American Economic Journal: Economic Policy 2020, 12(4): 275–297
https://doi.org/10.1257/pol.20180420

Reported MPC and Unobserved Heterogeneity†

By Tullio Jappelli and Luigi Pistaferri*

Panel data on reported marginal propensity to consume in the 2010 and 2016 Italian Survey of Household Income and Wealth uncover a strong negative relationship between cash on hand and MPC. Even though the relationship is attenuated when using regression methods that control for unobserved heterogeneity, the amount of bias is moderate. MPC estimates are used to evaluate the effectiveness of revenue-neutral fiscal policies targeting different parts of the distribution of household resources. (*JEL* E21, E62, G51)

An important parameter for evaluating the effectiveness of fiscal policy and for distinguishing between competing models of consumption behavior is the marginal propensity to consume (MPC). Most literature measures the MPC using structural models or quasi-experiments (see Jappelli and Pistaferri 2017, chapter 9, for a survey). A new wave of papers rely instead on a more direct measurement. The main advantage of this approach is that it does not require to take a stand on specific income processes or consumption models.

In particular, Shapiro and Slemrod (1995, 2003) pioneered the idea of eliciting the MPC from transitory income shocks using survey questions. Their approach is to ask respondents how they reacted to *actual* income changes induced by tax stimulus programs. A complementary approach is to use survey questions asking respondents to report their MPC in response to *hypothetical* income changes as in Jappelli and Pistaferri (2014). Another difference between these two approaches is that while Shapiro and Slemrod (1995, 2003) rely on quantitative but coarse responses (ranging from "mostly spend" to "mostly save"), in Jappelli and Pistaferri (2014), people report numerical information about the MPC (the percentage spent and saved). Recent contributions further distinguish between reported MPC in response to positive and negative transitory income shocks and between shocks of different magnitude (Christelis et al. 2019; Fuster, Kaplan, and Zafar 2018; Bunn et al. 2018).

Models of consumption behavior make strong predictions regarding MPC and their relation to household resources. For example, contrary to the standard permanent income hypothesis (PIH), which predicts a linear consumption function, models

* Jappelli: University of Naples Federico II, Via Cinzia, Monte S. Angelo, Napoli 80146, Italy, Center for Studies in Economics and Finance, and CEPR (email: tullio.jappelli@unina.it); Pistaferri: Stanford University, 579 Jane Lathrop Way, Stanford, CA 94305, Stanford Institute for Economic Policy Research, NBER, and CEPR (email: pista@stanford.edu). Matthew Shapiro was coeditor for this article. Jappelli acknowledges financial support from the Italian Ministry of Education, University and Research (PRIN Grant 2017RHFXK4). We thank three anonymous referees, seminar participants at the European Central Bank, and Marteen van Rooij, Dimitris Christelis, and Dimitris Georgarakos for useful comments.

† Go to https://doi.org/10.1257/pol.20180420 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

with precautionary saving and liquidity constraints generate a concave consumption function (Deaton 1991, Carroll 2001). One important implication of these models is that MPC is heterogeneous across households, and in particular, rich households should have smaller MPC than poor households. Another key prediction of these models is that the MPC from negative shocks is larger than the MPC from positive shocks because liquidity-constrained households can partially overcome the constraint if the income change is large enough.

A common finding of the papers estimating an MPC, regardless of approach, is indeed strong evidence for MPC heterogeneity. The few papers that distinguish between shocks of different sign also confirm the theoretical prediction that the MPC for negative shocks is larger than the MPC from positive shocks. The findings about the relation between MPC and household resources are more mixed, however. Bunn et al. (2018); Fuster, Kaplan, and Zafar (2018); and Christelis et al. (2019) all find that the MPC with respect to windfall losses declines with cash on hand but find little to no relation between MPC with respect to windfall gains and household resources. In this paper, we explore this relation further, drawing on Italian panel data with direct measures of the MPC. As we shall see, the question is identical over the years, and the sample spans a relatively long period (2010–2016).

One major issue with existing evidence is that the direct survey approach is based on cross-sectional data, where respondents are asked only once about actual or hypothetical income changes. In principle, both MPC heterogeneity and the negative association between MPC and household resources might be consistent with models with a linear consumption function and unobserved preference heterogeneity. To see this point, suppose that the consumption function of each individual is linear but that there is unobserved heterogeneity in discount rates (or, alternatively, heterogeneity in the propensity to leave bequests).[1] This would imply that people with low discount rates have a flatter consumption function (a lower but constant MPC) than people with high discount rates (a higher but still constant MPC). At the same time, people with low tastes for current consumption relative to future consumption have accumulated more wealth in the past and therefore have higher cash on hand (defined as current income plus wealth) than people with high discount rates, other things being equal. This combination of preference and resource heterogeneity generates a negative relation between MPC and cash on hand in the cross section even when the consumption function of each individual is linear.

To identify the shape of the consumption function while controlling for unobserved heterogeneity, one needs panel data on reported MPC and cash on hand. In this paper, we achieve this goal by relying on the rotating panel structure of the Italian Survey of Household Income and Wealth (SHIW) (Bank of Italy 2020). In 2010, survey respondents reported how much they would consume of a hypothetical, unanticipated, and transitory income change equivalent to a one-month increase in disposable income. Crucially, a group of households interviewed in 2010 were also reinterviewed in 2016, thus offering longitudinal data on MPC, cash on hand, and other demographic variables. The panel structure is the key advantage of our

[1] For instance, models with income risk and quadratic or exponential utility (allowing negative consumption) imply a linear relation between consumption and cash on hand.

data. While in previous work different questions are posed to the same person about the size and direction of income shocks, the unique feature of the SHIW is that the same question is asked to the same person at two different points in time.

Using cross-sectional data, we find that the MPC declines quite significantly with cash on hand. For example, moving from the tenth to the ninetieth percentile of the cash on hand distribution is associated with a reduction of the MPC by about 16 percentage points. Next, we use the panel structure of the SHIW to estimate the sensitivity of MPC with respect to cash on hand, controlling for unobserved heterogeneity. We find that OLS exaggerates the negative relationship between the two variables by around 20 percent. To follow up on the same example, going from the tenth to the ninetieth percentile of the cash on hand distribution would reduce the MPC by 12 percentage points.

The difference between the estimates obtained by cross-sectional and panel data supports the idea that unobserved factors correlated with cash on hand account for part of the relationship. However, it can also be noticed that the difference in point estimates is not large. Overall, the paper suggests that the amount of bias is relatively small and that cross-sectional survey data are broadly informative about the MPC, at least in the Italian context. This main finding is robust to various sensitivity checks regarding the specific functional form of the relation between MPC and cash on hand, sample selection, additional covariates, and quality of the interviews.

The paper proceeds as follows. In Section I, we summarize the literature that uses direct survey questions to measure the MPC. In Section II, we describe our panel data and compare the MPC distribution in 2010 and 2016. In Section III, we present the estimate of the relationship between MPC and cash on hand with cross-sectional and panel data. Section IV explores the robustness of the results, while Section V uses our estimates to calculate the impact of several revenue-neutral redistributive fiscal policies on aggregate consumption, showing that the impact calculated using cross-sectional or panel data is similar across fiscal experiments. Section VI summarizes the evidence and concludes.

I. The Direct Survey Approach

The direct survey approach to evaluate the impact of fiscal shocks on consumption consists of asking direct questions on how consumers have reacted to *actual* income changes or asking them to report how they would respond to *hypothetical* income changes. Shapiro and Slemrod (1995) pioneered this approach, asking direct questions in the University of Michigan Survey of Consumers. These questions elicited, in a quantitative but coarse format ("mostly spend," "mostly save"), the consumer response to the Bush administration's 1992 change in tax withholding. Subsequent work used a similar type of questions focusing on spending in response to the various tax rebates and tax credit interventions taking place in the United States in the past two decades (Shapiro and Slemrod 2003, 2009; Sahm, Shapiro, and Slemrod 2010, 2012, 2015). These studies find that consumers differ in reported MPC along many margins; however, the relationship between MPC and measures of household resources is typically nonmonotonic, and many households appear to use rule-of-thumb behavior to respond to fiscal policy.

