



The consumption response to positive and negative income shocks[☆]



Philip Bunn^a, Jeanne Le Roux^a, Kate Reinold^{a,*}, Paolo Surico^b

^a Bank of England, UK

^b London Business School, CEPR, UK

ARTICLE INFO

Article history:

Received 11 January 2017

Revised 14 November 2017

Accepted 17 November 2017

Available online 21 November 2017

JEL codes:

E21

E52

D12

Keywords:

MPC asymmetry

Household balance sheet

Debt

Heterogeneity

Transmission mechanism

ABSTRACT

A set of newly-added questions in the 2011–2014 Bank of England/NMG Consulting Survey reveals that British households tend to change their consumption by significantly more in reaction to temporary and unanticipated *falls* in income than to *rises* of the same size. Household balance sheet characteristics such as high debt-to-income ratios and small liquidity buffers, concerns about credit market access and higher subjective risk of lower future income account for a sizable share of this spending asymmetry. Our findings have important implications for predicting the response of aggregate consumption to expansionary and contractionary macroeconomic policies.

© 2017 The Bank of England. Published by Elsevier B.V. All rights reserved.

1. Introduction

The marginal propensity to consume (MPC) out of temporary income shocks plays an important role in the transmission of monetary and fiscal policy. According to the permanent income hypothesis (PIH), households only respond to shocks that alter their lifetime resources and therefore a non-permanent change in disposable income should only trigger a small change in current consumption. Yet, there is abundant evidence of significant and large estimates for the MPC, with liquidity constraints, income risk and precautionary motives being prominent drivers of violation of the PIH.

Despite a large empirical literature on the spending response to temporary income windfalls, there are far fewer works on what happens to household expenditure following a temporary fall in income. In the light of the more scant evidence on the latter, it is then unsurprising that there are even fewer contributions (if any) which are able to look – within the same dataset – at the actual consumption response to positive and negative income changes. This limitation is potentially

[☆] We are grateful to Urban Jermann and one anonymous referee for excellent comments and suggestions. The views in this paper are those of the authors and do not necessarily reflect the views of the Bank of England, the Monetary Policy Committee, the Financial Policy Committee or the PRA. Surico gratefully acknowledges financial support from the [European Research Council](http://www.eur-research-council.org/) (Consolidator Grant Agreement 647049).

* Corresponding author.

E-mail addresses: philip.bunn@bankofengland.co.uk (P. Bunn), jeanne.leroux@bankofengland.co.uk (J. Le Roux), kate.reinold@bankofengland.co.uk (K. Reinold), psurico@london.edu (P. Surico).

serious because it poses the external validity concern that the estimates obtained for income changes of a certain sign may not provide reliable inference for the consumption response to income changes of the opposite sign.

In this paper, we exploit a set of newly-designed questions in the Bank of England/NMG survey of British households to provide an unprecedented evaluation of the spending response to positive and negative income changes on actual data. Households were asked to report not only the sign *and* the size of any unanticipated change in income over the previous year but also the sign *and* the size of any associated change in spending. The answers to these questions allow us to compute household-specific MPCs for positive and negative income shocks. In addition, information on demographics and balance sheet positions allows us to relate these characteristics to the individual MPCs at the household-level, both within and across income changes of opposite sign. As with all survey approaches to estimating MPCs, the reliability of the inference one can draw depends on the ability of households to respond accurately to questions about changes in their income and consumption. Accordingly, we subject our findings to an extensive range of sensitivity analyses to ameliorate concerns on any possible recall bias, to shock identification, to different treatments of non-responses and to using a more qualitative indicator of the spending response.

Our main results can be summarised as follows. First, the average MPC out of a negative income shock is between 0.46 and 0.68 across alternative cuts of the data but the range is only 0.07–0.17 following a positive income shock. Second, the MPCs of borrowers are systematically larger than those for savers, with the asymmetry between the responses to positive and negative income changes being less pronounced among the latter. Third, households with little or no savings buffer, facing some form of liquidity constraints and more pessimistic about their income prospects tend to exhibit a significantly larger MPC, especially after a fall in income. Fourth, households' balance sheet characteristics, and in particular low liquid wealth relative to income and high debt-to-income ratios, are strong predictors of a larger spending response to a negative shock of a certain size than to a positive shock of the same magnitude. Fifth, among households who received a negative income shock, the MPC was significantly larger for those reporting to be more concerned about their future income, consistent with a precautionary motive linked to income risk. Our inference is robust to using a propensity score method to control for differences in the ex-ante probability of receiving a shock of the opposite sign based on observables, to a number of estimation strategies handling some forms of nonlinearity and measurement errors, as well as to using the responses to a set of questions that asked households about their consumption behaviour following a *hypothetical* income shock.

Most of the household balance sheet measures in the Bank of England/NMG survey refer to either past or present conditions. But household behaviour may also reflect expectations about future access to credit markets. To evaluate this hypothesis, we consider the consumption response to a temporary income shock in a partial equilibrium model with occasionally binding borrowing constraints. Despite its simplicity, the model is able to replicate well the consumption asymmetry observed across positive and negative changes as well as across households with different debt positions, thereby providing further evidence for a significant role played by household balance sheets in the transmission mechanism.

As for policy implications, the MPC asymmetry documented in this paper implies that monetary and fiscal policies aimed at stimulating the economy may exert a smaller impact on consumption than contractionary interventions of the same magnitude. While this is reminiscent of the “pushing on a string” conjecture often attributed to John Maynard Keynes, it also suggests that expansionary policy measures need to be bolder than contractionary measures in order to generate a change of the same magnitude in consumer spending.

Related literature. The theoretical literature identifies several mechanisms that can motivate an asymmetric response to positive and negative transitory income shocks, including – among others – imperfect access to credit markets for those with low holdings of assets or sizable, but illiquid, assets (Deaton, 1991; Kaplan et al., 2016; Hedlund et al., 2016), precautionary savings (Caballero, 1990; Carroll, 1992, 1994, 2001) and loss aversion (Kahnemann and Tversky, 1979). The empirical literature on the consumption responses to unanticipated income changes can be broadly divided into three groups.¹ The first group exploits abrupt policy changes as quasi-natural experiments, including income tax rebates (Johnson et al., 2006; Parker et al., 2013; Agarwal and Qian, 2014; Misra and Surico, 2014), minimum wage hikes (Aaronson et al., 2012) and credit card limit increases (Gross and Souleles, 2002; Aydin, 2016). The second strand is based on a statistical decomposition of the income process into predictable and unpredictable movements, which are then correlated to consumption (Hall and Mishkin, 1982; Attanasio and Davis, 1996; Blundell et al., 2008).² Finally, another branch of research focuses on survey-based responses to *intended* (as opposed to *actual*) changes in income, (Shea, 1995; Bracha and Cooper, 2013) or hypothetical increases in household resources (Jappelli and Pistaferri, 2014). With these earlier contributions we share an emphasis on the characteristics that drive the heterogeneity in the distribution of the MPCs across households. Unlike previous literature, however, we are able to provide evidence on the consumption responses to *actual* positive shocks as well as *actual* negative shocks within the same household survey dataset. The novelty of our investigation comes from the fact that the survey we consider contains self-reported, unanticipated income changes of either sign.

Structure of the paper. Section 2 describes the survey questions and the distribution of the income shocks. Section 3 reports the main finding of the paper on asymmetric MPCs out of actual income changes. Section 4 presents results from questions about a hypothetical income shock. Section 5 offers a regression analysis of the determinants of variation in the

¹ For a survey, see Jappelli and Pistaferri (2010).

