View metadata, citation and similar papers at core.ac.uk
brought to you by CORE
provided by Hochschulschriftenserver - Universität Frankfurt am Main

CFS WORKING PAPER

No. 2013/14

Fiscal Policy and MPC Heterogeneity

Tullio Jappelli, Luigi Pistaferri

Center for Financial Studies
Goethe University ■ House of Finance
Grüneburgplatz 1 ■ 60323 Frankfurt am Main

Tel: +49 69 798-30050
Fax: +49 69 798-30077
www.ifk-cfs.de ■ info@ifk-cfs.de

CFS Working Paper Series

The *Center for Financial Studies*, located in Goethe University's House of Finance in Frankfurt, is an independent non-profit research center, funded by the non-profit-making organisation Gesellschaft für Kapitalmarktforschung e.V. (GfK). The CFS is financed by donations and by contributions of the GfK members, as well as by national and international research grants. The GfK members comprise major players in Germany's financial industry. Established in 1967 and closely affiliated with the University of Frankfurt, it provides a strong link between the financial community and academia. CFS is also a contributor to policy debates and policy analyses, building upon relevant findings in its research areas.

The CFS Working Paper Series presents the result of scientific research on selected topics in the field of money, banking and finance. The authors were either participants in the Center´s Research Fellow Program or members of one of the Center´s Research Projects.

If you would like to know more about the Center for Financial Studies, please let us know of your interest.

Prof. Michalis Haliassos, Ph.D. Prof. Dr. Jan Pieter Krahnen Prof. Dr. Uwe Walz

Center for Financial Studies
Goethe University ■ House of Finance
Grüneburgplatz 1 ■ 60323 Frankfurt am Main

Tel: +49 69 798-30050
Fax: +49 69 798-30077
www.ifk-cfs.de ■ info@ifk-cfs.de

Fiscal Policy and MPC Heterogeneity

Tullio Jappelli
University of Naples Federico II, CSEF and CEPR
e-mail: tullio.jappelli@unina.it

Luigi Pistaferri
Stanford University, NBER, CEPR and SIEPR
e-mail: pista@stanford.edu

Abstract

We use responses to survey questions in the 2010 Italian Survey of Household Income and Wealth that ask consumers how much of an unexpected transitory income change they would consume. We find that the marginal propensity to consume (MPC) is 48 percent on average, and that there is substantial heterogeneity in the distribution. We find that households with low cash-on-hand exhibit a much higher MPC than affluent households, which is in agreement with models with precautionary savings where income risk plays an important role. The results have important implications for the evaluation of fiscal policy, and for predicting household responses to tax reforms and redistributive policies. In particular, we find that a debt-financed increase in transfers of 1 percent of national disposable income targeted to the bottom decile of the cash-on-hand distribution would increase aggregate consumption by 0.82 percent. Furthermore, we find that redistributing 1% of national disposable income from the top to the bottom decile of the income distribution would boost aggregate consumption by 0.33%.

Keywords: Marginal Propensity to Consume, Fiscal Policy, Consumption Heterogeneity.
JEL Code: E21, D91.

Acknowledgements. We thanks seminar participants at the Central Bank of Ireland, the Meeting of the ECB-HFCN group in Vienna, the MPC Conference at the Center for American Progress (Washington D.C.), the Macroeconomic Dynamics Workshop at LBS, Michalis Haliassos, the editor, two anonymous referees, and especially Fabrizio Perri for comments. Yichen Su provided valuable research assistance. We acknowledge financial support from the Italian Ministry of Universities and Research (Jappelli) and the European Research Council starting grant 284024 (Pistaferri).

1. Introduction

Knowledge on how consumers respond to income shocks –the Marginal Propensity to Consume (MPC) – is crucial for evaluating the macroeconomic impact of tax and labor market reforms, and for the design of stabilization and income maintenance policies. Distinguishing whether consumption responds differently to transitory or permanent income changes is equally important. In this paper, we use information on how much consumers would spend of an unexpected windfall gain equal to their average monthly income. Hence, we can characterize empirically the distribution of MPC in response to an unexpected transitory variation in income.

In a standard life-cycle permanent income model, consumption is proportional to lifetime disposable resources (or permanent income), and hence all consumers respond in a similar way to income shocks, that is, there is no heterogeneity in the MPC. Models with precautionary savings, liquidity constraints or bequest motives, in contrast, suggest that consumers respond differently to changes in their economic resources. Indeed, the MPC of prudent individuals falls with household resources; liquidity constrained consumers exhibit higher MPC than households who can access credit markets to smooth consumption; and if bequests are luxury goods, rich individuals consume a smaller fraction of their lifetime resources.[1] In addition to these level effects, the composition of household resources may also matter. For instance, households burdened with large debts might react to a positive change in income by reducing their debt rather than spending (Dynan, 2012; Mian and Sufi, 2010). Moreover, if most of the wealth is locked into illiquid assets, households must cut consumption even in the face of a negative transitory income shock (Kaplan and Violante, 2011).

The importance of the consumer response to fiscal policies is attracting renewed attention given the large fiscal stimulus packages enacted by governments on both sides of the Atlantic to counteract the Great Recession. One of the major problems for policymakers is to assess the effectiveness of such policies as debt-financed fiscal packages or redistributive programs that maintain the public deficit unchanged. For instance, fiscal reforms that increase the tax burden of the rich are often advocated as a way to finance unemployment insurance schemes or other welfare programs. As mentioned above, in standard life-cycle models, these revenue-neutral redistributive policies produce no consumption effects in the aggregate because of homogeneity in the MPC. In these models, the consumption of the rich is no more than a scaled-up version of the consumption of the poor and any increase in the latter's consumption is matched exactly by a decline in the consumption of the former. However, more realistic models with precautionary savings or liquidity

[1] See, e.g., Dynan, Skinner and Zeldes (2004).

constraints feature MPC heterogeneity. Since aggregate consumption depends on the distribution of MPC within the population, redistributive fiscal policy potentially might boost national income.

A major problem in estimating the MPC is isolating the exogenous shocks to income which can be used to track consumption behavior after a shock. The literature has considered three approaches. A first method identifies episodes in which income changes unexpectedly, and evaluates in a quasi-experimental setting how consumption reacts to such changes (due to, say, unemployment or disability, see Browning and Crossley, 2001, and Stephens, 2001). A second approach is to rely on a statistical decomposition of income shocks and the covariance restrictions imposed by the theory on the joint behavior of income and consumption, and use long panel data to relate income shocks to consumption changes (Blundell, Pistaferri and Preston, 2008). Survey questions containing responses to hypothetical income changes represent a third alternative and also the strategy adopted in this study, which relies on a unique question in the 2010 Italian Survey of Household Income and Wealth (SHIW).[2] The survey is designed to elicit information on how much people would consume or save were they unexpectedly to receive a reimbursement (or transfer) equal to their average monthly income. The responses to this question provide a sample distribution of the MPC that we can relate to observable characteristics and, most importantly, compare with the predictions of intertemporal consumption models and use to perform policy analyses.

In our sample, we find that the average MPC is 48 percent, at the high-end of current estimates based on survey data on observed consumption and income changes.[3] Most importantly, we find quite substantial heterogeneity in people's responses, ranging from zero to 1 across the entire distribution of the household's resources. In particular, households with low cash-on-hand exhibit a much higher MPC than affluent households, which is in agreement with models where income risk or liquidity constraints play an important role. The empirical distribution of the MPC can be used to perform some simple policy simulations designed to predict the consumption response to tax interventions. We consider the cases of both government enacting a transfer policy financed by issuing debt and government redistributing income from rich to poor in a revenue-neutral scheme. We find that the response of aggregate consumption to these policies is substantially higher than in a benchmark case when the MPC is assumed to be the same for each all household.

Our approach is based on the assumption that there is no systematic discrepancy between the intention to spend and actual spending behavior. Given this assumption, a first advantage of our approach is that it solves the econometric problems of identifying genuine income shocks and isolating their effect on consumption, from other confounding forces. A second advantage is that,

[2] See Shapiro and Slemrod (1995; 2003) and Sahm, Shapiro and Slemrod (2009).
[3] Sahm, Shapiro and Slemrod (2009) find an MPC of about one-third.

relative to experimental approaches and statistical decompositions of income shocks, where absence of long panels forces researchers to make strong distributional assumptions about the relationship between consumption growth and shocks to households' resources, no such assumptions are needed in our context. A third advantage is that the survey question delivers the MPC for each household, while in most studies in the literature, the sample sizes are too small to obtain precise estimates of MPC for specific population groups defined by age, education, or household resources - hence most focus on average MPC. This means that researchers are unable to evaluate the effect of redistributive policies on aggregate consumption.

The paper is organized as follows. Section 2 discusses the theoretical predictions about the MPC and reviews various methods proposed for its estimation, focusing particularly on studies that rely on direct survey questions. Section 3 describes the data and the question used to elicit the MPC. Section 4 presents the descriptive analysis and regression results relating the MPC to demographic variables, household resources, and proxies for credit constraints. In the same section we compare the empirical distribution of the MPC with the one that results from calibrating a dynamic heterogeneous agent model. Section 5 provides a series of policy experiments to illustrate how MPC heterogeneity affects the aggregate consumption response to tax reforms and redistributive fiscal policy. Section 6 concludes.

2. Theoretical predictions and empirical approaches

The standard life-cycle model and the permanent income hypothesis suggest that consumption should react to unanticipated income shocks, and that the response should be stronger for permanent than for transitory income shocks.[4] Our survey question refers to a transitory income shock, and hence we would expect a relatively low consumption response.

In models with quadratic utility, the MPC is equivalent to the annuity factor and is independent of household resources (the sum of current assets and the present discounted value of income). In the infinite version of the model, the annuity factor is (approximately) equal to the real interest rate, while in the finite-life version of the model MPC increases with age because the elderly have a shorter horizon over which to smooth consumption.

In contrast to models with quadratic utility, models with prudent individuals predict that MPC will depend on the level of household resources. Carroll and Kimball (1996) show that adding income uncertainty to the standard optimization problem with preferences characterized by

[4] The same models suggest that consumption does not respond to anticipated income changes.

prudence, produces a concave consumption function in which MPC from cash-on-hand declines with the level of wealth. The intuition is that consumers with less wealth have less ability to protect their consumption against income shocks. This means that as their wealth reduces, their fear of low income realization increases, which intensifies the strength of their precautionary saving motive.

