

Another look at whether the rich save more: evidence from a survey experiment

**Mike Brewer, Jon Burton, Thomas F. Crossley, Paul Fisher, Alessandra Gaia, Annette Jäckle,
Joachim Winter¹**

Preliminary and incomplete version: do not cite without permission

Draft of July 2018 prepared for the 2018 IARIW conference, Copenhagen, Denmark

Abstract: A long-standing question in empirical economics is whether more affluent households save a larger fraction of current income than poorer households. Much of the literature has used IV approaches to deal with any measurement error or transitory fluctuations in income, but it is challenging to find valid instruments to deal both with transitory income shocks and the systematic under-reporting of expenditure and misreporting of income that (we think) pervades household surveys. We use data from an experimental design that collected information on income, spending, and “net saving”, and then prompted respondents to make adjustments if these were not in balance. Those households who do, after prompting, report consistent data on their household finances have much lower saving rates, below 5% for low income households, and rising slowly to 25% for the richest. On the other hand, the majority of households do not report consistent data on their household finances even after a prompt. Our results highlight the potential for measurement error or transitory shocks to income and spending to distort seriously estimates of the saving function.

JEL codes:C81, D12

Keywords: household budget, income, spending, saving rate, reconciliation, survey methods, permanent income

¹ Correspondence to mbrewer@essex.ac.uk. Brewer, Burton, Crossley, Fisher, and Jäckle are at the University of Essex, Gaia is at City University, and Winter is at the University of Munich. We are grateful for funding from the ESRC grants which support the Understanding Society Innovation Panel, grant numbers RES-586-47-0001 and RES-586-47-0002; and a grant from the ESRC Transformative Research scheme and the National Centre for Research Methods (ES/N006534/1). Brewer and Crossley also acknowledge support from the ESRC-funded Research Centre on Micro-Social Change (grant number ES/L009153/1). We are grateful for comments made at seminars at the University of York, and at the Understanding Society 2017 conference and the 2018 Research Methods Festival. We thank Cormac O’Dea for Stata code and helpful conversations.

Introduction

A long-standing and important question in empirical economics is whether more affluent households save a larger fraction of current income than poorer households (see, for example, Dynan et al., 2004 and Alan et al. 2015 and references therein). This is relevant, for example, to distinguishing between alternative macroeconomic models, and for assessing the incidence of consumption taxes, and for thinking about the intergenerational transmission of advantage. In cross-sectional survey data, from many periods and countries, high income households appear to save a larger share of their income than lower income households. On the other hand, there has no broad increase in aggregate savings rates over time even though real per capita incomes have risen.

There are both data and conceptual reasons that make assessing the slope of the saving function difficult. One problem is the absence of data that, for a given household, records income, expenditure and the change in financial position (or “net saving”) over the same period of time. As an accounting identity, income minus expenditure must be the change in financial position (that is, additions to or withdrawals from net wealth) but, although data on any two of these concepts will imply the third, in recent years there has been considerable attention to the possibility that in household surveys income is under reported among the poor (see, e.g., Sabelhaus and Groen (2000) and Meyer and Sullivan (2003, 2011) for the US, and Brewer et al. (2017) for the UK) and consumption expenditure is under-reported among the rich (for the UK, for example, see Appendix B of Crossley and O’Dea (2010)). Either of these phenomena would also contribute to over-estimating the slope of the savings function.

The conceptual issues were articulated over 60 years ago, when Friedman (1957) famously emphasized that the upward sloping cross-sectional savings function (i.e. saving divided by income plotted as a function of current income) could be the spurious result of either transitory fluctuations or measurement error in income: if consumption expenditure depends on long-run rather than current

income, a positive shock to income will raise both apparent saving and measured current income, and a positive measurement error in income will have the same effect. In light of these concerns, the literature has employed econometric methods to deal with measurement error and transitory fluctuations in income. For example Dynan et al. (2004) and Alan et al. (2015) use instrumental variable methods in an attempt to separate a long-run component in current incomes from transitory fluctuations and measurement error. The results have been mixed, with Dynan et al. (2004) concluding (using several US datasets) that the rich do in fact save more, and Alan et al. reporting that the relationship between savings and a long-run measure of income (in Canadian data) is flat above the first quintile. However, there are reasons to be sceptical of these IV approaches, as each is a valid instrument only under particular assumptions. For example, using income lagged 1 year as an instrument is valid only if the measurement error and the transitory shocks to income show no persistence at all over time. Importantly, the IV approaches described above will not help solve these sort of measurement error issues if there is persistent under-reporting of expenditure or over reporting of income that is related to the level of income or expenditure.

Our contribution is to revisit the relationship between the saving rates and income in the context of an experimental data collection exercise that aims to reduce directly some of the measurement error (through a balance check) rather than econometrically correcting for it (using instruments). The experiment involved collecting information on the full household budgeting identity in a single household survey, with the goal of reducing measurement error in both income and consumption expenditure.² The important feature of the survey design is that respondents are invited (but not required) to adjust their answers if the information they provide on their income, their spending and their net saving over the same period is not consistent (i.e. if income less spending less net saving

² The Townsend Thai Monthly Survey (Samphantharak and Townsend 2010) has pioneered the collection of information about all components of the household accounting identity in a developing country, and demonstrated its advantages, both in terms of data quality and substantive research possibilities.

does not sum to zero). This builds on evidence on the use of a “balance edit” in a Canadian survey of household spending (Brzozowski and Crossley, 2011) and in experiments conducted by the BLS (Fricker et al., 2015). We show how this budget reconciliation both alone and in conjunction with the instrumental procedures of the previous literature, affects the estimated relationship between affluence and saving.

Our findings are as follows. First, as is common, the raw saving function in our data (income less spending as a fraction of income plotted against current income) slopes steeply upwards, and implies saving rates in excess of 25% for the top half of the income distribution. Second, like the previous literature, we have mixed success in finding convincing instruments for long-run income, but instrumenting current income with lags and leads of income removes most of the slope of the saving function. Third, giving respondents the opportunity to amend inconsistent data on their household finances leads to lower implied saving rates (because they revise upwards their spending). Fourth, even after prompting, a majority of households do not present consistent data on their household finances, with the mean error consistent with overreporting income, underreporting spending or underreporting directly net saving (indeed, reported net saving is much lower than implied net saving, being a tightly-estimated zero at the median for the bottom 60% of the current income distribution). Restricting attention to those households who do, after prompting, report consistent data on their household finances, suggests that saving rates are generally fairly low, being well below 5% for the bottom half of the income distribution, before rising slowly to about 25% for the richest decile; this is little affected by whether or not we instrument income. How much weight we should place on what is an unrepresentative sample in the absence of any external validation data is, though, hard to tell.

In the rest of the paper, Section 1 explains the data that we use, and the definition of our key concepts, and Section 2 presents our estimated saving functions. We conclude in Section 3. Annexes give more detail on the nature of our data, and the impact of the balance edit.

1. Description of survey design, the experimental variation, and the analysis sample

1.1 Overview of survey and experiment

Our data comes from the 9th wave of the Innovation Panel (known as IP9), part of the UK household longitudinal study (UKHLS, known as “Understanding Society”). This sweep contained a special module, known as the Benefit Unit module, and designed by the authors of this paper, to collect information on income, spending, and “net saving” (the change in net financial assets) over the past month. The survey varied experimentally the nature of the questions eliciting the amount of net saving in the period³, and the allocated survey mode (face-to-face (CAPI) or on-line (CAWI), although respondents allocated to CAWI could elect to do the survey face-to-face.⁴) The intent of the Benefit Unit module was to collect information on the full household budget constraint (i.e. on income, spending and net saving all over the same period), and then to prompt, but not force, respondents to make adjustments if their responses were not in balance (we refer to this as a “balance edit”).⁵

The module takes its name from the fact that data on income, spending and net savings were collected at the level of the “benefit unit”, where a benefit unit is defined as a single adult without a co-resident partner or spouse plus any dependent children (i.e. aged under 19), or a co-resident couple (whether married or cohabiting) plus any dependent children. This concept is used in other

³ There were few statistically-significant differences between the two methods and so we ignore this feature of the experiment from hereon. See Gaia et al (2017) for some more results.

⁴ There are some telephone respondents (CATI) in IP9, but they are omitted entirely from the analysis in this paper as the Benefit Unit Module was not administered to them.

⁵ Alan et al. (2015) point out that this data collection process is likely to result in data where the measurement errors in income, expenditure and net saving are likely to be correlated with each other.

UK household surveys and corresponds to the unit of assessment for most means-tested benefits, tax credits and welfare payments in the UK. Annex 1 contains more information on IP9 and the questions in the Benefit Unit Module.

1.2 Overview of analysis sample and key outcomes

1,044 BUs were given the questions in the Benefit Unit module, or 63 percent of all benefit units in the 1,277 households that responded to the full IP9 survey (Annex 2 analyses the reason why 37% of BUs did not provide information). We make no correction for differential non-response in this paper, and so our sample is representative of the population who would have participated in IP9 and in the Benefit Unit module, with 66% of the BUs being single adults (possibly with children). Most of our analysis is done on the sample that initially reported non-negative values of income and spending before the balance check (880 BUs). We include BUs with adults of all ages, but our measures of income-rank are always calculated within 10-year age-bands (under 25, 25-34, etc).

1.3 Construction of outcomes and analysis methods

For each BU, the key measures of interest are *income*, *spending*, *implied net saving*, *reported net saving* and *imbalance*.⁶ They are related in the following way:

$$\textit{income} - \textit{spending} = \textit{implied net saving}$$

$$\textit{balance} = \textit{income} - \textit{spending} - \textit{reported net saving}$$

$$= \textit{implied net saving} - \textit{reported net saving}$$

We make no adjustments for inflation and report all financial values in cash terms (interviews took place in 2016). We equalise all financial measures using the Modified OECD scale, reflecting the

⁶ Income and spending are restricted to be non-negative, and the other concepts can take any real value.

composition of the BU and taking a single adult with no children as the reference BU with a weight of 1. We say that a BU reports “consistent” information if $balance \approx 0$, i.e. if the difference between its reported income and its reported spending is (nearly) identical to the amount it reports to have saved in that same period.⁷ A non-zero value of $balance$ means that a household has reported “inconsistent” data. We refer to the values initially given as “before reconciliation”, and the adjusted valued as those “after reconciliation” (and these are identical if respondents made no changes). We construct measures of the saving rate by dividing net saving (whether reported or implied) by current income (meaning that this is measure is n/a for those reporting zero income).⁸

To rank BUs, we use either current income or a proxy for long-run income. The issue with using current income on the x-axis of the saving function is that both measurement error (if saving is measured as income less spend) and transitory income shocks will bias upwards the estimated slope of the saving function. The idea, then, is to find instruments that are correlated with long-run income but not the transitory parts of current income, nor measurement error in current income, and also that have no direct impacts on the saving rate (including through any measurement error in the saving rate). The literature has suggested using the education level of the adults, lags and leads of income, and non-durable consumption. Alan et al. (2015) suggest that the education level of the adults is likely to be an invalid instrument because of a direct link between adults’ amount of education and their preferences for saving. We make use of the longitudinal nature of our data, instrumenting current income with 4 lags and 1 lead of income.^{9,10}

⁷ Following Brzozowski and Crossley (2011), we define a balanced BU (or one reporting consistent information) as one for whom the absolute value of “balance” is smaller than 10% of the mean of its income and spending; other BUs are said to be “out of balance”.