² In independent work, Arellano, Blundell and Bonhomme (2017) report estimates of asymmetric expenditure effects using household level data for the US and a permanent/transitory income decomposition based on a nonlinear model.

MPCs across positive and negative shocks. Section 6 assesses the consumption responses to positive and negative income shocks through the lens of a model with occasionally binding borrowing constraints. Concluding remarks and some policy implications are drawn in the last section.

2. Identifying unanticipated income changes using the Bank of England/NMG survey

In this section, we briefly describe the Bank of England/NMG survey and discuss the novel questions that allow us to elicit the MPCs out of positive and negative income shocks. Particular emphasis is devoted to the statistical properties of the distribution of those income shocks as well as to the strengths and limitations of the approach based on recollection of actual events.

2.1. Survey questions

Every year since 2003, the Bank of England has commissioned NMG Consulting to conduct a survey of British households.³ Over the years in which the questions on unanticipated changes were asked, the sample size increased from 1000 in 2011 to 6000 in 2014. The survey contains a variety of questions on households' characteristics, such as income, and balance sheet positions. This information allows us to compare the MPCs of different types of households and the traits of households that received different types of income shocks.

The NMG survey is designed to be representative of the UK population in terms of age, gender, region, housing tenure and employment status. Anderson et al. (2016) conduct a detailed evaluation of the NMG survey relative to aggregate data and other UK household surveys. As is the case with most household surveys, average levels of income in the NMG survey tend to be lower than implied by aggregate data, but it follows broadly similar trends to aggregate data and other surveys in tracking changes in income and financial positions.⁴ In the next part of this section, we will show that indeed the distribution of income shocks in the NMG survey displays similar statistical properties relative to the distribution of actual income changes from administrative data. Given our focus on income changes, these cross-checks are important to draw reliable inference from our analysis.

The exact questions used to generate households' MPCs are reported in Appendix A. Households were asked about whether their income differed from what they expected a year ago, and if so, by how much. Then, they were asked how they had adjusted their spending over the previous year in response to the unexpected change in income.⁵ Where households reported the sign and size of their spending response and the sign and size of the income shock, we were able to compute their MPC as the ratio of the two. Given the novelty of the question design upon which the MPCs are estimated, we verify the robustness of our findings to both an alternative income shock identification strategy based on hypothetical shocks (similar to Jappelli and Pistaferri, 2014) and a more qualitative indicator of the spending response (closer to that used by Sahm et al., 2010).

To minimize the effects of outliers, we excluded the largest 2.5% of positive and negative shocks relative to household income. We also excluded MPCs that were either negative or implausibly large (i.e. above 1.5).⁶ Overall, we discarded around 9% of observations but we have verified that none of our results hinge upon these restrictions.

2.2. The distribution of positive and negative income shocks

Around a third of households in our sample reported experiencing an income shock (see Table B1 in Appendix B). This share fell in the later years of the survey, dropping from 38% in 2011 to 28% in 2014. More households reported experiencing a negative rather than positive income shock, which may reflect the state of the UK economy during our sample period. The share of negative income shocks did, however, fall in each year of the survey (from 70% in 2011 to 54% in 2014), possibly consistent with the improvement in the macro economy.


Reported income shocks tended to be fairly small relative to overall income (Fig. 1). The median size was around £2000 (6% of average household income) for both positive and negative shocks. There were, however, long tails of some very large income shocks, which led to mean positive and negative shocks of £5000 and £4200 respectively. Fig. 1 shows that the distribution of shock to income ratios as a whole is leptokurtic – that is, small and very large shocks are far more

³ The survey has been running since 2003, usually in September, but prior to 2010 it was conducted face-to-face. In 2011, a second parallel online trial was launched (in which our questions on positive and negative unanticipated income changes were first included). Since 2012, the survey has been fully conducted online. This has allowed the increase in the sample size. Since 2014, an additional survey has been carried out during April. The survey is primarily a repeated cross-section, but has also included a small panel element since moving online. See Anderson et al (2016) for more details on the NMG survey.

⁴ Average gross income per house hold in the NMG survey averaged 74% of the equivalent measure from aggregate data between 2005 and 2014. The average for other UK household surveys ranged between 77% and 99% (87% excluding the Wealth and Assets Survey, which is specifically designed to oversample wealthier households).

⁵ The NMG survey questions refer to changes in 'spending' and so are likely to incorporate a wider range of categories over and above non-durable goods and services. The survey questions only ask about the spending response in the year that the shock took place, and so do not capture the response over a longer period.

⁶ All our findings are robust to excluding MPCs above one.



a) Ratio of annual income shock to annual pre-tax income. Distribution cut at ± 2.5 for presentational purposes.

Fig. 1. Shock to income distribution.^a

frequent than the normal distribution would predict – in line with new evidence on actual income changes reported by Guvenen et al. (2015) using administrative data for the US.⁷ The income shock distribution also closely matches the distribution of unexplained changes in income from the British Household Panel Survey (BHPS, see Figure B1 in Appendix B).⁸

In the 2012 survey, households were asked about the reason why their income was not what they had expected. Important explanations for positive income shocks were new jobs, higher than expected pay rises and higher pension income (see Table B2 in Appendix B). Common reasons for negative income shocks were loss of benefits, higher taxes, lower than expected wages and lower interest income.

In Table 1, we report descriptive statistics on the characteristics of households who received positive and negative income shocks. Households who experienced a windfall appeared to be similar to those experiencing negative shocks along several dimensions: annual income,⁹ unsecured and mortgage debt, mortgage debt service ratios, loan-to-value (LTV) ratios. On the other hand, households who experienced positive income shocks were, on average, younger than those experiencing a negative shock and they held more net liquid assets.¹⁰

The discrete characteristics reported at the bottom of Table 1 show that a significantly greater proportion of households experiencing a negative income shock reported: (i) to be credit constrained; (ii) that a fall in their income was likely; (iii) to be somewhat or very concerned about debt; or (iv) that they were worse off than they might have expected before the occurrence of the financial crisis. In contrast, households hit by a positive shock were more likely: (a) to have a buffer of savings to draw on in an emergency; and (b) to perceive the income shock to be less persistent.

Non-respondents sample. The survey questions have multiple layers to them, and as such there is potential for households to answer some of the questions but not all. Indeed, there was a relatively high degree of non-response as the questions progressed: while 33% of all households in the sample reported having experienced an income shock, only 15% gave information on both the size of the income change and the size of the spending change for us to compute an MPC.¹¹ Whilst there were some differences in the characteristics of those who provided all the information to allow us to compute an MPC and those who did not, the analysis presented in Section 5.1, which is based on a Heckman-type selection model, reveals that these non-respondents are unlikely to distort inference on the MPC asymmetry or our analysis of its determinants.¹²

⁷ The kurtosis of the shock to income distribution is around 18, compared to 3 for the normal distribution.

⁸ The BHPS provides panel data for the level of annual income. In order to make them more comparable with our income shocks, we regress income changes on household characteristics and use the residuals from that equation as a measure of unexplained changes in income.

⁹ The annual income shock is reported on a post-tax disposable income basis, while annual income is reported on a pre-tax basis. We use a simple approach to adjust the self-reported post-shock annual income to reflect pre-shock annual income, simply adding (subtracting) the income shock to (from) reported post-shock annual income. To the extent that positive and negative shocks are randomly distributed across tax rate bands, the fact that this approach ignores the impact of tax paid on the income shock should not influence the results.

¹⁰ Net liquid assets are defined as less unsecured debt (excluding credit card balances which the household intends to pay in full over the month).

