in MPCs, we can investigate what observable characteristics explain the estimated MPC distributions that we recover. We tackle this issue in the next section.

[28]Our definition of durables follows Coibion et al. (2017). Parker et al. (2013), instead, define durable expenditures as the difference between total and nondurable expenditures. Using this categorization delivers very similar results.

4.3 What drives MPC heterogeneity?

Our approach uncovers the 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 MPCs correlate, *ex post,* with observable characteristics. We start by examining how observables are individually correlated with MPCs. We then turn to investigate the joint relationship between the estimated MPCs and various household characteristics. As such, we contribute to the literature in two ways. First, we show that, with our approach, a large number of statistically significant individual correlations between MPCs and observable drivers emerge. This is true despite the fact that we use a dataset and an identification strategy which previously failed to find statistically significant relationships (see Parker et al. (2013)). We believe that this contrast underscores the power of our approach. Second, we show how the distribution of MPCs is jointly correlated with observable characteristics, and can be confident that any lack of significant correlations is not due to loss of statistical power introduced by progressive interactions.

Table 2 reports individual correlations. Our estimated weighted MPCs (column (1)) are positively correlated with salary and non-salary income, the mortgage interest-to-income ratio, the average propensity to consume (APC), and log liquid wealth; however, they are negatively correlated with age.[29] As shown in columns 2-4, these findings barely change when considering the MPC associated with an individual's maximal weight, or the MPC distribution estimated via TSLS.[30] We also find that homeowners have larger MPCs, a result that echoes findings in Parker et al. (2013). Moreover, having a mortgage is associated with an even higher propensity to consume, as shown in Figure 4.

Table 3 reports multivariate regressions of the estimated MPCs on the same observables. Since the inclusion of liquid assets reduces the number of observations, we also estimate the MPC drivers without taking liquidity into account. The absence of a correlation between the MPC and liquid assets is in line with findings by Parker et al. (2013), who suggest it might be driven by non-response bias.[31] Our results are robust across specifications and to using alternative MPC definitions.[32]

[29] Additional relationships hold unconditionally. For instance, we find that households that put money into a tax-deferred or tax-free educational savings plan have a significantly higher MPC. Moreover, MPCs increase with education. All these relationships, however, are insignificant when tested jointly with other observables as in Table 3.

[30] In unreported results we also find that correlations with observables are robust to the exclusion of observations associated with statistically insignificant MPCs for total expenditures.

[31] Considering the ratio of liquid assets to total income leaves the coefficients unaffected. Moreover, this variable is unconditionally uncorrelated with the marginal propensities to consume.

[32] All of the regression results are broadly unaffected if using the MPC associated with an individual's

Table 2: Individual correlations with the MPC for total expenditures

	(1)	(2)	(3)	(4)
log salary income	0.11***	0.09***	0.10***	0.09***
log non-salary income	0.19***	0.17***	0.18***	0.18***
mortgage interest/income	0.06***	0.07***	0.05**	0.04*
APC	0.05*	0.05**	0.02	0.00
age	-0.04**	-0.04*	-0.04**	-0.03*
log liquid assets	0.12***	0.11***	0.14***	0.13***

Notes: Table 2 shows the correlations between MPC estimates listed in columns and observables listed in rows. Column (1) uses the weighted MPC from the WLS specification, while column (2) uses the MPC associated with the maximal weight from the WLS specification. Columns (3) and (4) use the weighted and maximal weight MPCs, respectively, based on the TSLS specification, respectively. All logged variables take a value of log(0.001) when the raw value is 0 or negative. *, ** and *** denote significance of the correlation at 10, 5 and 1% levels, respectively.

Importantly, only two explanatory variables remain statistically significant after the inclusion of additional covariates: non-salary income and the average propensity to consume, both of which are positively correlated with the marginal propensity to consume. We expand on these two drivers in the remainder of the section.