Another way to elicit the MPC is to confront consumers with hypothetical scenarios in which income changes unexpectedly. Jappelli and Pistaferri (2014) use Italian survey data from the 2010 SHIW, where consumers were asked to report the fraction of a positive income shock (a hypothetical unanticipated tax rebate) that they would consume or save. They find considerable heterogeneity in the reported MPC and a strong negative relation between MPC and cash on hand.

The literature has extended in at least three directions. First, some papers have asked how reliable are reported MPC in predicting behavior in response to actual income changes. A second important issue discussed in the literature is whether heterogeneity in MPC is a spurious reflection of failure to control for unobserved preference heterogeneity. Finally, a handful of papers have explored how reported MPC varies in response to income changes of different sign and different magnitude.

One way to validate the informational content of MPC based on hypothetical questions is to see if planned consumption decisions are confirmed by actual consumption choices. Graziani, van der Klaauw, and Zafar (2016) study reported use of the extra income accruing from the 2011 US payroll tax cuts. Workers were surveyed in early 2011, just after the tax cut was first signed into law, and then in mid-December 2011, close to the expiration of the initial tax cuts. The first survey asked respondents how they intended to spend the extra funds, while the second survey inquired about ex post usage of the funds. The paper finds that workers intended to spend 14 percent of the extra income but ex post reported spending 36 percent of the funds.

Parker and Souleles (2019) use an alternative strategy and compare the "revealed preference" approach (in which inference is based on actual data) with the "reported response" approach, which consists of asking people to report their choices. They find that households reporting that they "mostly spent" their economic stimulus payments in 2008 had indeed spent twice as much as those reporting that they used their payments "mostly to save or pay down debt." Furthermore, the quantitative reported-response estimate of the average propensity to spend is close to the average revealed-preference propensity to spend. Reported spending, however, is unrelated to cash on hand (liquidity or income).

An important finding of the literature is that MPC varies considerably with personal traits and not only from temporary income shocks combined with precautionary savings or borrowing constraints. In other words, persistent characteristics such as preferences or behavioral traits can be an important driver of MPC heterogeneity. Gelman (2019) provides evidence on the two channels using panel data from a personal finance app with data on spending, income, and liquid assets. He finds that within-individual variation in cash on hand results from temporary income shocks, while across-individual variation in cash on hand results from differences in persistent characteristics.

Parker (2017) evaluates theoretical explanations for the propensity of households to increase spending in response to the arrival of a one-off stimulus payment, using households in the Nielsen Consumer Panel. Parker finds that low liquidity is a strong predictor of large spending responses, but this does not appear to be due to the extra income received but rather is a persistent characteristic of low-income households. Since the consumption response to income changes varies considerably with personal traits, the implication is that controlling for unobservable but persistent household characteristics is important.

Carroll et al. (2017) offer a possible explanation for MPC heterogeneity. The authors solve a macroeconomic model with a household-specific income process and heterogeneity in discount rates. Their model matches the wealth distribution and implies an aggregate MPC of around 0.2. Furthermore, it suggests that the aggregate MPC can differ greatly depending on how the shock is distributed across households. For example, low-wealth and unemployed households have much larger spending propensities than high-wealth and employed ones.

However, the evidence does not always point to a negative relation between wealth and MPC. For instance, Kueng (2018) finds high MPCs from large predetermined payments from the Alaska Permanent Fund and that the MPC is monotonically increasing with income. He advances a behavioral explanation, pointing out that deviations from consumption smoothing are more costly for poor households, while high-income households suffer only small losses from excess sensitivity. One explanation for high MPCs among households with relatively high liquid wealth is therefore that their consumption decisions are nearly rational.[2]

A handful of recent papers rely on direct survey questions similar to Jappelli and Pistaferri (2014) to study asymmetric and size-based responses, i.e., whether the consumption response to a hypothetical income shock varies with the sign and magnitude of the shock itself. Christelis et al. (2019) use a representative sample of Dutch households from the CentER internet panel. Respondents are asked to report how much their consumption would change in response to unexpected, transitory income shocks of different sign (positive and negative). The Dutch questionnaire also distinguishes between relatively small income changes (a one-month increase or drop in income) and relatively larger ones (a three-month increase or drop in income). These data indicate that consumers react more to negative income changes than to positive changes. Furthermore, Christelis et al. (2019) find a negative association between MPC and cash on hand for negative income shocks but essentially no relation for positive shocks. Bunn et al. (2018) use a set of questions in the Bank of England/NMG Consulting Survey and find that British households tend to change their consumption by significantly more in reaction to temporary and unanticipated falls in income than to increases in income of similar magnitude. They also find that low liquid wealth relative to income is associated with higher MPC in response to negative shocks than positive shocks. Fuster, Kaplan, and Zafar (2018) use data from the Federal Reserve Bank of New York's Survey of Consumer Expectations. In this survey, respondents report how they would adjust their spending over the next quarter in response to receiving or losing dollar amounts ranging from \$500 to \$5,000. As Bunn et al. (2018) and Christelis et al. (2019), they also find smaller consumption responses from positive income shocks and little relationship between

[2] After a careful review of the literature, Campbell and Hercowitz (2019) conclude that in the United States the responses to tax rebates indicate that the MPC is high (relative to the PIH benchmark) even for households with high wealth. They explain this puzzle by pointing out that people save in anticipation of major expenditures such as home purchases and college education. Adding such savings to the standard precautionary-saving model allows them to generate high MPCs even for households with liquid wealth.

wealth and MPC for positive shocks, even though they find strong relationships for negative shocks.

Limiting ourselves to studies relying on direct survey questions, this finding is in contrast with Jappelli and Pistaferri (2014), who use Italian data to show a strong negative relation between MPC and household resources from positive unexpected income shocks.[3] There are several potential explanations for these differences. First, the wording of the question differs across studies. While in Italy and the United Kingdom durables and nondurables are lumped together, in the Dutch and US data they are kept separate. Moreover, in the Dutch questions the horizon is explicit (12 months), while in Italy and the United Kingdom it is not. The US and UK surveys ask how people would respond to a fixed dollar amount, while in the Italian and Dutch surveys the hypothetical change is proportional to income. Finally, in the Italian dataset used by Jappelli and Pistaferri (2014), respondents are only asked about their reaction to positive income shocks, and hence no test of asymmetric behavior (in absolute or marginal terms) is possible.[4]

A different explanation is that in Italy, liquidity constraints may induce a stronger concavity of the consumption function than in other countries. Using theoretical simulations of an intertemporal model with income risk and liquidity constraints, Jappelli and Pistaferri (2014) are able to reconcile the Italian evidence with theoretical models using either a high fraction of rule-of-thumb consumers or an implausibly low discount factor. On the other hand, the United Kingdom, the United States, and the Netherlands feature more developed households' credit markets, which might induce a lower concavity of the consumption function and therefore a milder relation between MPC and cash on hand.

A final way to reconcile the findings of Jappelli and Pistaferri (2014) with those of the rest of the literature is to argue that the relation between MPC and household resources that they estimate may be biased by failure to account for unobserved heterogeneity. In principle, the data used by Christelis et al. (2019); Bunn et al. (2018); and Fuster, Kaplan, and Zafar (2018) allow them to control for within-person heterogeneity (since different MPC questions are asked to the same respondent). However, the role of unobserved taste for saving can differ depending on the sign of the hypothetical income change.[5] This implies that comparing responses to positive and negative income shocks for the same person does not necessarily "difference out" unobserved preference heterogeneity. In contrast, looking at the MPC from positive income shocks across two time periods is not context dependent and eliminates the bias (as long as taste for saving is stationary). This is precisely the approach taken in the present paper, which is unique in its availability of panel data on reported MPC, household resources, and other observable characteristics. As we

[3] There is a large literature based on observational data that has looked at how MPC out-of-income shocks differ by cash on hand. See Jappelli and Pistaferri (2010) for a survey.

[4] Sample sizes in the UK, US, and Dutch cases are also considerably smaller than in the Italian survey, potentially affecting test power.

[5] To consider a sharp example, assume a group of households with different tastes for saving is at a binding liquidity constraint where consumption equals income. Due to the inability to borrow, an unexpected income decline would change consumption by the same amount for everyone, independently of heterogeneity in tastes for saving. In contrast, a positive income change that overcomes the constraint generates different consumption responses that depend on how strong tastes for saving are.

shall see, we continue to estimate a strong, negative relation between cash on hand and MPC even controlling for unobserved heterogeneity.

