

ASYMMETRIC CONSUMPTION EFFECTS OF TRANSITORY INCOME SHOCKS*

*Dimitris Christelis, Dimitris Georgarakos, Tullio Jappelli,
Luigi Pistaferri and Maarten van Rooij*

We use the responses of a representative sample of Dutch households to survey questions that ask how much their consumption would change in response to unexpected, transitory income shocks (positive or negative). The questionnaire also distinguishes between relatively small income changes (a one-month increase or drop in income), and relatively larger ones (equal to three-months' income). The results are broadly in line with models of intertemporal choice with precautionary saving, borrowing constraints and finite horizons.

Evaluating the effect of a broad set of policy interventions, including fiscal and monetary policies, on household and aggregate consumption requires reliable estimates of the consumption response to income shocks, i.e., the marginal propensity to consume (MPC). Distinguishing whether consumption responds differently to positive and negative income changes, and whether the response depends on the size of the shock are equally important questions.

To address these issues, we use the responses from a representative sample of Dutch households to survey questions that ask how much they would consume of an unexpected, transitory and positive income change, and by how much they would reduce their consumption in response to an unexpected, transitory and negative income change. In addition, the survey questionnaire allows respondents to distinguish between relatively small income changes (an increase or reduction equivalent to roughly one-month's income) and relatively larger ones (equivalent to three-months' income).

The survey hence allows us to characterise empirically the distribution of the MPC in response to shocks of different sign and size and compare the findings with the predictions of intertemporal consumption models. Specifically, we test whether the consumption response to income shocks declines with the level of economic resources; whether the MPC is smaller if the consumer has a relatively long time horizon; whether consumption responds differently to positive income shocks and negative income shocks; and whether the response is stronger for larger income shocks. The main advantage of using our survey is that it allows us to compare the responses of *the same household* to hypothetical income shocks, hence replicating a quasi-experimental setting or one in which it is possible to control for unobserved heterogeneity. In contrast, a *realised* income shock is either positive or negative, small or large, and therefore, comparing the consumption

* Corresponding author: Tullio Jappelli, Department of Economics, University of Naples Federico II, Via Cintia, Monte S. Angelo, I-80126 Naples (NA), Italy. Email: tullio.jappelli@unina.it

This paper was received on 15 February 2017 and accepted on 24 June 2018. The Editor was Morten Ravn.

We are grateful to Morten Ravn (the Editor) and two anonymous referees for their constructive comments. We would like to thank seminar participants at the 2017 NBER group on 'Consumption: Micro to Macro', the 2017 ASSA meetings, the European Central Bank, De Nederlandsche Bank, the 2017 CSEF-IGIER Conference on Economics and Institutions, the University of Birmingham and CSEF for helpful suggestions and comments. We are also grateful to Jian Li and Yichen Su for providing excellent research assistance. Christelis and Jappelli acknowledge financial support from the Italian Ministry of University and Research. The opinions expressed in the article are those of the authors and do not necessarily reflect the views of the European Central Bank or of De Nederlandsche Bank. The data and programmes used to generate all results can be found at the journal website.

responses to shocks of different sign or size also reflects the different characteristics (observed and unobserved) of the selected sample that is subject to the shocks (and in most cases, the business cycle context in which the shocks occur).

Our empirical findings are broadly in line with models of intertemporal choice with precautionary saving, borrowing constraints and finite horizons. The average MPC is in the 15–25% range; it is larger for negative income shocks, it is larger at low levels of cash-on-hand (at least in the case of negative shocks), and it increases with age, particularly for the oldest group. We also find that the MPC distribution is in line with two predictions of models with liquidity constraints. First, as shown by simple simulation analysis, in the presence of income risk and liquidity constraints the MPC from negative income shocks is larger than that from positive shocks. Second, in the presence of liquidity constraints the size of the shock also matters. In the case of negative income shocks the MPC is large, and increases with the size of the shock, particularly for households with low cash-on-hand. In the case of positive shocks, liquidity-constrained consumers are more likely to overcome the constraint when the shock is large, and therefore the MPC from large positive income shocks should be lower than that from small positive shocks. The survey allows us to test these important and as yet unexplored implications of liquidity constraints, and our empirical results are in line with both of them, the only exception being that the MPC does not vary much with the size of the negative shocks.

We find considerable heterogeneity in the MPC distribution, with about 40% of respondents exhibiting symmetric MPCs in response to shocks of different size and direction, as in the standard permanent income hypothesis (PIH). About 40% exhibit asymmetric responses consistent with models with liquidity constraints, and about 20% exhibit asymmetric MPCs that are possibly consistent with behavioural models or lack of financial sophistication.

We contribute to the literature on MPC estimation based on income shocks. One of the difficulties affecting the estimation of the MPC is isolating the exogenous income shocks affecting consumption over time. The literature suggests three approaches to deal with this issue (see Jappelli and Pistaferri, 2011, for a survey). The first approach identifies episodes in which income changes due to exogenous events such as unemployment, disability or tax rebates, and evaluates in a quasi-experimental setting how consumption reacts to such changes: see, for instance, Souleles (1999), Agarwal *et al.* (2007) and Misra and Surico (2014). The second approach relies on the statistical decomposition of income shocks and the covariance restrictions imposed by the theory on the joint behaviour of income and consumption, in combination with long panel data to relate income shocks to consumption growth (Blundell *et al.*, 2008). Survey questions measuring the responses to actual or hypothetical income changes are the third approach.¹ Shapiro and Slemrod (2003) and Sahm *et al.* (2010, 2015) asked U.S. households to report how their consumption had changed in response to tax rebates, tax credits and payroll tax changes in the previous 15 years. Jappelli and Pistaferri (2014) analyse how a hypothetical tax rebate affects consumption and find an inverse relation between MPC and cash-on-hand, which is consistent with models with liquidity constraints and precautionary saving.

Using this approach, a few studies compare the consumption response with positive and negative income changes. Bracha and Cooper (2014) report that taxpayers in the United States are much more likely to reduce spending in response to the payroll tax increase than they are to increase spending based on their anticipated tax refund. Bunn *et al.* (2018), using the Bank

¹ Parker and Souleles (2017) compare reported preferences for spending in response to various tax policies with actual follow-up spending behaviour and find that the two are well aligned. Smith *et al.* (2014) find that subjects' reported preferences over a set of food items is a good predictor for their follow-up actual food choice.

of England Survey, find that British households change their consumption more in reaction to negative transitory income shocks than in reaction to positive ones. One important advantage of our survey over previous studies is that we ask the same set of questions to all households. Second, we can analyse, for each household, the consumption response to shocks of different sign as well as different size. Third, we study if these asymmetries depend on differences in cash-on-hand, age and other demographic variables. Fuster *et al.* (2018) use the same approach, asking respondents from the New York Fed Survey of Consumer Expectations to 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.