⁸ Dynan et al. (2004) find their conclusions to be insensitive to the choice of denominator when calculating the saving rate.

⁹ We regressed BU level equivalised current income on functions of equivalised income from the previous 4 waves and the 1 following wave the IP9 survey (each wave is approximately 12 months apart). For each year, we estimated a 4th order polynomial, plus indicators for income in that year being below £250/month and being missing, and all these were interacted with whether the BU has 1 or 2 adults in wave 9 (the measure of lagged or lead of income is defined as the

All of the figures shown in this paper are non-parametric median regressions (when we look at median saving rates) or non-parametric linear regressions (when we look at mean saving rates; when we do this, we weight each observation by the BU-level income and we drop the two BUs for whom the absolute value of saving rate exceeds 2000%). The identity that income less spend equals net saving only holds at the mean, but looking at median saving rates removes the influence of (usually measurement-error induced) extreme values of the saving rate, and is more common in the literature.

2. Results

In this section, we show various non-parametric estimates of the saving function, that is saving as a function of current income plotted against some measure of income or long-run income, as described in Section 1.

2.1 A naïve estimate of the saving function using data only on current income and current spending

Figure 1 presents the (median) implied saving rate as a function of current income. This is precisely how savings functions are typically estimated if one has access only to a household budget survey, such as the CEX in the US, or the LCFS and its predecessors in the UK, and if one uses current income as a proxy for long-run income.

Our estimated saving function has three notable properties. First, saving rates rise monotonically with income. Second, saving rates are negative for those with the very lowest incomes (in our

sum of the lags or leads of the income of the adults present in the BU in wave 9). We also control for the mean age of the adults in the BU (in bands). We then use the predicted values from this regression as the dependent variable in the non-parametric median regression. Reported CIs do not (yet) account for the two-stage nature of this estimate.

¹⁰ We lack information on non-durable consumption, but Instrumenting current income with current spending in our data led to a downward sloping saving function, which Dynan et al. (2004) note can occur if the measure of spending contains sufficiently large amounts of transitory spending or measurement error. This could reflect that spending was collected in IP9 with a single-shot question (although Al Baghal (2014) argues that the one-shot expenditure question can be useful) or that the monthly period over which respondents were asked to report spending was so short that transitory elements of spending are important.

sample, this is the case for the poorest 5%, echoing Brzozowski and Crossley (2011) and Brewer et al. (2017)). Third, saving rates are high, exceeding 25% for incomes at about the 45th centile of within-age-band incomes; these are higher than estimates in Crossley and O’Dea (2010) based on the UK’s household budget survey (see their Figure 6.4), or estimates of the aggregate saving rate in the UK, and also higher than estimates based on US household data (see e.g. Figure 1 of Dynan et al. (2004)). Figure 1 also shows the same but for mean saving rate: the values are very similar to median saving rates, except at the extremes of the income distribution, leading to a saving function with a more pronounced S-shape.

As discussed earlier, either transitory income shocks (if these are saved at a higher rate than permanent income) or measurement error in income will serve to produce an estimate of the saving function whose slope is upward-biased. We now explore the impact of making corrections to this standard saving function estimate.

2.2 Using a better proxy for long-run income

Figure 2 shows the saving rate plotted against the estimated rank of long-run income, where we instrument current income with 4 lags and 1 lead. The main impact of instrumenting income is to considerably lessen the slope in the saving function, mostly by increasing the estimated saving rate of BUs at the bottom of the income distribution. If our instruments are valid, then it would suggest that the median saving rate rises from just over 25% to just under 35%., and that mean saving rates are just over 25% for all levels of income.

The finding of a marked flattening of the saving function is in line with Alan et al. (2015)’s findings with data from the Canadian FAMEX and with Dynan et al. (2004)’s findings from the CEX when they use non-durable consumption as an instrument, but it is not consistent with Dynan et al. (2004)’s findings when they use various lags of income as instruments (using other datasets).

However, there are reasons to be sceptical of most of these IV approaches (including ours), as each

is a valid instrument only under particular assumptions. For example, using income lagged one year as an instrument is valid only if the measurement error and the transitory shocks to income show no persistence at all over time. More generally, none of the instruments will help if there is persistent under-reporting of expenditure or over reporting of income that is related to the level of income/expenditure. This is part of the motivation for trying to eliminate some of the measurement error (through a balance check) rather than just econometrically correct for it (using IV).

2.3 Using on saving and income data after a balance edit

Having answered questions about income, spending and the change in net assets, respondents to the IP9 survey were presented with an image of the household budgeting identity and had the opportunity to adjust any of the values that they had previously given. They were not required to make any adjustments, though, and the survey did not enforce that the household budget was balanced. (Annex 3 shows how respondents adjusted their answers during the balance edit). Figure 3 repeats Figure 1 but using the values of income and spend after the balance edit.¹¹ The balance edit reduces implicit saving rates, and by more in the bottom half of the income distribution than at the top: median (mean) saving rates now reach 25% at the 45th (55th) centile of income. (Annex 3 shows that this occurred chiefly through respondents revising upwards their estimated monthly spending rather than reducing their income.)

Figure 4 shows the estimated saving function after the balance edit and with instrumenting current income, as in Section 2.2.¹² This leads to an estimated saving function with a small upwards slope, with saving rising from about 20% to about 30% over the income distribution (the rise is more

¹¹ The balance edit can change the estimated saving function by adjusting any of: the estimated amount of savings (in £), the value of current income used as the denominator in the saving rate, and the way that households are ranked along the x-axis: we do not attempt to decompose the overall changes.

¹² The horizontal axis here is now of predicted-post-reconciliation current income, although Annex [X] shows that income was hardly changed by the reconciliation.

pronounced with median rates). As before, instrumenting current income (i.e. comparing Figures 3 and 4) flattens the saving function. The impact of using post-reconciliation values of income and spending, conditional on instrumenting income (i.e. comparing Figures 2 and 4) is, as before, to lower saving rates, particularly at the bottom.

2.4 Using data on active saving

As discussed earlier, IP9 asked BUs to provide a direct measure of net saving, and they were prompted if this was not consistent with their answers on income and spending. Figure 5 reports the estimated saving function using this data on directly reported net saving (using instrumented income on the x-axis: analysis with current income on the x-axis is very similar and available on request); if this data were correct, then it would imply that median saving rates are a well-estimated zero for the bottom 60% of the income distribution, and never rise above 10%. Confidence intervals for mean saving rates are wider, but we can always reject that saving rates are as high as 25%. On the other hand, as Annex 3 confirms, for only 38% of the sample is the data on reported saving equal to the difference between reported income and saving: the mean imbalance amongst the other 62% is consistent with over-reporting of income or under-reporting of spending or net saving (see Annex Figure A6). It is this degree of imbalance that explains why Figure 3 and 5 are not identical to each other, even after households have had an opportunity to correct their responses.

2.5 Restricting to BUs that are “in balance” (i.e. provide consistent information on their household budget constraint)

The Figures above have used data on all BUs, regardless of whether they were in balance (i.e. regardless of whether or not they provided consistent data on their income, spending and net saving). Figure 6 plots the estimated saving function only for those BUs who are in balance after the balance

edit.¹³ Plotted against current income, saving rates are considerably lower for this group than are the implied saving rates for the whole sample, but are higher than are the directly reported saving rate for the whole sample: for this sample of balanced BUs, median saving is zero for the lowest incomes, and rises slowly with income to reach 25% at the 80th income centile. Mean saving is also lower, being negative for the bottom 25% of the income distribution, before reaching 20% at the highest incomes.

Instrumenting income with lags and leads for this sample leads, as before, to lower estimated saving rates and a slightly less sloped saving function, with the median saving rate retaining an upwards slope from close to 0% for the bottom quarter to 25% for the top 20%. The estimated function for mean saving rate has a slight U shape, with mean saving being positive for the poorest quarter, very close to zero for the next quartile, and above zero but some way off 25% for the richer half.

3. Conclusion

Ever since Friedman emphasized that the upward sloping cross-sectional savings function could be the spurious result of transitory fluctuations or measurement error in income, assessing whether or not the fraction of income that is saved varies with long-run income has proved empirically challenging.

A considerable part of the literature has employed IV approaches to deal with the measurement error and transitory fluctuations in income. But it is challenging both to find instruments that deal with transitory fluctuations in income but do not affect saving directly, and to find instruments that will help with the sort of persistent under-reporting of expenditure or over reporting of income that we think pervades household surveys.

¹³ Ranks are defined within the full sample of BUs, not just those who are balanced.

The data used in this paper comes from an experimental design, designed by the authors, that collected information on income, spending, and “net saving” over the same period of time, and then prompted (but did not force) respondents to make adjustments if their responses were not in balance. The aim was to reduce directly some of the measurement error through better survey instrument design, rather than econometrically correcting for it using IV.

As is common, we find that the raw saving function (income less spending as a fraction of income plotted against current income) slopes steeply upwards, and implies saving rates in excess of 25% for the top half of the income distribution. Although we remain unconvinced of their validity, instrumenting current income with lags and leads of income removes most of the slope of the saving function, but in doing so implies (implausibly?) high rates of saving across most of the income distribution. The novel feature of our survey instrument – giving respondents the opportunity to amend inconsistent data on their household finances – leads to lower implied saving rates, chiefly because they revise upwards their spending. However, even after prompting, a majority of households do not present consistent data on their household finances, with the average reported net saving being considerably smaller than mean income less mean spending (i.e. the errors are consistent with overreporting income, or underreporting spending or underreporting the direct measure of net saving). This is consistent with the Hansen et al. (2017), who report the results of (ex post) cognitive testing and qualitative interviewing on the questions in the Benefit Unit module. A common finding is that people think of saving as being an active thing, often associated with moving money from a current (checking) account into a savings account or similar. On the other hand, fluctuations in the balance of their current (checking) account would not be thought of as (dis)saving; as a result, people had difficulty understanding why income less spend had to equal net saving.