Binding liquidity constraints also change some of the implications of the standard models. If people want to borrow but face binding liquidity constraints, the current marginal utility of consumption is high relative to future periods. Therefore an unexpected positive income shock raises current consumption one-to-one, regardless of the persistence of the shock. This assumes that the income shock is not large enough to overcome the credit constraint, which may be a reasonable assumption in the context of the survey question we exploit. [5]

Any test of the theoretical predictions regarding the size of the MPC involves the problem of identifying the type of income change faced by consumers. Jappelli and Pistaferri (2011) survey the methodologies used by applied researchers to identify exogenous and *unexpected* income changes: (a) quasi-experimental settings in which researchers use episodes when income changes exogenously in an unexpected way due, for example, to tax reforms; (b) making specific assumptions about the income process in order to distinguish the distribution of shocks from actual income realizations. [6] A useful alternative to these approaches is to sidestep the problem of identifying income changes, and to ask direct survey questions about the response of consumption to hypothetical or actual changes in income. In this paper we use a survey question to estimate the response of consumption to an unanticipated transitory income change.

Shapiro and Slemrod (1995; 2003) and Sahm, Shapiro and Slemrod (2009) have used this approach extensively in the US context. In particular, they rely on survey data to measure individual responses to actual or hypothetical tax policies. Household responses to tax rebates are coded according to three categories: mostly spend, mostly save, mostly pay off debt. [7] Given the qualitative nature of the survey questions, they rely on distributional assumptions to translate mostly-spend rates into an aggregate MPC. Shapiro and Slemrod (1995) examined the effectiveness of President Bush's temporary reduction in income tax with-holding in 1992.[8] One month after the tax change was implemented, they surveyed about 500 taxpayers and asked: (a) whether they were

[5] The bequest motive for saving is a further cause of MPC heterogeneity. A bequest typically is modeled as a luxury good, which leads to consumption being a concave function of cash-on-hand. Hence, when cash-on-hand increases by a given income amount, the poor mostly consume it, while the rich mostly save it to contribute to a bequest.

[6] Identifying such shocks from income realizations is difficult because people often have information that is not observable by the econometrician. For instance, they may know in advance that they will face a temporary change in their incomes (e.g. seasonal lay-off, or a tax rebate). When this change is realized, the econometrician measures it as a shock while, in fact, it is an expected event.

[7] E.g., the question in the Michigan survey was: "Thinking about your family's financial situation this year, did the tax rebate lead you mostly to increase spending, mostly to increase saving, or mostly to pay off debt?"

[8] The change was transitory since it was offset by a planned smaller tax refund in 1993.

aware that income tax withholding had decreased, and (b) what they were planning to do with the extra money in their paychecks, i.e., mostly save it or mostly spend it. They found that 40 percent of those interviewed planned to spend the extra take-home pay, suggesting that even a temporary tax change can be effective for increasing household spending.

Analysis of the 2001 income tax rebate reports a lower estimate of the marginal propensity to consume (22% of interviewed households reported plans to spend the tax rebate), and little evidence of myopia or liquidity constraints (Shapiro and Slemrod, 2003). Sahm, Shapiro and Slemrod's (2009) analysis of the 2008 tax stimulus provides intermediate findings (an MPC of about one-third) and also that the lowest income group and the group with no stock ownership have slightly higher mostly-spend rates than the higher-income groups and stock owners but the differences are small and statistically insignificant. Their final conclusion is that less-well-off households are not more likely than rich households to spend a tax rebate.[9]

It is worth pointing out that a quite different branch of the literature estimates the effect of *anticipated* tax changes on consumption using quasi-experimental approaches, as opposed to the response to *unanticipated* income changes we study. The best known applications use features of the tax system, such as social security payroll caps (Parker, 1999), tax refunds (Souleles, 1999), pre-announced tax cuts (Souleles, 2002), and the recent tax rebates or stimulus payments (Johnson, Parker and Souleles, 2006; Parker, Souleles, Johnson and McClelland, 2011; Agarwal, Liu and Souleles, 2007; and Misra and Surico, 2012). These studies typically find evidence of excess sensitivity of consumption to anticipated income changes induced by tax reductions, with relatively larger responses among low-wealth and low-income households, which is consistent with liquidity constraints. Given the small sample sizes and short panels, the evidence on heterogeneity in these studies is limited to group comparisons and rarely is precisely estimated. Moreover, any evidence of heterogeneity obtained in the "anticipated income change" context does not necessarily reproduce the extent of response heterogeneity in the "unanticipated income change" setting that we are studying.

3. The survey question

To characterize the MPC, we rely on the following question posed to respondents to the 2010 SHIW (the original Italian wording is reported in Appendix 1):

[9] Sahm, Shapiro and Slemrod (2009) also compare the pattern of answers provided in response to other US surveys that address the same question in slightly different ways. They compare 12 surveys conducted by nationally recognized polling organizations, and find that they are relatively consistent in relation to frequency of responses in the mostly spend/mostly save categories.

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.

The survey covers a cross-section of 7,951 households and provides detailed information on demographic variables, income, consumption, wealth (broken down into real assets, financial assets, and various debt components). We will see that it contains also useful variables to measure credit constraints, attitudes to risk, financial literacy, and other characteristics that potentially might affect the MPC.

The SHIW covers a representative sample of the Italian resident population.[10] 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 wherever possible. The unit of observation is the family, which is 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.

Note that, in contrast to surveys that ask for qualitative information (mostly save/mostly spend) about how people spend temporary tax rebates, the question we analyze is hypothetical and quantitative (people are asked what percentage of the reimbursement they would spend). The advantage of quantitative survey responses is that they overcome the problem of comparing responses across individuals who might interpret the statement "mostly spend/mostly save" in different ways. Another advantage is that by asking for a numerical value for the MPC, one avoids the problem of obtaining an aggregate MPC from the aggregation of qualitative responses. Finally, the design of the survey question addresses the potential problem that consumers' responses to income changes depends on the size of the change. The survey question ties the amount of the transfer received to the monthly income as opposed to asking how the respondent would spend a fixed sum of money (i.e., a $500 tax rebate).[11]

[10] The sample design is similar to the Labour Force Survey conducted by ISTAT (the Italian national statistics agency). Sampling is carried out in two stages: the first covers the selection of municipalities, the second the selection of households. Municipalities are categorized into 51 strata, defined by 17 regions and 3 classes of population size (over 40,000, 20,000-40,000, less than 20,000). All municipalities in the first group are included; those in the second and third groups are selected randomly with a probability proportional to their population size. In the second stage households are selected randomly from registry office records (Bank of Italy, 2012).

[11] Questions about the US 2008 tax rebate may suffer from a "size effect," if the magnitude of the rebate was small relative to the incomes of many households. Though the US rebates have been mainly lump sum, Parker et al. (2011) and Sahm et al. (2010) try to tease out the "size effect" by looking at rebates relative to income.

Some caveats are needed related to four issues: (a) consumption vs. spending, (b) external validity, (c) quality of interviews, and (d) time period for the spending decision. First, the question refers to the marginal propensity to spend, not to consume, and the two differ by the propensity to use the refund to purchase durable goods. If people's response is that they intend to spend a large part of the refund, this could refer to a plan to buy a durable good rather than to spend on non-durable goods. In the first case, people "mostly save" the refund despite declaring that they would spend 100 percent of it, making it more correct to use the term "marginal propensity to spend" rather than "marginal propensity to consume". While acknowledging the importance of this distinction, we prefer the acronym MPC and its conventional use.[12] Second, most of the interviews took place between April and July 2011, when the economy had just recovered from the 2007-09 recession, but also heading towards a further slump with GDP growth of -2.4% in 2012. Hence it might not be obvious how to extrapolate the distribution of the MPC estimated during a given year to other periods. A third caveat, common to all research eliciting subjective expectations or behavior in hypothetical scenarios, is that respondents might have little incentive to answer the questions truthfully, might have trouble understanding the wording of the questions, or might in practice display quite different behavior from their reported behavior. A final caveat concerns the fact that the question offers no period of reference for the planned expenditure.

At the end of the next section we discuss in detail the sensitivity of our results to the framing and context of the survey question. In particular, in Section 4.4 we consider whether the pattern of responses is affected by measures of financial literacy available in the survey, while in Section 4.5, we look at the issues of durables vs. nondurable spending, external validity, and timing of planned spending.

4. Empirical evidence

4.1. Descriptive evidence

Figure 1 plots the distribution of the responses to the survey question on MPC. It shows a heavy concentration of responses in three categories. Around 22 percent of respondents reported that they would spend nothing out of the reimbursement (MPC=0), around 16 percent reported that they would spend the entire reimbursement amount (MPC=1), and around 24 percent said they

[12] The importance of distinguishing between non-durable and total spending is highlighted by Parker, Souleles, Johnson and McClelland (2012), who find that households spent 12-30% of their 2008 US stimulus payments on non-durable goods, but this rose to 50-90% when durable goods are included.

would spend 50 percent of it. These responses are therefore reminiscent of the "mostly save", "mostly spend" and "partly save/partly spend" coding used by Shapiro and Slemrod in their qualitative assessment of the MPC. Figure 1 also shows a "heaping" at rounded values (10%, 20%, etc.).

The sample mean of the MPC is 48 percent, substantially higher than the predictions of standard consumption models on the impact of a transitory shock.[13] A possible explanation for this large deviation from the theoretical benchmarks is that the question does not distinguish between durable and non-durable consumption. Another explanation is that many households belong to population groups for whom liquidity constraints or myopia are important. For these reasons, we focus on the relation of the MPC with observable characteristics and in particular on the difference in the MPC across the distribution of household resources.

We start by exploring the relationship between MPC and cash-on-hand, defined as the sum of household disposable income and financial wealth, net of consumer debt. Figure 2 shows that average MPC declines sharply with cash-on-hand, from around 65 percent in the lowest cash-on-hand percentile to some 30 percent for the richest households.[14] Above we noted a large concentration of responses in the MPC=0 and MPC=1 categories. Figures 3 and 4 plot the fraction of respondents reporting MPC values at the two extremes of 0 and 1, by cash-on hand percentiles. The fraction of MPC=0 is only 10 percent for the poorest households, and increases steadily to 40 percent in the richer segment of the sample (Figure 3). The fraction reporting MPC=1 declines from 40 percent to 10 percent across the cash-on-hand distribution (Figure 4). This descriptive evidence is indicative of two aspects: first, there is wide MPC heterogeneity; second, MPC declines sharply with household resources.