The article is organised as follows. Section 1 discusses the theoretical predictions related to the MPC and presents a simple simulation analysis of the effect of positive and negative income shocks on consumption in a model with income risk and liquidity constraints. Section 2 describes the data and discusses advantages and limitations of the questions used in our survey to elicit the MPC. Section 3 presents the regression results obtained when relating the MPC to demographic variables and household resources. Section 4 compares the distribution of positive and negative income shocks for the same individuals and uses information on the size of the shock to draw implications about the prevalence of liquidity constraints in the sample. Section 5 concludes.

1. Theoretical Predictions

In a standard life-cycle permanent income model with perfect credit markets, quadratic utility and an infinite horizon, consumption is proportional to lifetime disposable resources, and hence all consumers respond in the same way to transitory income shocks; that is, there is no heterogeneity in the MPC. Models with a finite horizon introduce a first important source of heterogeneity: the MPC is larger for households with short horizons (typically, older households).

Departing from quadratic utility, models in which the utility function exhibits prudence predict that the MPC depends on the level of household resources. Indeed, utility functions characterised by prudence produce a concave consumption function in which the MPC declines with the level of cash-on-hand (Carroll, 1996). The intuition is that consumers with less wealth have lower ability to protect their consumption against income shocks. Therefore, an unanticipated increase in income, by increasing cash-on-hand, has a smaller effect on consumption than a reduction in income. Adding liquidity constraints, that is, a set of constraints that prevent wealth from being negative each period, increases the sensitivity of consumption to income shocks, producing further asymmetries in the MPC.

To gauge the importance of the responses of consumption to income shocks, we simulate the MPC in a version of the life-cycle PIH with an exogenous borrowing constraint. We assume that agents solve the following problem:

$$\max \sum_{t=0}^T \beta^t \frac{c_t^{1-\gamma} - 1}{1-\gamma},$$

subject to (for all t):

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

$$a_{t+1} \geq 0.$$

$$a_T = 0.$$

Income follows the process $y_t = \exp(z_t + \varepsilon_t)$, with $z_t = \rho z_{t-1} + \eta_t$, and 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. From the budget constraint, we have that :

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

Hence, the MPC with respect to a proportional change in income can be obtained directly from the decision rule as:

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

We assume that people earn income from age 20 to age 80, at which point they die. To solve the model, we use standard parameter values. The interest rate is 4%, the discount factor 0.95, relative risk aversion equals 2, the AR(1) parameter is 0.98, the standard deviation of the persistent shock is 0.03, and the standard deviation of the transitory shock is 0.01. We then calculate the optimal consumption rule. To mimic a hypothetical income windfall equivalent to one-month of income (as in the survey questions we describe in Section 2), we normalise median annual labour income to 1, set the transitory shock to 1/12, and compute the distribution of MPCs with respect to this shock. We repeat the exercise setting the transitory shock to $-1/12$. Finally, we change the size of the shock to 3/12 and $-3/12$ to mimic a three-month income shock.

Figure 1 plots the MPC from positive (left graph) and negative (right graph) income shocks. In both cases, we fix age at 40, and plot the consumption response for different levels of cash-on-hand, defined as disposable income (labour income plus capital income) plus any financial assets minus non-mortgage debt. The MPC in response to positive income shocks range from around 11% (at low levels of cash-on-hand), to approximately 7% (for levels of cash-on-hand about three times larger than median annual labour income). The left graph in Figure 1 also shows that the MPC from a small income shock is larger than that from a large shock, and particularly so at low levels of cash-on-hand. This is because a large income shock makes it more likely that the potential liquidity constraint is no longer binding.

The MPCs in response to negative income shocks (right graph) is considerably larger for both small and large shocks, again especially at low levels of cash-on-hand. Moreover, the MPC from a large negative shock is greater than that from a small negative shock because the liquidity constraint is more likely to be binding in the former case. Finally, comparison of the two graphs in Figure 1 shows that at low levels of cash-on-hand the MPC in response to a negative shock is much larger than that in response to a positive shock, regardless of the size of the shock.

All in all, we conclude that there is substantial heterogeneity in the MPC for income changes of different signs and different size, particularly at low levels of cash-on-hand. On the other hand, at high levels of wealth the differences in MPC tend to be negligible, and to resemble the case without liquidity constraints. Indeed, for a given consumer's horizon, simulation of a model without liquidity constraints (where borrowing is only constrained by the terminal condition on wealth) reveals an essentially constant MPC (6% at age 40) across the cash-on-hand distribution, regardless of the direction and size of the shock. In other words, the absence of a period-by-period liquidity constraint leads to both a less heterogeneous MPC and a less concave consumption function.

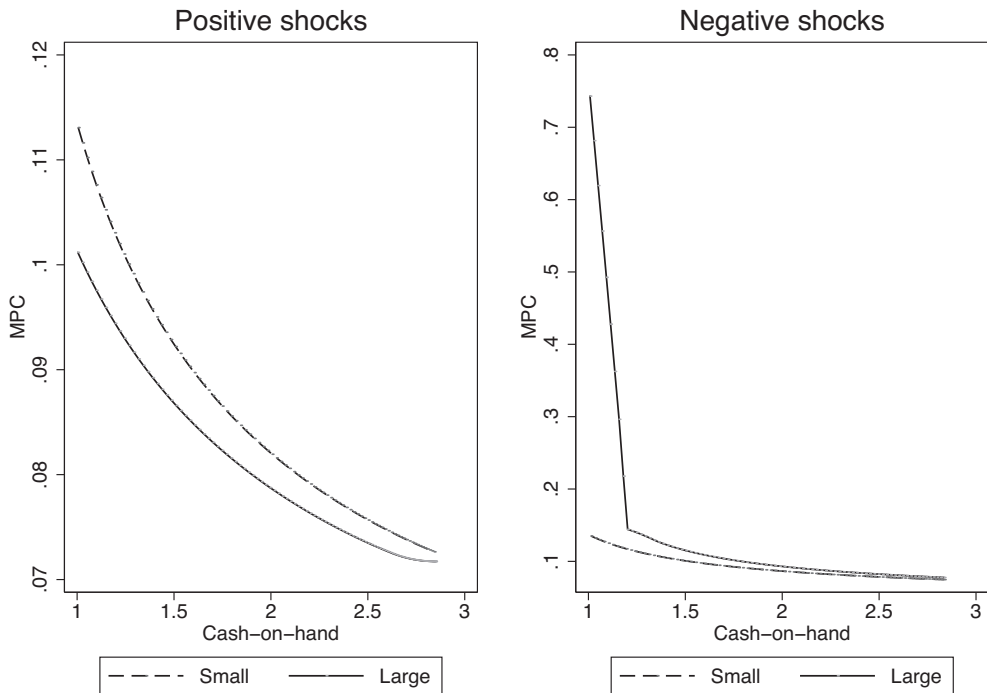


Fig. 1. *The Response of Consumption to Income Shocks.*

Notes: Both graphs plot the MPC for individuals who are 40 years old. Cash-on-hand values (normalised by median income) refer to pre-income shock levels.