II. The Data

The SHIW is a biannual representative sample of the Italian resident population. The surveys cover 7,950 households in 2010 and 7,416 households in 2016 and provide detailed information on demographic variables, income, consumption, and wealth (broken down into real assets and various components of financial assets and debt). The survey also has a rotating panel component: each year close to 50 percent of the sample is composed of households interviewed in the previous wave, while 50 percent represents new interviews.

For the present study, in particular, 2,138 households interviewed in 2010 were also interviewed in 2016.[6] To make sure that the question on hypothetical income change is answered by the same person, our panel sample further selects households with a stable demographic structure (the same household head and no change in marital status across the two waves). We end up with an estimation panel sample of 1,618 households.

To estimate the relation between MPC and cash on hand, we rely on the following question posed to respondents in the 2010 and 2016 SHIW:

> Imagine you unexpectedly receive a reimbursement equal to the amount your household earns in a month. How much of it would you save and how much would you spend? Please give the percentage you would save and the percentage you would spend.

While the dictionary meaning of "reimbursement" is a sum paid to cover money that has been spent or lost, in our context the question stresses more precisely that the reimbursement is received "unexpectedly." Thus, we assume that people interpret the question as referring to an unanticipated windfall gain similar to a tax rebate.[7] In Jappelli and Pistaferri (2014), we use the 2010 wave and discuss pros and cons of the survey question. The main advantage is that it provides a quantitative estimate of the MPC at the individual level, refining the quantitative but coarse approach of Shapiro and Slemrod (1995), which relies on a "mostly spend/mostly save" scale. However, several caveats are also in order: (i) the question does not distinguish between consumption and spending; (ii) the 2010 survey was fielded during a deep recession, and responses may be different during normal times or expansions; (iii) it may be hard for some people to answer this type of question, and actual MPC may differ from the reported ones; and (iv) the survey question offers no period of reference for the planned spending (i.e., 12 months, etc.).

[6] The survey was also conducted in 2012 and 2014, but MPC questions are comparable only for the 2010 and 2016 waves. Data are collected through personal interviews. Questions concerning the whole household are addressed to the household head or the person most knowledgeable about the family's finances. Questions on individual incomes are answered by the individual household member. The unit of observation is the family, defined as including all persons residing in the same dwelling who are related by blood, marriage, or adoption. Individuals described as "partners or other common-law relationships" are also treated as family members.

[7] The Italian wording for "unexpected reimbursement" is "rimborso inatteso."

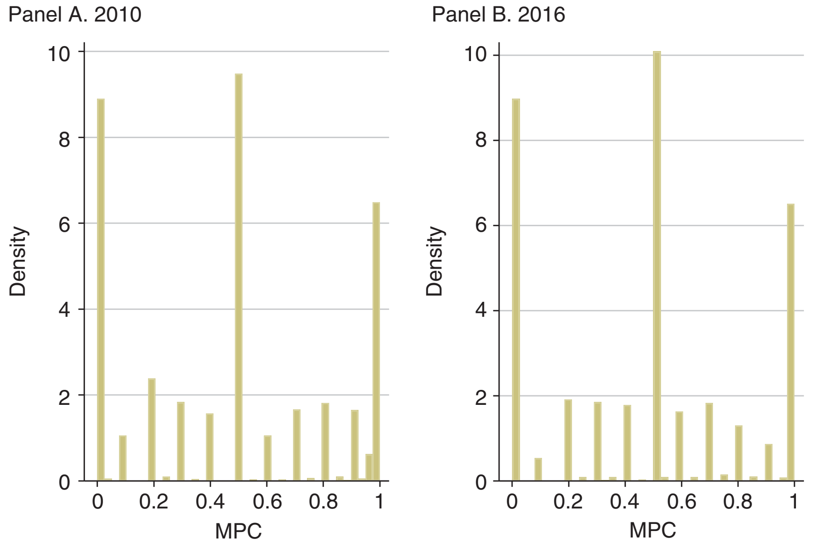

Figure 1. Histogram of the Distribution of Reported MPC, 2010 and 2016

Figure 1 plots the histogram of the cross-sectional distribution of reported MPC in the waves using all sample observations (7,950 households in 2010 and 7,416 in 2016). The figure shows that the two distributions are remarkably similar, supporting the reliability and information content of the data. The sample averages of the individual MPC are 48 percent in 2010 and 47 percent in 2016. Both distributions exhibit heaping at 0 percent, 50 percent, and 100 percent. In particular, in 2010, heaping at these three values is 22 percent, 24 percent, and 16 percent, respectively; in 2016, the values are slightly larger, at 24 percent, 27 percent, and 17 percent. Heaping and rounding can reflect uncertainty about responses or measurement error. We deal with these important issues in Section IV.

Table 1 provides descriptive statistics on the cross-sectional and panel samples we use in the regression analysis below, separately for 2010 and 2016. To conform to the survey question (which refers to a one-month income change), we define cash on hand as the sum of monthly income and the stock of financial assets (transaction accounts, mutual funds, stocks, outstanding claims, and corporate and government bonds), net of consumer debt. This definition of cash on hand is in line with Kaplan and Violante (2014), who argue that consumption in the short run is more strongly related to the liquid portion of total wealth since real estate can be liquidated only by incurring in high transaction costs.

Monetary variables are expressed in 2016 euros using the CPI. Table 1 shows that the cross-sectional sample does not differ appreciably from the longitudinal sample in basic demographic characteristics such as age, gender, etc. Households in the panel sample have slightly more schooling and are more likely to live in the North, which likely drives the difference in economic resources (cash on hand, income, and financial assets). Respondents also report whether they have been turned down for credit or were discouraged from applying for credit in the past 12 months. We

Table 1—Descriptive Statistics

Sample statistics	2010, All		2010, Panel		2016, All		2016, Panel	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
MPC	0.48	0.36	0.49	0.36	0.47	0.35	0.45	0.35
Age	58.37	15.76	58.43	13.82	62.17	15.67	64.43	13.82
Male	0.55	0.50	0.55	0.50	0.53	0.50	0.55	0.50
Married	0.62	0.49	0.65	0.48	0.53	0.50	0.65	0.48
Years of education	9.58	4.59	9.81	4.51	9.68	4.45	10.04	4.47
Resident in the South	0.32	0.47	0.37	0.48	0.33	0.47	0.37	0.48
Family size	2.49	1.26	2.60	1.28	2.22	1.21	2.38	1.21
Large city	0.09	0.29	0.06	0.23	0.08	0.28	0.06	0.23
Cash on hand (€1,000)	34.40	106.24	39.29	108.75	33.36	133.73	38.63	101.16
Income (€1,000)	2.95	2.19	3.10	2.18	2.54	1.90	2.83	2.07
Financial assets (€1,000)	31.45	105.21	36.19	107.55	30.82	133.06	35.80	100.19
Unemployed	0.04	0.18	0.04	0.20	0.06	0.23	0.04	0.20
Liquidity constrained	0.05	0.21	0.03	0.18	0.02	0.14	0.02	0.13
Observations	7,950		1,618		7,416		1,618	

Note: The table reports sample statistics for 2010 and 2016 SHIW and for the subsamples used in panel estimation.

use this information to construct an indicator of liquidity constraints. In the 2010 wave, which was conducted in the middle of a deep recession, 5 percent report to be liquidity-constrained as opposed to 2 percent in 2016.

Figure 2 starts delving in the relationship between MPC and cash on hand, again separately for the 2010 and 2016 waves. We allocate households to percentiles of the cash on hand distribution and plot the average MPC for each percentile together with a univariate regression line. The MPC declines quite significantly with cash on hand in each cross section. In 2010, a move from the tenth to the ninetieth percentile of the cash on hand distribution is associated with a reduction of the MPC of about 25 percentage points. In 2016, the same move is associated with an 18 percentage point decline in the MPC. In the next section, we use a regression framework to estimate the sensitivity of the MPC to cash on hand, using both the pooled cross sections as well as the panel sample.