Restricting attention to those households who do, after prompting, report consistent data on their household finances, implies that saving rates are generally low, being well below 5% for the bottom

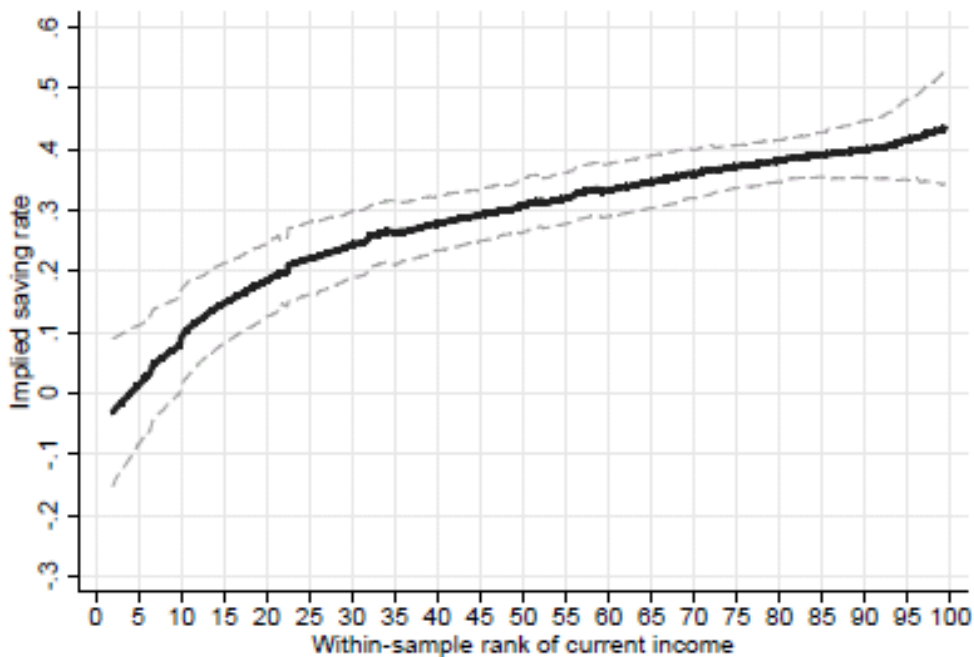
half of the income distribution, before rising slowly to about 25% for the richest decile. On the other hand, this data, although internally consistent, may not be accurate, let alone representative of the saving behaviour of the population.

Our experiment is clearly not the final word on this issue. On average, we consider that the balance edit did improve the accuracy of data on BU-level income and spending at the aggregate level, because the mean difference (i.e. the aggregate saving rate) after reconciliation seems a little less implausible. On the other hand, there is clearly more work to do if the majority of households are not able to report consistent data on their household finances even after a prompt (and it would have been extremely valuable to have been able to validate directly any of the data supplied by households). In this respect, our results highlight the potential for measurement error or transitory shocks to income and spending – which are the source of most estimates of the saving function with micro-data – to distort seriously estimates of the saving function.

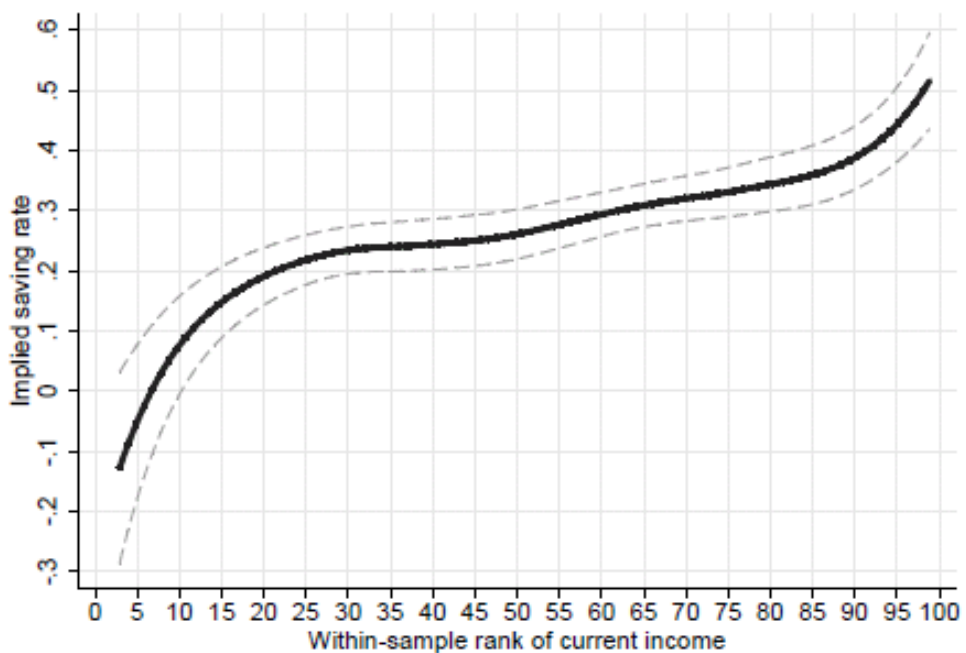
References

- Al Baghal, Tarek (ed.), Nick Allum, Katrin Auspurg, Margaret Blake, Cara Booker, Thomas F. Crossley, Joanna D'Ardenne, Malcolm Fairbrother, Maria Iacovou, Annette Jäckle, Olena Kaminska, Peter Lynn, Cheti Nicoletti, Zoe Oldfield, Steve Pudney, Sebastian Schnettler, S.C.N. Uhrig, and Joachim Winter (contributors). 2014. "Understanding Society Innovation Panel Wave 6: Results from Methodological Experiments." in *Understanding Society Working Paper 2014-4*. University of Essex.
- Alan, S., Atalay, K. and Crossley, T. (2015) "Do the Rich Save More?: Evidence from Canada", *Review of Income and Wealth*, 61(4), 739-758, 10.1111/roiw.12131
- Brewer, Mike, Ben Etheridge, and Cormac O'Dea. (2017). "Why Are Households That Report The Lowest Incomes So Well-Off?", *Economic Journal*, 127, F24-F49,
- Brzozowski, Matthew, and Thomas F. Crossley. 2011. "Measuring the Well-being of the Poor with Income or Consumption: A Canadian Perspective." *Canadian Journal of Economics* 44(1):88-106.
- Crossley, T. and O'Dea, C. (2010), *The Wealth and Saving of UK families on the eve of the crisis*, IFS report 71, doi: 10.1920/re.ifs.2010.0071.
- Dynan, K., Skinner, J. and Zeldes, S. (2004), "Do the Rich Save More?", *Journal of Political Economy*, 112(2), pp397-44, <https://doi.org/10.1111/eoj.12334>
- Friedman, M. (1957), *A Theory of Consumption Function*, Princeton University Press.
- Fricker, S., Kopp, B. and To, N. (2015), "Exploring a Balance Edit Approach in the Consumer Expenditure Quarterly Interview Survey" in C. Carroll, T. Crossley and J. Sabelhaus (eds), *Improving the Measurement of Consumer Expenditures*, Chicago: University of Chicago Press, DOI:10.7208/chicago/9780226194714.003.0013.
- Gaia, A. (ed). (2017). "Understanding Society Innovation Panel Wave 9: results from methodological experiments", *Understanding Society Working Paper 2017-07*, <https://www.understandingsociety.ac.uk/research/publications/524448>
- Hanson, T., Matthews, P. and Willis, D. (2016), "Understanding Society Innovation Panel Wave 9 Technical Report", mimeo.
- Hanson, T., Westwood, P., Burton, J. and Fisher, P/. (2017) "Improving household finances data with joint interviewing and a balance edit: cognitive testing of a 'Benefit Unit Finance' module", *Understanding Society Working Paper 2017-14.*, <https://www.understandingsociety.ac.uk/research/publications/524734>
- Meyer, B. D. and J. X. Sullivan, (2003), "Measuring the Well-Being of the Poor Using Income and Consumption," *The Journal of Human Resources*, 38, 1180-1220.
- Meyer, B. D. and J. X. Sullivan, (2011), "Viewpoint: Further results on measuring the well-being of the poor using income and consumption," *Canadian Journal of Economics*, 44, 52-87.
- Samphantharak, Krislert, and Robert M. Townsend. 2010. *Households as Corporate Firms*. New York: Cambridge University Press.

Figure 1. Implied saving rate vs age-adjusted rank of current income (top=median, bottom=mean)



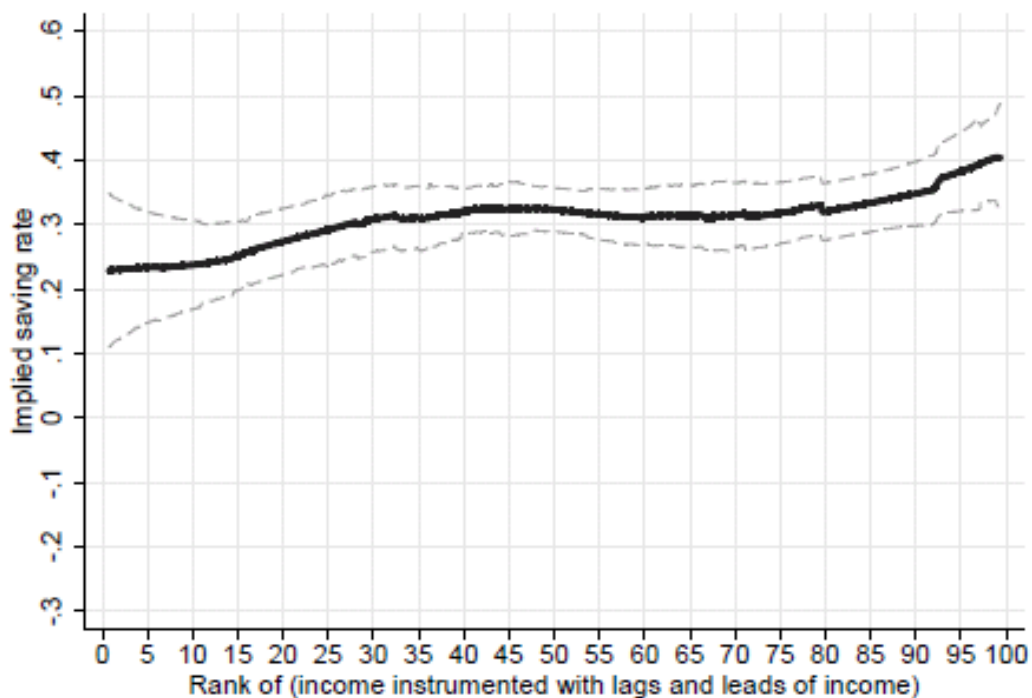
Notes:



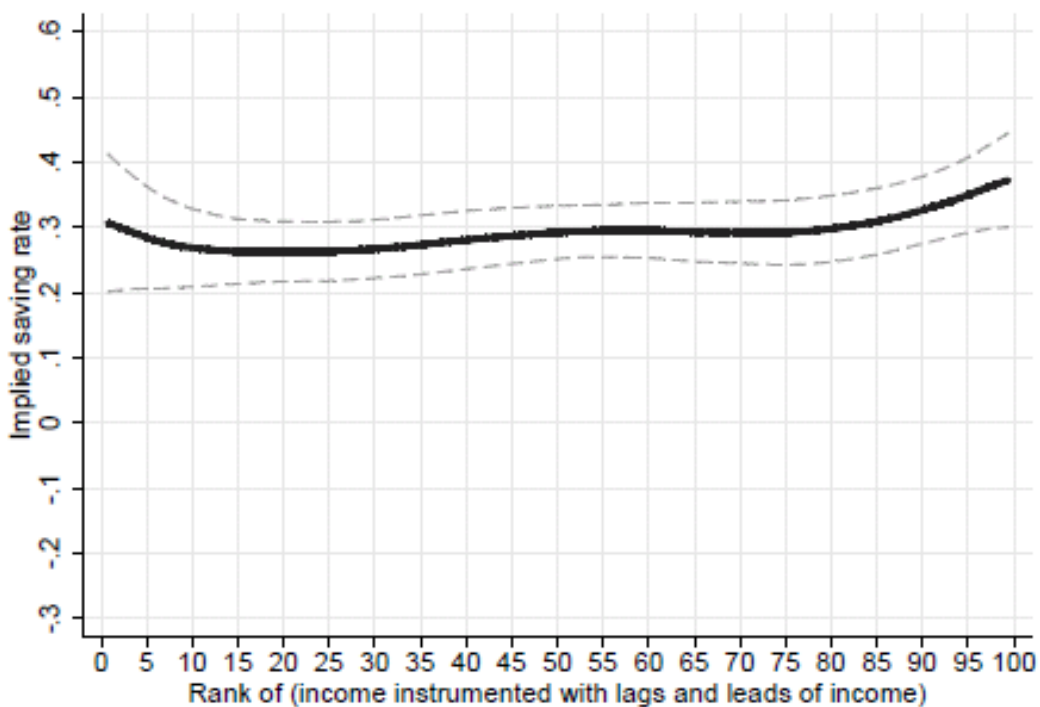
Notes:

Notes: sample is BUs reporting non-negative values of income and spending. Point estimates and 95% CI based on non-parametric median / linear regressions with a quartic (biweight) kernel.

Figure 2. Implied saving rate vs age-adjusted rank of current income instrumented by lags and lead of income (top=median, bottom=mean)



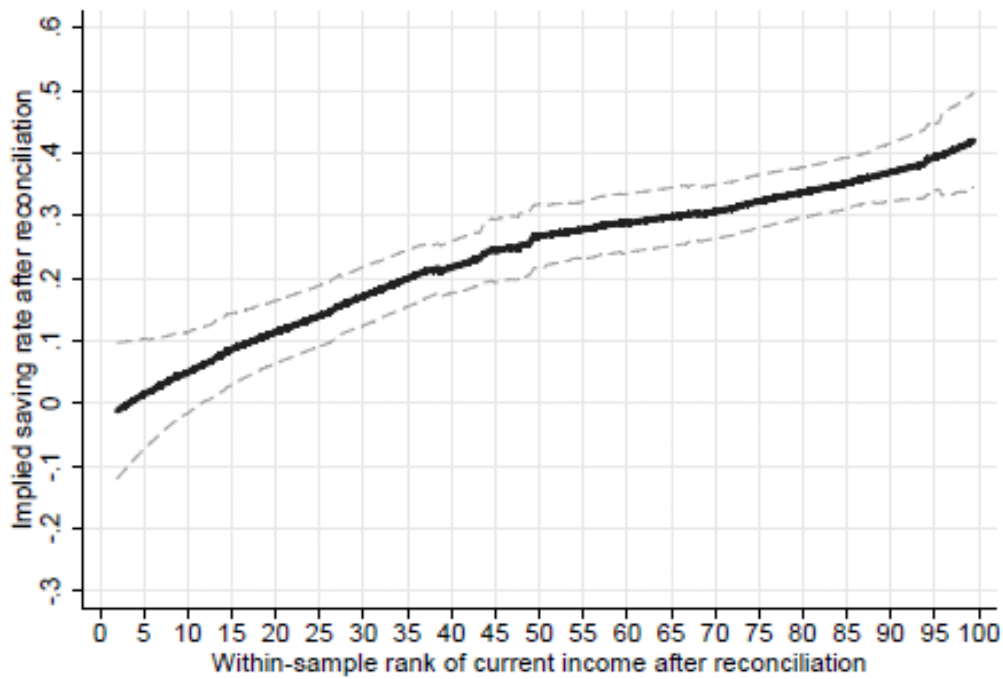
Notes:



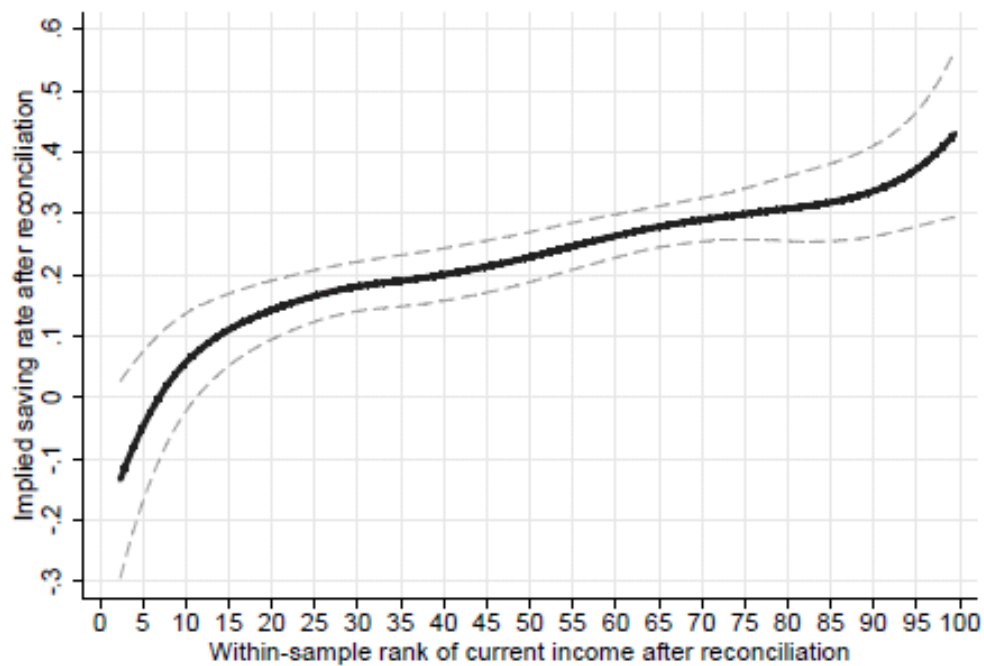
Notes:

Notes: sample is BUs reporting non-negative values of income and spending. Point estimates and 95% CI based on non-parametric median / linear regressions with a quartic (biweight) kernel.

Figure 3. Implied saving rate vs age-adjusted rank of current income, after reconciliation (top=median, bottom=mean)



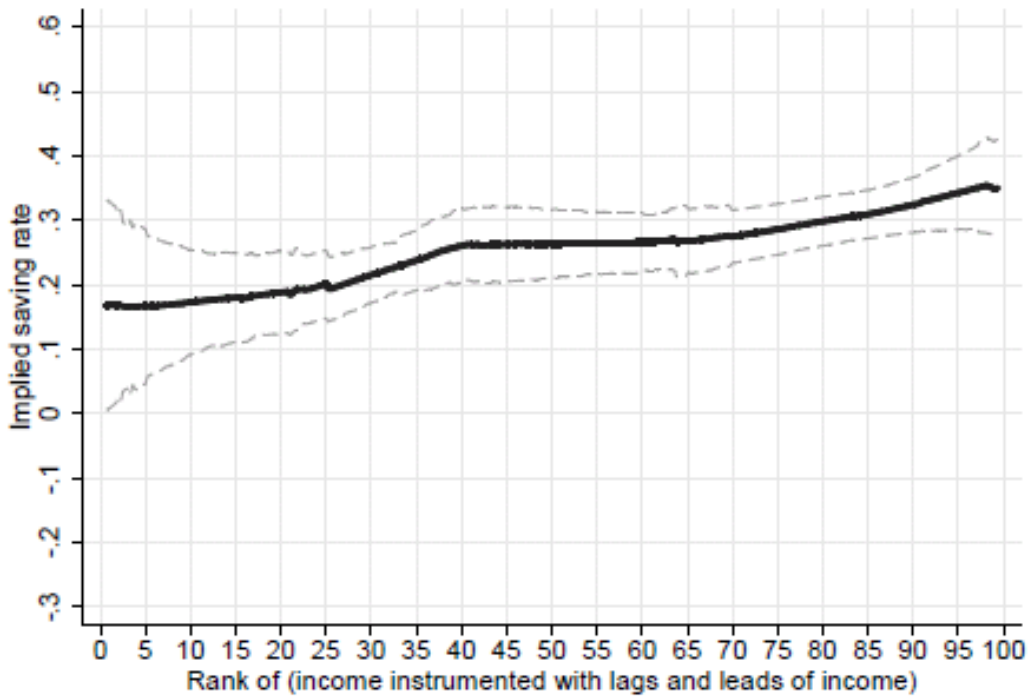
Notes:



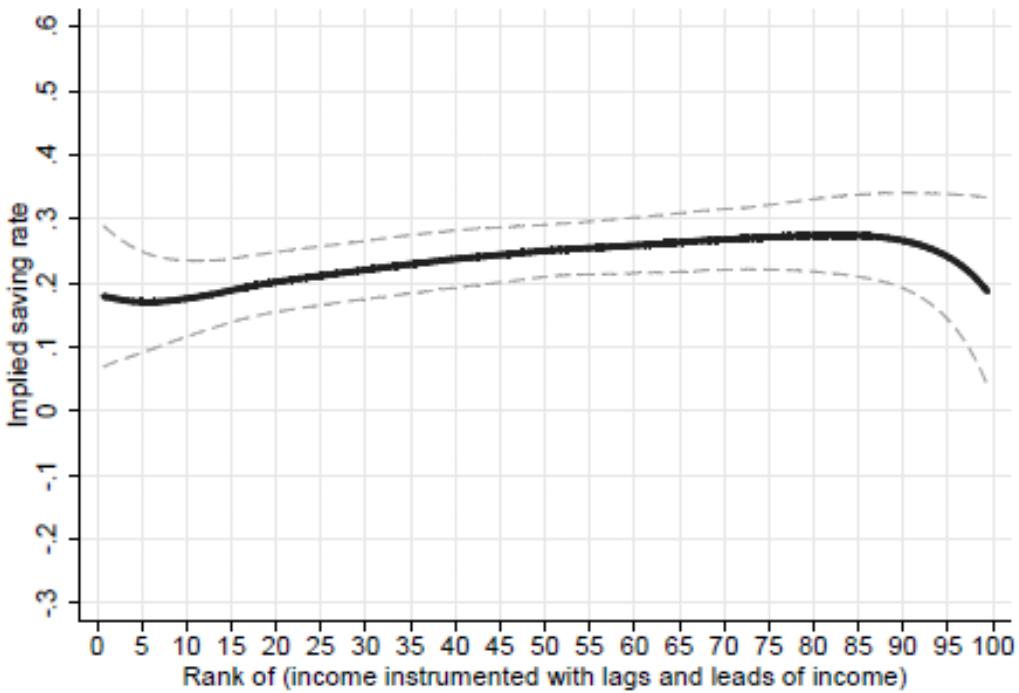
Notes:

Notes: sample is BUs reporting non-negative values of income and spending. Point estimates and 95% CI based on non-parametric median / linear regressions with a quartic (biweight) kernel.