The two graphs in Figure 2 plot the MPC distribution from positive and negative income shocks at different ages, for a fixed value of cash-on-hand (normalised by median labour income) equal to 1.25. For positive shocks, the MPC is 9–10% at age 30 (the youngest age) and increases steadily at older ages (to 12% at age 70). This happens because older people have less time to smooth consumption after the shock. Note also that the MPC from small shocks is larger than that from large shocks, as in the left graph of Figure 1, across the entire age distribution. The right graph of Figure 2 plots the MPC from negative income shocks at various ages. The shapes of the MPC distributions are similar to the positive income shocks, steadily increasing across the age distribution. As in Figure 1 (right graph), the MPC from large negative shocks is greater than the MPC from small negative shocks across the entire age distribution.

To check the robustness of our results, we also simulate the model with a more realistic income profile, allowing real income to grow until retirement at the constant rate of 1.5% per year, and remaining constant in real terms after age 65 at 60% of the pre-retirement income. The differences in MPC by direction, size of the shock, age and cash-on-hand are similar to the baseline case. For brevity, the results of these simulations are available on request.

To summarise, our simulations generate several predictions about MPC heterogeneity: (1) the MPC in response to a negative income shock is larger than that from a positive shock; (2) in the case of positive income shocks, the MPC from small shocks is greater than that from large shocks; (3) in the case of negative income shocks, the prediction is the opposite; (4) the MPC increases




Fig. 2. *The Response of Consumption to Income Shocks, by Age.*

Notes: The graphs plot the MPC at five different ages (30, 40, 50, 60 and 70), holding pre-income shock cash-on-hand (normalised by median income) constant at about 1.25.

with age, regardless of the sign and size of the shocks, particularly for the older households. Our specially designed set of questions allows us to test these important, and still unexplored, implications of models with precautionary saving and liquidity constraints.

Recent literature shows that the composition of household resources may also matter, introducing further sources of MPC heterogeneity that our simple model does not address. Households burdened with large debt amounts may react to a positive change in income by reducing their debt rather than spending (Mian and Sufi, 2010; Dynan, 2012). Moreover, if household wealth is locked into illiquid assets, spending may be more responsive to net liquid assets, rather than total wealth (Kaplan and Violante, 2014). In particular, households may reduce their consumption significantly even when a negative shock is transitory. The model does not distinguish between liquid and illiquid assets, but we address these issues empirically, exploring whether MPC heterogeneity depends also on the composition of households' resources, and not just cash-on-hand and consumers' horizons.

2. The Data

We use data from the CentER Internet panel, a project sponsored by the Dutch National Bank and maintained by CentERdata at Tilburg University. The baseline survey is conducted once a year via the Internet and collects detailed information on a range of demographics and asset holdings for a representative sample of Dutch-speaking households in the Netherlands. In addition to

the baseline survey, households may be asked, during the year, to participate in special-purpose surveys.

We designed a special-purpose survey that included questions aimed at measuring the MPC in response to positive and negative income changes, and to relatively small and relatively large income changes. Specifically, we characterise the MPC based on four separate questions asked to the financial respondent (i.e., the person responsible for the household's finances) in each household participating in the CentER survey.

In July 2015, we administered the first survey, which included two questions asking how people would respond to positive and negative income shocks of a relatively small size, respectively. To avoid influencing the respondents' reports, in October 2015 we administered to the same households a follow-up survey that asked how people would respond to positive and negative income shocks of a relatively larger size. To minimise framing concerns, we placed the questions referring to positive and negative changes in different parts of the survey questionnaire. Finally, to enforce the 'budget constraint' the survey questions ask respondents to allocate the hypothetical income change into its four possible uses (non-durable spending, durable spending, debt repayment, and saving). In this article we focus on the response to non-durable spending for a number of reasons. First, most of the literature focuses on the MPC with respect to non-durables (and not the marginal propensity to spend). Second, the predictions of the model with borrowing constraints discussed above are more appropriate in cases in which loans are non-collateralised (as in the case of loans to finance non-durable spending). Finally, if consumers report that they plan to allocate the extra income they receive to the purchase of durables, we ignore whether they plan to use the extra income as a down payment for a large durable good (e.g., a car), or to purchase a small durable good with no recourse to credit (e.g., jewellery). This has obvious consequences for the measurement of the marginal propensity to consume or spend. No such ambiguity exists when focusing on non-durables.

The two questions on positive income changes refer to a one-off bonus received from the government:

Imagine you unexpectedly receive a one-time bonus from the government equal to the amount of net income your household earns in (one-month / three-months). In the next 12 months, how would you use this unexpected income transfer? Distribute 100 points over these four possible uses:

1. *Save for future expenses [0, . . . ,100]*
2. *Repay debt [0, . . . ,100]*
3. *Purchase within 12 months durable goods (cars, home improvement, furniture, jewellery, other durable good) that you otherwise would not have purchased or that you would have purchased later [0, . . . ,100]*
4. *Purchase within 12 months non-durable goods and services that do not last in time (food, clothes, travel, vacation, etc.) [0, . . . ,100]*

[] Do not know

The two questions for negative changes refer to a one-off tax:

Imagine you unexpectedly have to pay a one-time tax to the government equal to the net income your household earns in (one-month / three-months). In the next 12 months, how would you react

to this unexpected reduction in your net income? Distribute 100 points over these four possible actions:

1. Reduce your saving for future expenses [0, . . . ,100]
2. Borrow more money or repay less debt [0, . . . ,100]
3. Cancel or postpone the purchase of durable goods (cars, home improvement, furniture, jewellery, other durable goods) that you otherwise would have purchased in the next 12 months [0, . . . ,100]
4. Reduce spending in the next 12 months on non-durable goods and services that do not last in time (food, clothes, travel, vacation, etc.) [0, . . . ,100]

[] Do not know

The survey is a cross-section of 1,543 households. It also requests information about demographics, household income, and wealth (broken down into real assets, financial assets and debt). Note that, in contrast to questions that elicit qualitative information ('mostly save/mostly spend') on how people spend temporary tax rebates, the responses to the questions we posed provide quantitative metrics for a proposed scenario (people are asked what percentage of the bonus they would spend, and what they would save). Similar to the 'mostly save/mostly spend' questions posed in Shapiro and Slemrod (1995, 2003), our questions refer to a bonus, or to a tax, and thus, reflect real-life situations.

The advantage of quantitative survey responses is that they overcome problems related to comparing responses across individuals who might interpret the statement 'mostly save/mostly spend' in different ways. The design of the survey questions also addresses the following potential problem: if the magnitude of the rebate is small relative to the incomes of many households, asking how the respondent would spend a fixed sum of money (e.g., a 500-euro tax rebate) may be subject to a size effect. In particular, respondents may not take optimal decisions if the change in their circumstances is trivially small. To overcome this issue, the survey question ties the amount of the transfer received to monthly income.²

It should be noticed that we compare the empirical MPC with the theoretical predictions of the model in Section 1, where people receive a *proportional* change in income. One could also be interested in how consumers respond to a *fixed* change in income as opposed to a *proportional* change in income.³ In general, it is not easy to derive the relation between the MPC in the two cases. Assuming that the fiscal expenditure is the same, and that in the fixed income case all taxpayers receive the same income change, simple simulations show that, given the concavity of the consumption function, the MPC from proportional income changes is larger than that for fixed income changes at the bottom of the distribution and *vice versa* at the top (and of course the two MPCs are equal when the consumption function is linear).