4.2. Regression evidence

To characterize better the determinants of MPC variability in our sample we rely on regression analysis. Summary statistics for the main variables used in our estimation are presented in Table 1; Table 2 reports the baseline regressions. Since the dependent variable is censored from

[13] Kaplan and Violante (2010) investigate this issue simulating a life-cycle model in which consumers have isoelastic preferences, face income risk during their working lives, and retire at age 65. They conclude that the MPC with respect to permanent shocks is considerably larger than the MPC with respect to transitory shocks (0.7 vs. 0.06). Blundell, Pistaferri and Preston (2008) provide empirical estimates of the MPC with respect to income shocks that are not too dissimilar from these theoretical benchmarks.

[14] Using a completely different approach based on matching actual and predicted consumption transition probabilities, Jappelli and Pistaferri (2004) estimate an MPC of 0.4 for the low educated and 0.10 for the high educated, confirming a negative relation between measures of permanent income (measured here as education) and the MPC.

above and below, we use a two-limit Tobit estimator.[15] We present the main results, then focus on credit constraints and finally perform robustness checks.

The first specification in column (1) includes only demographic variables: age dummies, gender, marital status, education, family size, dummies for city size, and residence in the South. Except for family size, all coefficients are statistically different from zero. In particular, the MPC is roughly constant throughout the working life, declining after retirement by about 7 percentage points (the omitted category is the 60+ age group). This pattern contrasts with the predictions of standard consumption models that MPC, with respect to transitory shocks, increases with age. One possible explanation is that bequest motives, survival risk or risk of large medical expenses lead elderly households to save a larger fraction of their windfall income than younger households. Married couples show a slightly lower MPC, while more educated households display slightly lower MPC. Regional and city size dummies signal that the reported MPC tends to be higher for households in the South and in larger cities.[16]

However, some of these differences might be explained by correlation with omitted measures of household resources, as indicated graphically by the pattern in Figure 3. In column (2) we add cash-on-hand quintile dummies to the list of regressors. The results confirm a strong negative correlation between MPC and cash-on-hand. The coefficients are precisely estimated and decline monotonically with the quintile dummies. In particular, going from the first to the fifth cash-on-hand quintile is associated with a 30 percentage point decline in the MPC. Interestingly, the strongest decline occurs at low levels of wealth (an 11 percentage point decline between the first and second quintiles). Adding household resources changes the impact of education and age and reduces the size and significance of other coefficients. In particular, the age dummies are smaller and show a slightly increasing pattern during the working life, although there is still a decline after retirement. The third specification, in column (3), adds a dummy for unemployed household heads. While all other coefficients are unaffected, we find that the MPC is 7 percentage points higher for the unemployed, perhaps an indication of binding borrowing constraints or higher incidence of debt.[17]

It may be more appropriate to focus on people in the labor force, who face rather different constraints and shocks to their resources (e.g., income and unemployment shocks) with respect to the elderly, for whom health shocks, bequest motives, and survival risk play more important roles. Table 3 repeats the estimation, but excluding households older than 60. The results are largely

[15] The results presented below are unchanged if we use a simple OLS estimator.
[16] None of the results change if we replace the South dummy with a full set of regional dummies.
[17] Some people have volatile monthly income and therefore the answers to the survey question may depend on the particular month they have in mind when answering the question. To control for this effect, we add employment and sector dummies in the regression, and find similar results. Dropping the self-employed also does not change our results.

confirmed: (1) a strong negative correlation between MPC and cash-on-hand, (2) a roughly constant age pattern of MPC during the working life,[18] and (3) a positive, albeit less precisely estimated, coefficient of the unemployment dummy.

The effect of cash-on-hand on MPC may be due in part to changes in disposable income and in part to changes in financial assets, net of debt. Table 4 presents a breakdown of the cash-on-hand distribution into its three components, adding dummies for income and financial assets quintiles, and dummies for homeownership and positive debt.[19] The results in Table 4 column 1 show that each of the three components of cash-on-hand are strongly negatively correlated with the MPC. Going from the first to the fifth income quintile results in an 11 percentage point decline in the MPC, while going from the first to fifth quintile of financial assets results in a 26 point decline. Again, the steepest declines occur at low levels of income and of financial wealth (5 and 11 points, respectively). On the other hand, the coefficient of the dummy for homeownership is rather small and not statistically different from zero. When we introduce the unemployment dummy, in column 2 we find that the coefficient is slightly lower than in Table 2 (0.056 vs. 0.07) and is less precisely estimated.

The survey question does not allow us to distinguish between households who intend to repay debt rather than save or spend. In Italy, only 24 percent of households have any debt (including mortgage debt), the average debt-income ratio is 0.3, and only 3 percent have a debt/income ratio greater than 3. This issue is then probably less relevant in our Italian sample than in countries where households are highly leveraged. To further analyze the role of debt, we include in our specification a dummy for positive debt. The coefficient is negative and precisely estimated. Households with debt have an MPC 9 points lower than the rest of the sample. A plausible interpretation is that these households would use part of the reimbursement to pay off their debts, rather than increasing their spending.[20] As a final check, in column (3) of Table 4 we drop observations with positive debt (which hence makes the interpretation of the survey question unambiguous relative to the debt issue) and find no appreciable changes in any of our results.

4.3. Credit constraints

[18] Figure 1 in Meghir and Pistaferri (2010) shows that in models with quadratic utility and finite horizon the MPC with respect to transitory shocks only rises substantially after around retirement, so data may not have the power to detect small increases over the working stage of the life cycle.

[19] We also introduce the value of real assets: the coefficient is close to zero and imprecisely estimated.

[20] Adding to the specification the level of debt itself does not change our results, and the debt level variable is statistically insignificant.

As already pointed out, the strong negative correlation between MPC and cash-on-hand might be explained by precautionary saving, liquidity constraints, or a combination of the two. Liquidity constrained individuals increase spending one-to-one if they receive additional income. In practice, cash-on-hand might be strongly correlated with access to credit, and lack of it might exclude people from the credit market.

Similarly, the response of prudent individuals to an additional unit of income depends on cash-on-hand. Poor consumers live on the edge of (precautionary) fear, and hence an extra unit of income leads to a higher MPC. Rich consumers, instead, have already accumulated resources to smooth consumption, and hence have a lower MPC.

As noted by Deaton (1991), it is hard to distinguish empirically between borrowing constraints and precautionary savings. To isolate the relation between cash-on-hand and MPC, we can control directly for indicators of access to credit available from the SHIW. We consider three indicators: (1) whether people have been turned down for credit in the past, or have been discouraged from borrowing (the wording of the question is reported in Appendix 1); (2) whether they have access to an overdraft facility; (3) whether they have a credit card.

While the effect of current credit constraints is to increase the MPC, the effect of our proxy (being turned down for credit) is less clear-cut. On the one hand, being turned down when applying for credit may be correlated with the presence of current constraints, and therefore may be associated with a high MPC. On the other hand, having been turned down for credit in the past signals a high probability of being credit constrained in the future, which increases the expected variability of consumption growth, and therefore induces a low MPC due to a combination of borrowing constraints and precautionary saving.

The results are reported in Table 5. The first regression adds to the baseline the indicator for being turned down. The estimated coefficient suggests that the MPC is lower among those who report having been turned down for credit or discouraged from borrowing. Since this variable signals also a high probability of being credit constrained in the future, the negative coefficient is consistent with fear of future borrowing constraints in combination with a precautionary motive for saving. Indeed, using the panel section of the SHIW we find evidence that being turned down for/discouraged from applying for credit is persistent over time, providing support for the explanation that people who were unable to access credit in the previous year might fear being excluded from credit in the future.[21]

[21] SHIW contains indicators for access to credit between 1991 and 2010. We calculate the coefficient of correlation between contemporaneous and lagged access to credit in the panel section of the survey. We find that being turned down for credit in year t increases the likelihood of being turned down in year t+1 by about 10 percentage points.

The coefficients of the other indicators for access to credit (credit line and credit card) are not statistically different from zero (Table 5 column 2). The coefficient of the turned down question is higher in absolute value (-0.17) if we exclude the elderly (columns 3 and 4) from the sample.

4.4. Measurement Error

One concern with subjective expectations and hypothetical questions similar to the one exploited in this paper is people's understanding of the survey question. Our question on MPC requires the respondent to understand that the reimbursement is an amount equal to their monthly income (not an absolute amount), and to be able to compute a percentage (planned consumption relative to reimbursement).

To address this, Table 6 reports three robustness checks. In the first experiment in column (2), we check whether understanding of the financial questions is correlated to the MPC. We rely on an indicator of financial literacy constructed by summing correct answers to three questions on interest compounding, risk diversification, and mortgage contracts (the indicator ranges from 0 to 3). The wording of the literacy questions, proposed in Lusardi and Mitchell (2011), is set out in Appendix 1. The regression coefficient of financial literacy is close to zero, supporting the idea that people with relatively low financial literacy do not exhibit a tendency to under-report or over-report MPC.[22]

We also check whether particular values of MPC are associated with lower literacy (results not reported here for reasons of space, but available on request). We define dummy variables for each of three modes of the MPC distribution (0; 0.5; 1) and run probit regressions using the specification in Table 6, column 1. We find no evidence that lower level of literacy is associated with a higher likelihood of any of the three responses. The result for the probit regression which models the probability of responding MPC=0.5, is particularly interesting in light of the fact that when asked to pick a number between 0 and 1, some people might choose 0.5 because they are uncertain about the meaning of the question. A related concern is that low-wealth households may be more likely to report MPC=0.5 because of lower understanding of the survey question. The results, however, indicate that no variable of the baseline specification is systematically related to the probability of reporting MPC=0.5, including wealth. In particular, three of the cash-on-hand quintile coefficients are small and insignificant, while the coefficient on the top wealth quintile is significant but positive (0.04). As a further check of the sensitivity of the result to the inclusion of

[22] As a further robustness check we focus on the sample of individuals with relatively high literacy, that is, those responding correctly to at least two literacy questions. The pattern of results in column 2 is unchanged with respect to the baseline estimates.

the MPC=0.5 value, we drop these observations, which reinforces the relation between cash-on-hand and MPC (the difference between the bottom and top quintiles is 0.44).

As a final check, we focus on a sub-sample of individuals with high quality responses. At the end of each interview, SHIW interviewers evaluate the quality of the answers provided by respondents. The evaluation is performed on a 0-10 scale, representing the respondents' general understanding of the survey questions. In column (3) of Table 6 we focus on a sample of 5,479 high-quality interviewees, i.e. respondents who received a score of at least 8 from the interviewer.[23] The results are again confirmed: there is strong correlation between MPC and cash-on-hand, a relatively flat age profile of MPC until retirement, a positive but not always significant correlation with unemployment, and a negative effect of previous credit constraints on the MPC.