Figure 4. Implied saving rate vs age-adjusted rank of current income instrumented by lags and lead of income, after reconciliation (top=median, bottom=mean)



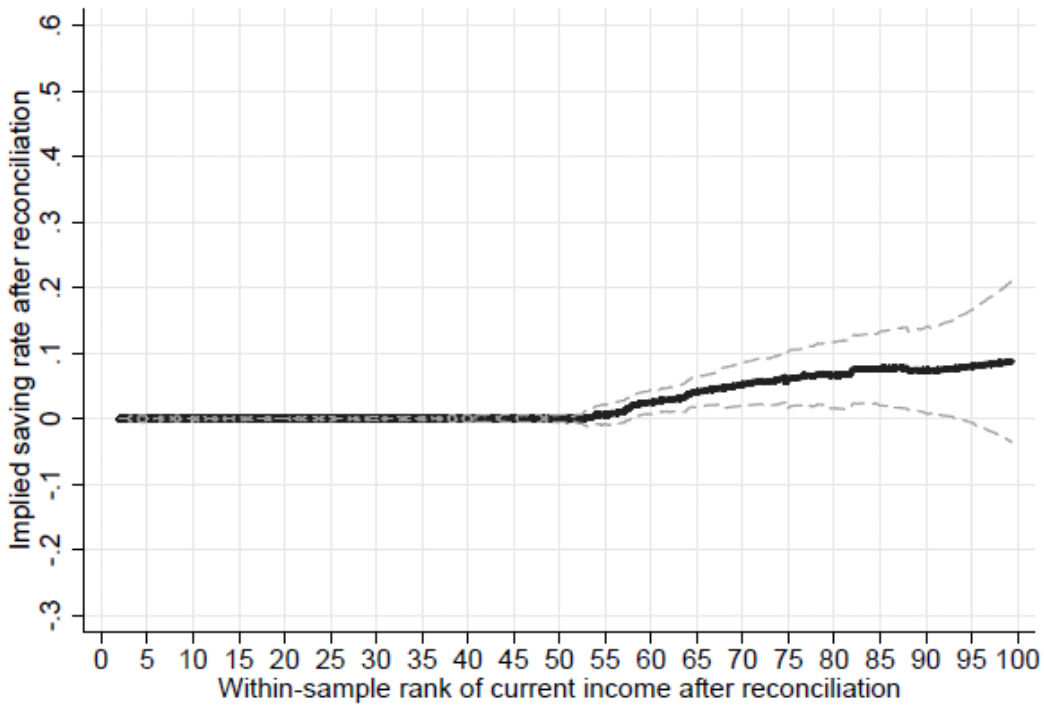
Notes:



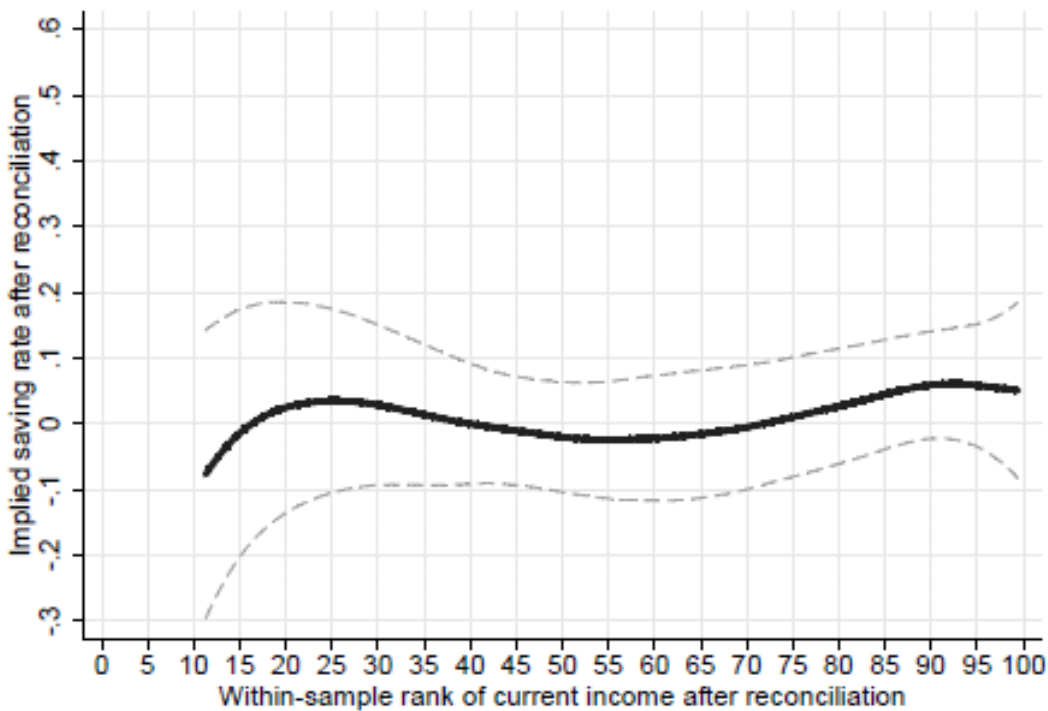
Notes:

Notes: sample is BUs reporting non-negative values of income and spending. Point estimates and 95% CI based on non-parametric median / linear regressions with a quartic (biweight) kernel.

Figure 5. Directly reported net saving rate vs age-adjusted rank of current income instrumented by lags and lead of income, after reconciliation (top=median, bottom=mean)



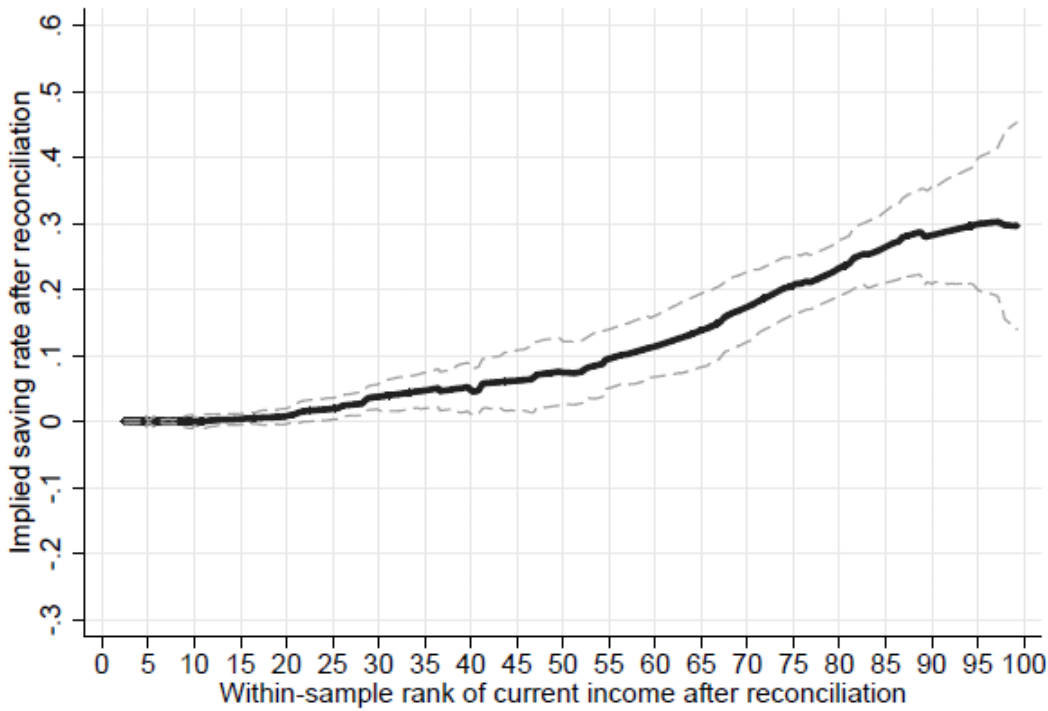
Notes:



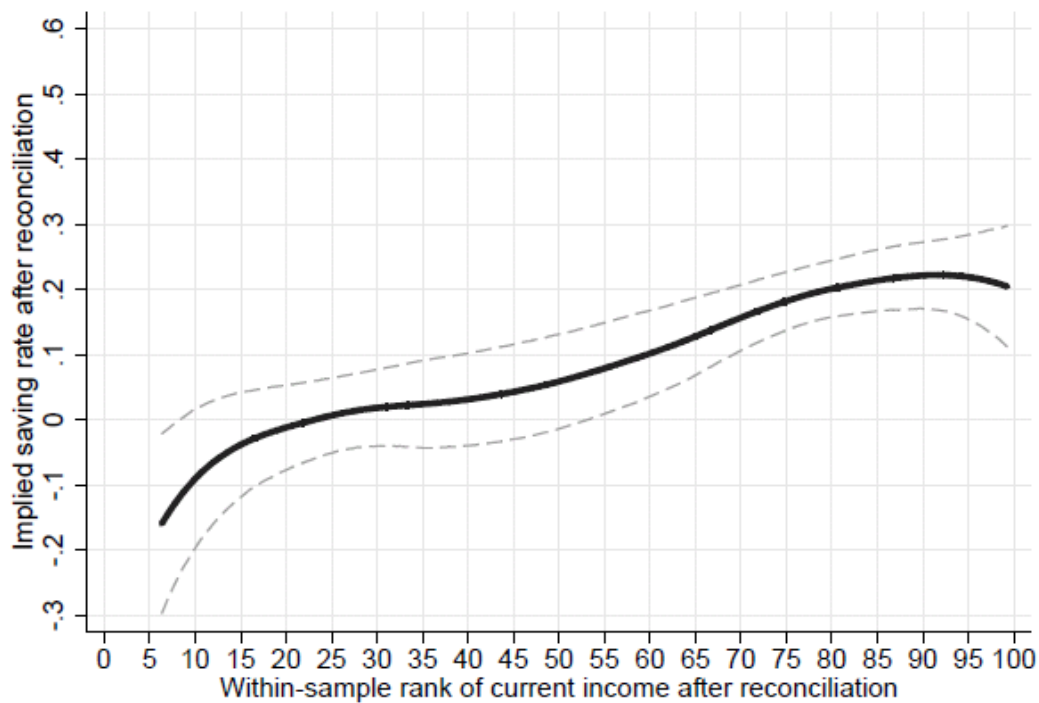
Notes:

Notes: sample is BUs reporting non-negative values of income and spending. Point estimates and 95% CI based on non-parametric median / linear regressions with a quartic (biweight) kernel.