² Parker *et al.* (2013) and Sahn *et al.* (2010) try to tease out the 'size effect' by looking at rebates relative to income. Here, we adopt the strategy of asking different questions for one-month and three-month income changes.

³ In some papers, researchers look at how people spend a fixed amount of money, as in the 2001 Bush tax rebate case, where (almost) all taxpayers received \$600 if married filing jointly (\$300 if single). See Johnson *et al.* (2006), and Fuster *et al.* (2018). But in other cases, researchers look at proportional income changes. For example, the 2010 Obama Temporary Employee Payroll Tax Cut was, effectively, an experiment similar to ours. This tax stimulus reduced the employee portion of the Social Security payroll tax from 6.2% to 4.2% on individual earnings up to the taxable maximum of \$110,100. This meant that most taxpayers received a positive, transitory income shock corresponding to 2% of their annual income (see Sahn *et al.*, 2015).

Finally, the survey allows us to characterise the MPC for positive and negative income shocks for the same household. Quasi-experimental data or retrospective data identify households who have experienced positive or negative shocks. This makes it difficult to compare the resulting two MPC distributions because the two samples are likely to represent segments of the population differing in terms of resources, socioeconomic characteristics and preferences. Thus, by asking hypothetical questions referring to both income increases and decreases to each respondent, the analysis in this article does not suffer from this problem.

Several features of the survey questions are noteworthy. First, the questions ask about consumption of non-durables and durables separately (questions on the latter mention cars, home improvements, furniture and jewellery). This allows us to identify the MPC without contamination arising from allocating some of the hypothetical income change to durables.⁴ This distinction might be especially relevant for the ‘three-month income changes’ questions, as a bonus equivalent to three-months’ income might allow the household to purchase more expensive durable goods (or to meet a down-payment constraint for purchasing an expensive durable good), while a tax equivalent to three months’ income might make it more likely that the household reduces or postpones planned expenditure on durable goods.

Second, consumers are asked by how much they would increase or cut spending ‘in the next 12 months’. This allows us to rule out that differences in the MPC arise from differences in the timing of planned spending. Each of the reported MPCs can be interpreted as the consumption response to an income change in the coming year. Of course, further adjustments in subsequent years cannot be ruled out. In principle, it would be useful to posit similar questions with other time horizons (e.g., how would consumption change in the second year after the shock) but this would increase the complexity of the questionnaire considerably.

Third, the questionnaire was administered in July and October 2015. In 2015, real GDP growth in the Netherlands was 2% and it was projected to grow by 1.7% in 2016 and 2% in 2017. In other words, the interviews took place several years after the financial crisis (GDP decreased by 4% in 2009) and the 2011–12 recession. Although business cycle effects can never be ruled out, the period in which the survey was administered should have weakened their impact.

Finally, responses to our questions might be affected by households’ financial sophistication. Thus, we check whether financial literacy affects our results and whether any part of the heterogeneity in MPCs can be attributed to lack of financial sophistication.

3. MPC Distributions

In this section, we report descriptive statistics of the distribution of responses to hypothetical income changes. We summarise the empirical correlations by employing regression analysis to examine how the MPC varies with certain household characteristics.

3.1. Descriptive Analysis

Table 1 reports summary statistics of the responses to the survey questions. It should be remembered, when evaluating responses, that the size of the income change is household-specific and

⁴ Parker *et al.* (2013) highlight the importance of distinguishing between non-durable and total spending and find that households spent between 12% and 30% of their 2008 U.S. stimulus payments on non-durable goods, and this rose to 50–90% when durable goods were included. This result is somewhat puzzling in light of a previous study which found that most spending goes to non-durables (Johnson *et al.*, 2006).

Table 1. *Summary Statistics.*

	Mean	Median	Standard deviation	Number of observations
One-month income change				
Income increase				
Increase non-durable consumption	19.59	10	23.01	1,319
Increase durable expenditures	19.24	10	22.87	1,319
Reduce debt	14.71	0	27.33	1,319
Increase saving	46.45	50	34.64	1,319
Income decline				
Reduce non-durable consumption	23.75	20	23.93	1,268
Reduce durable expenditures	25.76	20	24.71	1,268
Increase debt	6.98	0	17.61	1,268
Reduce saving	43.51	40	33.98	1,268
Three-month income change				
Income increase				
Increase non-durable consumption	14.34	10	16.28	1,484
Increase durable expenditures	22.28	20	22.81	1,484
Reduce debt	16.24	0	26.54	1,484
Increase saving	46.97	50	30.52	1,484
Income decline				
Reduce non-durable consumption	23.97	20	23.57	1,358
Reduce durable expenditures	26.99	25	25.02	1,358
Increase debt	7.30	0	18.94	1,358
Reduce saving	41.74	40	33.43	1,358

Notes: Mean and median refer to the percentage use of the income change.

that the average net monthly household income is 2,833 euro. Following a one-month income increase, the average respondent would allocate 19.6% of the additional income to non-durable consumption, 19.2% to durable consumption, 14.7% to debt reduction and save the remaining 46.5%. The distribution for a one-month income decline indicates a stronger consumption response: 23.8% of the income drop is absorbed by non-durable consumption, 25.8% by durables, 7% by a debt increase and 43.5% by reduced saving. Focusing on the MPC for non-durables, the median MPC from positive income changes is 10%, while it is 20% for negative changes. This pattern provides qualitative support for the insights from the simulations in Section 1, suggesting that the MPC in response to negative income shocks is higher than the MPC in response to positive shocks.

The MPC distributions for larger income changes highlight some interesting features: an assumed three-month rise in income is associated with an MPC on non-durables of 14.3% while the MPC associated with an equally sized income decline is 24%. Therefore, the MPC gap between positive and negative income changes is wider for large changes, again supporting the insights from the model.

The magnitudes of the average MPCs on non-durables are broadly consistent with the estimated average consumption response out of transitory income shocks reported in Johnson *et al.* (2009), Parker *et al.* (2013) and, for the case of the Netherlands, Carroll *et al.* (2014). Note however, that comparison with previous studies is not always straightforward because papers use different definitions of MPC. Some focus on proportional income changes (as in our case), some focus on absolute income changes, and others report elasticities.




Fig. 3. MPC Distribution.

Notes: Figure 3 plots the cross-sectional distribution of the MPC on non-durables due to a one-month income increase (upper-left panel), a one-month income decline (upper-right panel), a three-month income increase (lower-left panel), and a three-month income decline (lower-right panel).

An average MPC on non-durables of 19.6% associated with a one-month income increase is higher than implied by a standard model of intertemporal choice. However, the average hides substantial heterogeneity among the responses and the median (10%) is more in line with the predictions of models where households smooth a large fraction of the shock.