III. Regression Evidence

To interpret the regression estimates, let's consider the following regression for the MPC:

$$MPC_{it} = \alpha + \beta X_{it} + f_i + \nu_{it}, \tag{1}$$

where X_{it} is cash on hand (or cash on hand percentile) of individual i in period t, f_i is unobserved heterogeneity potentially correlated with cash on hand, and ν_{it} an i.i.d. error term capturing classical measurement error in reported MPC. For simplicity, we omit exogenous and observable variables from equation (1), such as age, education, etc. However, we fully control for such characteristics in the regression analysis.

The relationship (1) nests several consumption models. In the PIH with quadratic utility and homogeneous preferences, the MPC is constant, and hence $\beta = 0$. In models with precautionary savings and/or liquidity constraints, the consumption

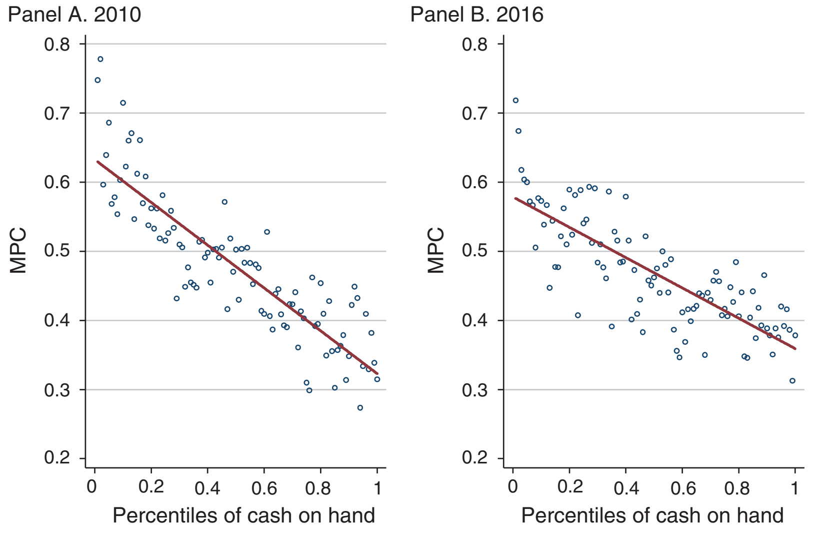

Figure 2. The Relationship between MPC and Cash on hand, 2010 and 2016

function is concave, and therefore the MPC is higher at low levels of economic resources, implying $\beta < 0$. A further reason for observing a negative relation is a nonhomothetic bequest motive, for instance treating intergenerational transfers as a luxury good in models where utility depends on terminal wealth. In support of the concavity of the consumption function, most papers (using cross-sectional data and OLS estimation) find $\hat{\beta}_{OLS} < 0$. In column 1 of Table 2, we confirm these findings, pooling data from 2010 and 2016, and obtain $\hat{\beta}_{OLS} = -0.27$. This coefficient estimate implies that a move from the tenth to the ninetieth percentile of the cash on hand distribution is associated with a 22 percentage point decline in the average MPC $(-0.27 \times (90 - 10))$.

However, in the presence of unobserved heterogeneity potentially correlated with cash on hand, the OLS estimate of β is biased and inconsistent. From regression (1), the bias can be inferred by computing the probability limit of $\hat{\beta}_{OLS}$:

$$p\lim \hat{\beta}_{OLS} = \beta + \frac{\operatorname{cov}(X_{it},f_i)}{\operatorname{var}(X_{it})}.$$

The expression above shows that the bias generated by unobserved heterogeneity (if it exists) depends on the sign and magnitude of the covariance term $\operatorname{cov}(X_{it},f_i)$. As discussed in the introduction, suppose that f_i represents unobserved differences in rates of time preference, implying that people with high values of f_i have high tastes for current consumption.[8] Since individuals with high rates of time preference have a tendency to report high MPC and may be more likely to have low cash on

[8] As standard, the validity of the fixed effects estimates rests on the assumption that the unobserved heterogeneity potentially correlated with the model regressors is time invariant. In our case, unobserved heterogeneity in MPC should reflect preference traits (such as discount rate, risk tolerance, etc.), which in the consumption literature are typically assumed to be permanent. It is possible that changes in economic circumstances may shift such traits. We control for employment status, family size, geographic mobility, etc. in the attempt to minimize this possibility.

Table 2—MPC Regressions Using Percentiles of Cash on Hand

Sample Estimation method	All OLS (1)	Panel Fixed effects (2)	All OLS (3)	Panel OLS (4)	Panel Fixed effects (5)
Percentiles of cash on hand	−0.266 (0.010)	−0.153 (0.047)	−0.197 (0.011)	−0.179 (0.025)	−0.147 (0.047)
Aged $\leq$30			0.052 (0.017)	0.067 (0.054)	0.037 (0.103)
Aged 30–45			0.036 (0.009)	0.020 (0.022)	0.001 (0.053)
Aged 45–60			0.030 (0.007)	0.033 (0.015)	0.025 (0.033)
Male			−0.005 (0.006)	−0.016 (0.013)	
Married			−0.017 (0.007)	−0.021 (0.017)	
Years of education			0.002 (0.001)	0.002 (0.002)	0.017 (0.012)
Resident in the South			0.113 (0.006)	0.117 (0.014)	
Family size			0.014 (0.003)	0.022 (0.007)	−0.001 (0.017)
City size >500,000			0.087 (0.010)	0.066 (0.026)	−0.061 (0.265)
Dummy for 2016			−0.004 (0.006)	−0.022 (0.012)	−0.035 (0.014)
Unemployed			0.064 (0.014)	0.036 (0.031)	−0.065 (0.054)
Credit constrained			−0.012 (0.015)	0.068 (0.038)	0.000 (0.053)
Constant	0.607 (0.006)	0.552 (0.026)	0.470 (0.010)	0.455 (0.022)	0.399 (0.129)
R^2	0.05	0.58	0.08	0.09	0.59
Observations	15,366	3,236	15,366	3,236	3,236

Note: We report standard errors in parentheses.

hand, we expect $\text{cov}(X_{it}, f_i) < 0$. Therefore, $\hat{\beta}_{OLS}$ will be greater (in absolute value) than the true β, and the OLS estimate will exaggerate the impact of cash on hand on the MPC. A policymaker who wants to forecast the impact of an expansionary fiscal policy targeting low-income households using $\hat{\beta}_{OLS}$ will predict larger effects than typically produced once the policy is in place.

With panel data, one can eliminate the bias by differencing the relationship (1) and hence estimate[9]

$$\Delta MPC_{it} = \beta \Delta X_{it} + \Delta \nu_{it}. \tag{2}$$

[9] Christelis et al. (2019) control for unobserved heterogeneity by considering within-person differences in MPC. This is because the same person responds to questions eliciting the MPC with respect to income changes of different sign and magnitude. Their approach can only identify *differences* in the sensitivity of MPC with respect to cash on hand across different scenarios (of income changes of different sign and size). However, the policy-relevant parameter (the *actual* sensitivity of MPC with respect to cash on hand) is not identified and can only be estimated using genuine panel data as we do in this paper.

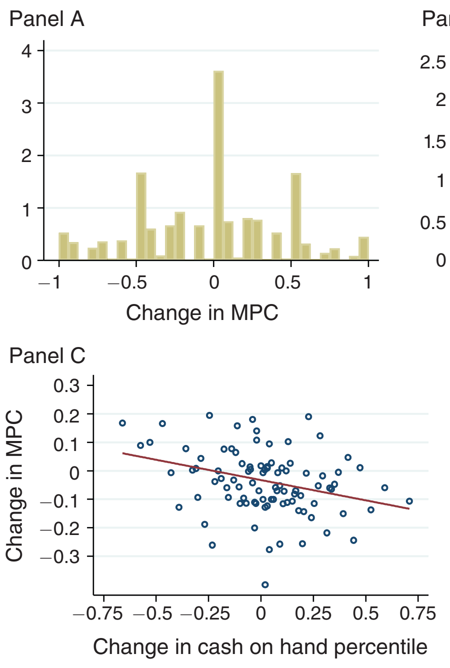

Figure 3. Panel Data Evidence: The Distribution of the Change in MPC and the Relation between the Change in the MPC and the Change in Cash on Hand

Panels A and B of Figure 3 report the histograms of the dependent and independent variables of equation (2), the change in the MPC (panel A), and the change in the percentile of cash on hand (panel B). There is much less heaping in the distribution of changes in MPC than in the level of MPC in the cross-sectional distribution of Figure 1. There is also considerable mobility in the cash on hand distribution, which is useful for identification purposes. In panel C of Figure 3, we plot the change in MPC against the change in the percentile of cash on hand together with a regression line, a way of describing graphically the relation in equation (2). The estimated coefficient (reproduced in column 2 of Table 2) is -0.15, implying that a move from the tenth to the ninetieth percentile of the cash on hand distribution is associated with a 12 percentage point reduction in the MPC $(-0.15 \times (90 - 10))$, significantly less than the 21 percentage point decline we found when using OLS. This suggests that unobserved heterogeneity may potentially account for part of the correlation between MPC and cash on hand estimated with cross-sectional data.