¹¹ See Table B1 of Appendix B for more details on non-response rates.

¹² Table B3 in Appendix B reports the characteristics of households who did not report the size of the income shock and/or spending response. This can be compared to Table 1 which contains the same information for households who did respond to all questions. The group that did not respond had lower educational attainment, lower net liquid assets, less unsecured debt and some slightly weaker self-reported balance sheet characteristics.

Table 1
Characteristics of households by sign of income shock.^a

	Positive shocks				Negative shocks				
Continuous variables:	Mean	Median	Interquartile range	Sample	Mean	Median	Interquartile range	Sample	Significant difference between positive & negative?
MPC	0.14	0	[0,0]	1080	0.64	0.8	[0.1,1]	1172	Yes***
Age	45	43	[30,61]	1080	49	50	[38,61]	1172	Yes***
Annual pre-tax income (£'000)	36.7	33.0	[20.8,45]	1060	36.7	30.5	[20.7,46.0]	1172	No
Unsecured debt (£'000)	8.2	4.3	[0.9,12.5]	734	8.7	4.3	[0.8,12.5]	758	No
Mortgage debt (£'000)	89.7	75.0	[35,115]	421	91.8	75.0	[40,120]	404	No
Mortgage debt service ratio	19.4	15.0	[9.4,21.4]	476	21	18.0	[12,26.5]	443	No
Loan to value ratio	0.4	0.4	[0.2,0.7]	421	0.5	0.5	[0.2,0.7]	404	No
Loan to income ratio	2.9	2.0	[1.0,2.9]	417	2.4	1.9	[1.1,2.9]	404	No
Net liquid assets (£'000)	43.8	5.5	[−1.8,45]	1040	29.8	1.2	[−4.0,27.1]	1118	Yes***
Annual income shock (£'000)	5.0	2.0	[0.8,5]	1080	−4.2	−2.0	[−5,−1]	1172	Yes***
Net liquid assets / Annual income	2.7	0.7	[0.1,2.3]	707	1.8	0.4	[0.1,2.3]	675	Yes***
Discrete variables:	Per cent of households	Standard deviation		Sample	Per cent of households	Standard deviation		Sample	Significant difference between positive & negative?
Credit constrained	19.0	0.40		1052	37.4	0.48		1145	Yes***
Risk of fall in income	48.4	0.50		954	82.7	0.38		958	Yes***
Concerned about debt	31.6	0.47		1025	52.1	0.50		1057	Yes***
Buffer stock of savings	68.4	0.47		988	38.8	0.49		995	Yes***
Worse off since 2006	26.4	0.44		814	82.1	0.38		759	Yes***
Temporary shock	36.0	0.48		968	18.3	0.39		995	Yes***
Education GCSE or lower	29.2	0.45		821	34.6	0.48		763	Yes**

^a Variation in sample sizes across rows reflects different subsamples (e.g. mortgagors only) and response rates to different survey questions. Significant difference indicates a difference in the mean or percent of households that is significant at the 1% (***), 5% (**) or 10% (*) level.

Recall bias. As the survey relies on the ability of households to accurately recall what happened to their income and spending, our results might be affected by the so-called recall bias.¹³ This type of survey bias has been extensively investigated in the psychology, sociology and epidemiology literature. Bradburn et al. (1987) and Hassan (2005), for instance, report that (i) about 20% of critical details are irretrievable one year after the event and (ii) emotional or personally important occurrences are associated with more vivid memories. In the NMG survey, all reported values for income shocks and spending changes refer to events that took place at most within one year of the interview (though we do not have information on their distribution over these twelve months). Furthermore, in Section 3.2, we will show that our main finding of MPC asymmetry survives when considering only income shocks associated with personally-salient events such as job changes, marriage, divorce and inheritance. Finally, exploiting income data from the small panel element of the survey, in Table B5 of Appendix B, we have verified that households reporting an income shock experienced an income change of the same sign and of a similar size to the reported shock, controlling for other factors such as demographics.¹⁴

Cross-checking the accuracy of the reported consumption responses to income shocks is more difficult because the NMG survey does not include any other data on consumption. But in Section 5.1 and Table B7 of Appendix B, we will show that the result of MPC asymmetry is insensitive to using probit regressions that simply distinguish between (the larger sample of) zero and non-zero spending responses among all households reporting an income shock.¹⁵ That relies only on respondents being able to recall whether or not they changed spending at all rather than remembering the exact amount that they changed spending by.

3. The marginal propensity to consume out of positive and negative income shocks

In this section, we present descriptive statistics about the distribution of the MPCs across a number of household and shock characteristics.

¹³ The desire to look financially responsible may be another source of survey bias.

¹⁴ The coefficients on the size of the unexplained income change for those experiencing positive and negative income shocks are never statistically different from one another (after accounting for their respective signs).

¹⁵ The regression sample is larger in Table B7 because many households only reported whether (and not also by how much) they changed their spending following an income shock.




Fig. 2. Mean MPC by year.

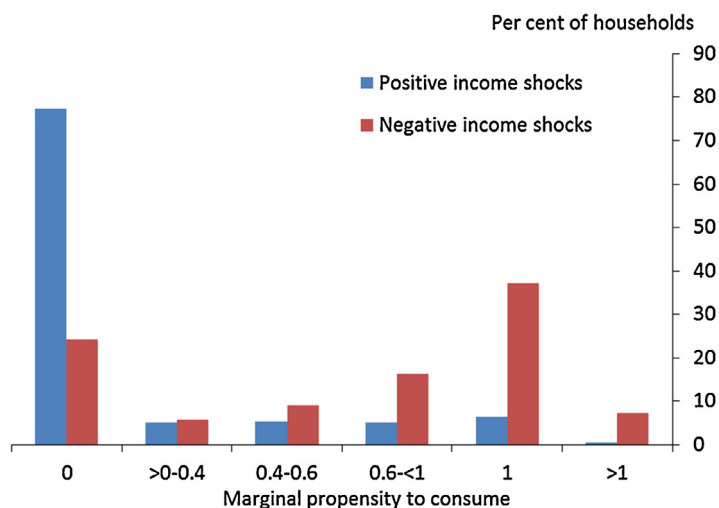


Fig. 3. Distribution of reported MPCs.

3.1. MPC distribution

A striking result across all years of the survey is that the MPC out of a negative income shock was considerably larger than that out of a positive income shock (Fig. 2), with very little variation in the values of MPCs between years. The mean MPC across the four waves of the survey was 0.64 for a negative shock, and 0.14 for a positive shock.

Fig. 3 shows the distribution of MPCs by direction of income shock. A notable feature of this distribution is the large share of MPCs with the value of zero for positive shocks (77% for the full sample). The distribution across MPC values was more even for negative income shocks, although there was still some ‘bunching’ of responses at zero and one. Jappelli and Pistaferri (2014) find a similar bunching of responses at zero and one (as well as a half), when asking about the spending response to a hypothetical windfall.

For households who experienced a positive income shock, the MPCs seemed to increase with the size of the shock. The opposite was true for households experiencing a negative income shock with the MPC falling as the shock size became larger.¹⁶

Temporary shocks. Unfortunately, the NMG survey contains no explicit questions about permanent shocks. A main reason for this design choice is that even salient (and potentially persistent) events such as a job loss or divorce are not necessarily

¹⁶ See Figure B7 in Appendix B.