Figure 3 plots the cross-sectional distribution of the MPC due to a one-month income increase (upper-left panel), a one-month income decline (upper-right panel), a three-month income increase (lower-left), and a three-month income decline (lower-right). The upper-left histogram in Figure 3 shows that 36% of respondents reported that they would not consume any of the bonus, and another 15% said they would consume 10% or less. Only 3% reported that they would consume more than 90% of the bonus, and only 2.8% said they would consume the entire bonus (MPC = 1). The histogram also shows a ‘heaping’ at rounded values (5%, 10%, etc). It is interesting that heaping is not concentrated in the ‘50%’ response, which often is interpreted as indicating respondent indecisiveness. We take this as an indication that the responses to the MPC questions are reliable.

The upper-right panel in Figure 3 reports the MPC distribution for one-month negative income changes. We noted that the average MPC corresponding to negative changes is higher (23.8%) than the average MPC corresponding to positive changes (19.6%). This higher average is due to a lower fraction of respondents reporting a low MPC (42% report that they would cut consumption by 10% of the income drop or less), and a higher fraction of households reporting that they would

cut consumption substantially (4% reported they would cut consumption by more than 90% of the income drop, and 3.4% reported an MPC equal to 1).

The lower two histograms in Figure 3 report similar distributions for larger income changes. The MPC distribution corresponding to a three-month negative income change is similar to the one-month change distribution. This result could be due to the fact that liquidity constrained households might exhibit very high MPCs irrespective of the size of the negative shock. Moreover, non-liquidity constrained households' MPC may not be sensitive to the size of the negative shock either, as suggested by the model simulation results discussed in Section 1.

In the case of positive income changes, a feature of the histogram worth noting is that only 1% of the sample reported an MPC from a three-month income increase of over 50%. In contrast, 7% of respondents reported an MPC from a one-month income change above 50%.

3.2. *Cash-on-hand and Age Profiles*

We next relate the MPC to household resources, which we measure empirically using cash-on-hand, defined as the sum of current income and financial wealth, net of consumer debt. This definition excludes real assets from cash-on-hand because it is unlikely that households will sell real assets when facing income shocks of the magnitude considered in our article. An additional reason for adopting this definition is that households likely face high transaction costs from selling illiquid assets in the time horizon described in our survey questions (one year). Therefore, low cash-on-hand households in our data include both those who have low net worth and those who are wealthy 'hand-to-mouth', as described by Kaplan and Violante (2014). Households of the latter type might have a large amount of illiquid assets (such as a house) but low levels of (liquid) financial assets.

Figure 4 plots the average MPC on non-durables by quartiles of cash-on-hand. There is no clear relation between the MPC and cash-on-hand for positive income changes, regardless of the size of the shock (one or three-month of income). In contrast, and consistently with theoretical predictions, the MPC in response to income declines is higher at low levels of cash-on-hand, for both one- and three-month income changes.

Figure 5 plots the MPC on non-durables against age (grouped in 10-year intervals). Theory predicts a positive relation between age and the MPC (as older people have a shorter horizon over which to smooth a transitory income change), and indeed we find that in all four graphs the relationships are upward sloping. For instance, the MPC in response to small positive income changes increases from 14% for the youngest age group (less than 30 years old) to 23% for the oldest group (over 80 years old). The MPC in response to a one-month income decline increases only for the oldest group. The age-MPC relation for three-month positive and negative changes is also upward sloping, as shown in the lower two graphs in Figure 5.

3.3. *Regression Analysis*

To properly characterise the various factors affecting the variability of the MPC, we rely on regression analysis. Summary statistics of the main variables used in the estimation are presented in Table 2. Table 3 presents the baseline OLS regression results for MPC on non-durables for each of the four scenarios: small negative and positive shocks (columns 1 and 2) and large negative and positive shocks (columns 3 and 4).

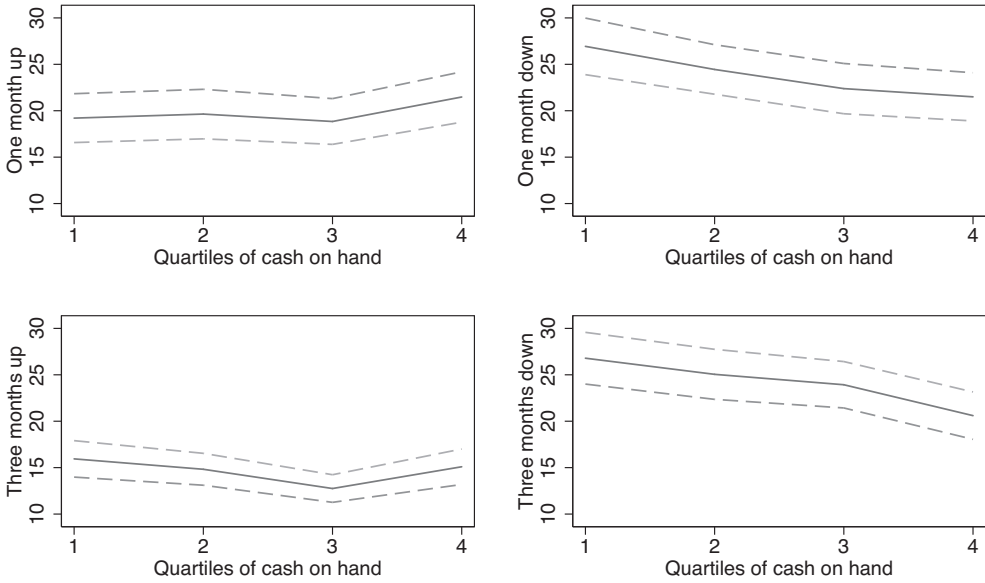


Fig. 4. MPC Distribution, by Cash-on-hand reports larger response and Quartiles.
 Notes: Dashed lines denote the lower and upper bounds of 95% confidence intervals.




Fig. 5. MPC Distribution, by Age.
 Notes: Dashed lines denote the lower and upper bounds of 95% confidence intervals.

Table 2. *Summary Statistics of Variables Used in the Empirical Analysis.*

Variable	Mean	Standard deviation	Number of observations
Age	56.75	14.57	1,543
Male	0.56	0.50	1,543
Family size	2.30	1.19	1,543
Cash-on-hand (median)	18,550	107,757	1,385
College degree	0.40	0.49	1,543
High school degree	0.33	0.47	1,543
Unemployed	0.03	0.18	1,474
Financial literacy	2.31	0.90	1,434

Table 3. *Regressions for the MPC on Non-Durable Consumption.*

Variables	(1) One-month down	(2) One-month up	(3) Three-months down	(4) Three-months up
35 ≤ Age < 50	0.335 (2.824)	4.571 (2.669)*	2.146 (2.640)	-0.688 (1.791)
50 ≤ Age < 65	-0.180 (2.756)	6.214 (2.601)**	3.682 (2.583)	0.584 (1.750)
Age ≥ 65	3.095 (2.838)	10.447 (2.694)***	5.499 (2.673)**	3.581 (1.816)**
Male	-1.369 (1.446)	-2.895 (1.371)**	-4.155 (1.384)***	-2.871 (0.929)***
Family size	0.244 (0.680)	-0.117 (0.631)	0.424 (0.627)	-0.446 (0.422)
II cash-on-hand quartile	-2.785 (2.054)	-0.271 (1.941)	-2.034 (1.920)	-1.267 (1.270)
III cash-on-hand quartile	-5.056 (2.048)**	-1.386 (1.924)	-3.167 (1.943)	-3.670 (1.285)***
IV cash-on-hand quartile	-5.932 (2.048)***	0.796 (1.947)	-6.487 (1.967)***	-1.383 (1.313)
Constant	26.493 (3.118)***	15.373 (2.951)***	24.932 (2.916)***	17.800 (1.962)***
R ²	0.01	0.02	0.02	0.02
N	1,160	1,208	1,230	1,332

Notes: We report standard errors in parentheses. *, **, *** indicate significance level at 10%, 5% and 1%, respectively.