As mentioned, some of the bias may be due to failure to control for observable characteristics correlated with cash on hand. In the remaining columns of Table 2, we provide estimates of β obtained after introducing in the regression a rich set of demographic and socioeconomic characteristics of survey participants. In particular, besides the percentile of cash on hand, we include age dummies, gender, marital status, years of schooling, residence in the South and a large city, family size, a dummy for unemployment, and an indicator for credit constraints.[10] Columns 3 and 4 report

[10] We also introduce dummies for retirement status and self-employment. The coefficients of these dummies are not statistically different from zero and are dropped from the baseline specification.

OLS estimates on the pooled cross-sectional sample (15,366 observations) and on the longitudinal sample (3,236 observations), respectively. The last column of Table 2 reports fixed effect estimates. Given that we have only two years of data, fixed effect estimates coincide with OLS first-difference estimates.

Column 3 of Table 2 indicates that the estimate of β is -0.20 and quite precisely measured. Comparison of columns 1 and 3 suggests that part of the association between MPC and cash on hand can be attributed to the omission of observables. The estimated coefficients indicate that the MPC is lower for married couples, higher for households with higher education, 11 percentage points higher for households living in the South, and 9 percentage points higher in large cities, and that it increases with family size. It is also significantly higher (6.4 percentage points) if the head of household is unemployed. As for age, we find that the MPC is negatively associated with it. The standard life cycle model predicts that the young should report a lower MPC since they have a longer horizon; however, there might be cohort effects working in the opposite direction, for instance, because younger generations might have lower discount factors. In general, the effect of age on the MPC is hard to interpret since it is not feasible to separate age and cohort effects in cross-sectional data. Finally, the coefficient of the credit constraint dummy is not statistically different from zero.

Column 4 replicates the specification of column 3 on the panel sample. The estimate of β is -0.18, which is again precisely estimated and not statistically different from the estimate in column 3. The pattern of the other coefficients is similar to the full sample estimates, but as expected, standard errors tend to be larger given the reduced number of observations.

In column 5, we report fixed effect estimates. The main coefficient of interest is $\hat{\beta}_{FE} = -0.15$. The first remarkable result is that the relation between MPC and cash on hand is negative and significant even controlling for unobserved heterogeneity. The second important result is that unobserved heterogeneity reduces the sensitivity of MPC with respect to cash on hand, although the bias appears to be moderate (a 25 percent change, or $(1 - (0.15/0.20))$).[11] The third result is that the gap between cross-sectional and panel estimates of β is consistent with $\mathrm{cov}(X_{it}, f_i) < 0$, namely that people with high taste for current consumption (as reflected in higher MPC) also tend to have relatively lower cash on hand.

Note that measurement error in the cash on hand variable is an alternative interpretation of the difference between cross-sectional and panel estimates. In fact, if measurement error is classical and cash on hand is positively correlated over time, panel data exacerbate the standard attenuation bias (Griliches and Hausman 1986).[12] One can show that the difference between cross-sectional and panel data estimates reflects the combined effect of preference heterogeneity and measurement error bias. However, since this difference is moderate, these two biases in isolation

[11] Part of this reduction is likely due to the different characteristics of the full sample and the panel sample. Comparing the two panel estimates (OLS of column 4 and fixed effects of column 5) shows that unobserved heterogeneity reduces the sensitivity of MPC with respect to cash on hand by 17 percent $(1 - (0.15/0.18))$.

[12] The degree of exacerbation in our context is likely reduced by the fact that the data cover a six-year difference.

Table 3—MPC Regressions Using log Cash on Hand and Cash on Hand Quintiles

Sample Estimation method	All OLS (1)	Panel OLS (2)	Panel Fixed effects (3)	All OLS (4)	Panel OLS (5)	Panel Fixed effects (6)
Aged ≤30	0.052	0.064	0.040	0.059	0.070	0.032
	(0.017)	(0.054)	(0.103)	(0.017)	(0.054)	(0.103)
Aged 30–45	0.035	0.016	0.004	0.040	0.022	−0.003
	(0.009)	(0.022)	(0.053)	(0.009)	(0.022)	(0.053)
Aged 45–60	0.029	0.031	0.024	0.032	0.035	0.023
	(0.007)	(0.015)	(0.033)	(0.007)	(0.015)	(0.033)
Male	−0.004	−0.016		−0.006	−0.017	
	(0.006)	(0.013)		(0.006)	(0.013)	
Married	−0.017	−0.020		−0.017	−0.022	
	(0.007)	(0.017)		(0.007)	(0.017)	
Years of education	0.002	0.003	0.017	0.001	0.002	0.016
	(0.001)	(0.002)	(0.012)	(0.001)	(0.002)	(0.012)
Resident in the South	0.114	0.115		0.114	0.120	
	(0.006)	(0.014)		(0.006)	(0.014)	
Family size	0.015	0.023	0.000	0.014	0.022	−0.002
	(0.003)	(0.007)	(0.017)	(0.003)	(0.007)	(0.017)
City size >500,000	0.088	0.070	−0.059	0.087	0.066	−0.053
	(0.010)	(0.026)	(0.265)	(0.010)	(0.026)	(0.265)
Dummy for 2016	−0.010	−0.028	−0.038	−0.004	−0.022	−0.036
	(0.006)	(0.012)	(0.014)	(0.006)	(0.012)	(0.014)
Unemployed	0.055	0.028	−0.074	0.065	0.038	−0.059
	(0.014)	(0.032)	(0.054)	(0.014)	(0.031)	(0.054)
Credit constrained	−0.014	0.063	−0.008	−0.013	0.069	−0.001
	(0.015)	(0.038)	(0.053)	(0.015)	(0.038)	(0.053)
Log cash on hand	−0.037	−0.036	−0.028			
	(0.002)	(0.005)	(0.009)			
II cash on hand quintile				−0.056	−0.047	−0.019
				(0.009)	(0.020)	(0.031)
III cash on hand quintile				−0.100	−0.083	−0.038
				(0.009)	(0.020)	(0.033)
IV cash on hand quintile				−0.129	−0.105	−0.068
				(0.009)	(0.021)	(0.035)
V cash on hand quintile				−0.157	−0.143	−0.110
				(0.010)	(0.022)	(0.040)
Constant	0.458	0.447	0.386	0.462	0.443	0.381
	(0.009)	(0.021)	(0.128)	(0.010)	(0.023)	(0.129)
R^2	0.08	0.09	0.58	0.08	0.09	0.58
Observations	15,303	3,230	3,230	15,366	3,236	3,236

Note: We report standard errors in parentheses.

are unlikely to be large (assuming, as intuition would suggest, that preference heterogeneity also imparts an attenuation bias).[13]

In Table 3, we use two different measures of cash on hand to check the robustness of our baseline estimates. In columns 1–3, we replace the percentile of cash on hand

[13] The most relevant forms of unobserved heterogeneity (such as high discount rates, high rates of risk tolerance, or high preferences for bequests) would all act in the sense of generating larger MPC and (assuming they are permanent traits) lower cash on hand.

with the log of cash on hand itself. The sensitivity of MPC with respect to log cash on hand is −0.04 in the pooled OLS estimates, essentially unchanged in the panel sample, and −0.03 (a 22 percent decline in absolute value) with the fixed effect estimator.