While higher income households have higher MPCs, it is mainly the non-salary component of income that drives this relationship.[33] This effect is partly the result of a particular category of households, such as entrepreneurs or investors (for example, those with a positive business or financial income), who have a significantly higher MPC. The intensive margin, however, seems to play the most prominent role. Putting the estimates together, we find that a 1% increase in non-salary income is associated with an increase in the MPC between 5 and 13 cents for each dollar of the rebate, depending on the specification. Put differently, a 2% increase in non-salary income predicts a 1 standard deviation increase in the MPC, in our preferred specification in column (2). The importance of business and financial income for the MPC might suggest the presence of wealthy hand-to-mouth households, as first posited by Kaplan and Violante (2014). Finally, the positive correlation between income and MPCs does not only hold for total expenditures, but also for nondurable expenditures, even when controlling for additional covariates.

maximal weight as the dependent variable.

[33]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). Business and financial income are positively correlated with the MPC.

Figure 4: MPCs by housing status

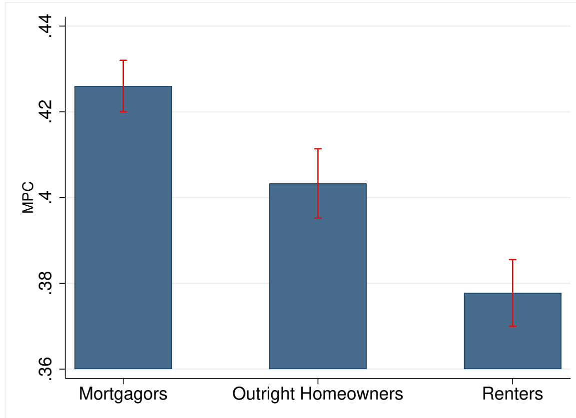

Notes: Figure 4 shows the average weighted MPC for total expenditures by housing status of households receiving a rebate. Error bars show 90% confidence intervals, computed with the standard deviation of weighted MPC within each group.

Some studies find 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. Other studies, however, find mixed results or even the opposite relationship, as we do. 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". 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. We instead find a monotonic relationship with income, as in Kueng (2018), who studies consumption responses to regular and predetermined payments from the Alaska Permanent Fund. Boutros (2020) finds that households whose 2008 rebate was a smaller fraction of their income – typically higher-income households – had a higher MPC. He explains this finding with 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. The theory of limited cognitive perception developed by Ilut and Valchev (2020) also delivers rich agents with high MPCs. 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. 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.

Table 3: Explanatory variables: weighted MPC for total expenditures

	(1)	(2)	(3)	(4)
dummy for no salary	-0.032	-0.125	-0.361	-0.379
	(0.108)	(0.097)	(0.282)	(0.235)
log salary income	0.001	-0.006	-0.016	-0.021
	(0.007)	(0.006)	(0.018)	(0.015)
log non-salary income	0.046^{***}	0.066^{***}	0.097^{***}	0.125^{***}
	(0.012)	(0.009)	(0.029)	(0.022)
mortgage interest/income	0.054	0.010	0.146	0.073
	(0.041)	(0.034)	(0.092)	(0.075)
APC	0.063^{***}	0.092^{***}	0.116^{***}	0.146^{***}
	(0.016)	(0.013)	(0.031)	(0.025)
homeowner dummy	0.032^{*}	0.024	0.026	0.022
	(0.015)	(0.012)	(0.039)	(0.031)
dummy for mortgage	-0.025	-0.013	-0.042	-0.026
	(0.015)	(0.012)	(0.036)	(0.028)
log liquid assets	-0.001		0.004	
	(0.003)		(0.006)	
dummy for liquid assets <= 0	-0.004		0.054	
	(0.040)		(0.089)	
N	723	1079	723	1079
R^2	0.101	0.112	0.101	0.092

Notes: Columns (1) and (2) use the weighted MPCs from our WLS specification. Columns (3) and (4) use the weighted MPCs from our TSLS specification. All logged variables take a value of log(0.001) when the raw value is 0 or negative. Non-salary income is positive for all observations. Standard errors are robust to heteroskedasticity and reported in parentheses. We control for marriage dummies, education dummies, number of children, age and age squared; those coefficients are not reported. Age and its square are controls in our FCP estimation. While this does not pose an issue for the point estimates shown in this table, it could potentially affect inference. However, we repeat the same regressions, excluding age and and age squared, and find that the remaining coefficients are unaffected. , *, and *** denote significance at the 10, 5, and 1% levels respectively.