Our baseline specification includes age dummies (the base category is the youngest age group), financial respondent's gender, family size, and dummies for cash-on-hand quartiles (the base category is the first quartile). The number of observations is not the same in each of the regressions, due to the different number of missing values in the responses to the four questions related to non-durable consumption.

The age coefficients in Table 3 are generally positive and statistically different from zero except for those in column 1, indicating that the youngest group (less than 35 years old, the base category) has a lower MPC than the oldest group (65 and over). In particular, for the specifications reflecting three-month income changes (both positive and negative), only those aged 65 and above exhibit significantly higher MPCs. For the specification reflecting a one-month positive change also the groups aged 35–49 and 50–64 exhibit significantly higher MPCs than the youngest group. For the specification reflecting a one-month negative change the age dummies are not statistically different from zero. All in all, these results are qualitatively congruent with the predictions of standard consumption models that the MPC in response to transitory shocks increases with age and with the unconditional relationships commented above.

As for the relationship of cash-on-hand with the MPC, we find a negative association for negative income changes. In particular, in column 1 the coefficient of the fourth quartile of cash-on-hand is -5.9% and is statistically different from zero at the 1% level (-6.5% in column 3 for the three-month income drop). On the other hand, there is no relationship between the MPC and cash-on-hand for positive changes. Thus, the pattern of the coefficients of the cash-on-hand quartile dummies confirms the descriptive analysis discussed above.⁵

To determine whether our results are sensitive to the omission of important variables and understanding of the survey questions, we perform various robustness checks. In the Online Appendix we report results including a richer set of education controls, distinguishing between income, real wealth, financial wealth and debt, and focusing on a sample that includes only younger households. We also control for knowledge of basic economic concepts and examine whether our estimates are robust to censoring or heaping of the MPC responses. We find that our results remain unchanged.

4. MPC Heterogeneity at the Individual Level

As discussed in Section 1, liquidity constraints are likely to have important effects on the MPC distribution. In particular, the distribution of negative income shocks is expected to stochastically dominate the MPC distribution of positive shocks, for both small and large shocks, as suggested by the simulation results of the theoretical model shown in Figure 1. Furthermore, in the case of positive income shocks, the MPC from relatively small (one-month) shocks should be greater than the MPC from relatively large (three-month) shocks. For negative shocks, we expect the opposite. Note that similar predictions would emerge in buffer stock models without liquidity constraints but with a positive probability of zero income in each period.

4.1. Means of MPC Differences

In this section, we provide direct evidence supporting some of these theoretical predictions. In the first three rows of Table 4 we test that household-level differences of the MPCs corresponding to different income change scenarios are statistically different. Since each household reports both MPCs, by taking differences of household MPCs we effectively eliminate the influence of household fixed unobservable effects (such as preferences or financial sophistication) which might affect both distributions. In column 1 the mean difference between the MPC due to a one-month negative income shock and its positive income shock counterpart is 4%; the test of equal means has a p -value well below 1%. The corresponding difference for three-month income shocks is even larger, 9.8%, and again statistically significant at the 1% level. We also find that the one-month positive income shock MPC is larger than its three-month counterpart by 5.6% (with the difference being again statistically significant at the 1% level). Finally, we find no difference between the MPCs due to one- and three-month negative income shocks, contrary to the predictions of the simulations.

In the other rows of Table 4 we move from unconditional to conditional means and try to understand what drives the household specific differences in MPCs across the various income change scenarios. In column 1 we focus on the determinants of the difference between small positive and small negative changes. The constant is large and statistically significant (as in the

⁵ Replacing cash-on-hand quartile dummies with income quartile dummies does not change the pattern of results, i.e., we still find a negative association between MPC and negative income changes, and no association with positive changes.

Table 4. *Test of Means and Regressions for MPC Differences at the Household Level.*

Variables	(1)	(2)	(3)	(4)
	One-month down minus one-month up	Three-months down minus three-months up	One-month down minus three-months down	One-month up minus three-months up
Test of means				
Mean	4.04	9.764	-0.367	5.552
Standard error	(0.77)	(0.712)	(0.870)	(0.648)
<i>p</i> -value test that the mean of MPC differences is 0	<0.01	<0.01	0.67	<0.01
Regression analysis				
35 ≤ Age < 50	-4.018 (3.211)	2.749 (2.945)	-1.799 (3.586)	5.587 (2.736)**
50 ≤ Age < 65	-6.319 (3.135)**	3.378 (2.883)	-3.566 (3.508)	5.912 (2.669)**
Age ≥ 65	-7.722 (3.234)**	1.913 (2.986)	-3.530 (3.615)	7.271 (2.761)***
Male	1.543 (1.642)	-1.499 (1.544)	3.502 (1.856)*	0.285 (1.408)
Family size	0.308 (0.769)	0.881 (0.698)	0.169 (0.867)	0.321 (0.644)
II cash-on-hand quartile	-1.791 (2.333)	-0.284 (2.143)	1.908 (2.663)	0.304 (1.985)
III cash-on-hand quartile	-3.812 (2.323)	0.990 (2.168)	0.190 (2.661)	1.328 (1.972)
IV cash-on-hand quartile	-6.009 (2.327)***	-5.029 (2.202)**	2.829 (2.663)	2.027 (2.001)
Constant	10.870 (3.544)***	6.952 (3.250)**	-1.401 (3.984)	-2.154 (3.018)
<i>R</i> ²	0.02	0.01	0.01	0.01
<i>N</i>	1,142	1,216	1,085	1,182

Notes: The first two rows of the table report the mean of the distribution of individual-specific MPC differences for income shocks of different size and sign, their standard errors, and the significance level of a test that the mean of the MPC differences is zero. The other rows of the table report regression coefficients using the MPC differences as the dependent variable. We report standard errors in parenthesis. *, **, *** indicate significance level at 10%, 5%, and 1%, respectively.

unconditional case). The effect decreases with age. Interestingly, the coefficient of the highest cash-on-hand quartiles is negative, indicating that the MPC in response to negative income changes has a stronger negative association with cash-on-hand than the MPC in response to positive changes, and thus tends to be larger at low levels of economic resources, a result in line with the simulations shown in Figure 1. In column 2 we focus on the difference of MPCs with respect to large positive and large negative income changes. The results are qualitatively similar to those we find for small income changes.