Table 2
MPC by discrete household characteristics.^a

	Mean MPC				Significant difference?
	Positive shock	Sample	Negative shock	Sample	
Housing Tenure					
Owner occupier	0.10	325	0.56	377	Yes***
Mortgagor	0.16	490	0.68	456	Yes***
Private tenant	0.15	196	0.65	180	Yes***
Social tenant	0.22	65	0.77	158	Yes***
Employment status					
Employed	0.16	656	0.64	606	Yes***
Self-employed	0.12	67	0.65	111	Yes***
Unemployed	0.07	12	0.79	37	Yes***
Retired	0.10	264	0.58	241	Yes***
Credit constrained					
No	0.13	853	0.58	718	Yes***
Yes	0.24	201	0.75	428	Yes***
Risk of fall in income					
No	0.15	492	0.49	166	Yes***
Yes	0.16	464	0.67	793	Yes***
Concerned about debt					
No	0.13	702	0.57	507	Yes***
Yes	0.19	325	0.72	551	Yes***
Buffer of savings					
No	0.21	312	0.74	609	Yes***
Yes	0.12	677	0.50	387	Yes***
Worse off than expected in 2006					
No	0.14	601	0.44	136	Yes***
Yes	0.16	215	0.67	624	Yes***
Positive net worth					
No	0.16	583	0.70	669	Yes***
Yes	0.14	501	0.58	510	Yes***
Salient shocks					
By reason (2012 only)	0.23	63	0.76	55	Yes***
Shock to income > 10%	0.18	351	0.64	503	Yes***
Shock size in top decile	0.16	144	0.54	139	Yes***

^a Significance tests compare mean MPC for positive income shocks to that for negative income shocks by characteristic and reports significance at 1% (***) 5% (**) or 10% (*) level. Net worth computed as net financial assets minus secured and unsecured debt. Salient reasons for income shocks include job changes, lifestyle changes (such as marriage, divorce, death of spouse, retirement, maternity/paternity leave), windfalls (such as inheritance and gifts).


absorbing states in the sense that they do not prevent new matches between new firms and workers or between new partners from taking place in the (possibly near) future. The NMG survey contains, however, a question about whether the shock was either transitory (which respondents may have potentially interpreted as lasting over the survey year only) or persistent (which might have been interpreted as potentially lasting beyond the survey year), though in the latter case we have no information on whether this could have been seen as two, three, five years or even lasting forever. Given our interest in isolating transitory shocks, we re-compute our main statistics excluding all shocks that were self-reported as persistent. This exercise confirms the strong asymmetry found in the full sample, with a MPC of 0.17 for positive shocks and 0.57 for negative shocks within the group of self-reported transitory income changes.¹⁷

3.2. MPC by household and shock characteristics

Table 2 documents the variation in MPC values by direction of the shock, conditional on a variety of discrete household characteristics. A standout feature of the results is the consistently low value of MPCs out of positive income shocks and high MPCs out of negative shocks, irrespective of household or shock characteristics. In general, the variation of MPCs across household characteristics was greater for negative income shocks than for positive shocks.

Outright owners responded less to both positive and negative income shocks relative to social tenants, perhaps reflecting the fact that they are less likely to face liquidity constraints. There was little difference between the MPCs of those working for somebody else and those self-employed. In contrast, households headed by someone unemployed reported a markedly higher MPC out of negative shocks but a smaller MPC out of positive shocks. This latter finding, however, is based on a very small sample.

¹⁷ Reported MPCs were higher for persistent negative shocks than for transitory negative shocks but there was no significant difference between temporary and persistent positive shocks. The similarity across positive shocks chimes with the evidence from [Sahm, Shapiro and Slemrod \(2010\)](#) that the response to the 2009 U.S. income tax rebate did not differ based on whether the policy change was perceived as temporary or persistent.



a) Total debt is the sum of any mortgage and unsecured debt.

Fig. 4. Mean MPC by total debt-to-income ratio.^a

MPCs were systematically higher for households reporting being credit constrained; having concerns about their debt; or reporting that they have an insufficient buffer of savings in the event of an emergency. That was true for both positive and negative shocks, although the differences in MPCs between households reporting the above characteristics and those not were sizably larger for negative shocks.

There is also suggestive evidence that a precautionary savings channel might be at play: among households who received a negative income shock, MPCs for those concerned about future falls in income were significantly higher (at 0.67) than those that were unconcerned (at 0.49). Meanwhile for positive income shocks there was no significant difference.

One factor potentially influencing our results is that the surveys were conducted in the wake of the financial crisis. This will have affected the financial positions of many households, but some are likely to have been more adversely affected than others. Households who reported being worse off at the time of interview, relative to what they might have expected prior to the crisis, had considerably higher MPCs out of negative income shocks than those who were not worse off.

As for balance sheet characteristics, borrowers (as defined by having negative net worth or no buffer of savings for emergencies) reported higher MPCs than savers out of both positive and negative income shocks, but the asymmetry in MPCs was clearly present for both groups. Fig. 4 shows that households with higher debt-to-income ratios reported higher MPCs out of negative shocks, but there was little difference for positive shocks. On the assets side, households with high net liquid assets reported lower MPCs out of negative income shocks, but there was little difference for positive shocks (Fig. 5).¹⁸

As for shock characteristics, in the 2012 survey, households were asked about the reason for the income shock that they had experienced. While we do not find evidence of significant differences in the MPC asymmetry across income shock reasons (Appendix B Table B2), the last panel of Table 2 restricts the attention to three alternative ways to categorise shocks that might be thought of as salient or easier for households to remember. The first row, based only on the 2012 survey, focuses on reasons such as job change, life style changes (e.g. marriage, divorce, etc), inheritance and gifts. The other two rows, based on all surveys, look only at relatively large shocks. According to all three partitions of the data, we find robust evidence of a spending asymmetry with MPCs for positive shocks ranging from 0.16 to 0.27 and MPCs for negative shocks varying from 0.54 to 0.76.¹⁹


4. The spending response to hypothetical income shocks

In the previous sections, we have shown that the reported income shock distribution is consistent with the properties of administrative data and that the finding of MPC asymmetry is robust to focusing only on salient shocks. Another way to assess the extent of any possible recall bias is to exploit another set of questions which, in the 2013 survey only, asked about the spending response to a hypothetical income shock of either sign (see Appendix A for the specific wording of the questions).²⁰ More specifically, households were randomised across the size of the hypothetical shock (between £500, £1000,

¹⁸ Figures B2 to B7 in Appendix B show some further breakdowns of MPCs along other dimensions.

¹⁹ In Appendix B Table B6, we obtain similar results if we focus only on transitory shocks (panel A) or exclude households for which the change of cash flows had a different sign from the income shock (panel B).

²⁰ This strategy has been used by Jappelli and Pistaferri (2014), but only for a hypothetical positive income shock, and we refer to their work for a discussion of the limitation of this alternative approach.



a) Net liquid assets are defined as directly held financial assets less unsecured debt (excluding credit card balances which the household intends to pay in full over the month).

Fig. 5. Mean MPC by net liquid assets to income ratio.^a

£2000 and £10,000) and were given categories from which to choose their spending response, implying MPCs bounded between zero and one by design.