4.5. Robustness

In this subsection we consider the importance of the framing and context of the survey question in explaining our findings.

In standard intertemporal consumption model with durables, the level of nondurable spending changes in the same proportion as the stock of durables. For instance, Bertola et al. (2005) develop a case with a Cobb-Douglas utility function that delivers this property even in the presence of income risk. With adjustment costs on durable purchases, the model is of the (S,s) variety and implies that consumers with high values of the durable stock/nondurable spending are the ones more likely, for given realization of the income shock, to hit one of the bands and adjust their durable stock. In our context, consumers of this type are precisely the ones who might report a high MPC because they are planning to purchase a durable good. To control for this source of MPC heterogeneity, we use information available in our survey about nondurable spending and the stock of durables, construct the durable stock/nondurable ratio, and add it to our baseline specification. We find that the coefficient on the ratio is insignificant and the difference in MPC between top and bottom quintiles remains unchanged.

Another issue is that the reported MPC could depend on the macroeconomic context, namely the question was asked in the second quarter of 2011 (a transitional year between a double-dip recession), potentially limiting the external validity of our results. To partially address the issue, we use information available in the survey of whether people's consumption choices for the survey year were unusually high or low relative to a "normal" period.[24] Restricting the sample to those

[23] Results focusing on respondents who received a score of at least 9 or exactly 10 are also similar.

[24] The question asked is: "You told me that your average monthly spending in 2010 was [actual amount]. Would you say that this spending was unusually high, unusually low, or equal to what your household would spend in a "normal" year?".

who report that their spending was similar to what they would do in a "normal" year, we find that results are very similar. This suggests that people respond to the MPC question mostly having in mind a "normal" situation.

A final issue is that the question does not ask consumers over which period they plan to spend or save the transfer (i.e., one month, one year, or the entire lifetime). Hence, in principle MPC heterogeneity may reflect different beliefs about the timing of planned spending. However, if this was the case, one would have to conclude that the poor, who report higher MPC, are systematically more likely to believe that the question is asking about spending over a long period of time (i.e., MPC is close to 1 if respondent interpret the survey question as asking about spending the transfer over the entire lifetime), while the rich are systematically more likely to believe that the question refers to short-term spending (i.e., MPC is close to zero if respondents interpret the survey question as asking about spending the transfer over the next month). Although we cannot rule out this possibility, we find it a highly implausible and contrived route to explain the pattern of MPC heterogeneity we find in the data.

4.6. Interpretation

Given the many caveats related to the interpretation of the survey question (timing, durable vs. non-durable spending, the role of debt, etc.) we think it is more interesting to focus on the heterogeneity of the MPC by household characteristics rather than on its level. The most striking source of heterogeneity in MPC uncovered by our empirical analysis is the 25-30 percentage point increase when moving from the top to the bottom quantiles of the cash-on-hand distribution. This result is quite robust to the particular specification used, inclusion of control variables (including the durable-nondurable ratio, financial literacy, and context variables), and omission of the 50 percent mode.

One may wonder whether this evidence is consistent with models of intertemporal consumption behaviour. In the absence of precautionary savings and liquidity constraint, the MPC with respect to a transitory shock is homogeneous across consumers. To get a sense of the ability of alternative models to replicate the patterns in the data, we consider the Ayiagari's (1994) model populated by heterogeneous agents with CRRA preferences, an exogenous borrowing constraint preventing wealth from being negative, and an income process featuring a stochastic component given by the sum of an AR(1) process and an i.i.d. transitory shock.[25] After solving the model using standard calibrated parameters (an interest rate of 4 percent, a discount factor of 0.95, risk aversion of 2, an AR parameter of 0.98, a standard deviation of the persistent shock of 0.03, and a standard

[25] We thank Fabrizio Perri for sharing with us the Matlab and Stata codes generating these results. Details are available in Appendix 2.

deviation of the transitory shock of 0.01), we find the optimal consumption rule. To mimic the hypothetical windfall gain equivalent to 1/12 of yearly income (as considered in the survey question), we normalize mean income to 1, set the transitory shock to 0.1, compute the distribution of MPC with respect to transitory shocks implied by the model and compare it with the distribution of MPC in the data. In the Ayagari's model, the consumption function is concave and the MPC is a decreasing function of cash-on-hand, so we know that the model will be at least qualitatively consistent with the data. A different question is whether a version of the model is able to replicate the quantitative pattern of MPC heterogeneity we observe in the data. In this respect, the most crucial moment to match is the 25-30 percentage point decline across the cash-on-hand distribution.

After exploring a wide range of parameter values, it turns out that the key parameter affecting the shape of the MPC distribution by cash-on-hand is the discount rate. If we set the discount factor to β=0.95, the model generates a 5 percentage point decline in the MPC across the wealth distribution (see the lower line in Panel B of Figure A1). Raising the level of impatience in models with precautionary savings produces what has come to be known as the buffer stock model (Deaton, 1991; Carroll, 1999). Only if we make agents extremely impatient (letting β=0.6) can the buffer stock model model reproduce the extent of heterogeneity we observe in the data (the upper line in Panel B of Figure A1). This value of the discount factor is commonly regarded as implausible. Furthermore, it generates a distribution of wealth that does not match what is observed in the data.

To generate an MPC distribution that is based on more plausible parameter values and replicates the shape observed in the data, one may considered a version of the Ayagari's model in which (1) a fraction of consumers follow the optimal consumption rule dictated by the model, (2) another fraction follows a rule-of-thumb in which consumption equals income in each period, and (3) the fraction of rule-of-thumb consumers declines with the level of cash-on-hand. Although the latter assumption may seem *ad-hoc*, it is supported by empirical evidence from Italy (Jappelli and Pistaferri, 2004), and it is isomorphic to a model in which the poor have higher discount rates than the rich (ex-ante heterogeneity).[26] Using a completely different approach based on matching actual and predicted consumption transition probabilities, Jappelli and Pistaferri (2004) estimated that the fraction of rule-of-thumb consumers is 40 percent among the low educated (which is a group that includes disproportionally low cash-on-hand consumers) and 10 percent among the high educated (a group characterized by high cash-on-hand).[27] In Appendix 2 we show that a version of this mixture

[26] For instance, Cagetti (2003) uses PSID data and finds that β=0.98 for the high educated (college degree) and 0.84 for the low educated (high-school dropouts).

[27] In our data, high-educated households have cash-on-hand that is almost three times larger than low educated households.

model is able to replicate quite well the shape of the relationship between MPC and cash-on-hand we observe in the data (see Panel C of Figure A1).

We stress that this exercise is illustrative and not designed to estimate "structural parameters" or to disentangle one model of consumption behavior from the other. Rather, the objective is to show that the degree of heterogeneity we observe in the data, while qualitatively consistent with models with incomplete markets and precautionary savings à la Ayagari, is still far from the one we observe in theoretical models, and that only by considering additional departures from the standard model (such as high discount rates, heterogeneity of discount rates across the income distribution or myopic behavior) can one match theoretical predictions with data evidence.

5. Fiscal policy with heterogeneous MPC

How important is MPC heterogeneity for assessing the effect of transitory income changes (such as tax rebates or stimulus packages) on the aggregate economy? In this final section we show that our data can be used to calculate the effect on aggregate consumption of a series of fiscal policies. In evaluating the consumption effects of these policies, for simplicity, we use the MPC sample distribution. Results using predicted MPC from our regressions (which control for a large set of variables) are qualitatively similar, and not reported here for reasons of space. It is worth stressing that the findings below should be seen more as a way to illustrate the direction in which redistributive policies or policies targeted to the poor might affect aggregate consumption, rather than at their face value.

We consider several experiments. In the first experiment, we assume that government enacts a transfer policy financed by issuing debt (no taxes are levied).[28] In particular, we study a policy in which government transfers 1 percent of national disposable income equally among all individuals in the bottom 10 percent of the income distribution. This policy is equivalent to a transfer of €3,308, or 120 percent of average monthly income. We next consider two scenarios: in one MPC=0.48 for all individuals (the sample average), and in the other, the MPC is heterogeneous across the sample distribution. The results of these two experiments are reported in Table 7 (cases (a) and (b)), where

[28] In all experiments, we abstract from any general equilibrium effects arising from changes in asset prices and distortions induced by taxes and transfers on labor supply. However, it is worth noting that labor supply distortions of taxes are reduced by the lump sum nature of the tax (absent the targeting, there would be no distortions). We also neglect network effects on aggregate consumption arising from interdependent preferences, which might induce changes in consumption in partial equilibrium even among individuals not directly targeted by the policy (De Giorgi, Frederikssen and Pistaferri, 2012). Note however that we do not need to assume that consumers are non-Ricardian (as long as the answers to the reimbursement question incorporate anticipation of future tax increases required to pay for current transfers).

we show the aggregate MPC (in column (1)) and the aggregate consumption growth resulting from the experiment (column (2)). In the homogeneous case (a), the aggregate MPC is obviously equal to 0.48, and aggregate consumption increases by 0.62 percent. If the MPC is heterogeneous, targeting transfers at the bottom 10 percent of the population results in a higher aggregate MPC (0.62) and higher aggregate consumption growth (0.82 percent). The difference between the two cases is due to the higher MPC prevailing in the bottom part of the cash-on-hand distribution (shown clearly in Figure 2). Note that if government were to implement a pro-rich transfer to the top 10 percent of the income distribution (case (c) in Table 7), the aggregate MPC and consumption growth would be significantly lower (0.36 and 0.47%, respectively).

Another experiment we consider is to transfer 1 percent of national disposable income equally among all households with at least one unemployed member (14% of the sample), see case (d) in Table 7. This is equivalent to an unemployment bonus of €2,400 (about 87% of average monthly income), roughly equal to three months of the unemployment insurance received by blue-collar workers. The quantitative impact of this policy will be to boost aggregate consumption by 0.76 percent, with an estimated aggregate MPC of 0.58. The reason for the similar effect relative to a transfer to the bottom 10 percent of the income distribution is that households with unemployed members are mostly concentrated among the poor. [29]

A different (and perhaps more compelling) type of experiment is a balanced-budget redistributive policy whereby the government finances a transfer to the bottom x percent of the income or cash-on-hand distribution (where $1 \leq x \leq 89$) by taxing the top 10 percent of the income distribution. In all experiments, as before, tax revenues equal 1 percent of national disposable income. We assume that this amount is obtained by imposing a lump sum equal-sized tax on the top income decile, and that the government transfers this amount equally among targeted households. The details of the calculation are in Appendix 3.