Figure 6. Implied saving rate vs age-adjusted rank of current income, sample of BUs who are balanced after the edit (top=median, bottom=mean)



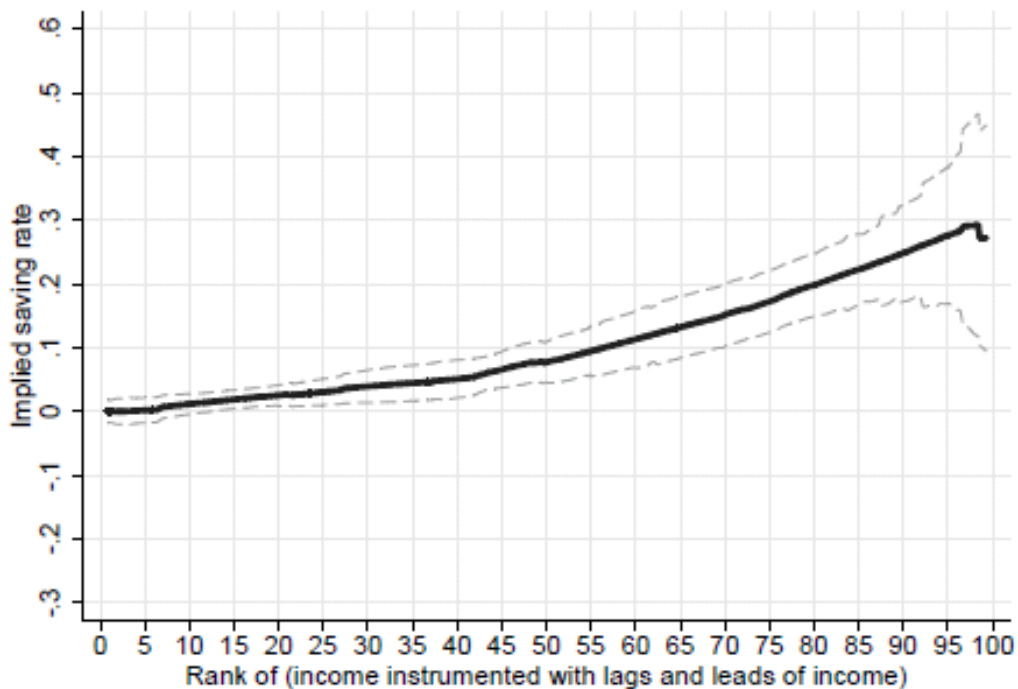
Notes:



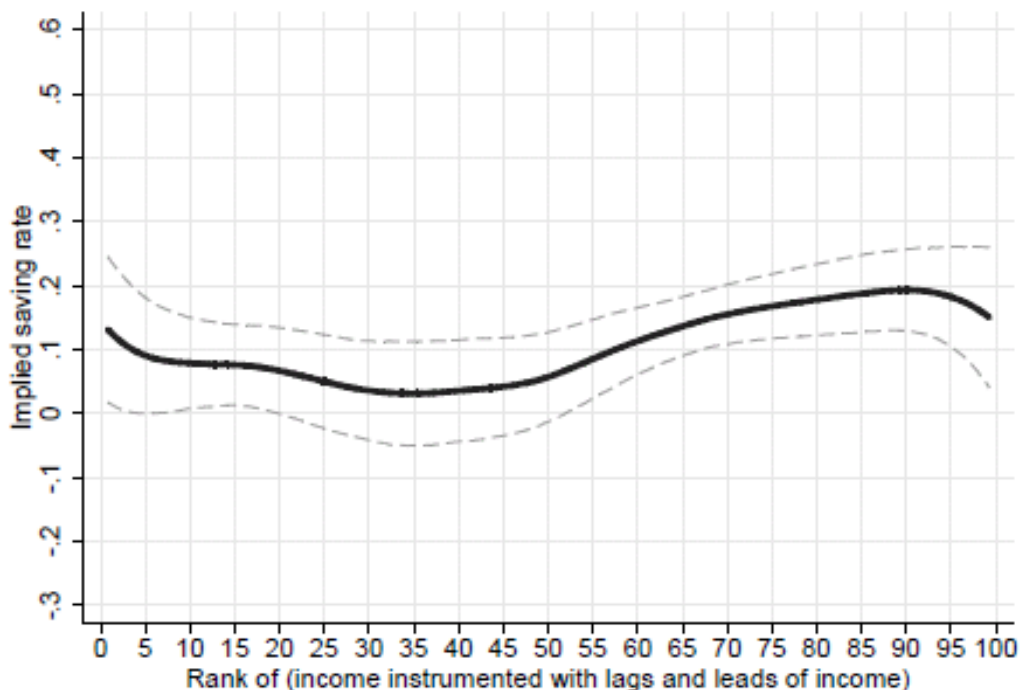
Notes:

Notes: sample is BUs reporting non-negative values of income and spending. Point estimates and 95% CI based on non-parametric median / linear regressions with a quartic (biweight) kernel.

Figure 7. Implied saving rate vs age-adjusted rank of current income instrumented by lags and lead of income, sample of BUs who are balanced after the edit (top=median, bottom=mean)



Notes:



Notes:

Notes: sample is BUs reporting non-negative values of income and spending. Point estimates and 95% CI based on non-parametric median / linear regressions with a quartic (biweight) kernel.

Annex 1. Overview of the Survey Experiment ¹⁴

Our data comes from IP9, the 9th wave of the Innovation Panel, part of the UK household longitudinal study (UKHLS), and known as “Understanding Society”.

In UKHLS, information on income, spending and change in net financial assets was collected and reconciled at the level of the benefit unit. A “benefit unit” is defined as a single adult without a co-resident partner or spouse plus any dependent children, or a co-resident couple (whether married or cohabiting) plus any dependent children. This concept is used in other UK household surveys – notably the Family Resources Survey – and it corresponds to the unit of assessment for most means-tested benefits and tax credits in the UK. Additional adults within a household, such as adult siblings, non-dependent children, or unrelated adults living together but not in a relationship, will form a separate benefit unit. The set of questions that reconciled income and spending are therefore referred to as the “Benefit Unit Module”.

Benefit units that contained two adults could be asked the Benefit Unit Module only if both adults agreed for the financial information they had given in their individual interview to be shared with their partner. If so, then the questions in the Benefit Unit Module were asked at the end of the second person’s interview. Single adult benefit units were asked the question at the end of their individual interview.

The intent of the Benefit Unit Module was to collect data on the benefit unit’s net income, spending and the change in net financial assets, all over the same period of time. For net income, the script calculated a benefit unit total of monthly net income from the amounts already reported in the individual interviews. Benefit Units were then asked to confirm or correct that amount, and the

¹⁴ Some of this draws on Gaia (2017) and Hanson et al. (2016).

analysis in this paper is based on the post-reconciliation values of BU-level income.¹⁵ Monthly total spending was elicited using a single question, based in part on an experiment in an earlier wave of the UKHLS Innovation Panel (see Al Baghal et al., 2014). The method of eliciting changes in net financial assets was varied experimentally. Responding households were randomly allocated to one of two ways of eliciting changes in net financial assets.¹⁶ Half were asked the ‘net flows’ model, in which outgoing money (spending, new savings and debts that have been paid off) was subtracted from incoming money (income, increases in debts and withdrawals from savings). Benefit units in the other half of households were asked the ‘change in stocks’ model. They were asked about the starting and ending balance in all of their financial assets (including debts), which were then aggregated to give a change in net financial assets. The difference between income and spending was then compared to the aggregate net change in the benefit unit’s financial assets. These alternatives are illustrated in Figure A1. We note that the “net flows” model would capture only active saving, and that the “change in stocks” method would also capture a change in the level of savings caused by e.g. changes in capital values.

With benefit unit-level data on total monthly net income, total spending, and additions and subtractions from wealth, the script calculated the difference between ‘Incoming Money’ (income plus any increase in borrowing or drawdown of savings) and ‘Outgoing Money’ (spending plus any reductions in debts or increases in savings). We refer to this difference as the “balance”, although according to the accounting identity discussed above it should be identically zero. Where the amounts did not balance (i.e. where ‘Incoming Money’ was not equal to ‘Outgoing Money’),

¹⁵ The corrections to BU-level income made at this stage are analysed further in [Paul and Alessandra’s paper]. Note that there was a separate experiment within IP9 that presented half of respondent households with an individual-level income screen: see Gaia (2017) for more details.

¹⁶ The randomization was at the household level (so multiple responding benefit units in the same household received the same treatment) and was stratified with respect to allocated mode.

respondents were asked to check the amounts and reconcile any differences. Respondents were not forced to make the amounts balance to zero.