Two main concerns with this hypothetical approach are that: (i) on actual data, there are differences between the groups that received positive and negative income shocks; and (ii) actual responses may be very different from hypothetical responses (see [Graziani et al., 2013](#)). One way to ameliorate the first concern is to compute the hypothetical responses of those households that received an actual income shock of the same sign: the average MPC out of a further hypothetical positive shock was 0.28 and for those that received an actual negative income shock, the MPC was 0.41. This asymmetry is narrower but still significant compared to the actual responses, which over the same year were 0.14 and 0.55 when we restrict the sample to values between zero and one for comparability with the MPCs based on the hypothetical questions.²¹

Both recall and hypothetical approaches to estimating MPCs are subject to misreporting: households may report hypothetical responses that are very different to how they would actually respond, and households may have difficulty recalling the actual size of the unexpected income change. Accordingly, we run a further exercise that combines information from both approaches. To the extent that the measurement/mis-reporting errors are uncorrelated across hypothetical and actual responses, the combined strategy could yield more accurate results. More specifically, we consider only households that gave a consistent MPC for both sets of questions, namely whether the MPC based on the hypothetical questions is within 15% of their reported actual MPC (about 60% of the 2013 sample).²²

In [Fig. 6](#), we report the mean actual MPC for 2013 (with an upper bound of one to be consistent with the hypothetical responses), and compare them to the mean actual MPC for the restricted sample of households with consistent hypothetical MPCs. This combined approach continues to reveal a large and significant asymmetry on actual MPCs, with an average value of 0.07 for positive income shocks and 0.46 for negative income shocks.


5. Explaining the response to positive and negative shocks

To explore more formally what drives our main finding of asymmetric spending behaviour with respect to the sign of the shock, in this section we relate the MPCs to household and shock characteristics. We begin with simple OLS regressions and then verify the robustness of our findings to using a propensity score matching approach.²³

²¹ We are also able to compare these hypothetical responses for those who did not give sufficient information to compute an MPC out of their actual shock (who were discussed in [Section 2.2](#)), and verify that the small differences between the hypothetical responses of the respondents and non-respondents (to the size of their actual income shock questions) are statistically insignificant (see Figure B9 in Appendix B).

²² Results below are robust to choosing 5%, 10%, 20% or 25% as cut-off for restricting the sample of actual MPCs based on the distance from the hypothetical MPC. Furthermore, we have verified that the characteristics of households reporting 'consistent' and 'inconsistent' MPCs across the two approaches do not vary in any particularly systematic way between positive and negative income shocks.

²³ The regression results are similar if estimated as tobit models.



a) Figure reports mean actual MPCs for households for the year 2013 (the year for which we also have hypothetical responses), conditional on an MPC between zero and one (for comparability with the buckets in the hypothetical questions).

Fig. 6. MPC distribution for actual income shocks where hypothetical response is within 15% of actual MPC.^a

5.1. Baseline results

Our OLS regressions use the MPC as the dependent variable and the characteristics of the shock as well as households' characteristics as regressors. We run six separate specifications for positive and negative income shocks (Table 3), all with HAC robust standard errors. In the first five columns, we project the MPC on a specific set of household characteristics reflecting (i) demographics, (ii) the nature of the shock, (iii) household balance sheets, (iv) net liquid wealth positions and (v) debt in order to assess the contribution of each set of variables in isolation. The final column combines all of the characteristics in our richest specification, which allows for a comparison of their relative merits.

Several characteristics are significant in explaining variation in the MPCs across households. Demographics appear to play a fairly small role with the youngest households associated with lower MPCs, especially after a negative income shock. In column 2 of Table 3, we report that MPCs are significantly higher out of large (top two quartiles by size) positive income shocks, whereas for negative shocks, the MPC out of income shocks in the top quartile is significantly lower.

An important determinant of the response to an income shock is likely to be the balance sheet position of a household, indicators of which are introduced in the third column. As theory predicts, credit constrained households were estimated to have higher MPCs out of both positive and negative shocks, +0.1 and +0.07 respectively. Households who reported some likelihood of a fall in income over the next year (an indicator of income risk) were estimated to have higher MPCs out of a negative shock than those unconcerned about a fall in income, but there was not a significant difference for positive shocks. Finally, respondents who had a buffer of savings for emergencies were estimated to have had significantly lower MPCs for both positive and negative income shocks (−0.09 and −0.18).

These are qualitative indicators of overall balance sheet health, so as a crosscheck we include more quantitative measures of financial assets (in the fourth column) and debt (in the fifth column). We consider the ratio of net liquid assets to household income,²⁴ which might be thought of as an indicator of whether the household is hand-to-mouth (in the traditional sense, or as the wealthy hand-to-mouth described by Kaplan et al., 2014). Indeed, households who have a high ratio of liquid assets to income tended to have lower MPCs for income shocks of both signs. The difference is particularly pronounced for negative shocks: households with a liquid asset to income ratio in the top tertile were estimated to have an MPC which was 0.2 lower than the average household in the bottom tertile. On the liabilities side, households who were concerned about their level of debt had a significantly higher MPC than those without concerns, for income shocks of both signs. For positive shocks, households with a loan-to-value ratio above 90% had MPCs 0.19 percentage points higher than those with a LTV below 75%, though the results are not monotonic across LTV bands.

The final column includes all of the regressors in our richest specification. As several regressors are correlated with each other, it should not come as a surprise that the size and significance of several of our individual coefficients are reduced. For positive income shocks, households facing larger shocks or who were credit constrained had a significantly higher MPC

²⁴ Net liquid assets are defined as directly held financial assets less unsecured debt (excluding credit card balances which the household intends to pay in full over the month)

Table 3

OLS regression results, dependent variable MPC.

VARIABLES	POSITIVE INCOME SHOCKS						NEGATIVE INCOME SHOCKS					
	Demographics ^a	Shock features ^b	Balance sheet dummies ^c	Net liquid wealth ^d	Debt ^e	All ^f	Demographics ^a	Shock features ^b	Balance sheet dummies ^c	Net liquid wealth ^d	Debt ^e	All ^f
Age 18–24 dummy	0.05 (0.03)					0.05 (0.04)	−0.13** (0.07)					−0.10 (0.07)
Age 45–64 dummy	0.02 (0.03)					0.02 (0.03)	0.02 (0.03)					0.06 (0.04)
Age 65 + dummy	−0.04* (0.02)					−0.01 (0.03)	−0.10** (0.04)					−0.02 (0.05)
Unemployed dummy	−0.10 (0.07)					−0.10 (0.08)	0.14** (0.06)					0.09 (0.07)
Shock size in 2nd quartile dummy		0.02 (0.03)				0.02 (0.03)	−0.01 (0.05)					0.05 (0.05)
Shock size in 3rd quartile dummy		0.05* (0.03)				0.08** (0.03)	0.02 (0.04)					0.05 (0.04)
Shock size in top quartile dummy		0.07** (0.03)				0.08*** (0.03)	−0.08** (0.04)					−0.04 (0.04)
Transitory shock dummy		0.02 (0.02)				0.02 (0.02)	−0.10*** (0.03)					−0.07* (0.04)
Credit constrained dummy			0.10*** (0.03)			0.09*** (0.03)			0.07** (0.03)			0.06 (0.04)
Income uncertainty dummy			−0.03 (0.02)			−0.04* (0.02)			0.13*** (0.04)			0.09** (0.04)
Has buffer of savings dummy			−0.09*** (0.03)			−0.11*** (0.03)			−0.18*** (0.03)			−0.12*** (0.04)
2nd third of liquid assets to income dummy				−0.03 (0.03)		0.00 (0.03)				0.03 (0.03)		0.06 (0.04)
Top third of liquid assets to income dummy				−0.06** (0.02)		0.01 (0.03)				−0.20*** (0.03)		−0.08 (0.05)
Mortgagor dummy					0.01 (0.02)	0.02 (0.03)					−0.01 (0.03)	−0.02 (0.04)
LTV > 90% dummy					0.19* (0.10)	0.17 (0.11)					0.07 (0.05)	0.09 (0.05)
LTV 75%–90% dummy					−0.07* (0.04)	−0.10** (0.04)					0.04 (0.06)	0.06 (0.07)
Concerned about debt dummy					0.06*** (0.02)	0.00 (0.03)					0.14*** (0.03)	0.02 (0.04)
2013 dummy	−0.03 (0.03)	−0.04 (0.03)	−0.04 (0.03)	−0.02 (0.02)	−0.03 (0.03)	−0.06** (0.03)	−0.09*** (0.03)	−0.09*** (0.03)	−0.08** (0.04)	−0.07** (0.03)	−0.09*** (0.03)	−0.08** (0.04)
2014 dummy	0.01 (0.03)	0.00 (0.03)	0.01 (0.03)	0.02 (0.02)	0.01 (0.03)	−0.01 (0.03)	−0.03 (0.03)	−0.02 (0.04)	−0.04 (0.04)	−0.01 (0.03)	−0.03 (0.03)	−0.03 (0.04)
Constant	0.16*** (0.03)	0.12*** (0.03)	0.23*** (0.04)	0.18*** (0.02)	0.13*** (0.02)	0.18*** (0.05)	0.71*** (0.03)	0.74*** (0.04)	0.62*** (0.05)	0.72*** (0.03)	0.62*** (0.03)	0.62*** (0.07)
Observations	1029	970	903	1084	1027	854	1069	996	875	1179	1058	817
R-squared	0.012	0.014	0.043	0.008	0.019	0.074	0.026	0.026	0.092	0.051	0.037	0.123