Marginal propensities to consume also increase with the average propensity to consume (APC). Empirically, we define the APC as the ratio between lagged consumption and lagged total income. We consider income as measured in the first interview for each household, which refers to the previous 12 months. We lag expenditures to avoid the possibility of a mechanical positive correlation with the MPC. To ensure stability of APCs, 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. Households that spent 1 percentage point more of their income before receiving the rebate spent 9 additional cents out of each rebate dollar. This effect is significant also for nondurable MPCs and conditional on a wide array of controls.[34]

In Figure 5 we show 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 weighted MPC for each quintile pair. The MPC increases with income, conditional on the APC, and vice versa. As the figure shows, our analysis uncovers 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, and vice versa, display intermediate MPCs. Third, the greatest marginal propensity to consume is found among households with a high APC and high total income. We label this group "rich spenders".[35]

We regard the results presented in this section as particularly relevant for disciplining macro models of household consumption. First, the average propensity to consume can easily be computed in a large number of micro datasets with minimal information. Second, these correlations can be directly tested in even the simplest of consumption/savings models. Existing models, however, predict very different relationships between MPC and APC. Hand-to-mouth, constrained agents, will typically have large MPC and APC. As they save towards their target level of wealth, both propensities fall. If agents are infinitely-lived, they eventually reach the target level of wealth, at which they stop saving (i.e.: APC = 1) and have an MPC equal to the annuity value of the transitory income shock. A life-cycle model, in contrast, can generate the empirically observed positive relationship between MPC and APC, as old households dissave, but also have a high MPC due to a low effective discount factor. This standard model, however, generates a clear

[34] A 1 percentage point increase in the APC for total expenditures predicts 2 additional cents per rebate dollar spent on nondurables. This effect goes up to 6 when considering the APC for nondurable expenditures only.

[35] We find similar relationships for MPC for nondurable goods, especially the presence of "rich spenders", as we show in Supplement (D.4).

Figure 5: The relationship between MPCs, APCs, and income

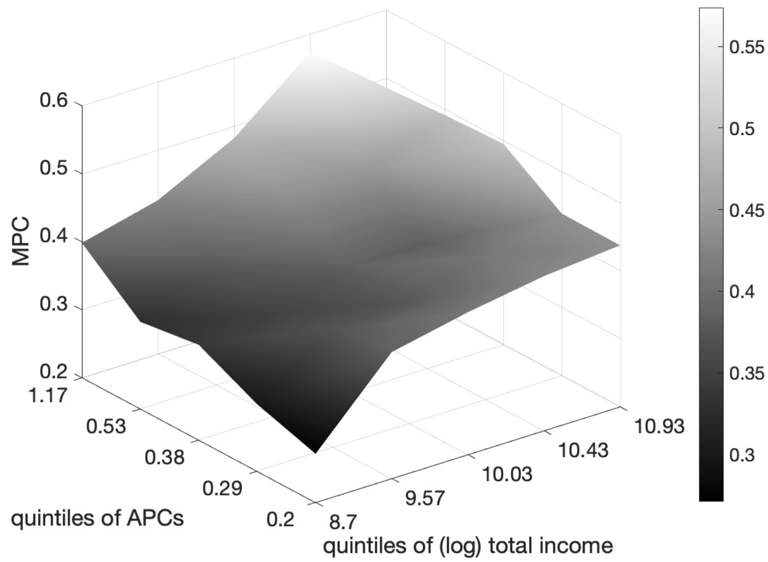

Notes: Figure 4 shows the average weighted MPC for total expenditures for pairs of quintiles of APC and log total income. The colorbar on the right represents the MPC.

relationship between the MPC and age, which our results do not bear out.[36]

All these characterizations are conditional on homogeneous preferences. Preference heterogeneity, in contrast, can break these relationships and rationalize some of our findings. Aguiar et al. (2019), for instance, highlight the importance of heterogeneity in the intertemporal elasticity of substitution in order to generate heterogeneous target levels of wealth; high-IES households have high MPCs and high APCs. Consistent with this, 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."