In columns 4–6, we break down cash on hand into quintiles to check for possible nonlinear effects of cash on hand on MPC. Cash on hand quintiles are also more resilient to measurement error than percentiles of cash on hand or log of cash on hand. The pattern of the coefficients suggests a monotonically declining relation, ranging from 0 (the excluded first quintile) to −0.16 (the top 20 percent group). There is mild evidence of nonlinearity, as the effect of cash on hand on MPC is stronger at low than at high levels of cash on hand. Moreover, the estimates are all significantly different from 0. The OLS estimates in the panel sample essentially mirror those in the whole sample. The fixed effect estimates confirm a monotonic relationship but weaker from both a statistical and quantitative point of view, with the estimates ranging from 0 (for the excluded first quintile) to −0.11 (the top quintile).

Finally, in Table 4 we break down cash on hand percentiles separately into income and financial wealth percentiles. Column 1 shows that there is a negative gradient between both income and financial assets and cash on hand. Comparing columns 1 and 3, it appears that both variables contribute to the exaggeration effect of cross-sectional OLS estimates; if one compares columns 2 and 3, the exaggeration effect is mostly due to income.

IV. Robustness Checks

One potential criticism of the reported MPC measure is that there is substantial heaping of responses as about two-thirds of respondents choose focal answers of 0, 50, or 100 percent. The high rate of 50 percent as a response (25 percent of the sample) is particularly concerning as it is often interpreted as a symptom of the respondent's epistemological uncertainty (Bruine de Bruin et al. 2002). Giustinelli, Manski, and Molinari (2018) show that survey respondents provide more refined responses in the tails of a 0–100 scale than in the center and that rounding/heaping practices are associated with observable respondent characteristics, such as personal finances, health, and macroeconomic events. Gideon, Helppie-McFall, and Hsu (2017) also find that survey responses to quantitative financial questions display strong patterns of heaping at round numbers and that rounding is more common for respondents with low ability and for more difficult questions. To address these important issues, we perform several experiments.

Table A1 in the online Appendix provides a number of robustness checks for the role of heaping in the MPC variable. First, we perform conditional logit regressions for the probability of reporting an MPC $= 0.5$. The results indicate that no variable of the baseline specification is systematically related to the probability of reporting MPC $= 0.5$, including cash on hand. The exception is the credit-constrained dummy, which is negatively associated with MPC $= 0.5$. This may reflect behavior (credit-constrained households are less likely to report an MPC $= 0.5$) or lower uncertainty (credit-constrained households are less uncertain about the meaning of the question).

TABLE 4—MPC REGRESSIONS DISTINGUISHING BETWEEN FINANCIAL ASSETS AND INCOME

Sample	All	Panel	Panel
Estimation method	OLS	OLS	Fixed effects
	(1)	(2)	(3)
Aged $\leq$30	0.048	0.046	0.038
	(0.017)	(0.054)	(0.103)
Aged 30–45	0.031	0.007	−0.000
	(0.009)	(0.022)	(0.053)
Aged 45–60	0.029	0.031	0.026
	(0.007)	(0.015)	(0.033)
Male	−0.003	−0.013	
	(0.006)	(0.013)	
Married	−0.011	−0.008	
	(0.007)	(0.017)	
Years of education	0.003	0.005	0.017
	(0.001)	(0.002)	(0.012)
Resident in the South	0.106	0.107	
	(0.006)	(0.014)	
Family size	0.019	0.029	0.003
	(0.003)	(0.007)	(0.018)
City size $>$500,000	0.088	0.069	−0.053
	(0.010)	(0.026)	(0.265)
Dummy for 2016	−0.002	−0.020	−0.032
	(0.006)	(0.012)	(0.014)
Unemployed	0.054	0.021	−0.068
	(0.014)	(0.031)	(0.054)
Credit constrained	−0.016	0.059	−0.001
	(0.015)	(0.038)	(0.053)
Percentile of financial assets	−0.138	−0.106	−0.115
	(0.011)	(0.025)	(0.041)
Percentile of disposable income	−0.103	−0.148	−0.073
	(0.015)	(0.033)	(0.068)
Constant	0.465	0.451	0.409
	(0.009)	(0.022)	(0.129)
R^2	0.09	0.09	0.59
Observations	15,366	3,236	3,236

Notes: Each regression includes a time dummy. We report standard errors in parentheses.

For these reasons, we focus next on heaping at 0 or 1. Note that both reported values should be more informative about actual behavior than other cutoffs. Indeed, MPC $=$ 1 is typically associated with rule-of-thumb behavior or binding liquidity constraints. On the other hand, MPC should be close to 0 for unconstrained PIH consumers with long horizons. To buttress the information contained in these values, we run conditional logit regressions for these two polar cases. We find that cash on hand is a strong (negative) predictor of MPC $=$ 1 and a strong (positive) predictor of MPC $=$ 0, as it should be if the consumption function is concave.

The heaping at 0, 50, and 100 suggests that our scale is closely related to the one used by Shapiro and Slemrod (2003, 2009). In their context, a "mostly spend" response can be interpreted as implying an MPC of 50 percent or higher, while a "mostly save or pay debt" response would imply an MPC of less than 50 percent. Accordingly, we estimate a conditional logit model for the "mostly spend" category, defined as

$MPC \geq 0.5$. The effect is identified from people switching from "mostly spend" to "mostly save" (and vice versa). The results are qualitatively similar to the ones where we use the continuous (albeit "heaped") variable. In fact, the marginal effect of the cash on hand percentile on the probability of "mostly spend" is -0.18. These results are reported in the last column of Table A1 of the online Appendix.

As a further check of the sensitivity of the result, in column 1 of Table A2 of the online Appendix, we drop the observations with an MPC $= 0.5$ value in one or both waves and find that our main results are confirmed. In columns 2 and 3 of online Appendix Table A2, we use information on the accuracy of survey responses. In particular, we use a set of indicators provided by the professional survey interviewer at the end of each one-to-one personal interview. First, we focus on a subsample of individuals with more reliable information on financial assets (measured by a dummy for whether the interviewer rates the quality of the financial asset information provided by the respondent with a score of 7 or above on a 1 to 10 scale). The results are similar to those of the baseline sample. Finally, we interact cash on hand with the indicator measuring high quality of the financial asset variable. The coefficient on the interaction term is not significant.[14]

The results are also robust to various sample splits and definition of the variables. The most relevant robustness checks are reported in Table A3 of the online Appendix. First, we replicate the regressions in Table 4, focusing on households that experienced (on average) an annual income change between -5 percent and $+5$ percent. This reduces the sample to about two-thirds of the original one. The coefficient on liquid financial assets is -0.10 (slightly lower than -0.12 in the original sample) and significant at the 5 percent level. Not surprisingly, the income coefficient is much noisier since there is now much less variation left to identify it.[15]

Next, we replace financial assets with financial asset net of nonmortgage debt. The results are again essentially unchanged (online Appendix Table A3, column 2). We also explore the role of permanent income as a possible driver of MPC heterogeneity. One possibility would be to compute an average of income in the years prior to the survey. While we could use the rotating panel structure of the SHIW to construct average past income for each household, it would be a fixed household characteristic and hence not identifiable in the panel regression. An alternative is to use consumption as a proxy for permanent income. We thus replace the income percentiles with the percentile of nondurable consumption. We confirm that liquid assets are the main driver of MPC, while permanent income (to the extent that consumption is a good proxy for it) plays no role (online Appendix Table A3, column 3).

An important implication that emerges from our analysis is that MPC heterogeneity arises from differences in cash on hand, controlling for differences in personal traits. As a further check of our findings, we perform the following test. If behavioral

[14] The results remain similar (and available on request) if we focus on a sample with greater ability to understand the survey questions (as assessed by the professional survey interviewers).

[15] To check that results are not affected by outliers in the income distribution, we also trim the top and bottom 1 percent of the sample, finding no remarkable changes. We check that our main findings are robust to sample selection, focusing on a sample of household heads younger than 60. We also control for real estate wealth and debt, which may create overhang effects (Dynan 2012). All these checks leave the pattern of results qualitatively unchanged: controlling for fixed effects attenuates the sensitivity of MPC with respect to cash on hand. These additional results are available upon request.

traits (for instance, impatience) predict MPC, one should expect that people with high MPC (that is, people with high impatience) accumulate fewer assets in the future. We thus regress growth of cash on hand in 2010–2016 on 2010 MPC, controlling for demographic variables and initial cash on hand. The coefficient of lagged MPC is -0.02 and not statistically different from 0. This suggests that initial MPC is not a good predictor of personal traits associated with asset accumulation. These additional results are reported in column 4 of online Appendix Table A3.