Robust standard errors in parentheses. Symbols indicate significance at the 1% (***), 5% (**) and 10% (*) level.

^a Benchmark household is aged 25–44 and employed, responding to the 2012 survey^b Benchmark household is one with a shock size in the lowest quartile, faces a persistent shock, responding to the 2012 survey^c Benchmark household is one with access to credit, not uncertain about income and without a buffer of savings, responding to the 2012 survey^d Benchmark household is in the bottom liquid assets to income third, responding to the 2012 survey^e Benchmark household is probably a renter or owner occupier who is unconcerned about debt, responding to the 2012 survey^f Benchmark household has all the characteristics of benchmark households in earlier regressions

Table 4Average treatment effect, alternative matching approaches.^a

VARIABLES	(1)	(2)	(3)	(4)	(5)
	Demographics	As (1), plus: Shock characteristics	As (2), plus: Balance sheet dummies	As (3), plus: Liquid assets	As (4), plus: Debt
Unmatched	−0.50*** (0.02)	−0.50*** (0.02)	−0.50*** (0.02)	−0.50*** (0.02)	−0.49*** (0.02)
Kernel matching (epanechnikov), Probit					
ATT	−0.48*** (0.02)	−0.47*** (0.02)	−0.37*** (0.03)	−0.37*** (0.03)	−0.36*** (0.03)
Nearest five neighbour matching, Probit					
ATT	−0.53*** (0.06)	−0.49*** (0.03)	−0.36*** (0.03)	−0.36*** (0.03)	−0.36*** (0.03)
Kernel matching (normal), Probit					
ATT	−0.53*** (0.06)	−0.49*** (0.03)	−0.36*** (0.03)	−0.36*** (0.03)	−0.36*** (0.03)
Kernel matching (epanechnikov), Logit					
ATT	−0.48*** (0.02)	−0.47*** (0.02)	−0.37*** (0.03)	−0.37*** (0.03)	−0.36*** (0.03)
Observations	2263	1966	1678	1678	1671

Standard errors in parentheses. Symbols indicate significance at the 1% (***) , 5% (**) and 10% (*) level.

whereas households who reported facing significant income risks, having a buffer of saving or an LTV between 75% and 90% had a lower MPC. Across negative income shocks, future income risks and having a buffer of savings remain significant predictors of the MPC in the augmented specification.

In Section 2.2, we explained that there were some differences in the characteristics of households who did and did not provide sufficient information to compute an MPC. One way of assessing the effect that may have on our results is to use a Heckman selection model, where the probability of being a non-respondent computed from a first-stage probit regression is included as a control in the second-stage regression of the MPC. These results are reported in Table B4 of Appendix B. The coefficient on the sample selection variable is insignificant and the other regression coefficients are little changed. An alternative approach for establishing robustness to non-response is to estimate a probit model with a dummy variable as the dependent variable which has a value of one if the MPC is greater than zero.²⁵ This substantially increases the sample size because households only needed to report the direction of the spending change, not the amount.²⁶ The direction and significance of the coefficients are consistent with the OLS regressions (Table B7 of Appendix B). Overall, these exercises suggest that the non-responses are not having a large effect on our results.

In summary, across all specifications and econometric methods, the characteristics that are consistently important in determining the response to a positive income shock are balance sheet characteristics (having a buffer of savings, credit constraints and income risk), larger income shocks and having a high LTV ratio. For negative income shocks, income risk and having a buffer of savings are consistently important across specifications.

5.2. A propensity score approach

A striking feature of our results is the sizable asymmetry between the MPC out of positive and negative income shocks. A possible concern with the interpretation of our findings is that households receiving a positive income shock may be fundamentally different along some dimensions from households subject to a negative shock.²⁷ To ameliorate such a concern, in this section we use a propensity score matching (PSM) framework. In the first stage, a probit or logit regression computes the likelihood that a household would have received a positive shock rather than a negative one based on certain characteristics (this is their propensity score). It then matches households from the positive and negative income shock groups based on their propensity scores, to find respondents who ex-ante would have faced a similar likelihood of receiving a positive rather than a negative income shock. This allows us to compute a 'counterfactual' MPC for households with a positive income shock, had they instead received a negative one. The difference between their actual MPC and their matched MPC can then be seen as an alternative measure of the asymmetric effect of receiving a positive shock rather than a negative shock.

Table 4 reports the results of this exercise. The top row gives the unmatched difference between MPCs when a household receives a positive rather than a negative shock and therefore represents the *unconditional* average difference between the

²⁵ Figure B8 of Appendix B shows that (as we'd expect) this statistic also displays a large asymmetry across the years of the survey: the share of households reporting an MPC of greater than zero for negative income shocks was 87%, compared to 24% for positive income shocks.

²⁶ The sample size increases to 5,397 versus a trimmed sample of 2,263 for actual MPCs. This exercise also makes the results more comparable to surveys that ask for more qualitative indications of how spending has changed (and so providing robustness to our quantitative approach) and would not be affected by any measurement errors in reporting the specific pounds change in consumption following the income shock.

²⁷ Prima facie evidence on this issue can be obtained by focusing on the small panel dimension of the survey and restricting attention to those households that, over two consecutive years, reported income shocks of different sign. Despite the limited number of observations (about one hundred), we still find a sizable asymmetry both in the average MPC and the share of households reporting a zero MPC across positive and negative shocks.

two groups. The remaining rows report the same average MPC difference but *conditional* on the characteristics in the column heading being held similar (or matched in the PSM jargon) across the groups of positive and negative shocks. The columns of the table give the results where different regressors are included in the production of the propensity score in the first stage. This implies that, for each column, the difference between the entry in the first row and each of the remaining rows could be interpreted as an indication of the marginal contribution of the added set of variables (in that column) to explaining the unmatched average difference between the MPCs of positive and negative shock groups in the first row. The remaining rows of Table 4 differ according to the specification of the regression to produce the propensity score (probit or logit) and the method by which households are matched. It is worth noting that the interpretation of our PSM exercise relies on the presumption that if a set of variables is a significant driver of the MPC asymmetry, then by adding them to the first stage it should reduce the MPC gap between positive and negative income shocks, conditional on this set of characteristics.