An additional finding underscores the importance of unobserved heterogeneity. All the observable drivers mentioned in this section — 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 of weighted MPCs on observable characteristics delivers an R^2 of 11%, at best. While other recent studies also find that most observable households characteristics are jointly

[36]Moreover, most incomplete markets 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).

uncorrelated with MPC heterogeneity, we can exclude that this finding is solely due to loss of statistical power.[37] Moreover, through the R^2 we can formally obtain a statistical measure of the portion of the variance in the MPC distribution explained by observable characteristics.

A low R^2 could be partly explained by non-linear relationships that are either difficult to parametrize, or are not captured by variables in our dataset.[38] For example, the CEX contains only sparsely populated information on wealth. In Supplement D.4, we show the relationship between the MPC and liquid wealth, aware of the potential nonresponse bias highlighted by Parker et al. (2013), but we refrain from showing any relationship with total wealth, given the lack of reliable data. While such unmeasured characteristics could potentially explain some of the variation in MPCs, our results strongly suggest the presence of latent drivers, since some of those unobserved characteristics may give rise to the observables we analyze in the first place, such as the APC.

4.4 Robustness to spurious heterogeneity

In this section we show that neither the MPC heterogeneity we uncover, nor its correlation with observables, is a spurious product of our estimation approach. For this exercise, we generate data using estimates from the homogeneous regression, with errors drawn from a Gaussian distribution with the empirical variance. We then obtain FCP estimates under the faulty assumption that five groups are present, and repeat the same analysis for 250 Monte Carlo samples. Table 4 shows that imposing spurious heterogeneity on a homogeneous distribution does not significantly bias moments of the distribution.

Limited spurious heterogeneity, however, arises towards the tails, driven by particularly noisy draws in the simulation. To show that these spuriously estimated MPCs do not invalidate our headline results, we regress the estimated weighted MPCs for each sample on the array of observable predictors used in specification (2) of Table 3. On average across samples, all the estimated correlations are almost exactly 0. For illustrative purposes, Figure 6 displays the distribution of the $t-$statistic for the coefficient on the APC, across samples. In only 6.8% of the samples is there significant evidence of a relationship between MPC heterogeneity and APC at the 5% level, a size distortion within the scope of Monte Carlo error. The same is also true for coefficients on all other observables, with even lower shares of significant coefficients. Further, a common concern when

[37] See, for instance, Table 6 in Fagereng et al. (2016).

[38] As discussed, our results are robust to different sets of explanatory variables. We also run a linear Lasso for the selection of the array of predictors. Regressing the MPC on the selected right hand side delivers the same R^2.

Table 4: Over-fitting G: average quantiles of the MPC distribution across simulated samples

	Average	p25	p50	p75
truth:	0.31	0.31	0.31	0.31
$G = 1$	0.31	–	–	–
	(0.11)			
$G = 5$	0.30	0.25	0.30	0.35
	(0.27)	(0.27)	(0.27)	(0.28)

Notes: Table 4 reports the average of various summary statistics of the distribution of (weighted) MPCs across Monte Carlo samples. The first row reports the truth, which is 0.31 for all statistics, since the distribution is homogeneous. The second row corresponds to a correctly-specified homogeneous regression in repeated samples (with the standard deviation across samples below in parentheses) and the fourth to FCP incorrectly assuming the presence of five distinct groups. pxx denotes the xx^{th} percentile. The final row reports the standard deviation of each moment for the $G = 5$ specification across simulated samples. Results are roughly unchanged computing medians across simulated samples, as well as distributions of the MPC associated with an individual's maximal weight.

using the CEX data is the role of measurement error. These exercises also serve to show that grouping on noise alone – like measurement error – does not lead to either the type of distribution that we recover or the correlations with observable characteristics that we estimate.