In columns 3 and 4 of Table 4 we focus on size, rather than sign changes. In column 3 we look at the difference of MPCs with respect to small versus large negative income changes; in column 4 we look at the difference of MPCs with respect to small versus large positive income changes. In the case of negative shocks (column 3) all the relevant effects appear imprecisely estimated. On the other hand, for the case of positive shocks (column 4), the MPC out of a small shock becomes increasingly larger than that out of a large shock as age increases.

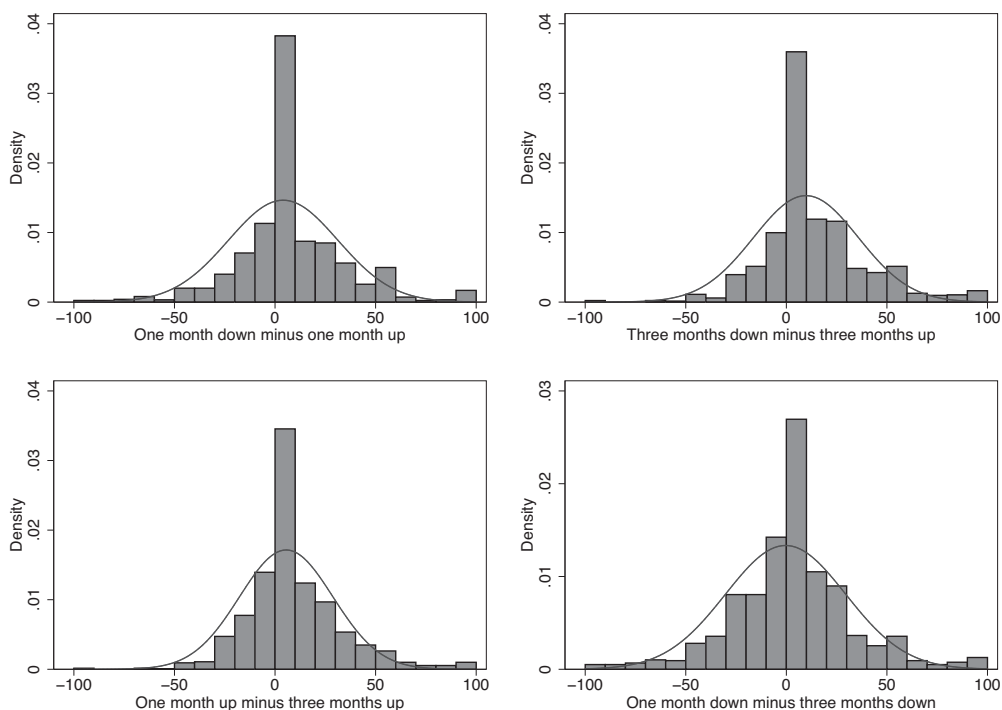


Fig. 6. *Plots of the Differences in MPC at the Individual Level.*

Notes: Figure 6 plots the cross-sectional distribution of differences in the MPC for changes in sign (upper panels) and size (lower panels) of the income shocks. The lines denote fitted normal distributions.

We also compare MPC distributions using quantile-quantile (QQ) plots and find that the results of these comparisons are consistent with those discussed above. Due to space constraints, we discuss these findings more extensively in the Online Appendix.

4.2. *Distribution of MPC Differences*

In this section we provide further evidence on the distribution of the household-level difference in the MPCs by plotting (in Figure 6) the histograms of the differences and superimposing a normal curve for comparison. Consider the first distribution for the difference between one-month negative income shock and one-month positive income shock (the upper-left graph in Figure 6). It is apparent that the distribution has a right tail that is thicker than the left one (implying a positive overall mean, as in Table 4). But it is also clear that there is considerable heterogeneity across the distribution. About 41.3% of individuals report differences close to zero (within plus or minus 5%), which is congruent with standard intertemporal models (such as the permanent income model without liquidity constraints), in which the MPC is symmetric and the difference between positive and negative income shocks equals zero. Another 33.4% of households reports MPC differences higher than 5%, a pattern that could be rationalised by models with liquidity constraints, or by other models that generate buffer stock behaviour. A third group of households (25.3%) reports a larger response for positive income shocks than for negative shocks and exhibits

MPCs that are not easy to reconcile with models with rational agents. One possible explanation for this behaviour could be loss aversion, which can induce households to adjust their behaviour more strongly to positive rather than negative shocks (Bowman *et al.*, 1999).

The histogram for the difference between three-month negative shocks and three-month positive shocks (upper-right in Figure 6 and second column in Table 4) exhibit a similar pattern: for 38.5% of the sample the difference is close to zero, for 43.1% it exceeds 5%, but for 18.4% it is negative and inconsistent with standard models, which could be due, as noted in the one-month case above, to loss aversion.

As already pointed out, when comparing responses to one and three-month negative shocks we expect the left tail of the distribution to be greater than the right tail, but the distribution is in practice symmetric, with about one-third reporting differences in MPCs close to zero, one-third reporting substantial negative MPC differences and one-third reporting differences exceeding 5%. Overall, however, the mean difference is not statistically different from zero.

Finally, in the distribution for differences between one-month and three-months positive shocks, 39.2% of the sample reports values around zero (within a tolerance limit of 5%), 37% positive values exceeding 5%, and 23.8% substantial negative values. Such negative values could be due to households being more responsive in situations in which the stakes are large (the ‘big peanuts’ phenomenon: see van den Assem *et al.*, 2012).

Altogether, our results indicate that about 38% of our sample exhibits behaviour that is consistent with the existence of liquidity constraints or with buffer stock saving behaviour. Another 40% of the sample has MPCs that appear independent of the size of the shock, which is consistent with a standard permanent income model. Finally, about 22% of the sample reacts more to positive than to negative shocks, which could be explained by behavioural theories as outlined above.

An additional explanation for these seemingly inconsistent responses is lack of financial sophistication. To check this hypothesis, we define a dummy variable equal to one for those who report inconsistent MPC values in the various scenarios, where we vary the size and sign of the shock. We then regress the dummy on education, financial sophistication and other demographic controls. We find that higher education and financial sophistication are associated with a lower probability of reporting inconsistent MPCs (results are available upon request).

5. Conclusions

We use a representative survey of the Dutch population to characterise empirically the distribution of the MPC in response to unexpected transitory income changes (positive and negative; small and large) and check several predictions of intertemporal consumption models. We find that the consumption response to income shocks declines with economic resources, and that the MPC is smaller if consumers have relatively long horizons. Most importantly, we detect significant asymmetries between the MPC in response to positive and negative income shocks. The main advantage of the survey questions is that they allow us to compare the responses to a hypothetical positive and negative income shock for the *same* household. In contrast, in real-life studies the income shock is either positive or negative. Thus, results obtained comparing the consumption responses of those facing positive shocks with the consumption responses of those facing negative shocks may confound genuine MPC heterogeneity with the heterogeneity of households that are subject to different types of shock.