The Average Treatment effect on the Treated (ATT), which is the actual MPC for a household with a positive shock minus the average matched MPC, reported in Table 4 declines monotonically across the columns as additional variables are included in the propensity score regression. The greatest effect comes when the balance sheet dummy variables are included in the first stage regression. This suggests that being credit constrained, having an uncertain income or a buffer of savings are prominent determinants of the observed MPC asymmetry, with a contribution of around one quarter of the unconditional MPC gap between positive and negative shocks of 0.50 (i.e. $100 * [-0.50 + 0.37] / [-0.50] = 26\%$).²⁸

6. The role of borrowing constraints: insights from a simple model

In the previous sections, we have shown that household balance sheet characteristics and the presence of debt/lack of liquidity are significant drivers of the MPC variation observed both across households and across positive and negative income changes. It is worth noting, however, that the financial questions in the Bank of England/NMG Consulting survey refer to either *past* or *current* access to credit markets. As such, these questions are silent on the extent to which current expenditure decisions may also be influenced by the likelihood that some financial circumstances (and borrowing constraints in particular) may change in the *future*, possibly as a result of a negative income shock.

This section explores the extent to which possible changes in future credit conditions may account for the asymmetry in MPCs reported earlier on, using a partial equilibrium model with an occasionally binding borrowing constraint. A representative household alternates between two regimes. In the first regime, the household is unconstrained and thus is able to borrow sufficient funds to smooth consumption optimally. In the second regime, the household is constrained in the borrowing amount that they can lever relative to income and so behave in a hand-to-mouth fashion. The probability of hitting the second regime depends on the size of the shock relative to the distance from the borrowing constraint.

We consider a very simple specification which sets out the decision faced by the household. The representative consumer in the model maximises its lifetime utility at a discount rate of β subject to a budget constraint,

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \log(c_t) \quad (1)$$

$$c_t + Rb_{t-1} = y_t(1 - H) + b_t \quad (2)$$

The resources available to households consist of income, y_t , net of a subsistence level of consumption made mostly by housing costs and a small share of essential non-durables, Hy_t , plus any possible one-period loan, b_t . These are spent on consumption goods, c_t , and repaying the debt from the previous period, b_{t-1} .

The representative agent uses debt, b_t , to smooth consumption intertemporally, which they repay at an interest rate, R , in the next period. But they may face a constraint which ensures that the debt that they acquire each period does not exceed a proportion M of their income.

$$b_t \leq My_t \quad (3)$$

The model incorporates two regimes according to whether the borrowing constraint restricts household behaviour or not. In the case where the borrowing constraint is slack, the Lagrange multiplier on the Euler condition, λ_t , is set to zero and the household optimally smoothes their consumption across time.

$$\lambda_t = \frac{1}{c_t} - \frac{\beta R}{c_{t+1}} \quad (4)$$

When the borrowing constraint is binding (i.e. $\lambda_t \neq 0$), households are unable to optimally allocate consumption and borrow at their maximum loan-to-income value. The current period consumption is limited by the borrowing constraint.

Finally, household income in the model is determined by an exogenous autoregressive income process, where ε_t is an i.i.d. normally distributed shock with mean zero and unitary standard deviation while σ represents the standard deviation

²⁸ An alternative approach for understanding the effects of receiving shock of a different sign would be to compute the fitted values for those who received positive income shocks using the negative regressions described in Section 5.1, and vice versa. The mean fitted values would be 0.19 and 0.52, suggesting a contribution of around a third ($100 * [-0.5 + 0.33] / [-0.5] = 34\%$).

Table 5
Model-implied MPC across debt position.

	Lower Debt	Baseline	Higher Debt
Positive income shock	0.12	0.13	0.13
Negative income shock	0.60	0.66	0.72

Note: MPCs to positive and negative income shocks for $M = [0, 0.3, 0.6]$ in the model (1)–(5).

of income shocks.

$$\log(y_t) = \rho \log(y_{t-1}) + \sigma \varepsilon_t \quad (5)$$

We solve the model using the occasionally binding constraints toolkit of [Guerrieri and Iacoviello \(2015\)](#). This toolkit delivers a piecewise linear approximation to the non-linear solution, which is shown to approximate the full solution well.²⁹

The model is calibrated at a quarterly frequency. In particular, β is set to 0.9925 and $R = 1.0075$ corresponding to an annual real interest rate of 3%. In the baseline parametrization, the loan-to-income ratio M is fixed to 0.3, consistent with the average unsecured debt-to-income ratio reported by the Bank of England/NMG survey respondents who described themselves as credit constrained. The share H is set to 0.4 whereas the process of income is characterized by an autoregressive parameter of 0.9 and a shock standard deviation of 0.02, consistent with the evidence from both survey and administrative data on household income (see also Table B5 in Appendix B).

As the parameter M pins down the steady state level of household debt-to-income ratio in the model, we consider two alternative calibrations relative to the baseline case. The first scenario of lower debt is represented by $M=0$. The second scenario corresponds to an environment of higher debt and is exemplified by $M=0.6$. The grid below summarizes the cumulated MPCs over the first four quarters after the shock and can be seen as the theoretical counterpart of the variation in empirical MPCs recorded by [Fig. 4](#) across debt positions ([Table 5](#)).³⁰

Two main results are worth emphasizing from this exercise. First, despite the model's simplicity, the baseline case in the middle column produces MPCs of 0.13 for positive income shocks and 0.66 for negative income shocks, which are remarkably close to the estimates of 0.14 and 0.64 from the Bank of England/NMG survey. Second, the MPCs to positive shocks display virtually no variation across columns whereas the MPCs to negative shocks increase monotonically with the degree of household indebtedness, consistent with the evidence reported in [Figs. 4](#) and [5](#). The intuition for these results comes from noting that a positive shock tends to relax the borrowing constraint irrespective of the household debt position. In contrast, a negative income shock not only makes the constraint more likely to bind but also tighten credit conditions, as any loan is now levered upon a reduced income.

7. Concluding remarks

A large empirical literature has reported significant and heterogeneous consumption responses to positive income changes, as generated for instance by tax rebates, minimum wage hikes or hypothetical/anticipated increases in household resources. Far less, however, was known about the consumption responses to negative, unanticipated income changes. In this paper, we have contributed to filling in this important gap in the literature by using a set of newly-added questions to the Bank of England/NMG survey of British households. These questions ask about the size and sign of any unanticipated income change during the previous year as well as the share of that change spent.

Our main contribution is to highlight how the sign of an income shock affects the spending response: MPCs out of negative income shocks are estimated to be much larger than for positive shocks. Between 2011 and 2014, an unexpected rise in income is associated with MPCs in the range of 0.07 to 0.17 whereas an unpredicted fall in income leads to MPCs between 0.46 and 0.68. These average effects, however, mask significant heterogeneity, with households holding higher debt, more likely to be facing credit constraints, liquidity shortages and a higher risk of lower income in the future typically reporting significantly higher MPCs. Sensitivity analyses suggest that the findings of asymmetric spending responses and MPC heterogeneity across households are robust to various strategies to deal with survey non-responses, recall bias, income shock identification and measurement errors.

The findings in this paper are consistent with models that feature income risk, liquidity constraints and precautionary savings. An extensive regression analysis along a large number of observable characteristics suggests that variations in the strength of the household balance sheet, debt position and income risks can account for a significant share of the observed MPC asymmetry between positive and negative income shocks.