V. Fixed Effects at Work in a Simulated Fiscal Experiment

The value of a fiscal stimulus depends crucially on the characteristics of the policy change. Elmendorf and Furman (2008) summarize the evidence on the effectiveness of fiscal policy and provide principles and examples for formulating effective stimulus. They point out that policymakers should implement policies that ensure that each dollar of tax cuts or higher spending maximizes short-run output and that money ends up in the pockets of families that are most vulnerable in a weakening economy. They conclude that "these two goals are complementary, because the families that most need the money are also the most likely to stimulate the economy by spending it quickly" (Elmendorf and Furman 2008, 5). The fiscal policy experiments that we present in this section follow their insight and provide quantitative evidence on the aggregate consumption effect of fiscal redistribution in the presence of MPC heterogeneity.

Consider a policymaker trying to forecast the effect on aggregate consumption of a fiscal policy that transfers an amount Δ to each household in the population. To see a concrete example, consider an approximation to a concave consumption function:

$$C_{it} \approx \phi + \alpha X_{it} + \frac{\beta}{2} X_{it}^2 + \epsilon_{it}, \tag{3}$$

with $\alpha > 0$ and $\beta < 0$. The individual MPC is $MPC_{it} = \alpha + \beta X_{it}$, which shows that its heterogeneity comes only from differences in cash on hand. To consider a case with preference heterogeneity, let's write a different approximation of the consumption function:

$$C_{it} \approx \phi + \left(\alpha + f_i\right) X_{it} + \frac{\beta}{2} X_{it}^2 + \epsilon_{it}, \tag{4}$$

with $E(f_i) = 0$. The individual MPC associated with the consumption function (4) is now

$$MPC_{it} = \left(\alpha + f_i\right) + \beta X_{it},$$

i.e., a version of regression equation (1) omitting the stochastic term v_{it}. The latter may be added to reflect measurement error in reported MPC. Note that if f_i captures tastes for current consumption (and hence higher values of f_i are associated to lower values of cash on hand, ceteris paribus), one may find a negative correlation between MPC and cash on hand even if the true consumption function is linear

($\beta = 0$ but preferences are heterogeneous), the example we discussed in the introduction. It is immediate to show that the effect on aggregate consumption of the fiscal policy considered above depends on the value of the average MPC, which is directly proportional to $\beta \bar{X}_t$, highlighting the importance of obtaining an estimate of the causal effect of a change in cash on hand on the MPC (β). Estimating this causal effect hinges crucially on the ability to control for unobserved heterogeneity (f_i), which may be potentially correlated with resources. One route is to run an experiment in which the policymaker increases cash on hand by an amount Δ while keeping everything else constant. Alternatively, one can use panel data and difference out the f_i term across two periods. This route relies on the assumption that tastes do not change over time and that the econometrician can control for the observable characteristics that may have shifted over time—our empirical approach.

To show the importance of unobserved heterogeneity when evaluating the macroeconomic impact of a policy, we simulate a revenue-neutral fiscal reform under different scenarios. In particular, we consider a policy that transfers the equivalent of 1 percent of national income (in equal amounts) to the bottom x percent of the cash on hand distribution. The policy is financed by taxing the top 10 percent of the cash on hand distribution (in the form of a lump-sum tax).[16] The design of the experiment provides a useful benchmark case: in models with a homogenous MPC, such as the PIH with certainty equivalence and no liquidity constraints ($\beta = 0$), this redistributive policy has no aggregate effects (absent labor supply and general equilibrium effects).

Suppose that the goal is to measure the impact of the policy using an estimate of the actual relationship between MPC and cash on hand. A policymaker may simply multiply the average reported MPC at each cash on hand percentile (i.e., the estimated relationship between MPC and cash on hand percentile from column 1 of Table 2) by the transfer received/tax paid and aggregate the corresponding consumption change. This is the calculation reported in the first block of Table 5 (column 1), where transfers are targeted to households below a given percentile of the cash on hand distribution (tenth, twenty-fifth, fiftieth, seventy-fifth, and ninetieth). The largest effects are found when the policy targets the poorest 10 percent. In this case, the revenue-neutral policy boosts aggregate consumption by 0.33 percent (a result coming from the bottom decile reporting much higher MPC than richer households). The smallest impact of the policy (0.14 percent) is when one targets all households (except of course those that finance it).

However, cash on hand correlates with many variables so that one should consider that differences in MPC by cash on hand partly reflect such correlation. To isolate the effect of cash on hand on MPC, controlling for *observable* characteristics,

[16] Unlike the example discussed above, the transfer is positive for some households ($\Delta_i > 0$) and negative for others ($\Delta_i < 0$). However, it is still true that the effect of the policy on aggregate consumption crucially depends on the value of β. Note that in each of the scenarios we consider, we make the implicit assumption that the MPC of the rich is symmetric with respect to positive and negative income changes (even though we only have reported MPC with respect to positive income changes). We believe that this is reasonable given that for the top 10 or 5 percent of the cash on hand distribution one should not expect liquidity constraints (the most likely source of asymmetric responses) to be important.

TABLE 5—THE EFFECT OF A REDISTRIBUTIVE FISCAL POLICY ON AGGREGATE CONSUMPTION

Policy	Unconditional MPC (1)	Conditional MPC, OLS (2)	Conditional MPC, fixed effects (3)
Transfers targeted to cash on hand distribution below			
10th percentile	0.33	0.23	0.17
25th percentile	0.28	0.21	0.16
50th percentile	0.21	0.18	0.14
75th percentile	0.17	0.15	0.11
90th percentile	0.14	0.13	0.10
Transfer randomly to cash on hand distribution below			
90th percentile	0.14	0.13	0.10
Transfer targeted to permanent income distribution below			
10th percentile	0.24	0.18	0.13
25th percentile	0.20	0.16	0.12
50th percentile	0.16	0.13	0.10
75th percentile	0.13	0.11	0.08
90th percentile	0.11	0.10	0.07
Transfer targeted to current income distribution below			
10th percentile	0.26	0.17	0.13
25th percentile	0.21	0.15	0.11
50th percentile	0.16	0.12	0.09
75th percentile	0.13	0.10	0.08
90th percentile	0.11	0.09	0.07

Notes: The table reports the growth in aggregate consumption corresponding to various redistributive policies. Total transfers are the same in each of the policy simulations. They are revenue neutral with financing coming from taxing people in the top decile of relevant distributions (by an equal amount such that total revenues raised equal 1 percent of the national income). Column 1 uses the OLS estimate of the relationship between MPC and cash on hand estimated from column 1 of Table 2. Column 2 uses the OLS estimate from column 3 of Table 2, and column 3 uses the panel data estimate from column 5 of Table 2.

one should perform the experiment using the predicted MPC obtained from the OLS regression reported in column 3 of Table 2. This is what we do in column 2 of Table 5, showing that there is substantial attenuation of the aggregate effect of the redistributive policy. For instance, targeting the transfer to households below the tenth percentile of the cash on hand distribution would boost aggregate consumption by only 0.23 percent (down from 0.33 percent for the unconditional estimates). There is a similar pattern if the transfer is more diffuse. For instance, targeting households below the twenty-fifth percentile would increase aggregate consumption by 0.21 percent (down from 0.28).

Still, the conditional correlation between MPC and cash on hand may be affected by *unobserved* heterogeneity (such as preferences) as argued above. For the final experiment, one should rely on the estimates obtained from the fixed effect regression reported in column 5 of Table 2. The results, reported in column 3 of Table 5, show that the aggregate consumption effect of the redistributive policy is further attenuated with respect to the case in which unobserved heterogeneity is ignored. For instance, the boost in aggregate consumption is 0.17 percent for the most concentrated transfer policy that targets households in the bottom tenth percentile of the cash on hand distribution.