Figure A1a: “Net Flows” model

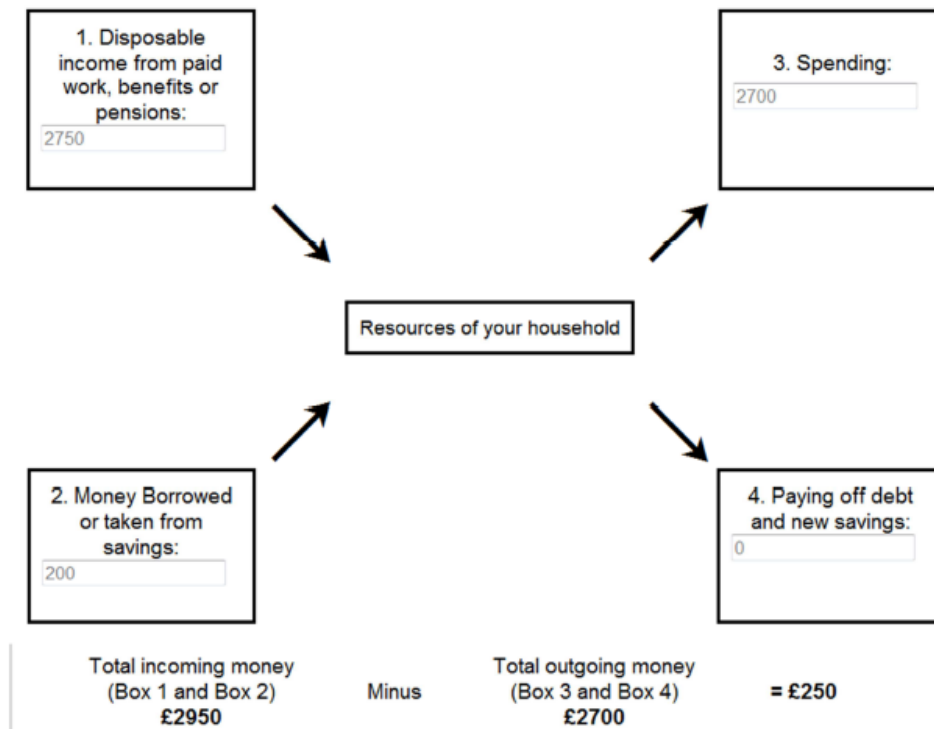
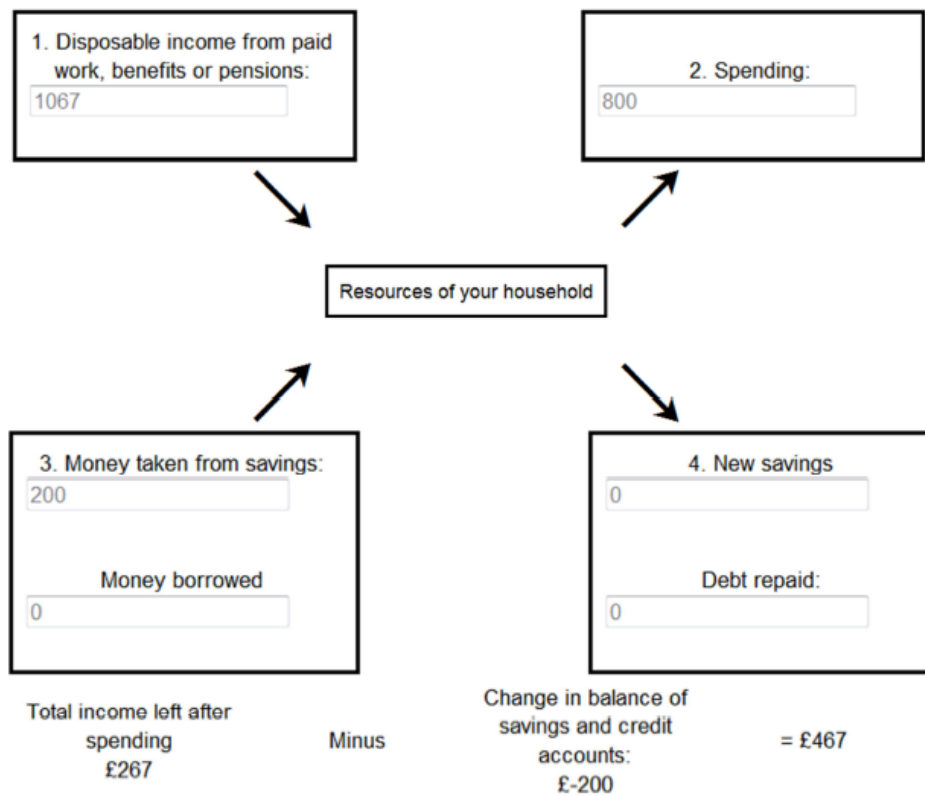


Figure A1b: “Change in stocks” model



Annex 2. Sample selection

Table A1 shows the reasons why the not BUs in households responding to IP9 were asked the Benefit Unit Module. The requirement for both adults in a couple to consent is the biggest reason why the sample who responded to the Benefit Unit Module questions is slanted towards single adults: of the 1,174 adults in couples asked for consent, 77% gave consent, but only in 66% of couples did both adults give consent. There was little difference in completion or consent rates by survey mode (other than those that used the phone were not eligible for the module at all).¹⁷ Of the 1,044 BUs who responded to the Benefit Unit Module, 880 initially gave non-negative answers to income and spending.¹⁸

¹⁷ There are 36 BUs that gave consent but do not appear in `i_bufinance.dta`, 35 of which are couples, and there is also 1 BU which gave consent and appears in `i_bufinance.dta` but where the type of benefit unit is not coded consistently across the survey; none of these appears in our analysis sample. We have the full characteristics of those who completed IP9 but did not give consent to the Benefit Unit module, and it would be possible to see which characteristics help predict consent.

¹⁸ The survey does not make it easy to distinguish zero values from responses that might otherwise be coded as missing. Our best guess is that 11% of BUs in the net flows allocation, and 5% of BUs in the change in gross stocks allocation, refused to answer, or did not know the answers to, the initial questions on how their debts and savings balances had changed. Just over 10% of BUs did not give a valid initial answer to the question on spending. However, before calculating the final balance, the CAPI and CAWI routines set any incomplete or missing values for spending or for changes in assets or for spending to zero, and this is what would have been presented to BUs during the reconciliation. There are almost no unusable values of income in the data-set at this stage, but net income is £0 in 9% of BUs, a far greater fraction than suggested by other households surveys, suggesting that some of these zeroes might reflect item non-response; on the other hand, all respondents had already had the opportunity to review and amend their income, and the reconciliation would have proceeded assuming that their income was zero.

Table A1. Derivation of final analysis sample for Benefit Unit Module

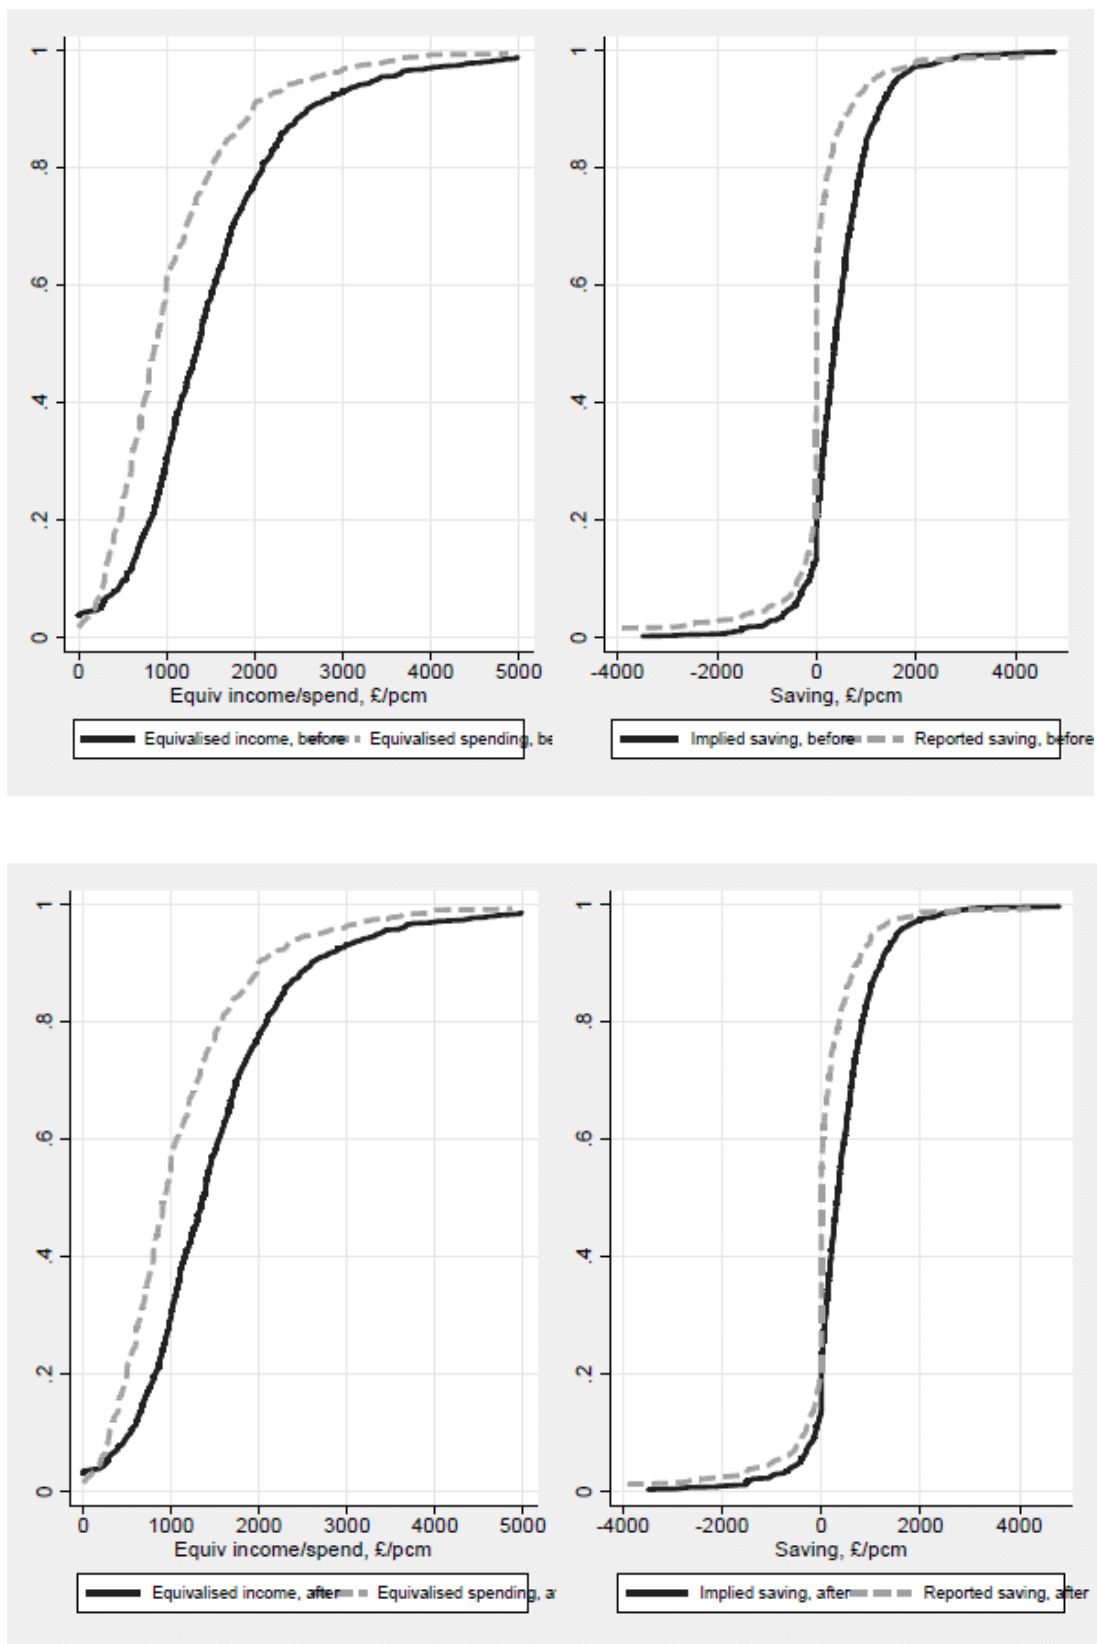
Sample	N adults aged 16+ and not in f-t education	N benefit units	Of which, share single adults	N households	Of which, share allocated to web
Respondent households	2,421	1,646	0.516	1,277	0.649
...where all adults provided individual interviews	2,014	1,373	0.533	1,146	0.646
...and no interviews in the BU were by telephone, or proxies	1,886	1,290	0.538	1,096	0.649
...and, if adults in couples used different modes, the second one was not web	1,868	1,281	0.542	1,090	0.651
...and both adults in a couple consented to share information with each other	1,471	1,081	0.639	930	0.628
...and who appear in i_bufinance.dta	1,399	1,044	0.660	895	0.636

Annex 3. Changes in the distributions of income, spending and change in net assets

Figure A3 shows percentiles of the distributions of income, spending, change in assets (defined both as “income minus spending”, and as provided by the respondents’ answers to explicit questions about changes in the stocks of assets and debts) and balance, before and after the reconciliation.