Our results are broadly in line with models of intertemporal choice with precautionary saving, borrowing constraints and finite horizons. The average MPC corresponding to non-durable consumption is in the 15–25% range, it increases with age, and it is larger at low levels of economic resources. We also find that the MPC distribution is in line with two important predictions of models with liquidity constraints. The empirical estimates confirm the results from a simple simulation analysis of a model with income risk and precautionary saving showing that in the presence of liquidity constraints the MPC in response to a negative income shock is larger than the MPC in response to a positive shock. In addition, in the presence of liquidity constraints the size of the shock also matters, especially at low levels of economic resources. For large increases in income, consumers are more likely to overcome the constraint (and therefore, the MPC is lower than for small increases).

The ability to look at within-household MPC differences also allows us to classify households in relation to the predictions of models of consumption behaviour. We find that about one-third of respondents provide symmetric MPCs in response to shocks of different size and direction (as in the standard PIH); about 40% display asymmetric responses consistent with models with liquidity constraints; and about 25% reports asymmetric MPCs that appear inconsistent with standard models and could be attributed to behavioural decision-making or lack of financial sophistication.

Our findings have important implications for predicting consumption responses to a broad set of policy interventions that may change household incomes. Such interventions could range from changes in policy interest rates to direct government money transfers, tax reforms and other redistributive policies. Our results suggest that the outcome of a temporary policy intervention depends on the distribution of household resources, and that negative shocks are likely to have more pronounced effects on consumption than positive ones.

University of Naples Federico II, CSEF, CFS, CEPAR and Netspar
European Central Bank and CFS
University of Naples Federico II, CSEF and CEPR
Stanford University, SIEPR, NBER and CEPR
De Nederlandsche Bank and Netspar

Additional Supporting Information may be found in the online version of this article:

Online Appendix Replication Package

References

- Agarwal, S., Liu, C. and Souleles, N.S. (2007). 'The reaction of consumer spending and debt to tax rebates-evidence from consumer credit data', *Journal of Political Economy*, vol. 115(1), pp. 986–1019.
- Blundell, R., Pistaferri, L. and Preston, I. (2008). 'Consumption inequality and partial insurance', *American Economic Review*, vol. 98(5), pp. 1887–921.
- Bracha, A. and Cooper, D. (2014). 'Asymmetric responses to income changes: The payroll tax increase versus tax refund in 2013', *Economics Letters*, vol. 24(3), pp. 534–8.
- Bowman, D., Minehart, D. and Rabin, M. (1999). 'Loss aversion in a consumption-savings model', *Journal of Economic Behavior & Organization*, vol. 38(2), pp. 155–78.
- Bunn, P., Le Roux, J., Reinold, K. and Surico, P. (2018). 'The consumption response to positive and negative income shocks', *Journal of Monetary Economics*, vol. 96(June 2018), pp. 1–15.

- Carroll, C.D. (1996). 'A theory of the consumption function, with and without liquidity constraints', *Journal of Economic Perspectives*, vol. 15(3), pp. 23–46.
- Carroll, C.D., Slacalek, J. and Tokuoka, K. (2014). 'The distribution of wealth and the MPC: Implications of new European data', *American Economic Review Papers and Proceedings*, vol. 104(5), pp. 107–11.
- Dynan, K. (2012). 'Is a household debt overhang holding back consumption?', *Brookings Papers on Economic Activity*, vol. 43(1), (Spring 2012), pp. 299–362.
- Fuster, A., Kaplan, G. and Zafar, B. (2018). 'What would you do with \$500? Spending responses to gains, losses, news and loans', Working Paper No. 24386, National Bureau of Economic Research.
- Jappelli, T. and Pistaferri, L. (2011). 'The consumption response to income changes', *Annual Review of Economics*, vol. 2(1), pp. 479–506.
- Jappelli, T. and Pistaferri, L. (2014). 'Fiscal policy and MPC heterogeneity', *American Economic Journal: Macroeconomics*, vol. 6(4), pp. 107–36.
- Johnson, D.S., Parker, J.A. and Souleles, N.S. (2006). 'Household expenditure and the income tax rebates of 2001', *American Economic Review*, vol. 96(5), pp. 1589–610.
- Johnson, D.S., Parker, J.A. and Souleles, N.S. (2009). 'The response of consumer spending to rebates during an expansion: evidence from the 2003 child tax credit', Working Paper, The Wharton School.
- Kaplan, G. and Violante, G.L. (2014). 'A model of the consumption response to fiscal stimulus payments', *Econometrica*, vol. 82(4), pp. 1199–239.
- Mian, A.R. and Sufi, A. (2010). 'Household leverage and the recession of 2007 to 2009', *IMF Economic Review*, vol. 58(1), pp. 74–117.
- Misra, K. and Surico, P. (2014). 'Consumption, income changes and heterogeneity: evidence from two fiscal stimulus programmes', *American Economic Journal: Macroeconomics*, vol. 6(4), pp. 84–106.
- Parker, J.A. and Souleles, N.S. (2017). 'Reported preference versus revealed preference: Evidence from the propensity to spend tax rebates', Working Paper no. 23920, National Bureau of Economic Research.
- Parker, J.A., Souleles, N.S., Johnson, D.S. and McClelland, R. (2013). 'Consumer spending and the economic stimulus payments of 2008', *American Economic Review*, vol. 103(6), pp. 2530–53.
- Sahm, C.R., Shapiro, M.D. and Slemrod, J. (2010). 'Household response to the 2008 tax rebates: Survey evidence and aggregate implications', in (J.R. Brown, ed.), *Tax Policy and the Economy*, vol. 24, pp. 69–110, Chicago: University of Chicago Press.
- Sahm, C.R., Shapiro, M.D. and Slemrod, J. (2015). 'Balance-sheet households and fiscal stimulus: Lessons from the payroll tax cut and its expiration', Working Paper No. 21220, National Bureau of Economic Research.
- Shapiro, M.D. and Slemrod, J. (1995). 'Consumer response to the timing of income: evidence from a change in tax withholding', *American Economic Review*, vol. 85(1), pp. 274–83.
- Shapiro, M.D. and Slemrod, J. (2003). 'Did the 2001 tax rebate stimulate spending? Evidence from taxpayer surveys', in (J. M. Poterba, ed.), *Tax Policy and the Economy*, vol. 17, pp. 83–109, Cambridge, MA: MIT Press.
- Smith, A., Bernheim, B.D., Camerer, C.F. and Rangel, A. (2014). 'Neural activity reveals preferences without choices', *American Economic Journal: Microeconomics*, vol. 6(2), pp. 1–36.
- Souleles, N.S. (1999). 'The response of household consumption to income tax refunds', *American Economic Review*, vol. 89(4), pp. 947–58.
- van den Assem, M.J., van Dolder, D. and Thaler, R.H. (2012). 'Split or steal? Cooperative behavior when the stakes are large', *Management Science*, vol. 58(1), pp. 2–20.