Comparison of columns 1 and 3 of Table 5 across the size of groups targeted by the policy yields two insights. The first insight is that the bias induced by neglecting heterogeneity (both observed and unobserved) is higher when the targets are the bottom decile or quartile than when the policy is more diffuse. The reason is that people at the bottom of the cash on hand distribution are also more likely to report high MPC (as revealed by OLS regression estimates) given their characteristics: they are more likely to be unemployed, to live in large cities or in the South, or to be young. At the same time, people at the bottom of the cash on hand distribution are also more likely to have preferences for current consumption as revealed by the difference between cross-sectional and panel estimates. The second important insight from Table 5 is that the bias induced by neglecting unobservable characteristics is moderate as implied by the relatively small difference in the β estimated with OLS or fixed effects.

In the rest of Table 5, we perform three additional experiments. First, we transfer the revenues obtained from taxing the top decile of the cash on hand distribution to a random 10 percent of the households below the ninetieth percentile. Comparing this fiscal experiment with the first row of Table 5 (where we transfer income only to the bottom 10 percent of the cash on hand distribution) is interesting because it highlights the importance of targeting fiscal policy to specific population groups.

One objection with the experiments performed so far is that they target cash on hand, which unlike income may not be observed (and hence targeted) accurately by the government. The remaining experiments focus on redistribution on the basis of permanent income or current income instead of cash on hand. In the first experiment, we define as "permanent income" the average of 2010 and 2016 disposable income. In the second experiment, we simply use current income. The results are qualitatively similar to those of the first experiment. For instance, transfers targeted to households below the tenth percentile of permanent income increase aggregate consumption by 0.24 percent (using the unconditional MPC), 0.18 percent (using the conditional MPC from the OLS estimates), and 0.13 percent (using the conditional MPC with the fixed effects estimates). The numbers are similar using current income instead of permanent income. Note that the aggregate consumption effects are lower than the corresponding ones in the first row of Table 5 because the MPC is more negatively related to cash on hand (which also includes liquidity) than to current or permanent income.

It is worth stressing that none of these calculations include general equilibrium effects (deriving, e.g., from changes in interest rates). Hence, they are likely providing an upper bound to the true effects of redistributive fiscal policies. Finally, the results may not generalize to other samples and countries.

VI. Summary

We analyze reported MPC from hypothetical income change questions posed to participants of the 2010 and 2016 Italian Survey of Household Income and Wealth. We confirm some of the findings from the existing literature, such as considerable heterogeneity in MPC. Different from previous studies, controlling for observable characteristics, we uncover a strong negative association between MPC and cash on hand, consistent with models of consumption with precautionary savings and liquidity constraints.

One limitation of the studies that use survey-based reported MPC is that they rely on cross-sectional data. However, some of the association between MPC and cash on hand could be spurious and attributable to unobserved heterogeneity. A unique feature of the SHIW is that the same hypothetical MPC question is available in two waves (2010 and 2016) and that the survey itself has a sizable longitudinal component. This allows us to use standard panel data estimation methods to purge the effect of cash on hand on MPC by fixed unobserved heterogeneity.

Comparison of cross-sectional and panel data estimation reveals that unobserved heterogeneity exaggerates the sensitivity of MPC to cash on hand, but the amount of bias is moderate (roughly 20 percent). In Section V, we simulate the impact of several fiscal experiments to study the implications of such bias for the effectiveness of revenue-neutral redistributive fiscal policies. We find that the effectiveness of such revenue-neutral fiscal policies does not change much relative to a case in which both observed and unobserved heterogeneity are ignored, particularly for policies that target the bottom part of the distribution of household resources.

REFERENCES

Bank of Italy. 2020. "Survey of Household Income and Wealth." Banca D'Italia—Eurosistema. https://www.bancaditalia.it/statistiche/tematiche/indagini-famiglie-imprese/bilanci-famiglie/index.html?com.dotmarketing.htmlpage.language=1 (accessed January 16, 2020).

Bruine de Bruin, Wandi, Paul S. Fischbeck, Neil A. Stiber, and Baruch Fischhoff. 2002. "What Number Is 'Fifty-Fifty?': Redistributing Excessive 50% Responses in Elicited Probabilities." *Risk Analysis: An International Journal* 22 (4): 713–23.

Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico. 2018. "The Consumption Response to Positive and Negative Income Shocks." *Journal of Monetary Economics* 96: 1–15.

Campbell, Jeffrey R., and Zvi Hercowitz. 2019. "Liquidity Constraints of the Middle Class." *American Economic Journal: Economic Policy* 11 (3): 130–55.

Carroll, Christopher D. 2001. "A Theory of the Consumption Function, with and without Liquidity Constraints." *Journal of Economic Perspectives* 15 (3): 23–45.

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

Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten van Rooij. 2019. "Asymmetric Consumption Effects of Transitory Income Shocks." *Economic Journal* 129 (622): 2322–41.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59 (5): 1221–48.

Dynan, Karen. 2012. "Is a Household Debt Overhang Holding Back Consumption?" *Brookings Papers on Economic Activity* 42 (1): 299–362.

Elmendorf, Douglas W., and Jason Furman. 2008. "If, When, How: A Primer on Fiscal Stimulus." https://www.urban.org/sites/default/files/publication/31076/1001127-If-When-How-A-Primer-on-Fiscal-Stimulus.PDF.

Fuster, Andreas, Greg Kaplan, and Basit Zafar. 2018. "What Would You Do With $500? Spending Responses to Gains, Losses, News and Loans." NBER Working Paper 24386.

Gelman, Michael. 2019. "What Drives Heterogeneity in the Marginal Propensity to Consume? Temporary Shocks vs Persistent Characteristics." https://sites.google.com/site/michaelgelman/research/mpc_heterogeneity_gelman.pdf?attredirects=0.

Gideon, Michael, Brooke Helppie-McFall, and Joanne W. Hsu. 2017. "Heaping at Round Numbers on Financial Questions: The Role of Satisficing." *Survey Research Methods* 11 (2): 189–214.

Giustinelli, Pamela, Charles F. Manski, and Francesca Molinari. 2018. "Tail and Center Rounding of Probabilistic Expectations in the Health and Retirement Study." NBER Working Paper 24559.

Graziani, Grant, Wilbert van der Klaauw, and Basit Zafar. 2016. "Workers' Spending Response to the 2011 Payroll Tax Cuts." *American Economic Journal: Economic Policy* 8 (4): 124–59.

Griliches, Zvi, and Jerry A. Hausman. 1986. "Errors in Variables in Panel Data." *Journal of Econometrics* 31 (1): 93–118.

Jappelli, Tullio, and Luigi Pistaferri. 2010. "The Consumption Response to Income Changes." *Annual Review of Economics* 2: 479–506.

Jappelli, Tullio, and Luigi Pistaferri. 2014. "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics* 6 (4): 107–36.

Jappelli, Tullio, and Luigi Pistaferri. 2017. *The Economics of Consumption—Theory and Evidence*. Oxford: Oxford University Press.

Jappelli, Tullio, and Luigi Pistaferri. 2020. "Replication Data for: Reported MPC and Unobserved Heterogeneity." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E116762V1.

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

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

Parker, Jonathan A. 2017. "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment." *American Economic Journal: Macroeconomics* 9 (4): 153–83.

Parker, Jonathan A., and Nicholas S. Souleles. 2019. "Reported Effects versus Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates." *American Economic Review: Insights* 1 (3): 273–90.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod. 2010. "Household Response to the 2008 Tax Rebates: Survey Evidence and Aggregate Implications." In *Tax Policy and the Economy*, Vol. 24, edited by Jeffrey R. Brown, 69–110. Chicago: University of Chicago Press.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod. 2012. "Check in the Mail or More in the Paycheck: Does the Effectiveness of Fiscal Stimulus Depend on How It Is Delivered?" *American Economic Journal: Economic Policy* 4 (3): 216–50.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod. 2015. "Balance-Sheet Households and Fiscal Stimulus: Lessons from the Payroll Tax Cut and Its Expiration." NBER Working Paper 21220.

Shapiro, Matthew D., and Joel Slemrod. 1995. "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding." *American Economic Review* 85 (1): 274–83.

Shapiro, Matthew D., and Joel Slemrod. 2003. "Consumer Response to Tax Rebates." *American Economic Review* 93 (1): 381–96.

Shapiro, Matthew D., and Joel Slemrod. 2009. "Did the 2008 Tax Rebates Stimulate Spending?" *American Economic Review* 99 (2): 374–79.