Before the reconciliation, the measure of balance has a mean and median that exceed zero, meaning that households are more likely to be doing at least one of over-reporting income, under-reporting spending, or under-reporting net saving (i.e. that they report that the resources coming in to the BU exceed the resources they report leaving the BU) than they are to be doing the opposite. The reconciliation has little impact on the distribution of income, shifting the distribution very slightly to the left for the net flows variant, and (except at the top) for the gross flows variant. The reconciliation has a greater effect on the distribution of spending, shifting it to the right so that median spending rises by £100 / month (a 9% rise). The reconciliation generally shifts the distribution of reported net change in assets to the right (that is, respondents report larger increases in net saving after the reconciliation), except at there is also an increase in dissaving in the net flows variant.

Figure A3. CDFs of income, spend, implied saving and reported saving (all equivalised £/pcm), before (top) and after (bottom) balance edit



We then focus on whether BUs were able to report consistent information on their finances, that is, whether the BU's finances were "in balance", which is defined as having a balance level that is less than 10% of the average of income and spending, i.e. if

$$abs\left(\frac{2x(\text{income}-\text{spend}-\text{net increase in assets})}{\text{income}+\text{spend}}\right) < 0.1.$$

Table A2 reports this, splitting BUs by their experimental allocation, and Table A3 does the same by survey mode. Overall, 23% of BUs were in balance before the reconciliation, and 38% after. Table A2 shows that the "gross flows" method had slightly fewer BUs initially in balance, but then slightly more in balance after the reconciliation. But Table A3 shows that there was a much larger difference between the two survey modes (f2f and web), where the reconciliation increases the fraction in balance by 20 ppts for those doing f2f, and by 10 ppts for those on-line. The sample doing the survey online were more likely to be in balance before the reconciliation, but less likely to be in balance after the reconciliation. The Tables also show that, amongst all 4 of these sub-groups, "change in net financial assets" is the most likely concept to be adjusted, followed by spending, and then by income; this is unsurprising given that participants have already had at least 1 chance to review their income, and that the question about spending was a one-shot. In general, these findings - that reconciliation reduces, but far from eliminates, budget imbalance is consistent with a small-scale lab study reported in Fricker et al. (2015). Figures A4 to A6 shows the distribution of income, spending and the measures of saving before and after the balance check, for the whole sample.

Table A2. Number of benefit units who are in balance before and after household budget reconciliation, by experimental allocation

	Total		
	Of which	Net flows	Change in stocks
In balance before	196	90	106
In balance after	317	155	162
Total	838	402	436
	0		
Of those initially out of balance:	0		
balance changed	259	140	119
where income changed	99	47	52
where spending changed	134	75	59
where “change in assets” changed	148	76	72
abs(balance) fell	240	133	107
abs(balance) rose	19	7	12
Total	642	312	330

Notes: in balance defined as $abs\left(\frac{2x(\text{income}-\text{spend}-\text{net increase in assets})}{\text{income}+\text{spend}}\right) < 0.1$.

Sample: BUs initially reporting non-zero values of income and spending.

Table A3. Number of benefit units who reported non-zero values of income and spending who report to be in balance before and after household budget reconciliation, by allocated survey mode

	Total		
	Of which	Ftf	Web
In balance before	196	67	129
In balance after	317	132	185
Total	838	306	532
	0		
Of those initially out of balance:	0		
balance changed	259	123	136
where income changed	99	32	67
where spending changed	134	67	67
where “change in assets” changed	148	73	75
abs(balance) fell	240	118	122
abs(balance) rose	19	5	14
Total	642	239	403

Notes: in balance defined as $abs\left(\frac{2x(\text{income}-\text{spend}-\text{net increase in assets})}{\text{income}+\text{spend}}\right) < 0.1$. Those completing by phone were not eligible for the BUfinance module.

Sample: BUs initially reporting non-zero values of income and spending.

Figure A4 Distribution of income and spending, before and after the balance check

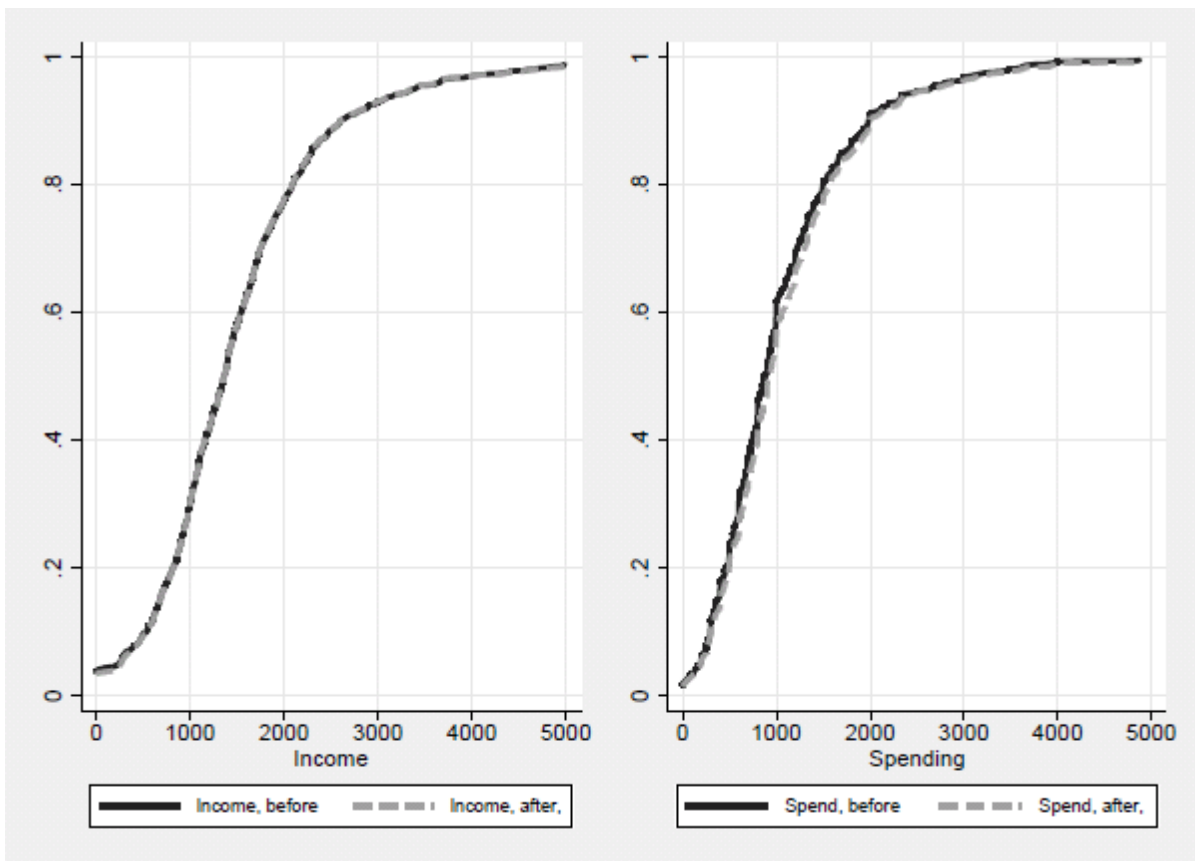


Figure A5 Distribution of implied and reported saving, before and after the balance check

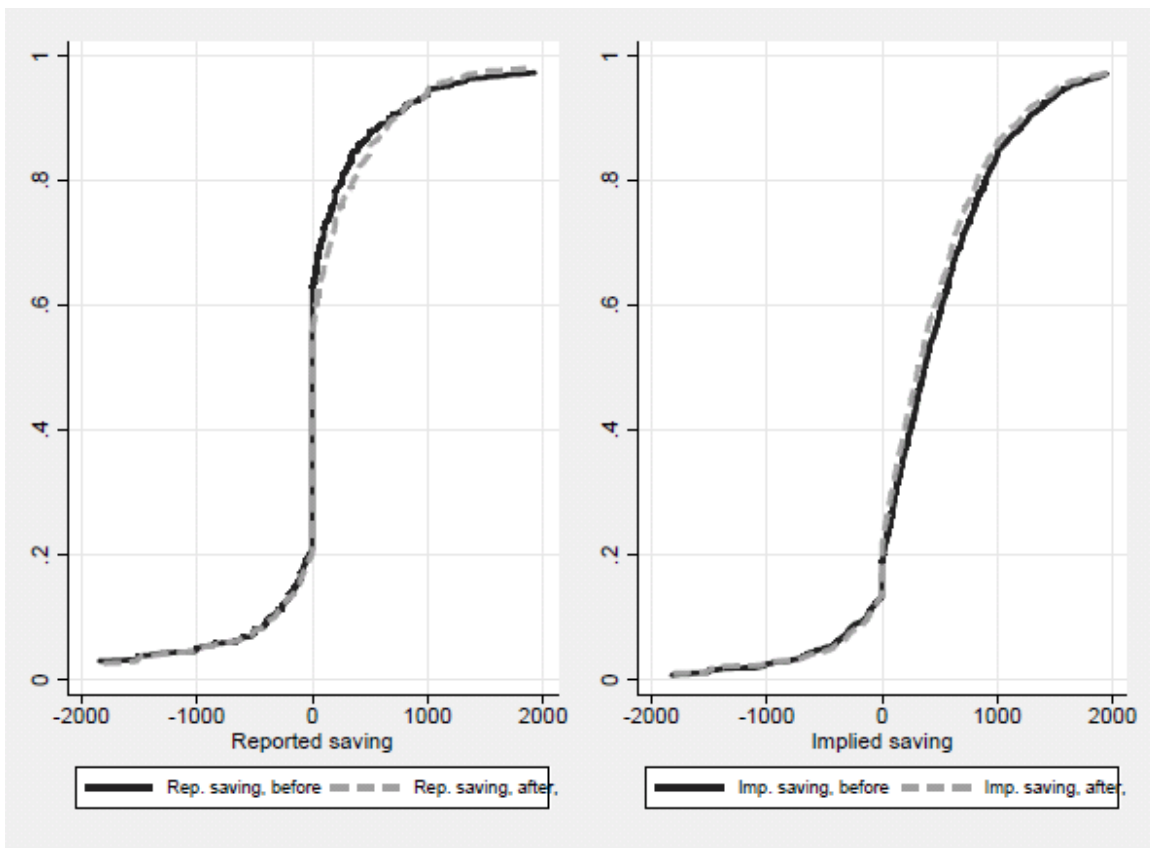


Figure A6. Distribution of imbalance as a fraction of income, before and after the balance check