²⁹ One mechanism that the piecewise linear solution is unable to account for is precautionary savings motives associated with the possibility that the constraint may bind in the future.

³⁰ Also in the Bank of England/NMG survey, households with higher debt tend to have systematically lower liquidity than households with lower debt.

The asymmetry in how spending responds to positive and negative income shocks has important implications for household responses to monetary and fiscal policies associated with significant income changes.³¹ To the extent that monetary policy influences households' resources (Cloyne et al., 2016) and engineers a redistribution between borrowers and savers, the asymmetry in MPCs that we have documented here implies that, at least in the short-term, a sufficiently large increase in interest rates would have a larger contractionary impact on aggregate spending than the expansionary effect from an equivalent reduction in rates.³² Some caution is needed in applying the MPC estimates from this study directly to changes in monetary policy, since the idiosyncratic income shocks in our sample are much larger than the typical policy rate move of 25 basis points. But they could be more relevant to rationalize the effects of large changes in monetary policy such as the 450 basis points cut in U.K. interest rates at the beginning of the financial crisis. An asymmetry in the consumption response to positive and negative income changes would suggest that the substantial loosening in interest rates of 2008/09 may have provided less stimulus than was expected on the basis of an "average" MPC for the whole sample, simply because the latter failed to recognize the very significant and sizable heterogeneity in the distribution of consumer spending responses across positive and negative income changes.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.jmoneco.2017.11.007](https://doi.org/10.1016/j.jmoneco.2017.11.007).

References

- Aaronson, D., Agarwal, S., French, E., 2012. The spending and debt response to minimum wage hikes. *Am. Econ. Rev.* 102 (7), 3111–3139.
- Agarwal, S., Qian, W., 2014. Consumption and debt response to unanticipated income shocks: evidence from a natural experiment in Singapore. *Am. Econ. Rev.* 104 (12), 4205–4230.
- Anderson, G., Bunn, P., Pugh, A., Uluc, A., 2016. The Bank of England/NMG survey of household finances. *Fiscal Stud.* 37 (1), 131–152.
- Arellano, M., Blundell, R., Bonhomme, S., 2017. Earnings and Consumption Dynamics: a nonlinear panel data framework. *Econometrica*. 85 (3), 693–734.
- Attanasio, O., Davis, S., 1996. Relative wage movements and the distribution of consumption. *J. Polit. Econ.* 104 (6), 1227–1262.
- Aydin, D., 2016. The Marginal Propensity to Consume Out of Liquidity. Stanford University, mimeo.
- Barnichon, R., Matthes, C., 2017. Functional approximations of impulse responses: new insights into the asymmetric effects of monetary policy. Unpublished.
- Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. *Am. Econ. Rev.* 98 (5), 1887–1921.
- Bracha, A., Cooper, D., 2013. Asymmetric responses to tax-induced changes in personal income: the 2013 payroll tax hike versus anticipated 2012 tax refunds. Federal Reserve Bank Boston Public Policy Briefs 13 (4).
- Bradburn, N., Rips, L., Shevell, S., 1987. Answering autobiographical questions: the impact of memory and inference on surveys. *Sci., New Ser.* 236 (4798), 157–161.
- Caballero, R.J., 1990. Consumption puzzles and precautionary savings. *J. Monet. Econ.* 25 (1), 113–136.
- Carroll, C.D., 1992. The buffer stock theory of saving: some macroeconomic evidence. *Brookings Pap. Econ. Act.* 2, 61–156.
- Carroll, C.D., 1994. How does future income affect current consumption? *Q. J. Econ.*, CIX (1) 111–148.
- Carroll, C.D., 2001. A theory of the consumption function, with and without liquidity constraints. *J. Econ. Perspect.* 15 (3), 23–45.
- Cloyne, J., Ferreira, C., Surico, P., 2016. Monetary Policy when Households Have debt: New Evidence On the Transmission Mechanism. Bank of England Staff Working Paper, No. 589.
- Deaton, A., 1991. Saving and liquidity constraints. *Econometrica* 59 (5), 1221–1248.
- Graziani, G., van der Klaauw, W., Zafar, B., 2013. A Boost in the Paycheck: Survey Evidence on Workers' Response to the 2011 Payroll Tax Cuts. Federal Reserve Bank of New York Staff Report 592.
- Gross, D., Souleles, N., 2002. Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data. *Q. J. Econ.* 117 (1), 149–185.
- Guerrieri, L., Iacoviello, M., 2015. OccBin: a toolkit to solve models with occasionally binding constraints easily. *J. Monet. Econ.* 70, 22–38.
- Guvens, F., Karahan, F., Ozkan, S., Song, J., 2015. What Do Data On Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk? NBER Working Paper 20913.
- Hall, R., Mishkin, F., 1982. The sensitivity of consumption to transitory income: estimates from panel data on households. *Econometrica* 50 (2), 461–481.
- Hassan, E., 2005. Recall bias can be a threat to retrospective and prospective research designs. *Internet J. Epidemiol.* 3 (2).
- Hedlund, A., Karahan, F., Mitman, K., Ozkan, S., 2016. Monetary policy, heterogeneity and the housing channel. 2016 Meeting Papers 663, Society for Economic Dynamics.
- Jappelli, T., Pistaferri, L., 2010. The consumption response to income changes. *Ann. Rev. Econ., Ann. Rev.* 2 (1), 479–506.
- Jappelli, T., Pistaferri, L., 2014. Fiscal policy & MPC heterogeneity. *Am. Econ. J.* 6 (4), 107–136.
- Johnson, D.S., Parker, J.A., Souleles, N.S., 2006. Household expenditure and the income tax rebates of 2001. *Am. Econ. Rev.* 96 (5), 1589–1610.
- Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47 (2), 263–291.
- Kaplan, G., Moll, B., Violante, G., 2016. Monetary Policy According to HANK. NBER Working Paper, No. 21897.
- Kaplan, G., Violante, G., Weidner, J., 2014. The wealthy hand-to-mouth. *Brookings Paper on Economic Activity*, Spring 2014.
- Misra, K., Surico, P., 2014. Consumption, income changes and heterogeneity: evidence from two fiscal stimulus programmes. *Am. Econ. J.* 6 (4), 84–106.
- Parker, J.A., Souleles, N.S., Johnson, D.S., McClelland, R., 2013. Consumer spending and the economic stimulus payments of 2008. *Am. Econ. Rev.* 103 (6), 2530–2553.
- Sahm, C.R., Shapiro, M.D., Slemrod, J., 2010. Check in the Mail or More in the Paycheck: Does the Effectiveness of Fiscal Stimulus Depend on How It is Delivered? NBER Working Paper, No. 16246.
- Shea, J., 1995. Union contracts and the life-cycle/permanent-income hypothesis. *Am. Econ. Rev.* 85 (1), 186–200.
- Tenreiro, S., Thwaites, G., 2016. Pushing on a string: US monetary policy is less powerful in recessions. *Am. Econ. J.* 8 (4), 43–74.

³¹ This is consistent with evidence from Tenreiro and Thwaites (2016) and Barnichon and Matthes (2017) that identify a similarly asymmetric effect from monetary tightenings versus loosening using time series models, aggregate data from national statistics and macro identification of policy shocks.

³² One (partial equilibrium) channel through which reducing interest rates can stimulate the economy is by redistributing income away from savers, who make relatively smaller reductions in spending following a change in resources, and giving to borrowers, who increase spending by more than savers reduce it. Stronger growth in the economy would then lead to further increases in income via general equilibrium channels, which in turn would lead to higher spending.