4.5 Aggregate partial equilibrium effects of the 2008 ESA

In this section, we estimate the partial equilibrium (PE) aggregate response to the 2008 tax rebate based on our estimated heterogeneous coefficients. For this exercise, we use a lagged specification that takes into account the possible persistent effects of rebate receipt, as in Parker et al. (2013). In particular, we estimate the following model:

$$\Delta C_j = \beta' W_j + \sum_{g \in G} \left(\phi_g \mathbf{1}\left[j \in g\right] R_j + \phi_g^{\text{lag}} \mathbf{1}\left[j \in g\right] R_j^{\text{lag}} + \alpha_g \mathbf{1}\left[j \in g\right] \right) + \epsilon_j, j = 1, \ldots, N, \tag{10}$$

where the coefficient ϕ_g^{lag} represents the lagged effect of the rebate for group g.[39] We do not force a household group membership for household i to be fixed across t, since we want to preserve flexibility. Even if individuals' preferences may be constant, the MPC may be time-varying, due, for instance, to changes in state variables such as income and wealth. To correctly estimate the cumulative response to the rebate, we therefore

[39] Kaplan and Violante (2014) suggest that the rebate coefficient might differ from the marginal propensity to consume because some households in the control group have already received the rebate, and some households might anticipate receiving the rebate in the future. Adding lagged rebate partially address this concern. See Supplement D.5 for further discussion of this specification.

Figure 6: (Lack of) spurious correlations with observables

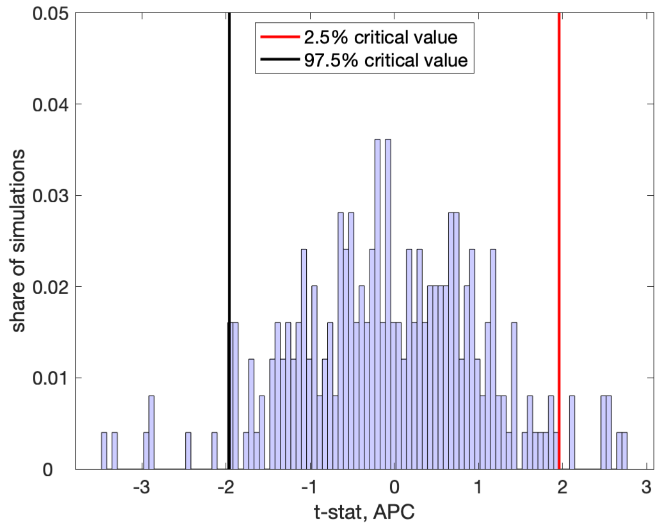

Notes: For each of 250 simulated samples, we regress the weighted MPCs for total expenditure estimated imposing spurious heterogeneity on the set of observables used in specification (2) of Table 3. Figure 6 plots the histogram (light blue bars) of the$t-$statistics for the coefficient on the APC. Total expenditures and baseline OLS specification. The red and black lines represent the critical values for a 5% test of equality with zero.

track individual weights over the two quarters following the rebate. We use these to construct the individual 2-quarter total effect of the rebate, by adding twice the weighted contemporaneous rebate coefficient to the weighted lagged coefficient.[40]

Figure 7 plots a histogram of this object among those who received the rebate. Relative to the baseline results depicted in Figure 1, the distribution spreads out, with some households having a total effect near zero. Moreover, as depicted in Figure 7, the estimated partial equilibrium effect of the tax rebate more than doubles relative to its homogeneous counterpart, from 0.28 to 0.69. Indeed, accounting for individual heterogeneity of treatment effects becomes even more important, even if the object of interest is simply the average response; as a share of 2008-Q1 aggregate consumption in our sample, the total effect is equivalent to 3.6%. Broda and Parker (2014) find that the stimulus increased consumption by 1.3% and 0.6% in Q2 and Q3, respectively.