Figure 5 plots the aggregate consumption growth generated by the policy. Of course, with a homogeneous MPC, a pure redistributive policy has no effect on aggregate consumption. However, with a heterogeneous MPC, the effect is positive and highest if the program targets the very poor. For instance, a transfer to the bottom 10 percent of the income distribution would raise aggregate consumption by 0.33 percent; if the same program targets people with below-median income, the boost in consumption would be around 0.2 percent.

We also consider a case where transfers are "means-tested" rather than being income-based. This case captures an income support program for the lowest x percent of the cash-on-hand

[29] We performed similar experiments giving the transfer to the lowest decile of the income (instead of the cash-on-hand) distribution, and obtain similar results, that is, a larger consumption effect for transfers to the bottom decile of the cash-on-hand distribution.

distribution. The shape of the curve in Figure 5 is similar to the income-based case, but the overall consumption effect is stronger. The larger effect for programs based on cash-on-hand transfers depends on the stronger negative correlation of MPC with financial assets (which of course are part of cash-on-hand) rather than income (as illustrated by the regression results in Table 4).[30]

One might wonder how much of these aggregate consumption effects are due to a correlation between the MPC and the income (or cash-on-hand) distributions as opposed to a correlation with other characteristics that vary across the income distribution, such as age, education, family size, etc. To control for these confounding factors, we expand the baseline regression in column 3 of Table 2, replacing the cash-on-hand quintiles with a set of percentile dummies (D_k). Thus, we run the regression:

$$MPC_i = X_i\beta + \sum_{k=1}^{100} \gamma_k D_{ik} + \varepsilon_i$$

We then use the predicted value $\hat{MPC}_i = \sum_{k=1}^{100} \hat{\gamma}_k D_{ik}$ to compute the consumption effect of the means-tested program that can be attributed solely to the MPC heterogeneity across the cash-on-hand distribution, controlling for demographic characteristics. The triangle scatter in Figure 5 ("means tested-based transfer, conditional MPC") shows that around 75% of aggregate consumption growth should be attributed to MPC heterogeneity across the income distribution, and about one quarter to correlation between MPC and other variables. To illustrate, consider the case of a redistributive policy targeted at the bottom 10 percent of the population. The figures shows that aggregate consumption would increase by 0.32 percent (the baseline scenario, or the red circle line). Of this increase, 0.27 per cent can be attributed to heterogeneity in the income distribution, and 0.05 percent to sources of heterogeneity that are unrelated to income.

6. Conclusions

We draw on survey questions in the 2010 Italian Survey of Household Income and Wealth that ask consumers how much they would consume of an unexpected transitory income change. We find that the marginal propensity to consume (MPC) is 48 percent on average, substantially higher

[30] While these experiments consider lump-sum transfers, we also experimented with a progressive transfer (assuming the transfer decreases linearly with the level of income). The effect on aggregate consumption growth is higher as expected, but not dramatically different, at least with the level of progressivity we chose. For example, a progressive transfer to the bottom 10% of the income distribution increases aggregate consumption by 0.36% as opposed to 0.32% in the lump-sum case.

than predicted by the standard intertemporal consumption choice model. This average masks very substantial MPC heterogeneity across households. Indeed, regression analysis uncovers four main facts: (1) a strong negative correlation between MPC and cash-on-hand, (2) a relatively flat age profile of MPC until retirement, (3) a positive but often imprecisely measured correlation with unemployment, and (4) a negative association between being turned down for credit and the MPC. Our interpretation of the evidence is that the MPC distribution reflects the importance of income risk and expectations of future constraints in shaping consumption decisions. However, comparison of the empirical distribution of the MPC with models with incomplete markets and precautionary savings à la Ayagari indicates that only by considering additional departures from the standard model (such as high discount rates, heterogeneity of discount rates across the income distribution or myopic behavior) can one match theoretical predictions with data evidence.

The results have important implications for evaluations of fiscal policy, and in particular predicting responses to tax reforms and redistributive policies. In particular, we find that a debt-financed increase in transfers of 1 percent of national disposable income targeted to the bottom decile of the cash-on-hand distribution would increase aggregate consumption by 0.82 percent. Furthermore, redistributing income from the top decile to the bottom decile of the income distribution would boost aggregate consumption by about 0.1 percent. One important caveat is that our calculations of the aggregate effects of fiscal policy are performed assuming no general equilibrium effects, and in particular that fiscal policy does not impact asset prices and that tax changes have no effect on labor supply. Hence, our calculations are likely to be an upper bound to the true effects of fiscal policy.

Appendix 1
Survey Questions

Propensity to spend

"*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*."[31]

Note that respondents only reply to the "how much would you spend" part of the question, as the survey taker enforces adding up to 100 of the two questions at the point of the interview.

Access to credit

1. Did the household contact a bank or financial company in 2010 with a view to obtaining a loan or mortgage?
2. Was the request granted in full, granted in part or refused?
3. Did your household later succeed in obtaining the amount needed, either from the same or from another financial intermediary?
4. During 2010 did you or a member of the household consider applying for a mortgage or a loan from a bank or financial company but later change your mind because you thought the request would be refused?

Financial literacy

1. Imagine leaving 1,000 euro in a current account that pays 1% interest and has no charges. Imagine that inflation is running at 2%. Do you think that if you withdraw the money in a year's time you will be able to buy the same amount of goods as if you spent the 1,000 euro today? The question is coded as: (1) Yes; (2) No, I will be able to buy less; (3) No, I will be able to buy more.

2. Which of the following investment strategies do you think entails the greatest risk of losing your capital? The question is coded as: (1) Investing in the shares of a single company; (2) Investing in the shares of more than one company; (3)

3. Which of the following types of mortgage do you think would allow you from the very start to fix the maximum amount and number of installments to be paid before the debt is extinguished? The question is coded as (1) Floating-rate mortgage; (2) Fixed-rate mortgage; (3) Floating-rate mortgage with fixed installments.

General understanding of the survey questions

At the end of the interview, interviewers are asked to evaluate the overall quality of the interview. The question we used was: "On a scale between 0 and 10, where 0 is the minimum and 10 the maximum, how would you evaluate the respondent's level of understanding of the survey questions?"

[31] In Italian the question is: "Supponga di ricevere improvvisamente un rimborso pari a quanto la sua famiglia guadagna in un mese. Di questa somma quanta parte ne risparmierebbe e quanta ne spenderebbe? Indichi la percentuale che verrebbe risparmiata e la percentuale che verrebbe spesa."

Appendix 2
The Aiyagari's model

We assume that agents solve the following problem:

$$\max \sum_t \beta^t \frac{c_t^{1-\gamma} - 1}{1-\gamma}$$

$$\text{s.t. } c_t + a_{t+1} \leq y_t + a_t(1+r)$$
$$a_{t+1} \geq 0$$

We assume the following income process:

$$y_t = \exp(z_t + \varepsilon_t)$$
$$z_t = \rho z_{t-1} + \eta_t$$

where ε_t and η_t are i.i.d. normal processes with mean zero and respective standard deviations of σ_ε and σ_η.

Let $c(a,z,\varepsilon)$ and $a'(a,z,\varepsilon)$ be the optimal decision rules. Normalizing income to 1, an increase of ε_t of 10 percent ($\bar{\varepsilon}$ =0.1) is akin to the rebate considered in the survey question. From the budget constraint, we have that

$$\left(c(a,z,\bar{\varepsilon}) - c(a,z,0)\right) + \left(a'(a,z,\bar{\varepsilon}) - a'(a,z,0)\right) = \exp(z)\bar{\varepsilon}$$

Hence, the model's equivalent of the MPC can be obtained directly from the decision rule as:

$$MPC(a,z)^* = \frac{c(a,z,\bar{\varepsilon}) - c(a,z,0)}{\exp(z)\bar{\varepsilon}}$$

For calibration purposes, we use the following parameter values: r=0.04, β=0.95, γ=2, ρ=0.98, σ_ε=0.01, σ_η=0.03.. In a second experiment, we set β=0.6. These parameter configurations generate the MPC distributions shown in Figure A1, together with the empirical distribution (Panels A and B). Unlike what presented in Figure 2, here we plot the MPC by the level of cash-on-hand (normalized by median income) rather than by percentiles of cash-on-hand. In particular, we allocate households in ten bins of the distribution of cash in hand (relative to median disposable income): (0-0.2), (0.2-0.4), (0.4-0.6), (0.6-0.8), (0.8-1), (1-1.8), (1.8-3.5), (3.5-6), (6-8.5), (8.5-10), and compute average MPC within each bin. The continuous lines are fractional polynomial fits, the shaded areas 95% confidence intervals.

In Panel C we consider a mixture model in which consumers have plausible discount rates (β=0.95) and they can either behave as in the Ayiagari's model above, or as rule-of-thumb consumers (consuming in each period their entire resources). We assume that the proportion of rule-of-thumb consumer declines with the level of cash-on-hand. In particular, we assume that the fraction of rule-of-thumb consumers is 75% in bins 1-3, 45% in bins 4-7, and 30% in bins 8-10.

Appendix 3
Calculation of the Fiscal Policy Experiments

In all experiments (both debt-financed and balanced-budget), we assume that total revenues raised by the government are equal to:

$$G = 0.01 \times \sum_i y_i$$

The net transfer received by household *i* in the debt-financed experiment *k* (see Table 7) is:

$$\tau_i(k) = \frac{G}{n^k} \mathbf{1}\{i \in g^k\}$$

where $\mathbf{1}\{i \in g^k\}$ is an indicator of whether *i* is a transfer recipient (which depends on which group is targeted by the policy) and n^k is the total number of transfer recipients.
The net transfer received by household *i* in the balanced-budget experiment *k* is:

$$\tau_i(k) = \frac{G}{\sum_{j=1}^{k} n_j} \mathbf{1}\{i \in p_1, ..., p_k\} - \frac{G}{\sum_{j=90}^{99} n_j} \mathbf{1}\{i \in p_{90}, ..., p_{99}\}$$

with $\sum_i \tau_i(k) = 0$, *k*=1,…,89, p_j indicates the j-th percentile of the relevant distribution (income or cash-in-hand), and n_j is the number of individuals in percentile *j*.