[40]For example, a household may be categorized to be in some group a in the period in which they receive the rebate, and then in some group b the period after they receive the rebate. For such an individual, we construct the individual 2-quarter total effect of the rebate by adding twice the contemporaneous rebate coefficient for group a to the lagged rebate coefficient of group b, since ϕ_g^{lag} captures the change in consumption *relative* to consumption in the period of rebate receipt.

Figure 7: Estimated distribution of total 2-quarter effect of the tax rebate

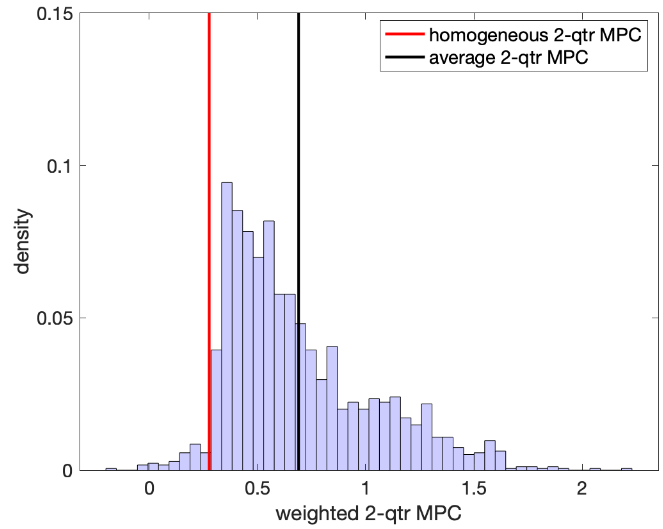

Notes: Figure 7 plots the histogram (light blue bars) of the estimated distribution of the total effect of the 2008 ESA for total expenditures, defined as in Parker et al. (2013), using the lagged specification in Equation (10). The homogeneous MPC (red vertical line) is estimated assuming a homogeneous contemporaneous and homogeneous lagged response to the tax rebate, as in Parker et al. (2013). For each household we compute the weighted MPC, weighted across groups $g \in G$. The black vertical line shows the average weighted MPC in our sample.

5 Conclusion

We develop a flexible clustering approach to uncover heterogeneity in the marginal propensity to consume. In particular, our fuzzy C-parameters methodology extends the fuzzy C-means algorithm to regression problems, is well-suited to cross-sectional settings and has a better simulation performance than competing clustering approaches and quantile regressions. Our approach has three main advantages, which allow us to obtain as many empirical contributions. First, we can estimate the full unconditional distribution of MPCs, which could be driven both by latent factors and observable characteristics, without having to take a stand on the latter. Second, we can investigate, *ex post*, how MPCs individually correlate with observables, without being subject to the loss of statistical power that is likely to affect existing methodologies. Third, we explore how MPCs are jointly correlated with observables, and obtain a formal estimate of the portion of the MPC variance they explain.

Taking FCP to the data on the 2008 Economic Stimulus Payments, we find that households display a considerable degree of heterogeneity in their marginal propensities to consume. Even after estimating the full unconditional distribution, households spent at least one fourth of the rebate within one quarter. Nondurable consumption, however, is characterized by smoothing for most households. We do not find evidence of individual-level substitution across consumption goods in response to transitory income shocks, but rather a very mild positive correlation. Finally, we explore what observables – individually and jointly – best predict different portions of the MPC distribution. We are able to retrieve various statistically significant relationships with the MPC, and show that only those with non-salary income and the APC survive the inclusion of additional drivers. Our findings suggest that there exists a tight relationship between marginal and average propensities to consume, which is easy to derive in many models of consumption behavior and yet has received relatively little attention. Moreover, since observable characteristics explain a minor portion of the estimated MPC heterogeneity, we posit that other latent factors, such as preference heterogeneity, might be important in determining marginal propensities to consume.

Finally, two caveats help to highlight possible avenues for future work. Importantly, we measure the distribution of MPCs to the 2008 tax rebate. This means our estimated distribution uses 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 windfalls, to which households may respond differently (Fuster et al. (2018)). However, the fuzzy C-parameters approach we develop 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.

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

Bezdek, J. (1981): *Pattern Recognition With Fuzzy Objective Function Algorithms*, New York: Plenum Press.

Bezdek, J. C. (1973): "Fuzzy Mathematics in Pattern Classification," *Ph.D. dissertation, Cornell University*.

Bezdek, J. C., R. Ehrlich, and W. Full (1984): "FCM: The Fuzzy C-Means Clustering Algorithm," *Computers & Geosciences*, 10, 191 – 203.

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): "Discretizing unobserved heterogeneity," *University of Chicago, Becker Friedman Institute for Economics Working Paper*.

Bonhomme, S. and E. Manresa (2015): "Grouped Patterns of Heterogeneity in Panel Data," *Econometrica*, 83, 1147–1184.

Boutros, M. (2020): "bounded Intertemporal rationality and the Marginal Propensity to Consume," Tech. rep.

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.

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*.

Cui, J., Q. Li, J. Wang, and D. Zong (2010): "Research on Selection Method of the Optimal Weighting Exponent and Clustering Number in Fuzzy C-means Algorithm," in *2010 International Conference on Intelligent Computation Technology and Automation*, vol. 3, 104–107.

Dunn, J. C. (1973): "A Fuzzy Relative of the ISODATA Process and Its Use in Detecting Compact Well-Separated Clusters," *Journal of Cybernetics*, 3, 32–57.

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," .

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.

JING, L., D. DENG, AND J. YU (2014): "Weighting Exponent Selection of Fuzzy C-Means via Jacobian Matrix," in *Knowledge Science, Engineering and Management*, ed. by R. Buchmann, C. V. Kifor, and J. Yu, Cham: Springer International Publishing, 115–126.

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.

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.

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

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

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

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.

PAL, N. R. AND J. C. BEZDEK (1995): "On Cluster Validity for the Fuzzy C-Means Model," *IEEE Trans. Fuzzy Systems*, 3, 370–379.

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.

ROUSSEEUW, P., E. TRAUWAERT, AND L. KAUFMAN (1995): "Fuzzy clustering with high contrast," *Journal of Computational and Applied Mathematics*, 64, 81 – 90.

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.

SŁOCZYŃSKI, T. (2020): "Interpreting OLS Estimands When Treatment Effects Are Heterogeneous: Smaller Groups Get Larger Weights," *The Review of Economics and Statistics*, forthcoming, 1–27.

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

TIBSHIRANI, R., G. WALTHER, AND T. HASTIE (2001): "Estimating the Number of Clusters in a Data Set via the Gap Statistic," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 63, 411–423.

WINKLER, R., F. KLAWONN, AND R. KRUSE (2011): "Fuzzy clustering with polynomial fuzzifier function in connection with m-estimators," *Special Issue on Fuzzy Set Theory and its Applications*, 10, 146–163.

WU, K.-L. (2012): "Analysis of Parameter Selections for Fuzzy C-Means," *Pattern Recognition*, 45, 407–415.

YANG, M.-S. (1994): "On Asymptotic Normality of a Class of Fuzzy C-Means Clustering Procedures," *International Journal of General Systems*, 22, 391–403.

YANG, M.-S. AND K. F. YU (1992): "On Existence and Strong Consistency of a Class of Fuzzy C-Means Clustering Procedures," *Cybernetics and Systems*, 23, 583–602.

YU, J., Q. CHENG, AND H. HUANG (2004): "Analysis of the Weighting Exponent in the FCM," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 34, 634–639.