We define pre- and post-experiment aggregate consumption as:

$$C = \sum_i c_i, C'(k) = \sum_i (c_i + \beta_i \tau_i(k))$$

where β_i is the individual MPC elicited in the survey. The aggregate marginal propensity to consume we report in column 1 of Table 7 is simply:

$$MPC = \frac{C'(k) - C}{G} = \frac{\sum_i \beta_i \tau_i(k)}{G}$$

Consumption growth generated by the policy is:

$$\frac{C'(k) - C}{C} = \frac{\sum_i \beta_i \tau_i(k)}{\sum_i c_i}$$

This is what is reported in column 2 of Table 7, and plotted in Figure 5 for the 89 different values of *k*. In the same graph we also plot the result of basing the experiment on the distribution of cash-on-hand rather than income. This is what we called the "means-tested" experiment. Finally, we plot the result of basing the experiment on the distribution of cash-on-hand controlling for demographic characteristics as explained in Section 5.

References

Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles (2007), "The Reaction of Consumer Spending and Debt to Tax Rebates-Evidence from Consumer Credit Data," *Journal of Political Economy*, 115, 986-1019.

Aiyagari, S. Rao (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", *The Quarterly Journal of Economics*, Vol. 109 (3), 659-684.

Bank of Italy (2012), "I bilanci delle famiglie Italiane nell'anno 2010," Supplemento al Bollettino Statistica, No. 6, January.

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

Cagetti, Marco (2003), "Wealth Accumulation Over the Life Cycle and Precautionary Savings", *Journal of Business and Economic Statistics*, 4(28), 15-57.

Carroll, Christopher D (1996), "A Theory of the Consumption Function, With and Without Liquidity Constraints", *Journal of Economic Perspectives*, 15(3), 23–46.

Carroll, Christopher D, and Miles Kimball (1996), "On the Concavity of the Consumption Function," *Econometrica*, 64, 981-92.

Dynan, Karen (2012), "Is A Household Debt Overhang Holding Back Consumption?" *Brookings Papers on Economic Activity* , Spring, 299-362.

Dynan, Karen, Jonathan Skinner, and Stephen P. Zeldes (2004), "Do the Rich Save More?," *Journal of Political Economy*, 112, 397-444.

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

De Giorgi, Giacomo, Anders Frederikssen, and Luigi Pistaferri (2012), "Consumption Network Effects." Stanford University, mimeo.

Jappelli, Tullio, and Luigi Pistaferri (2004), "Intertemporal choice and consumption mobility", *Journal of the European Economic Association* 4(1), 75-115, 2006.

Jappelli, Tullio, and Luigi Pistaferri (2011), "The Consumption Response to Income Changes," *Annual Review of Economics* 2, 479-506.

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

Kaplan Greg, and Giovanni L. Violante (2010), "How Much Consumption Insurance Beyond Self-Insurance?," *American Economic Journal: Macroeconomics* 2, 53-87.

Kaplan Greg, and Giovanni L. Violante (2011), "A Model of the Consumption Response to Fiscal Stimulus Payments," NBER Working Papers 17338.

Lusardi, Annamaria, and Olivia S. Mitchell (2011), "Financial Literacy Around the World: an Overview," *Journal of Pension Economics and Finance* 10, 497-508.

Meghir, Costas, and Luigi Pistaferri (2011), "Earnings, consumption and life cycle choices", in *Handbook of Labor Economics*, O. Ashenfelter and D. Card (eds), vol. 4B, 774-854.

Mian, Atif R., and Amir Sufi (2010), "Household Leverage and the Recession of 2007 to 2009," *IMF Economic Review*, 58(1), 74-117.

Misra, Kanishka and Paolo Surico (2012), "Consumption, Income Changes and Heterogeneity: Evidence from Two Fiscal Stimulus Programmes". London Business School, mimeo.

Parker, Jonathan A. (1999), “The Reaction of Household Consumption to Predictable Changes in Social Security Taxes”, *American Economic Review* 89, 959-7

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson and Robert McClelland (2011), “Consumer Spending and the Economic Stimulus Payments of 2008,” NBER Working Papers No. 16684.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod (2010). "Household Response to the 2008 Tax Rebates: Survey Evidence and Aggregate Implications," in Brown, Jeffrey R. ed., *Tax Policy and the Economy*, Vol. 24. National Bureau of Economic Research, pp. 69-110.

Shapiro Matthew D., Slemrod Joel (1995), “Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding”, *American Economic Review* 85, 274-283.

Shapiro Matthew D., Slemrod Joel (2003), “Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys”, in *Tax Policy and the Economy*. Vol. 17, ed. James M. Poterba, Cambridge, MA: MIT Press.

Stephens Jr. Mel (2001), The Long-Run Consumption Effects of Earnings Shocks. *The Review of Economics and Statistics*. 83(1):28-36

Souleles Nicholas S. (1999), “The Response of Household Consumption to Income Tax Refunds”, *American Economic Review* 89, 947-958.

Souleles Nicholas S. (2002), “Consumer Response to the Reagan Tax Cuts”, *Journal of Public Economics* 85, 99-120.

Figure 1. Self-reported MPC from transitory income shock

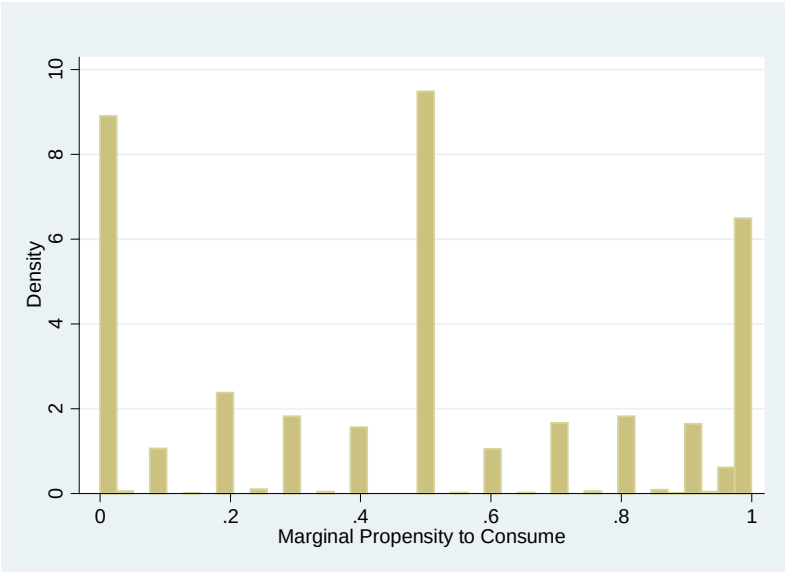

Figure 2. Average MPC by cash-on-hand percentiles

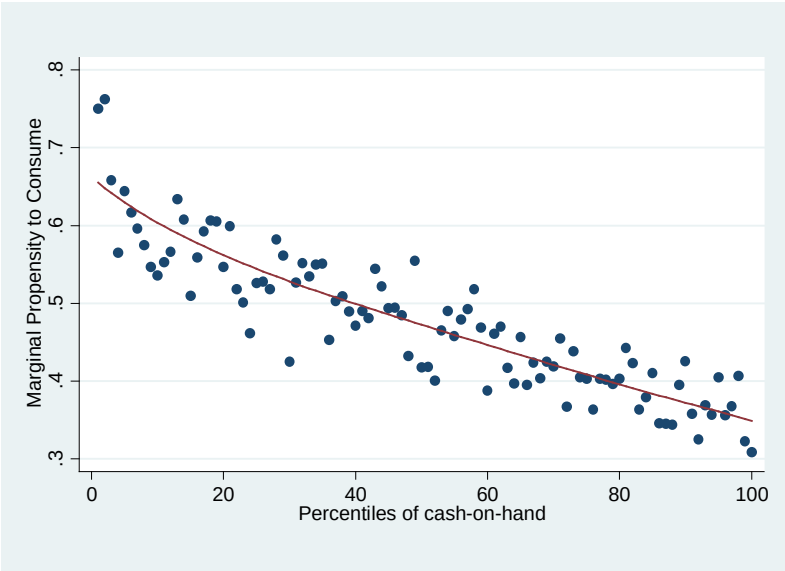

Figure 3. Fraction with MPC=0 by cash-on-hand percentiles

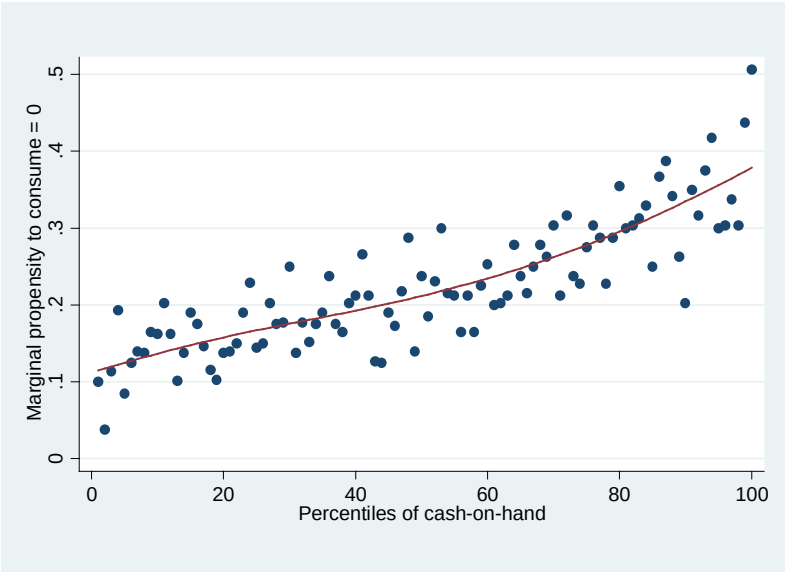

Figure 4. Fraction with MPC=1 by cash-on-hand percentiles

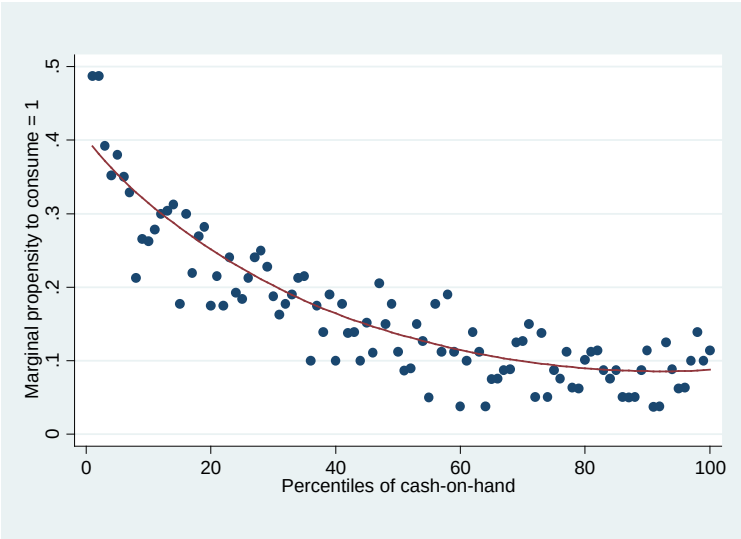

Figure 5: The effect of a redistributive transfer program

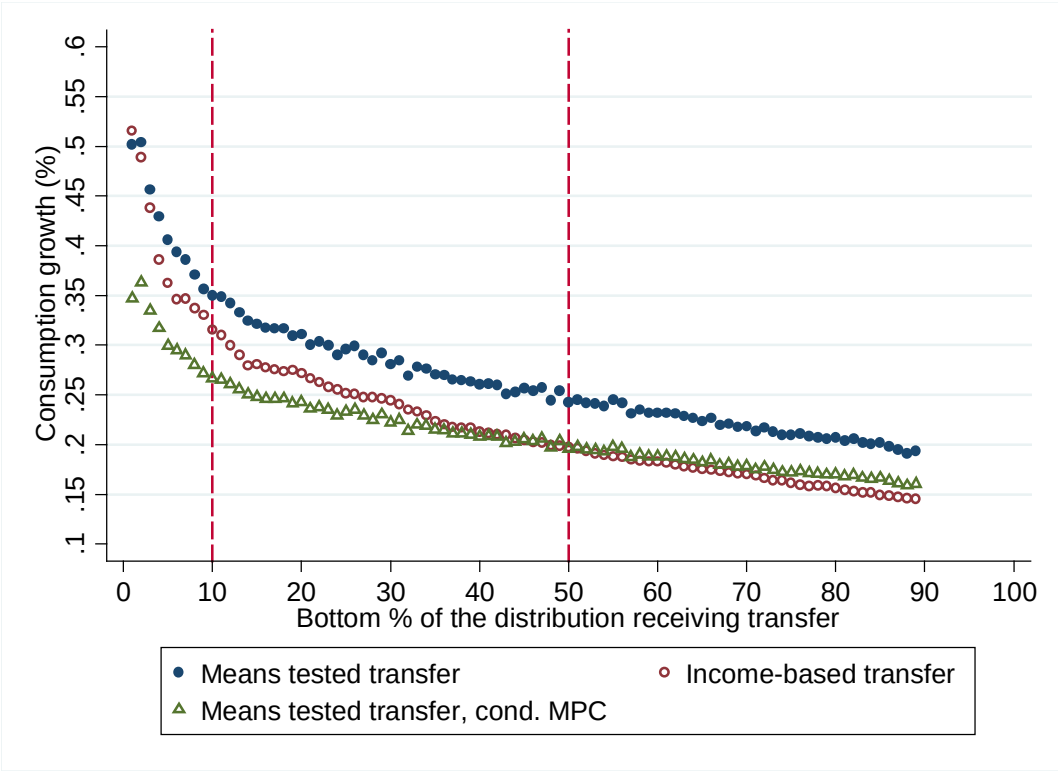

Table 1. Descriptive statistics

	Mean	*Standard deviation*	*Minimum*	*Maximum*
Marginal propensity to consume (MPC)	0.476	0.357	0.000	1.000
Age	58.374	15.761	18.000	99.000
Male	0.545	0.498	0.000	1.000
Married	0.623	0.485	0.000	1.000
Years of education	9.576	4.595	0.000	20.000
Resident in the South	0.324	0.468	0.000	1.000
Family size	2.495	1.257	1.000	12.000
City size less then 20,000	0.256	0.437	0.000	1.000
City size 20,000-40,000	0.182	0.386	0.000	1.000
City size 40,000-500,000	0.470	0.499	0.000	1.000
City size over 500,000	0.091	0.288	0.000	1.000
Cash (‘000 euro)	62.482	111.871	0.000	4381.469
Disposable income (‘000 euro)	33.089	24.530	0.000	587.784
Financial wealth (‘000 euro)	29.393	98.324	0.000	4154.771
Unemployed	0.035	0.185	0.000	1.000
Liquidity constrained	0.047	0.212	0.000	1.000
Overdraft	0.269	0.444	0.000	1.000
Credit card	0.320	0.467	0.000	1.000
Late payment	0.010	0.102	0.000	1.000

Note. Data are drawn from the 2010 SHIW. Total number of observations is 7,950.

Table 2. Baseline estimates

	(1)	(2)	(3)
Age 18-30	0.111	0.023	
	(0.034)***	(0.035)	
Age 31-45	0.106	0.042	0.040
	(0.018)***	(0.019)**	(0.019)**
Age 46-60	0.076	0.051	0.048
	(0.016)***	(0.016)***	(0.016)***
Male	-0.028	-0.016	-0.017
	(0.013)**	(0.013)	(0.013)
Married	-0.041	-0.013	-0.012
	(0.016)**	(0.016)	(0.016)
Years of education	-0.005	0.005	0.005
	(0.001)***	(0.002)***	(0.002)***
Family size	0.009	0.023	0.022
	(0.007)	(0.007)***	(0.007)***
Resident in the South	0.339	0.276	0.276
	(0.013)***	(0.014)***	(0.014)***
City size less than 20,000	-0.208	-0.193	-0.193
	(0.023)***	(0.023)***	(0.023)***
City size 20.000-40,000	-0.185	-0.174	-0.174
	(0.024)***	(0.024)***	(0.024)***
City size 40,000-500,000	-0.125	-0.120	-0.120
	(0.022)***	(0.021)***	(0.021)***
I cash-on-hand quintile		0.293	0.287
		(0.024)***	(0.024)***
II cash-on-hand quintile		0.186	0.184
		(0.021)***	(0.021)***
III cash-on-hand quintile		0.133	0.132
		(0.020)***	(0.020)***
IV cash-on-hand quintile		0.063	0.062
		(0.019)***	(0.019)***
Unemployed			0.070
			(0.034)**
N	7,950	7,950	7,950

Note. Estimation performed using Tobit. Standard errors are reported in parentheses. *indicates significance at the 10% level, ** significance at 5%, *** significance at 1%.

Table 3. Determinants of MPC: Age less than 60

	(1)	(2)	(3)
Age 18-30	0.036	-0.033	-0.035
	(0.034)	(0.034)	(0.034)
Age 31-45	0.032	-0.008	-0.008
	(0.017)*	(0.017)	(0.017)
Male	-0.045	-0.032	-0.032
	(0.016)***	(0.016)**	(0.016)**
Married	-0.050	-0.020	-0.018
	(0.022)**	(0.022)	(0.022)
Years of education	-0.010	0.001	0.001
	(0.002)***	(0.002)	(0.002)
Family size	0.005	0.019	0.019
	(0.008)	(0.008)**	(0.008)**
Resident in the South	0.338	0.269	0.269
	(0.018)***	(0.019)***	(0.019)***
City size less than 20,000	-0.176	-0.158	-0.158
	(0.030)***	(0.030)***	(0.030)***
City size 20.000-40,000	-0.178	-0.166	-0.166
	(0.031)***	(0.031)***	(0.031)***
City size 40,000-500,000	-0.099	-0.095	-0.095
	(0.028)***	(0.027)***	(0.027)***
I cash-on-hand quintile		0.321	0.315
		(0.031)***	(0.031)***
II cash-on-hand quintile		0.199	0.197
		(0.028)***	(0.028)***
III cash-on-hand quintile		0.123	0.122
		(0.026)***	(0.026)***
IV cash-on-hand quintile		0.070	0.070
		(0.025)***	(0.025)***
Unemployed			0.039
			(0.034)
N	4,315	4,315	4,315

Note. Estimation performed using Tobit. Standard errors are reported in parentheses. *indicates significance at the 10% level, ** significance at 5%, *** significance at 1%.

Table 4. Determinants of MPC: The role of disposable income, financial assets and debt

	(1)	(2)	(3)
Age 18-30	0.038	0.033	0.088
	(0.035)	(0.035)	(0.041)**
Age 31-45	0.057	0.055	0.090
	(0.019)***	(0.019)***	(0.023)***
Age 46-60	0.060	0.058	0.074
	(0.016)***	(0.016)***	(0.018)***
Male	-0.013	-0.013	-0.027
	(0.013)	(0.013)	(0.015)*
Married	-0.012	-0.011	-0.005
	(0.016)	(0.016)	(0.018)
Years of education	0.006	0.006	0.007
	(0.002)***	(0.002)***	(0.002)***
Family size	0.022	0.021	0.021
	(0.007)***	(0.007)***	(0.008)**
I income quintile	0.115	0.110	0.123
	(0.028)***	(0.028)***	(0.032)***
II income quintile	0.058	0.055	0.094
	(0.024)**	(0.024)**	(0.027)***
III income quintile	0.057	0.055	0.073
	(0.021)***	(0.021)**	(0.025)***
IV income quintile	0.032	0.031	0.056
	(0.020)	(0.020)	(0.023)**
I financial asset quintile	0.258	0.257	0.242
	(0.024)***	(0.024)***	(0.027)***
II financial asset quintile	0.146	0.146	0.144
	(0.022)***	(0.022)***	(0.024)***
III financial asset quintile	0.098	0.098	0.112
	(0.020)***	(0.020)***	(0.023)***
IV financial asset quintile	0.045	0.046	0.043
	(0.020)**	(0.020)**	(0.022)*
Homeowner	-0.003	-0.003	-0.002
	(0.015)	(0.015)	(0.017)
Positive debt	-0.090	-0.090	
	(0.015)***	(0.015)***	
Unemployed		0.056	0.017
		(0.034)*	(0.040)
N	7,950	7,950	5,967

Note. Regressions include dummies for South and City Size (less than 20,000 inhabitants, 20-40,000, and 40,000-500,000. Estimation performed using Tobit. Standard errors are reported in parentheses. *indicates significance at the 10% level, ** significance at 5%, *** significance at 1%.

Table 5. Determinants of MPC: The role of liquidity constraints

	(1)	(2)	(3)	(4)
Age 18-30	0.023	0.024	-0.033	-0.033
	(0.035)	(0.035)	(0.034)	(0.034)
Age 31-45	0.047	0.048	-0.003	-0.003
	(0.019)**	(0.019)**	(0.017)	(0.017)
Age 46-60	0.052	0.053		
	(0.016)***	(0.016)***		
Male	-0.015	-0.015	-0.030	-0.030
	(0.013)	(0.013)	(0.016)*	(0.016)*
Married	-0.014	-0.014	-0.023	-0.022
	(0.016)	(0.016)	(0.022)	(0.022)
Years of education	0.005	0.005	0.000	0.001
	(0.002)***	(0.002)***	(0.002)	(0.002)
Family size	0.022	0.022	0.019	0.019
	(0.007)***	(0.007)***	(0.008)**	(0.008)**
I cash-on-hand quintile	0.293	0.292	0.326	0.322
	(0.024)***	(0.025)***	(0.031)***	(0.033)***
II cash-on-hand quintile	0.185	0.184	0.200	0.198
	(0.021)***	(0.022)***	(0.028)***	(0.029)***
III cash-on-hand quintile	0.133	0.132	0.125	0.123
	(0.020)***	(0.021)***	(0.026)***	(0.027)***
IV cash-on-hand quintile	0.062	0.062	0.070	0.069
	(0.019)***	(0.019)***	(0.025)***	(0.025)***
Unemployed	0.075	0.075	0.044	0.044
	(0.034)**	(0.034)**	(0.034)	(0.034)
Turned down for credit or discouraged	-0.125	-0.126	-0.176	-0.176
	(0.029)***	(0.029)***	(0.032)***	(0.032)***
Overdraft		0.013		0.000
		(0.014)		(0.018)
Credit card		-0.011		-0.007
		(0.015)		(0.019)
N	7,950	7,950	4,315	4,315

Note. Regressions include dummies for South and City Size (less than 20,000 inhabitants, 20-40,000, and 40,000-500,000. Estimation performed using Tobit. Standard errors are reported in parentheses. *indicates significance at the 10% level, ** significance at 5%, *** significance at 1%.

Table 6. Determinants of MPC: Measurement error

	Total sample	Excluding financial literacy <2	Excluding quality of interview less than 8
	(1)	(2)	(3)
Age 18-30	0.029	-0.006	-0.008
	(0.035)	(0.042)	(0.043)
Age 31-45	0.055	0.036	0.038
	(0.019)***	(0.022)	(0.023)*
Age 46-60	0.059	0.034	0.047
	(0.016)***	(0.019)*	(0.020)**
Male	-0.012	-0.027	-0.020
	(0.013)	(0.015)*	(0.016)
Married	-0.010	0.004	-0.018
	(0.016)	(0.020)	(0.021)
Years of education	0.007	0.002	0.005
	(0.002)***	(0.002)	(0.002)**
Family size	0.021	0.010	0.011
	(0.007)***	(0.008)	(0.008)
I cash-on-hand quintile	0.258	0.296	0.307
	(0.024)***	(0.030)***	(0.031)***
II cash-on-hand quintile	0.165	0.143	0.167
	(0.021)***	(0.025)***	(0.026)***
III cash-on-hand quintile	0.120	0.121	0.119
	(0.020)***	(0.023)***	(0.023)***
IV cash-on-hand quintile	0.054	0.047	0.056
	(0.019)***	(0.021)**	(0.022)**
Unemployed	0.075	0.066	0.093
	(0.033)**	(0.040)*	(0.047)**
Turned down for credit or discouraged	-0.131	-0.067	-0.145
	(0.029)***	(0.036)*	(0.038)***
Financial literacy	-0.009		
	(0.006)		
N	7,950	5,292	5,479

Note. Regressions include dummies for South and City Size (less than 20,000 inhabitants, 20-40,000, and 40,000-500,000. Estimation performed using Tobit. Standard errors are reported in parentheses. *indicates significance at the 10% level, ** significance at 5%, *** significance at 1%.

Table 7. Effect of transfer policy financed by debt

Policy: Transfer equivalent to 1% of national disposable income	***MPC=ΔC/ΔY*** **(1)**	***Aggregate consumption growth*** **(2)**
Homogeneous MPC		
(a) Transfer to bottom 10%	0.48	0.62%
Heterogeneous MPC		
(b) Transfer to bottom 10%	0.62	0.82%
(c) Transfer to top 10%	0.36	0.47%
(d) Transfer to households with unemployed members	0.58	0.76%

Note. In column (1), we report the aggregate MPC, computed as: $\frac{\sum_i \beta_i \tau_i(k)}{G}$, where β_i is the individual MPC, τ_i (k) is the transfer received by household i in policy experiment k, and G are total government revenues raised. In column (2), we report aggregate consumption growth, defined as $\frac{\sum_i \beta_i \tau_i(k)}{\sum_i c_i}$, where c_i denotes household consumption. In experiments (a) and (b), transfers are distributed equally among members of the first decile of the cash-on-hand distribution; in experiment (c), transfers are distributed equally among members of the top decile of the cash-on-hand distribution; finally, in experiment (d) transfers are distributed equally among households with at least one unemployed member. See Appendix 3 for more details.

Figure A1: MPC Heterogeneity in the data and in two consumption models

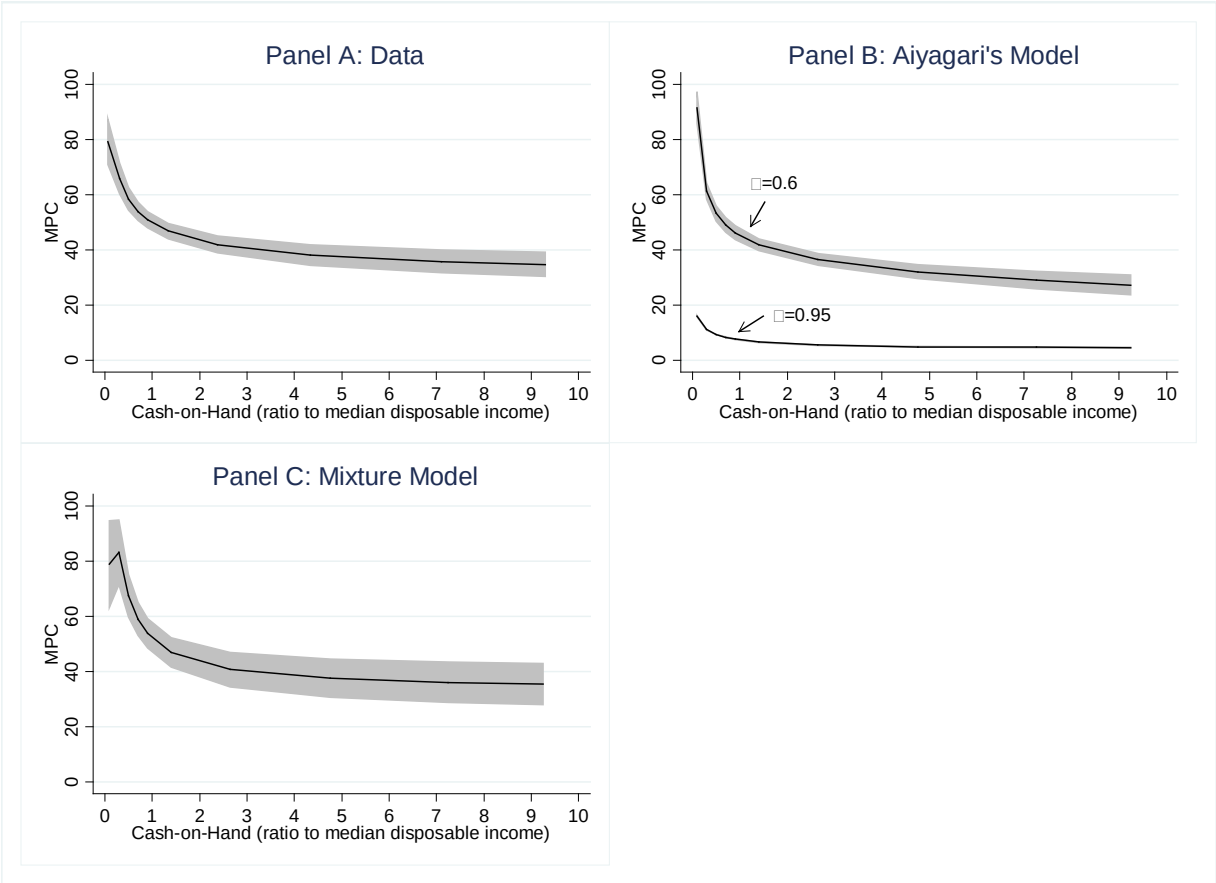

Note. Each graph plots average MPC for households in ten bins of the distribution of cash-on-hand (relative to median disposable income): (0-0.2), (0.2-0.4), (0.4-0.6), (0.6-0.8), (0.8-1), (1-1.8), (1.8-3.5), (3.5-6), (6-8.5), (8.5-10). The continuous lines are fractional polynomial fits, the shaded areas are 95% confidence intervals. In Panel B we solve Ayagari's model for two values of β (0.95 and 0.6). In Panel C we assume β=0.95 and that the fraction of rule-of-thumb consumers is 75% in bins 1-3, 45% in bins 4-7, and 30% in bins 8-10.

CFS Working Paper Series

No.	Author(s)	Title
2013/13	**Jeffrey V. Butler** **Luigi Guiso** **Tullio Jappelli**	Manipulating Reliance on Intuition Reduces Risk and Ambiguity Aversion
2013/12	**Tullio Jappelli** **Mario Padula**	Investment in Financial Literacy, Social Security and Portfolio Choice
2013/11	**Christiane Baumeister** **Lutz Kilian**	Forecasting the Real Price of Oil in a Changing World: A Forecast Combination Approach
2013/10	**Christiane Baumeister** **Lutz Kilian**	Do Oil Price Increases Cause Higher Food Prices?
2013/09	**Christiane Baumeister** **Lutz Kilian** **Xiaoqing Zhou**	Are Product Spreads Useful for Forecasting? An Empirical Evaluation of the Verleger Hypothesis
2013/08	**Hugh H. Kim** **Raimond Maurer** **Olivia S. Mitchell**	Time is Money: Life Cycle Rational Inertia and Delegation of Investment Management
2013/07	**Andreas Hubener** **Raimond Maurer** **Olivia S. Mitchell**	How Family Status and Social Security Claiming Options Shape Optimal Life Cycle Portfolios
2013/06	**Günter Coenen** **Anders Warne**	Risks to Prize Stability, the Zero Lower Bound and Forward Guidance: A Real-Time Assessment
2013/05	**Stefan Mittnik**	VaR-implied Tail-correlation Matrices
2013/04	**Hans-Joachim Dübel**	The Capital Structure of Banks and Practice of Bank Restructuring

All CFS Working Papers can be downloaded at *www.ifk-cfs.de/publications/working-papers*.